When and Why is Optimistic MWU Slow? The Geometry of Energy Dissipation

John Lazarsfeld[1] **Anas Barakat**[2] **Georgios Piliouras**[3]
Antonios Varvitsiotis[2,4,5] **Andre Wibisono**[6]

Abstract

This paper studies the convergence of the Optimistic Multiplicative Weights Update algorithm (OMWU) in two-player zero-sum games. Recent works have identified instances on which the last-iterate of OMWU can converge arbitrarily slowly, but understanding when and why this slow convergence occurs has remained open. In this work, we develop a new analysis framework that gives sharp, quantitative explanations for this behavior. Our analysis is based on viewing the algorithm's dual iterates as an *optimistic skew-gradient descent* with respect to an energy function. We prove over the dual iterates that energy is *dissipative*, and by establishing tight bounds on the magnitude of dissipation, our analysis quantifies the geometric bottlenecks that arise when the corresponding primal iterates are close to the simplex boundary. This further translates into a new linear last-iterate convergence rate in KL divergence on games with a unique and interior Nash equilibrium. Compared to prior work, this new rate contains a much sharper dependence on game-specific constants, and we prove this dependence is optimal. Moreover, these geometric insights further translate into new separations on *uniform* convergence rates for OMWU. On the one hand, we prove *constant lower bounds* on the uniform *best-iterate* convergence rate in KL divergence and total variation distance from Nash. On the other hand, we establish for the 2×2 setting a new $\widetilde{O}(T^{-1/2})$ best-iterate rate in duality gap, improving substantially over prior work. Together, this shows in general that uniform convergence rate guarantees do not transfer across different measures of distance to Nash.

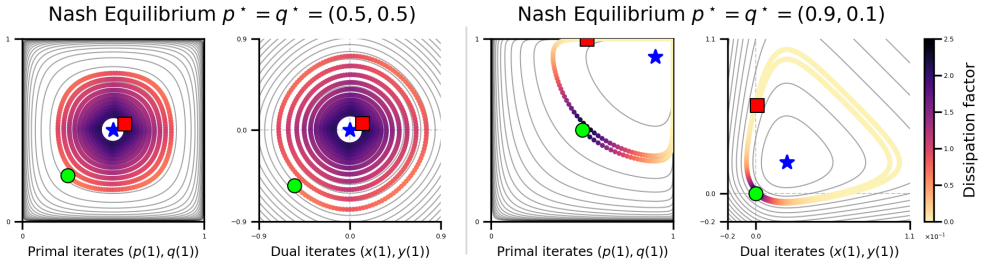

Figure 1: **The primal and dual iterates of OMWU** on two instances of a 2×2 zero-sum game. The algorithm is run for $T = 1500$ iterations with stepsize $\eta = 0.2$ in both instances. The Nash equilibrium (and corresponding dual equilibrium point) is denoted by the blue star; the OMWU initializations by the green circle; the T'th primal and dual iterates by the red square. In the left instance, the Nash equilibrium (NE) is $p^\star = q^\star = (0.5, 0.5)$ and the primal initialization is $(0.25, 0.25)$. In the right instance, the NE is $p^\star = q^\star = (0.9, 0.1)$ and the primal initialization is $(0.5, 0.5)$. Note that the algorithm's dual iterates are completely determined by the primal iterates. In both instances, the left subplot shows the trajectory of the **primal iterates overlaid on the levelsets of KL divergence from NE**, and the right subplot shows the trajectory of the **dual iterates overlaid on the levelsets of the log-sum-exp energy function**. In the right instance, the last-iterate convergence to NE is significantly slower due to the lower energy dissipation over the dual iterates.

Affiliations: [1] Georgia Tech [2] SUTD [3] Google DeepMind [4] NUS CQT [5] Archimedes/Athena RC [6] Yale
Contact: jlazarsfeld@gmail.com, anas_barakat@sutd.edu.sg, gpil@deepmind.com, antonios@sutd.edu.sg, andre.wibisono@yale.edu.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

Optimistic Multiplicative Weights Update (OMWU) is among the most prominent algorithms for online learning in games. While the standard Multiplicative Weights Update (MWU) algorithm diverges to the simplex boundary in two-player zero-sum games (Bailey and Piliouras, 2018), the optimistic variant *converges in last-iterate to a Nash equilibrium* (Daskalakis and Panageas, 2019; Wei et al., 2021). This desirable convergence property has led OMWU to gain increasing relevance in machine learning applications ranging from adversarial training (Daskalakis et al., 2018; Mertikopoulos et al., 2019), to zero-sum Markov games (Cen et al., 2023), and Nash learning from human feedback (Munos et al., 2024; Tiapkin et al., 2026).

However, recent work has identified instances on which the last-iterate convergence of OMWU can be arbitrarily slow (Cai et al., 2024), raising new doubts on the algorithm's reliability for fast learning in games. Moreover, this finding has highlighted the fact that, despite an extensive recent literature studying the *average-iterate* convergence behavior of OMWU (Syrgkanis et al. (2015); Daskalakis et al. (2021); Anagnostides et al. (2022a)), and despite the long history of optimistic algorithms in general,[1] obtaining a fine-grained and quantitative understanding of the *last-iterate* behavior of Optimistic MWU has still remained surprisingly elusive, even in the most fundamental constrained setting of two-player bilinear zero-sum games.

By way of background, the *asymptotic* last-iterate convergence of OMWU to a unique Nash equilibrium in zero-sum games was first proven by Daskalakis and Panageas (2019). This result was later strengthened by Wei et al. (2021), who proved a non-asymptotic *linear* last-iterate convergence rate for the same setting, but with a bound containing constant factors that depend on the game's smallest non-zero equilibrium mass δ. Such bounds are called *universal* convergence rates, and they imply that the algorithm's performance can vary dramatically depending on the magnitude of δ in a specific game instance. However, whether this potentially poor scaling in their analysis accurately captures the true iterate-to-iterate behavior of the algorithm was left unanswered.

Towards this direction, Cai et al. (2024) recently gave new evidence that OMWU *can indeed* converge arbitrarily slowly: for every time horizon T, they proved that the last iterate of OMWU on a 2×2 instance remains an absolute constant distance away from equilibrium in duality gap. While this confirms for fixed time horizons that the last-iterate of OMWU can remain far from Nash, it does not explain what *features* of a specific game instance and what *algorithmic properties* of OMWU govern its possibly slow iterate-to-iterate behavior. Moreover, whether the analysis of Wei et al. (2021) – which involves tracking distances between *primal* iterates – and its dependence on δ is quantitatively sharp still remains unresolved. Thus, the central focus of our paper is to address the following fundamental question:

When and why can the last-iterate convergence of OMWU be arbitrarily slow in zero-sum games?

1.1 Our Contributions and Techniques

Our main contributions provide sharp, quantitative answers to this question.

New Analysis Framework via Dual Energy Dissipation. For zero-sum games with a unique and interior Nash equilibrium, we develop a new framework for analyzing the last-iterate convergence of OMWU. We analyze the *dual iterates* of OMWU in the unconstrained space of payoff vectors, where we show the algorithm follows an *optimistic skew-gradient descent* with respect to the *log-sum-exp energy function*. This energy function is the convex conjugate of the negative entropy regularizer, and we show that minimizing energy over the dual iterates corresponds to minimizing KL divergence from Nash equilibrium over the primal iterates (Proposition 3.2 and Proposition D.11).

Under this dual perspective, we prove that *energy is dissipative* over the iterates of OMWU, which is in stark contrast to the behavior of the standard MWU variant. For OMWU, we obtain *tight* upper and lower bounds on the one-step dissipation in energy, the strength of which is quantified in the local geometry of the energy function Hessian (Lemma 4.1). Specifically, we relate the magnitude of dual energy dissipation to the primal KL divergence from Nash by proving a novel, general *skew-gradient domination* inequality (Proposition 4.4) that may be of independent interest. Together, these results

[1] OMWU is the canonical instantiation of the Optimistic Follow-the-Regularized Leader (OFTRL) family with entropic regularization (Rakhlin and Sridharan, 2013), and this family itself has its roots in optimistic and extragradient methods for solving monotone variational inequalities (see, e.g., Korpelevich (1976), Popov (1980), Nemirovski (2004), Nesterov (2007)).

	KL Divergence	TV Distance	Duality Gap
Universal Last-Iterate	$O(\exp(\frac{1}{\delta}) \cdot \exp(-\exp(\frac{-1}{\delta}) \cdot T))$ $^{\diamond}$	(via (2)) $^{\diamond}$	(via (2)) $^{\diamond}$
	$\Theta(\exp(-\exp(\frac{-1}{\delta}) \cdot T))$ * $^{\blacktriangle}$	(via (2)) *	(via (2)) *
Uniform Last-Iterate	$\Omega(1)$ †	$\Omega(1)$ †	$\Omega(1)$ †
Uniform Best-Iterate	$\Omega(1)$ **	$\Omega(1)$ **	$O(T^{-1/6})$ ††
			$\widetilde{O}(T^{-1/2})$ ***

Table 1: **The landscape of universal and uniform convergence rates for OMWU**, initialized from the uniform distribution (except for $^{\blacktriangle}$), and with constant stepsize η, on zero-sum games with a unique and interior Nash equilibrium $w^\star$ with minimum coordinate $\delta > 0$. Our results are highlighted in gray. $^{\diamond}$: Wei et al. (2021). *: upper bound in Theorem 4.7, and $^{\blacktriangle}$: lower bound in Theorem 4.10, using horizon-dependent games and a positive-measure set of nonuniform initializations. Both (*) and ($^{\diamond}$) imply rates in TV and duality gap due to (2). Compared with Wei et al. (2021), our bound improves the leading multiplicative factor by an exponential factor in the inverse minimum Nash coordinate δ. †: proven by Cai et al. (2024) in duality gap, which also implies the same lower bound in TV and KL due to (2). **: Theorem 5.1. ***: Theorem 5.2, which holds only for the 2×2 setting. ††: Cai et al. (2025), also only for the 2×2 setting.

give a new and more clear description of the geometric properties that lead to the algorithm's slow behavior: when the primal iterate is near the simplex boundary, the dual iterate moves in directions over which the energy function is extremely flat, thus leading to vanishingly small magnitudes of energy dissipation. As the decrease in energy is equivalent to decrease in KL divergence from Nash, a geometric bottleneck arises when the primal iterates spend many steps near the boundary of the simplex, a property that emerges when the game's equilibrium is also close to the boundary.

Tight Universal Last-Iterate Convergence Rates. Our results on energy dissipation further translate into a new and optimal *linear last-iterate convergence rate* in KL divergence for zero-sum games with a unique and interior Nash equilibrium (Theorem 4.7). This new universal rate has a significantly sharper dependence on δ (the minimum non-zero Nash equilibrium coordinate) compared to the prior rate of Wei et al. (2021). In particular, it improves by a leading multiplicative factor of $\exp(\frac{1}{\delta})$, a term that is unbounded as $\delta \to 0$. While our new bound can still decay slowly as δ grows small, we prove a lower bound on the last-iterate convergence rate with a matching dependence on δ (Theorem 4.10), establishing in general that our new rate and analysis are tight. Note also that upper bounds on convergence rates in KL divergence also imply convergence rates in duality gap (see expression (2)).

New Separations on Uniform Convergence Guarantees. Our geometric insights additionally yield new contrasting bounds on *uniform convergence rates* for OMWU. Uniform rates hold simultaneously over all instances of a fixed dimension and have no dependence on game-specific constants like δ. Importantly, a potentially slow *universal* last-iterate convergence rate does not preclude the possibility of obtaining *uniform* rates under weaker notions of convergence. While the result of Cai et al. (2024) confirms that no uniform *last-iterate* rates are attainable, their followup work proved for the 2×2 setting that OMWU does have a uniform $O(T^{-1/6})$ *best-iterate* rate in duality gap (Cai et al., 2025). Note that upper bounds on KL divergence imply upper bounds in duality gap (see the relationships of expression (2)), but the reverse direction does not necessarily hold. Thus, given our improved geometric understanding and new universal last-iterate analysis in KL, it is natural to ask:

Can we obtain any stronger uniform convergence guarantees for OMWU?

In particular, can we strengthen the guarantee of a uniform best-iterate convergence rate to also hold in KL divergence or total variation (TV) distance from Nash? And, can we obtain any *faster* rates than the current $O(T^{-1/6})$ bound?

Our work additionally gives answers to these questions. On the negative side, we prove a new *separation* between the uniform guarantees attainable in duality gap versus KL divergence and TV distance. Specifically, we prove that OMWU has *no uniform best-iterate convergence rate* in KL or TV (Theorem 5.1), meaning that on certain hard instances, *all* iterates of OMWU over a finite time horizon remain far from Nash under these measures. On the other hand, in duality gap, we prove a new $\widetilde{O}(T^{-1/2})$ best-iterate convergence rate for the 2×2 setting (Theorem 5.2), a bound that substantially improves over the prior rate of Cai et al. (2025) for this same regime. Together, these results indicate that uniform convergence guarantees do not translate between different measures of

distance to Nash, a distinction that arises due to the inherent relationship between OMWU and KL divergence, and the fundamental geometric differences between the level sets of KL divergence, TV distance, and duality gap over the simplex (see Section I.4.4 for a visual comparison).

Ultimately, our new techniques and results open the door to a broader, quantitative study of last-iterate learning in games that is both physics-informed and geometry-aware. We summarize our new rates in Table 1, and we defer a more detailed discussion of other related work to Section A.

2 Preliminaries

Notations. For $k \in \mathbb{N}$, let $[k] = \{1, \dots, k\}$. Let Δ_k denote the probability simplex in $\mathbb{R}^k$, and let $\mathrm{ri}(\Delta_k) = \{p \in \Delta_k : 0 < p(i) < 1\}$ denote its relative interior. For $p \in \mathrm{ri}(\Delta_k)$, we say that p is *interior* or *fully mixed*. For $u, v \in \mathbb{R}^k$, we write $\langle u, v \rangle = u^\top v = \sum_{i=1}^k u(i) \cdot v(i)$ to denote the ℓ_2 inner product. We write $\mathbf{1}_k \in \mathbb{R}^k$ to denote the vector of all ones.

Online learning in zero-sum games. We study the setting where two players repeatedly play a zero-sum game with payoff matrix $A \in \mathbb{R}^{m \times n}$, which yields the min-max optimization problem:

$$\min_{p \in \Delta_m} \max_{q \in \Delta_n} \langle p, Aq \rangle \ .$$

At each step $t \geq 0$ the players choose strategies $p_t \in \Delta_m$ and $q_t \in \Delta_n$, incur losses $\langle p_t, Aq_t \rangle$ and $-\langle q_t, A^\top p_t \rangle$, and then observe gradient feedback payoff vectors $Aq_t \in \mathbb{R}^m$ and $-A^\top p_t \in \mathbb{R}^n$, respectively. The players seek to converge to a Nash equilibrium (NE) of A. For this, let $\mathcal{W} := \Delta_m \times \Delta_n$. For $\epsilon \geq 0$, an (ϵ-approximate) Nash equilibrium $w = (p, q) \in \mathcal{W}$ of A satisfies

$$\max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \epsilon \leq \langle p, Aq \rangle \leq \min_{p' \in \Delta_m} \langle p', Aq \rangle + \epsilon. \tag{1}$$

Every payoff matrix A has at least one ($\epsilon = 0$) Nash equilibrium (von Neumann, 1928).

Distances to Nash equilibrium. This work focuses on games with an *interior* Nash equilibrium. Let $\mathrm{ri}(\mathcal{W}) := \mathrm{ri}(\Delta_m) \times \mathrm{ri}(\Delta_n)$. For an interior NE $w^\star = (p^\star, q^\star) \in \mathrm{ri}(\mathcal{W})$ of A, we measure convergence to $w^\star$ under several different distances. For any $w = (p, q) \in \mathrm{ri}(\mathcal{W})$, we define:

- **Duality gap (DG)**: $\mathrm{DG}(w) = \max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \min_{p' \in \Delta_m} \langle p', Aq \rangle$.
- **Kullback-Leibler divergence (KL)**: Define the components $\mathrm{KL}(p^\star, p) = \sum_{i=1}^m p^\star(i) \log\big(p^\star(i)/p(i)\big)$ and $\mathrm{KL}(q^\star, q) = \sum_{j=1}^n q^\star(j) \log\big(q^\star(j)/q(j)\big)$. Then $\mathrm{KL}(w^\star, w) = \mathrm{KL}(p^\star, p) + \mathrm{KL}(q^\star, q)$.
- **Total variation distance (TV)**: Define the components $\mathrm{TV}(p^\star, p) = \frac{1}{2}\|p^\star - p\|_1$ and $\mathrm{TV}(q^\star, q) = \frac{1}{2}\|q^\star - q\|_1$. Then $\mathrm{TV}(w^\star, w) = \mathrm{TV}(p^\star, p) + \mathrm{TV}(q^\star, q)$.

If $\mathrm{DG}(w) \leq \epsilon$, then w is an ϵ-approximate NE. Moreover, letting $a_{\max} := \max_{(i,j) \in [m] \times [n]} |A(i,j)|$ denote the largest absolute entry of A, the following relationships hold:

$$\tfrac{1}{(2a_{\max})^2} \cdot \mathrm{DG}(w)^2 \leq \mathrm{TV}(w^\star, w)^2 \leq \mathrm{KL}(w^\star, w) \ . \tag{2}$$

Thus, upper bounds on $\mathrm{KL}(w^\star, w)$ or $\mathrm{TV}(w^\star, w)$ also correspond to approximate Nash equilibria in the sense of (1). We give the proof of (2) and additional preliminaries in Section B.

Additional preliminaries. For a symmetric matrix $M = M^\top \in \mathbb{R}^{k \times k}$, we write $M \succeq 0$ when M is positive semi-definite (PSD) and $M \succ 0$ when M is positive definite (PD). For a (not-necessarily-symmetric) matrix M, let $\sigma_{\max}(M) := \max_{x \neq 0} \|Mx\|_2 / \|x\|_2$ denote its spectral norm. For a linear subspace $\mathcal{V}$, let $\Pi_{\mathcal{V}}(\cdot)$ denote the orthogonal projection operator onto $\mathcal{V}$, and for a matrix M, let $\sigma_{\min}(M, \mathcal{V}) = \inf_{v \in \mathcal{V} \setminus \{0\}} \|\Pi_{\mathcal{V}}(Mv)\|_2 / \|v\|_2$ denote its minimum singular value restricted to $\mathcal{V}$. A matrix $M \in \mathbb{R}^{k \times k}$ is *skew-symmetric* when $M^\top = -M$, which implies $\langle v, Mv \rangle = 0$ for all $v \in \mathbb{R}^k$. For an integer $k \geq 2$ and $x \in \mathbb{R}^k$, we write $\mathrm{softmax}_k(x) \in \mathrm{ri}(\Delta_k)$ to denote the distribution with coordinates $\mathrm{softmax}_k(x)(i) = \exp(x(i)) / (\sum_{j=1}^k \exp(x(j)))$ for all $i \in [k]$. For readability, we further use the notation $\mathrm{softmax}_k(x)(i) \propto \exp(x(i))$ for $i \in [k]$. For a distribution $p \in \mathrm{ri}(\Delta_k)$, the discrete entropy of p is given by $\mathrm{Ent}_k(p) = \sum_{i=1}^k p(i) \log(1/p(i))$. Finally, for continuously differentiable function $f : \mathcal{X} \to \mathbb{R}$, recall the Bregman divergence $D_f : \mathcal{X} \times \mathrm{ri}(\mathcal{X}) \to \mathbb{R}$ is given by $D_f(x', x) = f(x') - f(x) - \langle \nabla f(x), x' - x \rangle$ for all $x' \in \mathcal{X}, x \in \mathrm{ri}(\mathcal{X})$.

3 OMWU as Optimistic Skew-Gradient Descent

In this section, we show that the dual iterates of OMWU can be viewed as optimistic skew-gradient descent. We start with preliminaries on OMWU and defer a review of standard MWU to Section C.

Primal OMWU update rule. The algorithm is initialized from an arbitrary $w_0 = (p_0, q_0) \in \text{ri}(\mathcal{W})$. Then for $t \geq 1$, using a fixed stepsize $\eta > 0$, the iterate $w_t = (p_t, q_t) \in \text{ri}(\mathcal{W})$ has coordinates

$$\begin{aligned} p_t(i) &\propto p_{t-1}(i) \cdot \exp\big(-\eta\big(2Aq_{t-1} - Aq_{t-2}\big)(i)\big) \quad \forall i \in [m]\,, \\ q_t(j) &\propto q_{t-1}(j) \cdot \exp\big(\eta\big(2A^\top p_{t-1} - A^\top p_{t-2}\big)(j)\big) \quad \forall j \in [n]\,. \end{aligned} \tag{OMWU}$$

For notational convenience, we assume that $w_{-1} = w_0 \in \text{ri}(\mathcal{W})$.

Dual iterates and OMWU as Optimistic FTRL. OMWU can also be derived as an instantiation of the Optimistic Follow-the-Regularized-Leader (OFTRL) family using the negative entropy regularizer (Rakhlin and Sridharan, 2013; Syrgkanis et al., 2015). Specifically, under OFTRL, the primal iterates are chosen by applying a *regularized* best-response map to a sequence of *dual iterates*. In the zero-sum game setting, these primal and dual iterates have a compact representation that we describe here:

First, given $A \in \mathbb{R}^{m \times n}$, let $J \in \mathbb{R}^{(m+n)\times(m+n)}$ be the block skew-symmetric matrix

$$J = \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix} = -J^\top\,. \tag{3}$$

Under OFTRL, the primal iterates $\{w_t\}$ are initialized at $w_0 = (p_0, q_0) \in \text{ri}(\mathcal{W})$. For the dual iterates, choose $x_0 \in \mathbb{R}^m$ and $y_0 \in \mathbb{R}^n$ such that $\text{softmax}_m(x_0) = p_0$ and $\text{softmax}_n(y_0) = q_0$, for example, $x_0(i) = \log p_0(i)$ and $y_0(j) = \log q_0(j)$. Let $z_0 = z_{-1} = (x_0, y_0) \in \mathbb{R}^{m+n}$. At time $t = 1$, let $z_1 = z_0 - \eta J w_0$. Then at time $t+1 \geq 2$, the dual iterate $z_{t+1} = (x_{t+1}, y_{t+1}) \in \mathbb{R}^{m+n}$ is

$$z_{t+1} = z_0 - \eta\big(\textstyle\sum_{k=1}^{t} Jw_k + Jw_t\big)\,, \tag{4}$$

which is a sum that depends only on $w_0, \ldots, w_t$. The primal iterate at time $t+1$ is defined as follows: first, let $R_m : \Delta_m \to \mathbb{R}$ and $R_n : \Delta_n \to \mathbb{R}$ be a pair of strictly convex regularizers, and let R be the joint, separable regularizer defined by $R(w) = R_m(p) + R_n(q)$ for $w = (p, q) \in \mathcal{W}$. Then OFTRL sets

$$w_{t+1} = \text{argmin}_{w=(p,q)\in\mathcal{W}} \left\{ -\langle w, z_{t+1}\rangle + R(w) \right\}. \tag{5}$$

For OMWU, the regularizer R is instantiated by the *negative entropy* functions $R_m = -\text{Ent}_m$ and $R_n = -\text{Ent}_n$. Then, the first-order optimality conditions of (5) under this setting of R yield the closed-form expression $w_t = (p_t, q_t) = (\text{softmax}_m(x_t), \text{softmax}_n(y_t)) \in \text{ri}(\mathcal{W})$ for all $t \geq 1$. This recovers the update rule of (OMWU). We review the full derivation of this equivalence in Section C.

3.1 Dual OMWU Iterates as Optimistic Skew-Gradient Descent

Our analysis of OMWU is based on viewing the dual iterates $\{z_t\}$ as an *optimistic skew-gradient descent* with respect to an energy function F. We derive this connection here:

Energy function. Let LSE_m and LSE_n denote the m and n-dimensional *log-sum-exp* functions. For $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^n$, $\text{LSE}_m(x) = \log(\sum_{i=1}^m \exp(x(i)))$ and $\text{LSE}_n(y) = \log(\sum_{j=1}^n \exp(y(j)))$. Together, these functions define the convex, continuously differentiable, and separable *energy function* $F : \mathbb{R}^{m+n} \to \mathbb{R}$. For $z = (x, y) \in \mathbb{R}^{m+n}$, we let

$$F(z) = \text{LSE}_m(x) + \text{LSE}_n(y)\,. \tag{6}$$

Specifically, F is the *dual function* (convex conjugate) of the negative entropy regularizer R, and the gradient of F maps from the dual to the primal space. For $z = (x, y) \in \mathbb{R}^{m+n}$, define $\nabla F(z) = (\nabla_x F(z), \nabla_y F(z))$. Then for $z \in \mathbb{R}^{m+n}$ and $w = (p, q) \in \mathcal{W}$ such that $p = \text{softmax}_m(x)$ and $q = \text{softmax}_n(y)$, we have $\nabla F(z) = w$. The central consequence of this property is that the dual iterates of (4) can be written as a modified *skew-gradient descent* with respect to F. Formally:

Proposition 3.1 (Optimistic Skew-Gradient Descent)**.** *Let $\{w_t\}$ be the iterates of (OMWU) with stepsize $\eta > 0$ initialized from $w_0 \in \text{ri}(\mathcal{W})$. Let $\{z_t\}$ be the dual iterates of (4) initialized from $z_0 \in \mathbb{R}^{m+n}$ such that $\nabla F(z_0) = w_0$. Then for all $t \geq 0$, it holds that $w_t = \nabla F(z_t)$ and*

$$z_{t+1} = \underbrace{z_t - \eta J \nabla F(z_t)}_{\textit{skew-gradient descent}} \underbrace{- \eta(J\nabla F(z_t) - J\nabla F(z_{t-1}))}_{\textit{optimistic skew-gradient correction}}\,. \tag{OMWU Dual}$$

Proposition 3.1 (proof in Section C.2) shows that (OMWU Dual) consists of a *skew-gradient descent* step with an additional *optimistic corrective term*. As shown by Wibisono et al. (2022), the dual iterates following skew-gradient descent correspond to the standard MWU algorithm, which is a forward Euler discretization of the continuous-time skew-gradient flow $\dot{z}(t) = -J\nabla F(z(t))$. Under skew-gradient flow, energy F is conserved along the trajectory, while under the forward discretization, energy is non-decreasing (see Section C.2.2 for a review). In contrast, for OMWU, we prove the correction term leads to strict *energy dissipation*, which will yield last-iterate convergence to NE.

3.2 Primal Convergence in KL via Dual Dissipation of Energy

In this section, we show that the change in energy over the dual iterates is exactly the change in $\mathrm{KL}(w^\star, w_t)$ over the primal iterates when $A \in \mathbb{R}^{m\times n}$ has an interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$. This equivalence is due to a more general relationship between F and $\mathrm{KL}(w^\star, \cdot)$ that holds over the *effective* space of dual iterates. To state this relationship, we first require introducing the linear subspaces $\mathcal{Z}, \mathcal{S}, \mathcal{S}^\perp \subset \mathbb{R}^{m+n}$ and several additional preliminaries. First, we define the subspaces:

$$\mathcal{Z} = \mathrm{Span}(J\mathcal{W})\,,\ \mathcal{S} = \mathrm{Span}\left(\begin{pmatrix}\mathbf{1}_m\\0\end{pmatrix}, \begin{pmatrix}0\\\mathbf{1}_n\end{pmatrix}\right),\ \text{and } \mathcal{S}^\perp = \{v \in \mathbb{R}^{m+n} : \forall s \in \mathcal{S},\ \langle v, s\rangle = 0\}. \quad (7)$$

We refer to $\mathcal{Z}$ as the *effective dual space*, since (4) gives $z_t - z_0 \in \mathcal{Z}$ for all $t \geq 0$. In particular, $z_t \in \mathcal{Z}$ whenever $z_0 \in \mathcal{Z}$. Moreover, $\mathcal{S}$ is the subspace of constant shift directions, and $\mathcal{S}^\perp$ is its orthogonal complement.

Geometry of the effective dual space. A key property is that $\mathcal{Z}$ is *orthogonal* to interior Nash equilibria of A. This is because for any interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$, $Aq^\star$ and $A^\top p^\star$ are the constant vectors $Aq^\star = c\mathbf{1}_m$ and $A^\top p^\star = c\mathbf{1}_n$ for some $c \in \mathbb{R}$ (see Proposition B.1). Together with the skew-symmetry of J, this implies $\langle z, w^\star\rangle = 0$ for all $z \in \mathcal{Z}$. See Proposition D.6.

Primal-dual coupling despite lack of strict convexity. The energy function is not globally strictly convex, as ∇F is invariant under shifts $s \in \mathcal{S}$ (see Proposition D.3). However, the dual relationship of R and F and the structure of the gradient map ∇R still allows for establishing the following key relationship: for $z \in \mathbb{R}^{m+n}$ and $w = \nabla F(z)$, it holds that $\mathrm{KL}(\widetilde{w}, w) = R(\widetilde{w}) + F(z) - \langle z, \widetilde{w}\rangle$ for all $\widetilde{w} \in \mathrm{ri}(\mathcal{W})$. We state and prove this formally in Proposition D.2.

Together, these properties lead to the following relationship between $\mathrm{KL}(w^\star, \cdot)$ and F over $\mathcal{Z}$:

Proposition 3.2 (Equivalence between energy and KL differences)**.** *Let $A \in \mathbb{R}^{m\times n}$ have an interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$. Fix $z, z' \in \mathcal{Z}$, and let $w = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$ and $w' = \nabla F(z') \in \mathrm{ri}(\mathcal{W})$. Then it holds that*

$$\mathrm{KL}(w^\star, w') - \mathrm{KL}(w^\star, w) = F(z') - F(z)\ .$$

The proof of the proposition is in Section D.2. For the OMWU iterates, $z_{t+1} - z_t \in \mathcal{Z}$ and hence $\langle z_{t+1} - z_t, w^\star\rangle = 0$. Applying the preceding primal-dual identity at z_t and z_{t+1} gives $\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t)$ for every $t \geq 0$. Moreover, when $w^\star$ is unique and interior, we prove in Section D.3 that the map $\nabla F : \mathcal{Z} \to \mathrm{ri}(\mathcal{W})$ is *surjective* (Proposition D.9), and that $\mathrm{KL}(w^\star, w) = F(z) - \min_{z'\in\mathcal{Z}} F(z')$ for all $z \in \mathcal{Z}$ and $w = \nabla F(z)$ (Proposition D.11). Taken together, these propositions establish an exact equivalence between energy dissipation in the dual space and minimizing KL divergence to NE in the primal space.

4 Tight Bounds on Energy Dissipation

In this section, we establish tight bounds on energy dissipation under (OMWU Dual). Due to the relationships discussed in Section 3.2, this leads to a sharp last-iterate convergence rate in KL.

4.1 One-Step Change in Energy under OMWU

We show a dissipative one-step change in energy under (OMWU Dual) that is measured using the local seminorm induced by the Hessian of F. For this, we first introduce the following notation:

Local seminorms induced by the energy Hessian. The energy function F is convex and twice continuously differentiable, and thus for all $z \in \mathbb{R}^{m+n}$ its Hessian matrix $\nabla^2 F(z)$ is PSD and induces a degenerate inner product and seminorm. Specifically, we define $\langle u, v\rangle_z := \langle u, \nabla^2 F(z)v\rangle$ and $\|v\|_z := \sqrt{\langle v, \nabla^2 F(z)v\rangle}$ for $u, v \in \mathbb{R}^{m+n}$. We refer to $\|\cdot\|_z$ as a *local seminorm*, as the

spectrum of $\nabla^2 F(z)$ is dependent on the state z. By Proposition E.1(iii), this seminorm vanishes precisely on the shift subspace $\mathcal{S}$, and thus induces a norm on the quotient space $\mathbb{R}^{m+n}/\mathcal{S}$. We give further preliminaries on the energy Hessian in Section E.

Under the following constant stepsize condition, we prove our key energy dissipation lemma.

Assumption 1 (Stepsize). $0 < \eta \leq \frac{1}{4(54\sigma_{\max}+9)}$, where $\sigma_{\max} = \|A\|_2$.

Lemma 4.1 (Energy Dissipation). *For* (OMWU Dual) *with η satisfying Assumption 1, for all $t \geq 1$:*

$$-\tfrac{5}{4} \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \;\leq\; F(z_{t+1}) - F(z_t) \;\leq\; -\tfrac{1}{20} \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \,.$$

The significance of Lemma 4.1 is due to its *exact* characterization of the one-step change in energy, with matching upper and lower bounds. We make several further remarks on its interpretation:

Dependence on local geometry. The one-step change in energy is nonpositive and proportional to a *dissipation term* $\|J\nabla F(z_t)\|_{z_t}^2$. This term measures the magnitude of the skew-gradient payoff vector under the *local geometry* of $\nabla^2 F(z_t)$. When the energy Hessian has low curvature, the dissipation term is small and leads to a slow one-step change in energy. Note that the *flatness* of F is dual to the *steepness* of the entropy regularizer R, and thus dissipation will grow small when $w_t = \nabla F(z_t)$ approaches the boundary of the simplex, where R is steepest.

Primal characterization of dissipation. The dissipation term can also be expressed as the *variance* of the payoff vector Jw_t under $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$. Specifically, $\|J\nabla F(z_t)\|_{z_t}^2 = \mathrm{Var}_{p_t}(Aq_t) + \mathrm{Var}_{q_t}(A^\top p_t)$. This characterization stems from the structure of $\nabla^2 F$ and the relation $\nabla F(z_t) = w_t$. See Section E.1. Since w_t is interior, these variances vanish exactly when each player's actions have equal payoffs, or equivalently when w_t is a Nash equilibrium. Thus, for $t \geq 1$, Lemma 4.1 gives a strict decrease in energy whenever the primal iterate is not an equilibrium.

Asymptotic last-iterate convergence in KL. Together with Proposition 3.2, the nonpositive upper bound in Lemma 4.1 also immediately implies that OMWU has *asymptotic* last-iterate convergence to NE. Formally, we have the following result (see Section H for the short proof):

Theorem 4.2 (Asymptotic Last-Iterate Convergence). *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star$, and let $\{w_t\}$ be the iterates of* (OMWU) *on A with η satisfying Assumption 1. Then the sequence of iterates $\{w_t\}$ converges, and $\lim_{t\to\infty} w_t = w^\star$.*

4.2 Proof Overview of Lemma 4.1

We now give the main intuitions for proving the lemma. The full proof is developed in Section F.

Leading dissipation term. To see how the term $\|J\nabla F(z_t)\|_{z_t}^2$ arises, recall from (OMWU Dual) that

$$z_{t+1} - z_t = -\eta J\nabla F(z_t) - \eta J(\nabla F(z_t) - \nabla F(z_{t-1})).$$

Then using a first-order Taylor expansion, the difference of energy gradients in the corrective term can be expressed exactly as

$$\nabla F(z_t) - \nabla F(z_{t-1}) = \nabla^2 F(z_t)(z_t - z_{t-1}) + G_F(z_t, z_{t-1}) \,, \tag{8}$$

where $G_F(z_t, z_{t-1})$ is a remainder defined precisely in Proposition E.4. By using the dual update rule at time t and repeating a similar expansion, we can further express $z_t - z_{t-1} = -\eta J\nabla F(z_t) - Q_t$ for an error term Q_t (defined precisely in Proposition F.3). Substituting this difference into (8), and substituting (8) back into the dual update rule at time $t+1$, we can ultimately write

$$z_{t+1} - z_t = -\eta J\nabla F(z_t) + \eta^2 J\nabla^2 F(z_t) J\nabla F(z_t) + \eta W_t \,, \tag{9}$$

where $W_t = J\nabla^2 F(z_t) Q_t - J G_F(z_t, z_{t-1})$. Then, using the definition of the Bregman divergence D_F, we substitute (9) to quantify the one-step change in energy:

$$\begin{aligned} F(z_{t+1}) - F(z_t) &= \langle \nabla F(z_t), z_{t+1} - z_t \rangle + D_F(z_{t+1}, z_t) \\ &= -\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 + \eta \langle \nabla F(z_t), W_t \rangle + D_F(z_{t+1}, z_t) \,. \end{aligned} \tag{10}$$

Here, the final equality follows by the skew-symmetry of $J = -J^\top$ and by definition of $\|\cdot\|_{z_t}$.

Controlling the error terms. The expansion of $F(z_{t+1}) - F(z_t)$ in (10) is exact, and thus the remaining technical task is to control the magnitude of the latter two terms. In Propositions F.7, F.8, and F.9, we prove that $|\eta\langle\nabla F(z_t), W_t\rangle + D_F(z_{t+1}, z_t)| \le C\eta^2\|J\nabla F(z_t)\|_{z_t}^2$ for an absolute constant $C < 1$ when η satisfies Assumption 1. This allows the leading dissipation term to absorb the error terms, resulting in the statement of the lemma. Our analysis for controlling these error terms relies on the following *local Hessian stability* (LHS) property of the energy function F:

Proposition 4.3 (LHS Property)**.** *For any $z, z' \in \mathbb{R}^{m+n}$ and $\alpha > 0$, if $\|z - z'\|_\infty \le \alpha$, then:*

$$\exp(-2\alpha) \cdot \nabla^2 F(z) \preceq \nabla^2 F(z') \preceq \exp(2\alpha) \cdot \nabla^2 F(z) \; . \tag{LHS}$$

Roughly speaking, the (LHS) property is a local and relative notion of third-order smoothness. The property states that the spectra of $\nabla^2 F(z')$ and $\nabla^2 F(z)$ are similar, up to a multiplicative factor depending on the closeness of z, z'. The dual OMWU iterates all satisfy $\|z_{t+1} - z_t\|_\infty \le 3\eta\sigma_{\max}$ (see Proposition F.4), so for a small enough constant stepsize η, the (LHS) property allows for establishing the bounds on the error terms above. The (LHS) property is also related to *self-concordance*, with the former often derived as a consequence of the latter. Note that a recent work of Freund et al. (2026) established self-concordance of the log-sum-exp function (which implies Proposition 4.3 by definition of F). For completeness, we give a short direct proof of Proposition 4.3 in Section E.2.

4.3 Universal Linear Last-Iterate Convergence Rate

In addition to the asymptotic last-iterate convergence of Theorem 4.2, our new dual energy-based analysis leads to a sharp linear last-iterate convergence rate. This relies on a structural inequality relating $\|J\nabla F(z)\|_z^2$ to $\mathrm{KL}(w^\star, w)$. For this, let $\sigma_{\min} := \sigma_{\min}(J, \mathcal{S}^\perp) = \inf_{v \in \mathcal{S}^\perp \setminus \{0\}} \|\Pi_{\mathcal{S}^\perp}(Jv)\|_2 / \|v\|_2$ denote the minimum singular value of J restricted to $\mathcal{S}^\perp$. When A has a unique and interior NE, then $\sigma_{\min} > 0$ (see Proposition G.9). We then prove the following general inequality:

Proposition 4.4 (Non-uniform skew-gradient domination)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$. Fix $z \in \mathcal{Z}$, and let $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Let $p_{\min} = \min_{i\in[m]} p(i)$ and $q_{\min} = \min_{j\in[n]} q(j)$, and define $w_{\min} := \min\{p_{\min}, q_{\min}\}$. Then $\sigma_{\min} > 0$, and moreover*

$$\|J\nabla F(z)\|_z^2 \ge \sigma_{\min}^2 \cdot w_{\min}^2 \cdot \mathrm{KL}(w^\star, w) \; .$$

The proof of the inequality is given in Section G. Here, we make several remarks:

Remark 4.5 (Skew-gradient domination)**.** As $\mathrm{KL}(w^\star, w) = F(z) - \min_{z^\star \in \mathcal{Z}} F(z^\star)$ (Proposition D.11), note that Proposition 4.4 establishes a structural property akin to *gradient domination* in optimization, but with key differences arising from the non-Euclidean geometry induced by $\|\cdot\|_z$ and the concern of *skew*-gradients. In particular, the relationship is constrained by a non-uniform and state-dependent factor $w_{\min}^2$, a term that captures a bottleneck arising from the local geometry of F. Specifically, $w_{\min}$ is a lower bound on the smallest positive eigenvalue of $\nabla^2 F(z)$, and this bound deteriorates when the corresponding primal variable is near the simplex boundary. In these regions, the term $\|J\nabla F(z)\|_z^2$ can be small even when w is far from Nash in KL. In Proposition J.1, we give an *upper bound* on $\|J\nabla F(z)\|_z^2$ showing that this non-uniform dependence on $w_{\min}$ is necessary.

Remark 4.6 (Independence of bottlenecks)**.** A second bottleneck arises when J is ill-conditioned and has a small restricted singular value $\sigma_{\min}$. Intuitively, this means that large primal perturbations $w - w' \in \mathcal{S}^\perp$ lead to only small relative changes in $J(w - w')$, which makes OMWU less reactive to meaningful signal from payoff vectors. Importantly, $\sigma_{\min}$ is a property of J that is completely independent from the geometry of F and the proximity of the Nash $w^\star$ to the simplex boundary: we show in Propositions G.15 and G.16 that there are families of matrices where (i) $w^\star$ is uniform but $\sigma_{\min}$ is arbitrarily small, and (ii) where $\sigma_{\min} = \frac{1}{2}$, but $w^\star$ is arbitrarily close to a vertex. See Section G.5.

Recall by Proposition 3.2 that $\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t)$ along the OMWU iterates. Then combining the energy dissipation bound of Lemma 4.1 and the skew-gradient domination property of Proposition 4.4, we obtain a new, linear last-iterate convergence rate in KL:

Theorem 4.7 (Last-Iterate Convergence Rate in KL)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star = (p^\star, q^\star)$. Let $\{w_t\}$ denote the iterates of (OMWU) on A with η satisfying Assumption 1, initialized from $w_0 \in \mathrm{ri}(\mathcal{W})$. Let $\delta_p = \min_{i\in[m]} p^\star(i)$, $\delta_q = \min_{j\in[n]} q^\star(j)$, and $\delta := \min\{\delta_p, \delta_q\}$.*

For every t, let $p_{t,\min} = \min_{i\in[m]} p_t(i)$, $q_{t,\min} = \min_{j\in[n]} q_t(j)$, and $w_{t,\min} := \min\{p_{t,\min}, q_{t,\min}\}$. Let $\Lambda := \mathrm{CrossEnt}(w^\star, w_0) = \mathrm{KL}(w^\star, w_0) - R(w^\star)$. Then for all $t \geq 1$, the following hold:

$$\textit{(1)} \quad \mathrm{KL}(w^\star, w_{t+1}) \leq \mathrm{KL}(w^\star, w_t) \cdot \left(1 - \tfrac{1}{20} \cdot \eta^2 \sigma_{\min}^2 \cdot w_{t,\min}^2\right) ,$$
$$\textit{(2)} \quad \mathrm{KL}(w^\star, w_{t+1}) \leq 2\mathrm{KL}(w^\star, w_0) \cdot \exp\left(-\tfrac{1}{20} \cdot \eta^2 \sigma_{\min}^2 \cdot \exp\left(\tfrac{-3\Lambda}{\delta}\right) \cdot t\right) .$$

The proof of the theorem is given in Section H. Here, we make the following remarks:

Remark 4.8 (Trajectory-dependence)**.** The dependence on $w_{t,\min}$ in Part (1) of the theorem highlights the influence of the iterates' trajectory on the rate of convergence. The trajectory depends on the initialization w_0, and more importantly on the location of the NE. When $w^\star$ is close to the simplex boundary (meaning δ is small), the iterates $\{w_t\}$ are also more frequently near the boundary. In these regions $w_{t,\min}$ is small, meaning the one-step contraction in KL is also small. This compounds in slower overall convergence. Note that the sequence $\{w_{t,\min}\}$ is not monotonic. However, by establishing a *uniform lower bound* of $w_{t,\min}^2 \geq \exp(\frac{-3\Lambda}{\delta})$ (Lemma H.2) we obtain Part (2).

Remark 4.9 (Comparison with Wei et al. (2021))**.** The linear rate of Wei et al. (2021, Theorem 3) also depends on $\exp(\frac{-\Lambda}{\delta})$. However, explicitly tracking this dependence in their proof yields the much slower bound of $\mathrm{KL}(w^\star, w_t) \leq O(\exp(\frac{\Lambda}{\delta}) \cdot \exp(-\eta^2 \exp(\frac{-\Lambda}{\delta}) \cdot t))$. Compared with Wei et al., our bound improves the leading multiplicative factor by an exponential factor in the inverse minimum Nash coordinate δ. Moreover, our new energy-based proof technique offers a more transparent and quantitative understanding of the geometric bottlenecks that contribute to slow convergence. See Section H.4 for further comparisons.

Matching Last-Iterate Lower Bound in KL. We further prove that the rate in Theorem 4.7 is *optimal*. In particular, we prove a *lower bound* on the last-iterate convergence rate in KL with a matching dependence on δ, establishing that our new analysis is tight. Formally:

Theorem 4.10 (Lower bound in KL)**.** *Assume the definitions of Theorem 4.7. Fix any $T \geq 3$. Then there exists $A \in [-1, 1]^{2\times 2}$ with an interior NE $w^\star = (p^\star, q^\star)$ with $\delta := \delta_p = \delta_q$ such that, from a positive measure set of initializations w_0, each with $\mathrm{KL}(w^\star, w_0) \leq 6$, the OMWU iterates $\{w_t\}$ with η as in Assumption 1 satisfy*

$$\mathrm{KL}(w^\star, w_t) \geq \mathrm{KL}(w^\star, w_0) \cdot \exp(-40\eta^2 \exp(\tfrac{-1}{\delta}) \cdot t) \quad \textit{for all } t \in [T].$$

Together with Theorem 4.7, this lower bound indicates that the most severe geometric bottleneck in convergence arises when $w^\star$ is near a *vertex* of the simplex. We give the proof in Section J.

5 New Bounds on Uniform Convergence Rates

We use the geometric insights developed in Theorems 4.7 and 4.10 to additionally prove new bounds on *uniform* convergence rates under OMWU. For concreteness, fix the dimensions m and n, and recall that an algorithm has a *uniform last-iterate* or *uniform best-iterate* convergence rate of order $f(T)$ under a distance measure D if the following holds: *for all* $T \geq 1$ and *for all* $A \in [-1, 1]^{m\times n}$:

$$\begin{aligned} &\text{(uniform last-iterate)}: && D(w_T) \leq O(f(T)) \\ \text{or } &\text{(uniform best-iterate)}: && \min_{t\in[T]} D(w_t) \leq O(f(T)) . \end{aligned}$$

Here, we emphasize that f must be a function only of the time horizon T and the dimensions m and n, but with *no dependence* on the minimum coordinate δ of the NE in a particular instance. Proving such uniform rates thus allows quantitative guarantees to hold in an instance-agnostic manner (see also Cai et al. (2025) for further discussion on these different notions).

While Cai et al. (2024) proved for OMWU that no uniform last-iterate convergence rates are attainable, the follow-up of Cai et al. (2025) proved for the class of 2×2 games with an interior NE, that, when initialized at the uniform distribution, OMWU does have an $O(T^{-1/6})$ *uniform best-iterate convergence rate* in *duality gap*. Unfortunately, we prove that similar uniform best-iterate rates are *not* attainable under the stronger measures of TV distance and KL divergence.

Lower Bounds on Best-Iterate in KL and TV. We prove the following *constant lower bounds*:

Theorem 5.1 (Uniform Best-Iterate Lower Bounds)**.** *For every $T \geq 2$, there exists $A \in [-1, 1]^{2\times 2}$ with an interior Nash equilibrium $w^\star$ such that, for the iterates $\{w_t\}$ of (OMWU) on A with η*

satisfying Assumption 1, from a positive measure set of initializations that includes the joint uniform distributions, the following hold:

$$\min_{t\in[T]} \mathrm{TV}(w^\star, w_t) \geq \tfrac{1}{3} \quad \textit{and} \quad \min_{t\in[T]} \mathrm{KL}(w^\star, w_t) \geq \tfrac{1}{9} \ .$$

Note that these lower bounds in best-iterate also clearly imply a lower bound on last-iterate. We give the proof and a full discussion in Section K, where we also establish an additional uniform best-iterate lower bound in DG when the initialization is not the uniform distribution (see Theorem K.1).

Fast Best-Iterate Convergence Rate in Duality Gap for 2x2 Games. In contrast, for 2×2 games, we obtain in *duality gap* a new upper bound on *uniform best-iterate convergence rate* that substantially improves over the prior $O(T^{-1/6})$ rate of Cai et al. (2025) that was proven for the same setting:

Theorem 5.2 (Uniform Best-Iterate Convergence Rate in DG)**.** *Let $A \in [-1,1]^{2\times 2}$ be a zero-sum game with an interior Nash equilibrium $w^\star$, and let $\{w_t\}$ denote the iterates of running* (OMWU) *on A with a fixed absolute constant stepsize η satisfying Assumption 1, initialized at the uniform distributions. Then for any $T \geq 2$, it holds that*

$$\min_{t\in[T]} \mathrm{DG}(w_t) \leq O(T^{-1/2} \cdot \sqrt{\log T}) \ .$$

See Section L for the proof of the theorem and a full discussion. Notably, our proof relies on the new, universal last-iterate analysis in KL divergence from Theorem 4.7, along with an exact control of the OMWU trajectory in the 2×2 setting (see Section I). While we prove the result only for this low-dimensional regime, we conjecture that the same bound holds in higher dimensions. We expect our techniques will be useful for obtaining such a result, but we leave this for future work.

6 Conclusion

Our work obtains tight new guarantees on the iterate-to-iterate behavior of OMWU in two-player zero-sum games, and it offers sharp geometric explanations for the algorithm's potentially slow convergence. We believe our new energy-based analysis framework can be naturally extended to study the last-iterate convergence properties of other instantiations of Optimistic FTRL, in other settings (such as under bandit feedback), and for other more general classes of games. We leave these as directions for future work.

Acknowledgements. JL thanks JJ for inspiration. AV, AB, and JL were supported by the MOE Tier 2 Grant (MOE-T2EP20223-0018), the CQT++ Core Research Funding Grant (SUTD) (RS-NRCQT-00002), the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (Award Number: AISG2-RP-2020-016), and partially by Project MIS 5154714 of the National Recovery and Resilience Plan, Greece 2.0, funded by the European Union under the NextGenerationEU Program. AW is supported by NSF awards CCF-2403391 and CAREER CCF-2443097.

References

Abernethy, J., Lai, K. A., and Wibisono, A. (2021). Fast convergence of fictitious play for diagonal payoff matrices. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*.

Anagnostides, I., Daskalakis, C., Farina, G., Fishelson, M., Golowich, N., and Sandholm, T. (2022a). Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*.

Anagnostides, I., Panageas, I., Farina, G., and Sandholm, T. (2022b). On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning (ICML 2022)*.

Arora, S., Hazan, E., and Kale, S. (2012). The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, 8(1):121–164.

Azizian, W., Iutzeler, F., Malick, J., and Mertikopoulos, P. (2024). The rate of convergence of bregman proximal methods: Local geometry versus regularity versus sharpness. *SIAM Journal on Optimization*, 34(3):2440–2471.

Bach, F. (2010). Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4:384–414.

Bailey, J. and Piliouras, G. (2019a). Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. *Advances in Neural Information Processing Systems (NeurIPS 2019)*.

Bailey, J. P., Gidel, G., and Piliouras, G. (2020). Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Conference on Learning Theory (COLT 2020)*.

Bailey, J. P. and Piliouras, G. (2018). Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC 2018)*.

Bailey, J. P. and Piliouras, G. (2019b). Multi-agent learning in network zero-sum games is a Hamiltonian system. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2019)*.

Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

Cai, Y., Farina, G., Grand-Clément, J., Kroer, C., Lee, C., Luo, H., and Zheng, W. (2024). Fast last-iterate convergence of learning in games requires forgetful algorithms. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2024)*.

Cai, Y., Farina, G., Grand-Clément, J., Kroer, C., Lee, C., Luo, H., and Zheng, W. (2025). On separation between best-iterate, random-iterate, and last-iterate convergence of learning in games. *CoRR*, abs/2503.02825.

Cai, Y., Oikonomou, A., and Zheng, W. (2022). Finite-time last-iterate convergence for learning in multi-player games. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.

Cen, S., Chi, Y., Du, S. S., and Xiao, L. (2023). Faster last-iterate convergence of policy optimization in zero-sum markov games. In *The Eleventh International Conference on Learning Representations (ICLR 2023)*.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.

Chen, X. and Peng, B. (2020). Hedging in games: Faster convergence of external and swap regrets. *Advances in Neural Information Processing Systems (NeurIPS 2020)*.

Daskalakis, C., Fishelson, M., and Golowich, N. (2021). Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems (NeurIPS 2021)*.

Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. (2018). Training GANs with optimism. In *International Conference on Learning Representations (ICLR 2018)*.

Daskalakis, C. and Panageas, I. (2019). Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *10th Innovations in Theoretical Computer Science Conference, (ITCS 2019)*.

Facchinei, F. and Pang, J.-S. (2003). *Finite-dimensional variational inequalities and complementarity problems*. Springer.

Fasoulakis, M., Markakis, E., Pantazis, Y., and Varsos, C. (2022). Forward looking best-response multiplicative weights update methods for bilinear zero-sum games. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2022)*.

Feng, Y., Li, P., Panageas, I., and Wang, X. (2024). Last-iterate convergence separation between extra-gradient and optimism in constrained periodic games. In *The 40th Conference on Uncertainty in Artificial Intelligence (UAI 2024)*.

Freund, Y., Harvey, N. J., Portella, V. S., Qi, Y., and Wang, Y.-X. (2026). A second order regret bound for normalhedge. *arXiv preprint arXiv:2602.08151*.

Freund, Y. and Schapire, R. E. (1999). Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103.

Gao, B. and Pavel, L. (2017). On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*.

Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. (2019). A variational inequality perspective on generative adversarial networks. In *International Conference on Learning Representations (ICLR 2019)*.

Golowich, N., Pattathil, S., and Daskalakis, C. (2020). Tight last-iterate convergence rates for no-regret learning in multi-player games. *Advances in Neural Information Processing Systems (NeurIPS 2020)*.

Gorbunov, E., Taylor, A., and Gidel, G. (2022). Last-iterate convergence of optimistic gradient method for monotone variational inequalities. *Advances in Neural Information Processing Systems (NeurIPS 2022)*.

Katona, J. E., Wang, X., and Wibisono, A. (2026). A symplectic analysis of alternating mirror descent. *Journal of Machine Learning Research*, 27(44):1–61.

Korpelevich, G. M. (1976). The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756.

Lazarsfeld, J., Piliouras, G., Sim, R., and Skoulakis, S. (2025a). Optimism without regularization: Constant regret in zero-sum games. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025)*.

Lazarsfeld, J., Piliouras, G., Sim, R., and Wibisono, A. (2025b). Fast and furious symmetric learning in zero-sum games: Gradient descent as fictitious play. *Conference on Learning Theory (COLT 2025)*.

Lee, C.-W., Kroer, C., and Luo, H. (2021). Last-iterate convergence in extensive-form games. *Advances in Neural Information Processing Systems (NeurIPS 2021)*.

Legacci, D., Mertikopoulos, P., Papadimitriou, C., Piliouras, G., and Pradelski, B. (2024). No-regret learning in harmonic games: Extrapolation in the face of conflicting interests. *Advances in Neural Information Processing Systems (NeurIPS 2024)*.

Lei, Q., Nagarajan, S. G., Panageas, I., et al. (2021). Last iterate convergence in no-regret learning: constrained min-max optimization for convex-concave landscapes. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*.

Liang, T. and Stokes, J. (2019). Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *The 22nd International Conference on Artificial Intelligence and Statistics (AISTATS 2019)*.

Mertikopoulos, P., Lecouat, B., Zenati, H., Foo, C.-S., Chandrasekhar, V., and Piliouras, G. (2019). Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile. In *International Conference on Learning Representations (ICLR 2019)*.

Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. (2018). Cycles in adversarial regularized learning. In *Proceedings of the twenty-ninth annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)*.

Mokhtari, A., Ozdaglar, A., and Pattathil, S. (2020). A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*.

Munos, R., Valko, M., Calandriello, D., Gheshlaghi Azar, M., Rowland, M., Guo, Z. D., Tang, Y., Geist, M., Mesnard, T., Fiegel, C., Michi, A., Selvi, M., Girgin, S., Momchev, N., Bachem, O., Mankowitz, D. J., Precup, D., and Piot, B. (2024). Nash learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*.

Nemirovski, A. (2004). Prox-method with rate of convergence o (1/t) for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251.

Nesterov, Y. (2007). Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2):319–344.

Nocedal, J. and Wright, S. J. (2006). *Numerical optimization*. Springer.

Ota, T. and Fujimoto, Y. (2025). Hamiltonian of polymatrix zero-sum games. *arXiv preprint arXiv:2505.12609*.

Popov, L. D. (1980). A modification of the arrow-hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848.

Rakhlin, S. and Sridharan, K. (2013). Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems (NeurIPS 2013)*.

Rudin, W. (1991). *Functional Analysis*. McGraw-Hill, New York, second edition.

Shalev-Shwartz, S. et al. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.

Soleymani, A., Piliouras, G., and Farina, G. (2025). Faster rates for no-regret learning in general games via cautious optimism. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*.

Sun, T. and Tran-Dinh, Q. (2019). Generalized self-concordant functions: a recipe for newton-type methods. *Mathematical Programming*, 178(1):145–213.

Syrgkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. (2015). Fast convergence of regularized learning in games. *Advances in Neural Information Processing Systems (NeurIPS 2015)*.

Tiapkin, D., Calandriello, D., Belomestny, D., Moulines, E., Naumov, A., Rasul, K., Valko, M., and Menard, P. (2026). Proximal point nash learning from human feedback. *arXiv preprint arXiv:2505.19731*.

Tran-Dinh, Q., Li, Y.-H., and Cevher, V. (2015). Composite convex minimization involving self-concordant-like cost functions. In *Proceedings of the 3rd International Conference on Modelling, Computation and Optimization in Information Systems and Management Sciences (MCO 2015)*.

von Neumann, J. (1928). Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320.

Wang, G., Acharya, K., Lakshmikanthan, L., Ziani, J., and Muthukumar, V. (2026). Last-iterate convergence for symmetric, general-sum, 2x2 games under the exponential weights dynamic. In *37th International Conference on Algorithmic Learning Theory (ALT 2026)*.

Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2021). Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations (ICLR 2021)*.

Wibisono, A., Tao, M., and Piliouras, G. (2022). Alternating mirror descent for constrained min-max games. *Advances in Neural Information Processing Systems (NeurIPS 2022)*.

Table of Contents

1 Introduction **2**
2 Preliminaries **4**
3 OMWU as Optimistic Skew-Gradient Descent **5**
4 Tight Bounds on Energy Dissipation **6**
5 New Bounds on Uniform Convergence Rates **9**
6 Conclusion **10**
A Other Related Work **14**
B Preliminaries on Zero-Sum Games **15**
C Details on Optimistic MWU Algorithm **19**
D Details on Equivalence of KL Minimization and Energy Dissipation **22**
E Details on Energy Function Hessian Matrix **28**
F Details on Energy Dissipation Under OMWU **34**
G Details on Non-Uniform Skew-Gradient Domination **45**
H Details on Universal Last-Iterate Convergence in KL **54**
I Details on OMWU Dynamics in 2x2 Setting **59**
J Details on Lower Bound for Universal Last-Iterate Convergence in KL **69**
K Details on Uniform Convergence Rate Lower Bounds **73**
L Details on Best-Iterate Convergence Rate in Duality Gap for 2x2 Setting **77**

A Other Related Work

In this section, we give additional discussion on other related works.

Optimistic algorithms for min-max optimization and variational inequalities. The OMWU algorithm is more broadly related to optimistic and extra-gradient methods in monotone variational inequalities (Facchinei and Pang, 2003). Such algorithms date back to the 1970s and 80s (Korpelevich, 1976; Popov, 1980), and within the past decade have received growing theoretical interest, in part due to their applications in machine learning settings (see, e.g., Nemirovski (2004), Nesterov (2007), Daskalakis et al. (2018), Gidel et al. (2019), Mertikopoulos et al. (2019), Liang and Stokes (2019), Mokhtari et al. (2020), Golowich et al. (2020), Gorbunov et al. (2022), Cai et al. (2022), and references therein).

Optimistic MWU for online learning and learning in games. The Optimistic Mirror Descent and Optimistic Follow-the-Regularized-Leader families for adversarial online learning, of which OMWU is the canonical instantiation with negative entropy regularizer, were introduced by Rakhlin and Sridharan (2013). Since then, OMWU has emerged as a premier algorithm for online learning in games, and it has been proven to achieve optimal or near-optimal regret bounds (which correspond to time-average convergence guarantees) in many classes of games, including zero-sum and general-sum normal-form games (see, e.g., Syrgkanis et al. (2015); Chen and Peng (2020); Daskalakis et al. (2021); Anagnostides et al. (2022a); Soleymani et al. (2025), and references therein).

Last-iterate Convergence of OMWU in zero-sum games. The asymptotic last-iterate convergence of OMWU in two-player bilinear zero-sum games was first established by Daskalakis and Panageas (2019) on games with a unique Nash equilibrium, and Lei et al. (2021) generalized this result to hold also for convex-concave zero-sum games. Wei et al. (2021) later obtained a quantitative (universal) linear last-iterate convergence rate in the bilinear case. Unlike Optimistic Gradient-Descent-Ascent (OGDA), which has a uniform $O(T^{-1/2})$ last-iterate convergence rate in this setting (Gorbunov et al., 2022; Cai et al., 2022), OMWU was shown by Cai et al. (2024) to have no such uniform last-iterate convergence rate in duality gap (i.e., that does not depend on any game-dependent constant such as the minimum non-zero Nash equilibrium mass). Cai et al. (2024) attributed this behavior of OMWU to a notion they call *forgetfulness*, as the algorithm's update rule (as is the case for all OFTRL instantiations) inherently depends on the accumulation of all previous gradients. Our paper reinforces the main takeaways of their result: the last-iterate of OMWU (and, under some measures of distance to Nash, the best-iterate convergence, cf. Theorem 5.1) can be arbitrarily slow. However,

our results also give new and more precise quantitative explanations for this slow convergence that are related to geometric properties of the algorithm.

Energy-based perspective for learning in games. Several prior works have also relied on using a dual or energy-based perspective for analyzing the convergence of various algorithms for learning in games. See, e.g., Mertikopoulos et al. (2018), Bailey and Piliouras (2018), Bailey and Piliouras (2019a), Bailey and Piliouras (2019b), Bailey et al. (2020), Abernethy et al. (2021), Wibisono et al. (2022), Azizian et al. (2024), Legacci et al. (2024), Ota and Fujimoto (2025), Lazarsfeld et al. (2025b), Lazarsfeld et al. (2025a), Katona et al. (2026), and the references therein. However, our work is the first to prove quantitative last-iterate convergence rates for the bilinear zero-sum setting by proving bounds on *strictly dissipating* energy.

Other related works on learning in games. The extra-gradient variant of MWU has been studied in bilinear zero-sum games and proven to have asymptotic last-iterate convergence (Fasoulakis et al., 2022). In the unconstrained min-max setting, optimistic and extra-gradient methods have also been studied in time-varying and periodic games, where Feng et al. (2024) established a separation between the convergence guarantees of the extra-gradient and optimistic variants. Lee et al. (2021) proved last-iterate convergence for a variant of OMWU in extensive-form games. Anagnostides et al. (2022b) proved an $O(T^{-1/2})$ best-iterate convergence rate in duality gap using Optimistic Mirror Descent on games with non-negative sums of regret (which is satisfied by zero-sum games). However, this result assumes a smoothness condition on the regularizer that is not satisfied by the negative entropy function. Thus, the $\widetilde{O}(T^{-1/2})$ best-iterate convergence rate in duality gap that we prove in Theorem 5.2 for the 2×2 setting is the fastest-known best-iterate result for OMWU. Moreover, this result also adds to a growing body of work that has established various convergence results for algorithms in the setting of 2×2 games, which often still requires intricate proof techniques despite the low dimension (Bailey and Piliouras, 2019a; Chen and Peng, 2020; Cai et al., 2024; Lazarsfeld et al., 2025a; Cai et al., 2025; Wang et al., 2026).

B Preliminaries on Zero-Sum Games

This Section gives additional details on the preliminaries of zero-sum games introduced in Section 2.

Organization of Section. This section is organized as follows:

- **Section B.1** recalls several important properties of zero-sum games with an interior NE.
- **Section B.2** derives the relationships between distances to Nash introduced in expression (2).
- **Section B.3** states and proves two useful properties related to the spectrum of the payoff matrix A.

The claims of this section are mostly standard but included for completeness.

B.1 Properties of Zero-Sum Games with Interior Equilibria

Recall that for $A \in \mathbb{R}^{m \times n}$ and $\epsilon \geq 0$, an *ϵ-approximate Nash equilibrium (NE)* $w^\star = (p^\star, q^\star) \in \mathcal{W}$ for the min-max game $\min_{p \in \Delta_m} \max_{q \in \Delta_n} \langle p, Aq \rangle$ satisfies the inequalities (see also (1))

$$\max_{q \in \Delta_n} \langle q, A^\top p^\star \rangle - \epsilon \leq \langle p^\star, Aq^\star \rangle \leq \min_{p \in \Delta_m} \langle p, Aq^\star \rangle + \epsilon \,. \tag{11}$$

When $\epsilon = 0$, we say $w^\star$ is a Nash equilibrium.

If a game has an interior Nash equilibrium, then the following key properties hold:

Proposition B.1. *Fix $A \in \mathbb{R}^{m \times n}$. Then the following properties hold:*

(i) Suppose $w^\star = (p^\star, q^\star) \in \mathrm{ri}(\mathcal{W})$ is an interior NE of A with $\langle p^\star, Aq^\star \rangle = c \in \mathbb{R}$. Then $Aq^\star = c\mathbf{1}_m$ and $A^\top p^\star = c\mathbf{1}_n$.

(ii) Fix $w = (p, q) \in \mathcal{W}$. Suppose $Aq = d\mathbf{1}_m$ and $A^\top p = d'\mathbf{1}_n$ for $d, d' \in \mathbb{R}$. Then $d = d'$ and w is a Nash equilibrium of A.

Proof. We prove the two claims separately:

Claim (i): By the definition of an NE from (11), and since $\langle p^\star, Aq^\star \rangle = c$, we must have

$$c = \langle p^\star, Aq^\star \rangle \leq \min_{p \in \Delta_m} \langle p, Aq^\star \rangle \,. \tag{12}$$

We will first show that both $(Aq^\star)(i) \geq c$ and $(Aq^\star)(i) \leq c$ for all $i \in [m]$. For the first direction, suppose by way of contradiction that there exists $j \in [m]$ such that $(Aq^\star)(j) < c$. Then construct $p' \in \Delta_m$ as follows: set $p'(j) = 1$, and $p'(i) = 0$ for all other $i \neq j \in [m]$. Then observe that

$$\min_{p \in \Delta_m} \langle p, Aq^\star \rangle \leq \langle p', Aq^\star \rangle = \sum_{k=1}^m p'(k) \cdot (Aq^\star)(k) = (Aq^\star)(j) < c = \langle p^\star, Aq^\star \rangle \, .$$

However, this contradicts the fact that $\langle p^\star, Aq^\star \rangle \leq \min_{p \in \Delta_m} \langle p, Aq^\star \rangle$ from (12). Thus we must have $(Aq^\star)(i) \geq c$ for all coordinates $i \in [m]$.

For the other direction, suppose by contradiction that there exists $j \in [m]$ where $(Aq^\star)(j) > c$. Note that if $(Aq^\star)(j) > c$ for *all* coordinates $j \in [m]$, then clearly $\langle p^\star, Aq^\star \rangle > c$, which contradicts (12). Thus suppose instead there exists coordinate $i \neq j \in [m]$ with $(Aq^\star)(i) \leq c$. Now construct $p' \in \Delta_m$ as follows: set $p'(k) = p^\star(k)$ for all $k \in [m] \setminus \{i, j\}$, set $p'(j) = 0$, and set $p'(i) = p^\star(i) + p^\star(j)$. Observe by construction that this implies

$$p'(j)\big((Aq^\star)(j)\big) + p'(i)\big((Aq^\star)(i)\big) < p^\star(j)(Aq^\star)(j) + p^\star(i)(Aq^\star)(i) \, .$$

Then this further implies by construction that

$$\min_{p \in \Delta_m} \langle p, Aq^\star \rangle \leq \langle p', Aq^\star \rangle < \langle p^\star, Aq^\star \rangle = c \, ,$$

which contradicts (12). Thus we also have $(Aq^\star)(i) \leq c$ for all coordinates $i \in [m]$, and thus $Aq^\star = c\mathbf{1}_m$. Repeating identical arguments, we also have $A^\top p^\star = c\mathbf{1}_n$, which completes the claim.

Claim (ii). By the assumptions of the claim that $Aq = d\mathbf{1}_m$ and $A^\top p = d'\mathbf{1}_n$, observe for any $u \in \Delta_m$ that $\langle u, Aq \rangle = \langle u, d\mathbf{1}_m \rangle = d$, and for any $v \in \Delta_n$ that $\langle v, A^\top p \rangle = \langle v, d'\mathbf{1}_n \rangle = d'$. Then since $\langle p, Aq \rangle = \langle q, A^\top p \rangle$ and both $p \in \Delta_m$ and $q \in \Delta_n$, we must have $d = d' = \langle p, Aq \rangle$. To further show that (p, q) is a Nash equilibrium, observe that since $Aq = d\mathbf{1}_m$ and $A^\top p = d\mathbf{1}_n$, then

$$\max_{q' \in \Delta_n} \langle q', A^\top p \rangle = \langle q, A^\top p \rangle = \min_{p' \in \Delta_m} \langle p', Aq \rangle \, .$$

Thus by definition in (11), $w = (p, q)$ is a Nash equilibrium of A. ■

B.2 Distances to Nash Equilibrium

Recall from Section 2 that we consider several measures of distance to Nash; duality gap (DG), total variation distance (TV), and KL divergence (KL). We first restate these definitions. Let $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$ be a Nash equilibrium of A. Then for any $w = (p, q) \in \mathcal{W}$, we define:

- **Duality gap (DG)**: $\text{DG}(w) = \max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \min_{p' \in \Delta_m} \langle p', Aq \rangle$.
- **Kullback-Leibler divergence (KL)**: Define the components $\text{KL}(p^\star, p) = \sum_{i=1}^m p^\star(i) \log\big(p^\star(i)/p(i)\big)$ and $\text{KL}(q^\star, q) = \sum_{j=1}^n q^\star(j) \log\big(q^\star(j)/q(j)\big)$. Then $\text{KL}(w^\star, w) = \text{KL}(p^\star, p) + \text{KL}(q^\star, q)$.
- **Total variation distance (TV)**: Define the components $\text{TV}(p^\star, p) = \frac{1}{2}\|p^\star - p\|_1$ and $\text{TV}(q^\star, q) = \frac{1}{2}\|q^\star - q\|_1$. Then $\text{TV}(w^\star, w) = \text{TV}(p^\star, p) + \text{TV}(q^\star, q)$.

We also define the chi-squared divergence $\chi^2(w^\star, w)$ as follows:

- **Chi-squared divergence:** First define the components $\chi^2(p^\star, p) = \sum_{i=1}^m \frac{(p^\star(i) - p(i))^2}{p(i)}$ and $\chi^2(q^\star, q) = \sum_{j=1}^n \frac{(q^\star(j) - q(j))^2}{q(j)}$. Then $\chi^2(w^\star, w) = \chi^2(p^\star, p) + \chi^2(q^\star, q)$.

B.2.1 Bounded Duality Gap Implies Approximate Nash Equilibrium

A standard fact is that a bounded duality gap implies an approximate Nash equilibrium. Formally:

Proposition B.2 (Small DG implies Nash)**.** *Let $A \in \mathbb{R}^{m \times n}$ be a zero-sum game. Fix $w = (p, q) \in \mathcal{W}$ and $\epsilon \geq 0$. If $\text{DG}(w) \leq \epsilon$, then w is an ϵ-approximate Nash equilibrium of A.*

Proof. Using the definition of duality gap and adding and subtracting $\langle p, Aq \rangle = \langle q, A^\top p \rangle$, we can write and expand

$$\begin{aligned} \text{DG}(w) &= \max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \min_{p' \in \Delta_m} \langle p', Aq \rangle \\ &= \big(\max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \langle q, A^\top p \rangle\big) + \big(\langle p, Aq \rangle - \min_{p' \in \Delta_m} \langle p', Aq \rangle\big) \, . \end{aligned}$$

Observe that both terms in the above expression are non-negative. Thus, using the assumption that $\mathrm{DG}(w) \leq \epsilon$, it follows by rearranging that

$$\max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \epsilon \leq \langle q, A^\top p \rangle \quad \text{and} \quad \langle p, Aq \rangle \leq \min_{p' \in \Delta_m} \langle p', Aq \rangle + \epsilon \ .$$

Thus $w = (p, q)$ satisfies the inequalities in (11) and is therefore an ϵ-approximate NE of A. ■

B.2.2 Relationships between Duality Gap, KL, and TV

In this section, we prove a series of inequalities relating the distances DG, TV, KL, and χ^2. These relationships were introduced in expression (2) from Section 2. Here, we give a proof of each component separately, most of which are standard, but included for completeness. Further below, we restate the full sequence of inequalities from (2) in Corollary B.6.

Relating Duality Gap and TV Distance.

Proposition B.3 (Duality-Gap vs. TV distance)**.** *Let $A \in \mathbb{R}^{m \times n}$ be a zero-sum game with Nash equilibrium $w^\star = (p^\star, q^\star) \in \mathcal{W}$. Let $a_{\max} = \max_{(i,j) \in [m] \times [n]} |A(i,j)|$ denote the maximum absolute entry of A. Then for any $w = (p, q) \in \mathcal{W}$, it holds that:*

$$\mathrm{DG}(w) \leq 2a_{\max} \mathrm{TV}(w^\star, w) \ .$$

Proof. Let $\widetilde{w} = (\widetilde{p}, \widetilde{q}) \in \mathcal{W}$ be a pair of distributions such that $\max_{q' \in \Delta_n} \langle q', A^\top p \rangle = \langle \widetilde{q}, A^\top p \rangle$ and $\min_{p' \in \Delta_m} \langle p', Aq \rangle = \langle \widetilde{p}, Aq \rangle$. Then by definition of a Nash equilibrium $w^\star = (p^\star, q^\star)$ from (11), it follows that

$$\langle \widetilde{q}, A^\top p^\star \rangle - \langle q^\star, A^\top p^\star \rangle \leq 0 \text{ and } \langle \widetilde{p}, Aq^\star \rangle - \langle p^\star, Aq^\star \rangle \geq 0 \ .$$

Let $c = \langle p^\star, Aq^\star \rangle$, and thus the above inequalities imply $\langle \widetilde{q}, A^\top p^\star \rangle - c \leq 0$ and $\langle \widetilde{p}, Aq^\star \rangle - c \geq 0$. Then it follows by definition of $\mathrm{DG}(w)$, that we can write and further simplify

$$\begin{aligned}
\mathrm{DG}(w) &= \max_{q' \in \Delta_n} \langle q', A^\top p \rangle - \min_{p' \in \Delta_m} \langle p', Aq \rangle \\
&= \langle \widetilde{q}, A^\top p \rangle - \langle \widetilde{p}, Aq \rangle \\
&\leq \langle \widetilde{q}, A^\top p \rangle - (\langle \widetilde{q}, A^\top p^\star \rangle - c) - \langle \widetilde{p}, Aq \rangle + (\langle \widetilde{p}, Aq^\star \rangle - c) \\
&= \langle \widetilde{q}, A^\top (p - p^\star) \rangle + \langle \widetilde{p}, A(q^\star - q) \rangle \\
&\leq \|A\widetilde{q}\|_\infty \cdot \|p - p^\star\|_1 + \|A^\top \widetilde{p}\|_\infty \cdot \|q - q^\star\|_1 \ , \qquad (13)
\end{aligned}$$

where the final inequality comes from Hölder's inequality. As $(\widetilde{p}, \widetilde{q}) \in \mathcal{W}$, it follows that both $\|A\widetilde{q}\|_\infty, \|A^\top \widetilde{p}\|_\infty \leq a_{\max}$. Moreover, by definition of $\mathrm{TV}(w^\star, w)$, we conclude from (13) that

$$\mathrm{DG}(w) \leq a_{\max}(\|p^\star - p\|_1 + \|q^\star - q\|_1) = 2a_{\max} \mathrm{TV}(w^\star, w) \ ,$$

which concludes the proof. ■

Relating TV Distance and KL Divergence.

Proposition B.4 (Pinsker-like Inequality for TV and KL)**.** *Fix any $w^\star = (p^\star, q^\star) \in \mathcal{W}$ and any $w = (p, q) \in \mathcal{W}$. Then it holds that* $\mathrm{TV}(w^\star, w)^2 \leq \mathrm{KL}(w^\star, w)$.

Proof. Recall by our definition of TV and KL that $\mathrm{TV}(w^\star, w) = \mathrm{TV}(p^\star, p) + \mathrm{TV}(q^\star, q)$ and $\mathrm{KL}(w^\star, w) = \mathrm{KL}(p^\star, p) + \mathrm{KL}(q^\star, q)$. By Pinsker's inequality, we have for each component that

$$2 \cdot \mathrm{TV}(p^\star, p)^2 \leq \mathrm{KL}(p^\star, p) \text{ and } 2 \cdot \mathrm{TV}(q^\star, q)^2 \leq \mathrm{KL}(q^\star, q) \ .$$

Thus together this implies

$$\mathrm{KL}(w^\star, w) \geq 2 \cdot \big(\mathrm{TV}(p^\star, p)^2 + \mathrm{TV}(q^\star, q)^2\big) \geq \big(\mathrm{TV}(p^\star, p) + \mathrm{TV}(q^\star, q)\big)^2 = \mathrm{TV}(w^\star, w)^2 \ .$$

Here, the second inequality comes from the fact that $\sqrt{(a^2 + b^2)/2} \geq (a + b)/2$ and thus also $2(a^2 + b^2) \geq (a + b)^2$ for all $a, b \in \mathbb{R}$. ■

Relating KL Divergence and Chi-Squared Divergence.

Proposition B.5 (KL Divergence vs. χ^2 Divergence)**.** *Fix $w^\star = (p^\star, q^\star) \in \mathcal{W}$ and $w = (p, q) \in \mathcal{W}$. Then it holds that* $\mathrm{KL}(w^\star, w) \leq \chi^2(w^\star, w)$.

Proof. We have by definition of $\mathrm{KL}(p^\star, p)$ that

$$\mathrm{KL}(p^\star, p) = \sum_{i=1}^m p^\star(i) \cdot \log\left(\tfrac{p^\star(i)}{p(i)}\right) \leq \sum_{i=1}^m p^\star(i) \cdot \left(\tfrac{p^\star(i)}{p(i)} - 1\right) = \sum_{i=1}^m \tfrac{p^\star(i)^2}{p(i)} - 1 \,,$$

where the inequality comes from the fact that $\log(u) \leq u - 1$ for all $u > 0$. On the other hand, we have by definition of $\chi^2(w^\star, w)$ that

$$\chi^2(p^\star, p) = \sum_{i=1}^m \tfrac{(p^\star(i) - p(i))^2}{p(i)} = \sum_{i=1}^m \left(\tfrac{p^\star(i)^2}{p(i)} + \tfrac{p(i)^2}{p(i)} - \tfrac{2p(i)p^\star(i)}{p(i)}\right) = \sum_{i=1}^m \tfrac{p^\star(i)^2}{p(i)} - 1 \,.$$

Combining the two expressions then yields $\mathrm{KL}(p^\star, p) \leq \chi^2(p^\star, p)$. An identical calculation also gives $\mathrm{KL}(q^\star, q) \leq \chi^2(q^\star, q)$, which then implies by definition that $\mathrm{KL}(w^\star, w) \leq \chi^2(w^\star, w)$. ■

Full Sequence of Relationships. The following corollary summarizes the relationships between the measures of distance to Nash from Propositions B.3, B.4, and B.5:

Corollary B.6 (Relationships of Distances to Nash)**.** *Let $A \in \mathbb{R}^{m \times n}$ be a zero-sum game with a Nash equilibrium $w^\star = (p^\star, q^\star) \in \mathcal{W}$. Let $a_{\max} = \max_{(i,j) \in [m] \times [n]} |A(i,j)|$ denote the maximum absolute entry of A. Then for any $w = (p, q) \in \mathcal{W}$, it holds that:*

$$\tfrac{1}{(2a_{\max})^2} \cdot \mathrm{DG}(w)^2 \leq \mathrm{TV}(w^\star, w)^2 \leq \mathrm{KL}(w^\star, w) \leq \chi^2(w^\star, w) \,.$$

Note that, due to Proposition B.2, Corollary B.6 implies that upper bounds on TV, KL, and χ^2 each also imply an approximate Nash in the sense of (11).

B.3 Properties of Spectrum of Payoff Matrix

Here, we state and prove two useful properties related to the spectrum of the payoff matrix A.

Proposition B.7 (Relationship between max entry and max singular value)**.** *Fix $A \in \mathbb{R}^{m \times n}$. Let $a_{\max} = \max_{(i,j) \in [m] \times [n]} |A(i,j)|$ and let $\sigma_{\max} = \|A\|_2$. Then:*

$$a_{\max} \leq \sigma_{\max} \leq \sqrt{mn} \cdot a_{\max} \,.$$

Proof. For the first inequality, let $e_i \in \mathbb{R}^m$, $e_j \in \mathbb{R}^n$ for $i \in [m]$, $j \in [n]$ denote the standard m and n-dimensional basis vectors. Then by Cauchy-Schwarz and the definition of $\sigma_{\max}$, we have

$$\max_{i,j} |A(i,j)| = \max_{i,j} |\langle e_i, Ae_j \rangle| \leq \max_{i,j} \|e_i\|_2 \cdot \|Ae_j\|_2 \leq \sigma_{\max} \,.$$

For the second inequality, recall that the Frobenius norm of A is given by $\|A\|_F = (\sum_{(i,j) \in [m] \times [n]} A(i,j)^2)^{1/2}$, and that $\|A\|_2 \leq \|A\|_F$. Then we can further bound

$$\|A\|_2^2 \leq \|A\|_F^2 = \sum_{(i,j) \in [m] \times [n]} A(i,j)^2 \leq mn \cdot a_{\max}^2 \,.$$

Taking square roots yields the desired inequality. ■

In the next proposition, we show that $\|J\|_2 = \|A\|_2$, where J is the skew-symmetric matrix from (3).

Proposition B.8 (Spectral norms of J and A)**.** *Fix $A \in \mathbb{R}^{m \times n}$, and let $J = ((0, A), (-A^\top, 0)) \in \mathbb{R}^{(m+n) \times (m+n)}$ be the block skew-symmetric matrix from (3). Then $\|J\|_2 = \|A\|_2$.*

Proof. For a symmetric matrix $M \in \mathbb{R}^{k \times k}$, let $\lambda_{\max}(M)$ denote the largest eigenvalue of M. Then recall by definition of the matrix spectral norm that, for a matrix P, we have $\sigma_{\max}(P) = \|P\|_2 = \sqrt{\lambda_{\max}(P^\top P)}$. For the skew-symmetric J, we can compute

$$J^\top J = \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix} \begin{pmatrix} 0 & -A \\ A^\top & 0 \end{pmatrix} = \begin{pmatrix} AA^\top & 0 \\ 0 & A^\top A \end{pmatrix} \,.$$

Given the block structure of $J^\top J$, it follows that $\lambda_{\max}(J^\top J) = \max(\lambda_{\max}(AA^\top), \lambda_{\max}(A^\top A))$. However, $\lambda_{\max}(AA^\top) = \lambda_{\max}(A^\top A)$, and thus

$$\|J\|_2 = \sqrt{\lambda_{\max}(J^\top J)} = \sqrt{\lambda_{\max}(A^\top A)} = \|A\|_2 \,,$$

which yields the desired claim. ■

C Details on Optimistic MWU Algorithm

In this section, we review preliminaries on the OMWU algorithm that were introduced in Section 3.

Organization of Section. This section is organized as follows:

- **Section C.1** reviews the derivation of the primal and dual OMWU iterates (and also reviews the standard MWU algorithm). In particular, Section C.1.2 gives details on OMWU as the instantiation of Optimistic FTRL using the negative entropy regularizer.
- **Section C.2** gives the proof of Proposition 3.1, which establishes the perspective of the dual OMWU iterates as an optimistic skew-gradient descent on the log-sum-exp energy function.

C.1 Primal and Dual OMWU Iterates

C.1.1 Primal Update Rule

Standard MWU. We first recall the primal update rule of the standard Multiplicative Weights Update algorithm (MWU) (see, e.g., Freund and Schapire (1999), Arora et al. (2012)), which is the basis of the optimistic variant. The MWU iterates are initialized from $w_0 = (p_0, q_0) \in \mathrm{ri}(\mathcal{W})$. Then using a fixed stepsize $\eta > 0$, for $t \geq 1$, the MWU iterates $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$ are given by

$$\begin{aligned} & p_t(i) \propto p_{t-1}(i) \cdot \exp\big(-\eta(Aq_{t-1})(i)\big) \quad \text{for all } i \in [m] \\ \text{and} \quad & q_t(j) \propto q_{t-1}(j) \cdot \exp\big(\eta(A^\top p_{t-1})(j)\big) \quad \text{for all } j \in [n]\,. \end{aligned} \tag{MWU}$$

Optimistic MWU. The Optimistic MWU update rule is similar to the standard (MWU) update, but with an additional recency bias. This leads to increases in the weights of coordinates corresponding to smaller losses in the most recent gradient vector. Specifically, recall from (OMWU) that the primal iterates $\{w_t\}$ are initialized from $w_0 = (p_0, q_0) \in \mathrm{ri}(\mathcal{W})$. Then using a fixed stepsize $\eta > 0$, for $t \geq 0$, the iterates $w_{t+1} = (p_{t+1}, q_{t+1})$ in $\mathrm{ri}(\mathcal{W})$ are given by

$$\begin{aligned} & p_{t+1}(i) \propto p_t(i) \cdot \exp\big(-\eta(2Aq_t - Aq_{t-1})(i)\big) \qquad \text{for all } i \in [m] \\ \text{and} \quad & q_{t+1}(j) \propto q_t(j) \cdot \exp\big(\eta(2A^\top p_t - A^\top p_{t-1})(j)\big) \quad \text{for all } j \in [n]\,. \end{aligned} \tag{14}$$

We assume for notational convenience that $w_{-1} = (p_{-1}, q_{-1}) = w_0 \in \mathrm{ri}(\mathcal{W})$. This means at time $t = 1$ that $p_1(i) \propto p_0(i) \cdot \exp(-\eta(Aq_0)(i))$ for $i \in [m]$ and $q_1(j) \propto q_0(j) \cdot \exp(\eta(A^\top p_0)(j))$ for $j \in [n]$.

C.1.2 Dual Iterates and Optimistic FTRL Instantiation

Component-Wise Dual Iterates. We start by introducing the player-wise components of the dual iterates $\{z_t\}$ from (4). For this, let $x_0 = x_{-1} \in \mathbb{R}^m$ and $y_0 = y_{-1} \in \mathbb{R}^n$ be initial dual vectors satisfying $p_0 = \mathrm{softmax}_m(x_0)$ and $q_0 = \mathrm{softmax}_n(y_0)$. Let $\{w_t\}$ denote the sequence of primal (OMWU) iterates, where each $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$. Then for stepsize $\eta > 0$, the dual iterates $x_t \in \mathbb{R}^m$, $y_t \in \mathbb{R}^n$ are defined for all $t \geq 1$ as follows: first, at time $t = 1$, let $x_1 = x_0 - \eta A q_0$ and $y_1 = y_0 + \eta A^\top p_0$. Then for all $t \geq 2$, we define

$$\begin{cases} x_t = x_0 - \eta \cdot \big(\sum_{k=1}^{t-1} Aq_k + Aq_{t-1}\big) \\ y_t = y_0 + \eta \cdot \big(\sum_{k=1}^{t-1} A^\top p_k + A^\top p_{t-1}\big)\,. \end{cases} \tag{15}$$

Optimistic FTRL Using Negative-Entropy Regularizer. As introduced in Section 3, the primal OMWU iterates can also be derived as an instantiation of Optimistic FTRL (OFTRL) using negative entropy regularization (Rakhlin and Sridharan, 2013; Syrgkanis et al., 2015). For this, recall that we define the strictly convex regularizers $R_m : \Delta_m \to \mathbb{R}$ and $R_n : \Delta_n \to \mathbb{R}$ as the discrete negative entropy functions:

$$\begin{aligned} R_m(p) &:= -\mathrm{Ent}_m(p) = \textstyle\sum_{i=1}^m p(i)\log(p(i)) \\ \text{and} \quad R_n(q) &:= -\mathrm{Ent}_n(q) = \textstyle\sum_{j=1}^n q(j)\log(q(j))\,. \end{aligned} \tag{16}$$

Here and throughout, we use the standard convention that $0 \log 0 = 0$.

Under Optimistic FTRL, the primal iterates $p_t \in \Delta_m$, $q_t \in \Delta_n$ at time $t \geq 1$ are chosen via a regularized best-response map applied to the dual iterates $x_t \in \mathbb{R}^m$ and $y_t \in \mathbb{R}^n$, respectively.

Specifically, let the OFTRL primal iterates be initialized at $p_0 \in \text{ri}(\Delta_m)$ and $q_0 \in \text{ri}(\Delta_n)$, with the compatible dual initialization specified above. At time $t+1 \geq 1$, suppose $x_{t+1} \in \mathbb{R}^m$ and $y_{t+1} \in \mathbb{R}^n$ are the dual iterates from (15) defined in terms of $\{q_0, \ldots, q_t\}$ and $\{p_0, \ldots, p_t\}$, respectively. Then at time $t+1$, the primal iterates of OFTRL instantiated with the regularizers R_m and R_n are given by the objectives

$$\begin{aligned} p_{t+1} &:= \text{argmin}_{p \in \Delta_m} \left\{ -\langle p, x_{t+1} \rangle + R_m(p) \right\} \\ q_{t+1} &:= \text{argmin}_{q \in \Delta_n} \left\{ -\langle q, y_{t+1} \rangle + R_n(q) \right\} . \end{aligned} \tag{17}$$

Here, observe that the accumulated payoff terms in (x_{t+1}, y_{t+1}) from (15) are already scaled by the stepsize $\eta > 0$. Thus, smaller values of η correspond to more weight on the regularization terms in (17), and vice-versa. Moreover, under the negative entropy regularizers, the variables (p_{t+1}, q_{t+1}) have the following closed-form solution:

Proposition C.1 (Softmax update rule)**.** *For $t \geq 1$, let $x_t \in \mathbb{R}^m$ and $y_t \in \mathbb{R}^n$, and let $p_t \in \Delta_m$ and $q_t \in \Delta_n$ be defined as in (17). Then $p_t = \text{softmax}_m(x_t)$ and $q_t = \text{softmax}_n(y_t)$.*

This result is standard and is based on the first-order optimality conditions of the objective in (17), and thus we omit the proof; See, e.g., Shalev-Shwartz et al. (2012), Section 2.6.

Now, as previously mentioned, the update rules of (17) correspond exactly to the (OMWU) primal iterates. Due to Proposition C.1, this will follow as an immediate corollary of Proposition 3.1 from Section 3. We restate and give the proof of this latter result in the next subsection, which requires using the *concatenated* primal and dual variables $w_t = (p_t, q_t) \in \text{ri}(\mathcal{W})$ and $z_t = (x_t, y_t) \in \mathbb{R}^{m+n}$. This notation was introduced in expressions (3), (4), and (5) from Section 3, and we briefly review these preliminaries again here:

Concatenated Primal and Dual Iterates. Given $A \in \mathbb{R}^{m \times n}$, recall from expression (3) that $J = -J^\top$ is the block skew-symmetric matrix $J = ((0, A), (-A^\top, 0)) \in \mathbb{R}^{(m+n)\times(m+n)}$. Then observe for any $w = (p, q) \in \mathcal{W}$ that

$$Jw = \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} = \begin{pmatrix} Aq \\ -A^\top p \end{pmatrix} \in \mathbb{R}^{m+n} .$$

Thus, given the sequence of primal iterates $\{w_k\}_{k=1}^t$, where each $w_k = (p_k, q_k) \in \mathcal{W}$, it follows for all $t+1 \geq 2$ that the dual iterates $x_{t+1} \in \mathbb{R}^m$ and $y_{t+1} \in \mathbb{R}^n$ from (15) can be written jointly as

$$z_{t+1} = (x_{t+1}, y_{t+1}) = z_0 - \eta \cdot \left(\textstyle\sum_{k=1}^t Jw_k + Jw_t \right) \in \mathbb{R}^{m+n} ,$$

where $z_0 = z_{-1} = (x_0, y_0) \in \mathbb{R}^{m+n}$ is the dual initialization. It follows recursively that we can write:

$$\begin{cases} z_1 = z_0 - \eta J w_0 & \text{at } t = 1 \\ z_{t+1} = z_t - 2\eta J w_t + \eta J w_{t-1} & \text{at } t+1 \geq 2 . \end{cases} \tag{18}$$

Further recall from Section 3 that we define the joint negative entropy regularizer $R : \mathcal{W} \to \mathbb{R}$ to be $R = R_m + R_n = -(\text{Ent}_m + \text{Ent}_n)$, where for $w = (p, q) \in \mathcal{W}$, we have

$$R(w) = R_m(p) + R_n(q) = -(\text{Ent}_m(p) + \text{Ent}_n(q)) . \tag{19}$$

The separability of the regularizer R allows us to write the OFTRL objective functions from (17) using the concatenated dual variable $z_{t+1} = (x_{t+1}, y_{t+1}) \in \mathbb{R}^{m+n}$. Specifically, it holds for $t+1 \geq 1$ that $w_{t+1} = (p_{t+1}, q_{t+1}) \in \mathcal{W}$ is given by

$$w_{t+1} = \underset{w=(p,q)\in\mathcal{W}}{\text{argmin}} \left\{ -\langle w, z_{t+1} \rangle + R(w) \right\} = \begin{pmatrix} \text{argmin}_{p \in \Delta_m} \left\{ -\langle p, x_{t+1} \rangle + R_m(p) \right\} \\ \text{argmin}_{q \in \Delta_n} \left\{ -\langle q, y_{t+1} \rangle + R_n(q) \right\} \end{pmatrix} \in \mathcal{W} . \tag{20}$$

C.1.3 Energy Function and Gradient Map

Log-Sum-Exp Energy Function. We recall the definition of the energy function $F : \mathbb{R}^{m+n} \to \mathbb{R}$ from (6) in Section 3. First, we have that $\text{LSE}_m : \mathbb{R}^m \to \mathbb{R}$ and $\text{LSE}_n : \mathbb{R}^n \to \mathbb{R}$ are the m and n-dimensional log-sum-exp functions, where:

$$\text{LSE}_m(x) = \log \left(\textstyle\sum_{i=1}^m \exp(x(i)) \right) \quad \text{and} \quad \text{LSE}_n(y) = \log \left(\textstyle\sum_{j=1}^n \exp(y(j)) \right) .$$

Both LSE_m and LSE_n are convex and continuously differentiable. Together, they define the convex and separable energy function $F = \text{LSE}_m + \text{LSE}_n$, where for $z = (x, y) \in \mathbb{R}^{m+n}$ we have

$$F(z) = \text{LSE}_m(x) + \text{LSE}_n(y) .$$

Energy Gradients. By a straightforward differentiation, note that $\nabla \mathrm{LSE}_k(v) = \mathrm{softmax}_k(v)$ for any $k \geq 2$ and $v \in \mathbb{R}^k$. Thus, the gradient map $\nabla F : \mathbb{R}^{m+n} \to \mathbb{R}^{m+n}$ is given by

$$\nabla F(z) = \begin{pmatrix} \nabla_x F(z) \\ \nabla_y F(z) \end{pmatrix} = \begin{pmatrix} \mathrm{softmax}_m(x) \\ \mathrm{softmax}_n(y) \end{pmatrix} \in \mathrm{ri}(\mathcal{W}) \quad \text{for all } z = (x, y) \in \mathbb{R}^{m+n}. \tag{21}$$

Note also that the log-sum-exp energy function F and the negative entropy regularizer R are dual functions (convex conjugates). We give more details and implications of this duality in Section D.1.

C.2 OMWU as Optimistic Skew-Gradient Descent

The joint OMWU dual iterates $\{z_t\}$ can be written as an optimistic skew-gradient descent on the energy function F. This key property was stated in Proposition 3.1 in Section 3. Here, we restate the proposition and give the proof:

Proposition 3.1 (Optimistic Skew-Gradient Descent)**.** *Let $\{w_t\}$ be the iterates of (OMWU) with stepsize $\eta > 0$ initialized from $w_0 \in \mathrm{ri}(\mathcal{W})$. Let $\{z_t\}$ be the dual iterates of (4) initialized from $z_0 \in \mathbb{R}^{m+n}$ such that $\nabla F(z_0) = w_0$. Then for all $t \geq 0$, it holds that $w_t = \nabla F(z_t)$ and*

$$z_{t+1} = \underbrace{z_t - \eta J \nabla F(z_t)}_{\textit{skew-gradient descent}} \underbrace{- \eta (J \nabla F(z_t) - J \nabla F(z_{t-1}))}_{\textit{optimistic skew-gradient correction}} . \tag{OMWU Dual}$$

C.2.1 Proof of Proposition 3.1

For $t \geq 0$, let $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$ be the primal iterates from (OMWU) (equivalently, from (14)), and let $z_t = (x_t, y_t) \in \mathbb{R}^{m+n}$ be the dual iterates from (4) (equivalently, from (18)). Moreover, recall from (21) that $\nabla_x F(z) = \mathrm{softmax}_m(x)$ and $\nabla_y F(z) = \mathrm{softmax}_n(y)$ for any $z = (x, y) \in \mathbb{R}^{m+n}$. Thus, we first prove inductively that $w_t = (p_t, q_t) = (\mathrm{softmax}_m(x_t), \mathrm{softmax}_n(y_t))$ for all $t \geq 0$.

For the base case, we have by assumption of the proposition that $z_0 = (x_0, y_0) \in \mathbb{R}^{m+n}$ is such that $w_0 = (p_0, q_0) = \nabla F(z_0) = (\mathrm{softmax}_m(x_0), \mathrm{softmax}_n(y_0))$. Now suppose through time t that $w_t = (\mathrm{softmax}_m(x_t), \mathrm{softmax}_n(y_t))$. Then at time $t+1$, we have by definition of (OMWU) that for all coordinates $i \in [m]$:

$$\begin{aligned}
p_{t+1}(i) &= \frac{p_t(i) \cdot \exp(-\eta(2Aq_t - Aq_{t-1})(i))}{\sum_{k=1}^m p_t(k) \cdot \exp(-\eta(2Aq_t - Aq_{t-1})(k))} \\
&= \frac{\frac{\exp(x_t(i))}{\sum_{\ell=1}^m \exp(x_t(\ell))} \cdot \exp(-\eta(2Aq_t - Aq_{t-1})(i))}{\sum_{k=1}^m \frac{\exp(x_t(k))}{\sum_{\ell=1}^m \exp(x_t(\ell))} \cdot \exp(-\eta(2Aq_t - Aq_{t-1})(k))} \\
&= \frac{\exp(x_t(i) - \eta(2Aq_t - Aq_{t-1})(i))}{\sum_{k=1}^m \exp(x_t(k) - \eta(2Aq_t - Aq_{t-1})(k))} \\
&= \mathrm{softmax}_m\big(x_t - \eta(2Aq_t - Aq_{t-1})\big)(i) \, .
\end{aligned} \tag{22}$$

Here, the second equality follows by applying the inductive hypothesis at time t. By definition of x_{t+1} from (15), we have $x_{t+1} = x_t - \eta(2Aq_t - Aq_{t-1})$ for all $t \geq 0$, using $q_{-1} = q_0$ at $t = 0$. As (22) holds for all coordinates $i \in [m]$, we conclude that $p_{t+1} = \mathrm{softmax}_m(x_{t+1})$. By an identical calculation, we similarly find $q_{t+1} = \mathrm{softmax}_n(y_{t+1})$. Thus it holds inductively that $w_t = (\mathrm{softmax}_m(x_t), \mathrm{softmax}_n(y_t)) = \nabla F(z_t)$ for all $t \geq 0$.

For the second claim, we have by definition of $\{z_t\}$ from (18) that $z_1 = z_0 - \eta J \nabla F(z_0)$ at time $t = 1$, and at times $t + 1 \geq 2$:

$$\begin{aligned}
z_{t+1} &= z_t - 2\eta J \nabla F(z_t) + \eta J \nabla F(z_{t-1}) \\
&= z_t - \eta J \nabla F(z_t) - \eta J (\nabla F(z_t) - \nabla F(z_{t-1})) \, .
\end{aligned}$$

This concludes the proof. ■

Equivalence of (OMWU) and the OFTRL Instantiation. As a corollary of Proposition 3.1, we also have that the OFTRL iterates from (5) (equivalently (20)) are also identical to those of (OMWU). This follows from the fact that $\nabla F(z_t) = (\mathrm{softmax}_m(x_t), \mathrm{softmax}_n(y_t))$ for all $t \geq 1$ (from the first part of Proposition 3.1), and from the characterization of the iterates of (5) from Proposition C.1.

C.2.2 Skew-Gradient Flow and Standard Skew-Gradient Descent

Here, we briefly review the continuous-time *skew-gradient flow* and the standard forward discretization *skew-gradient descent*.

Skew-gradient flow. The skew-gradient flow dynamics is the continuous-time ODE:

$$\dot{z}(t) = -J\nabla F(z(t)) \ . \qquad \text{(Skew-gradient flow)}$$

By the chain rule, observe that energy is *conserved* under (Skew-gradient flow), as we can compute

$$\frac{d}{dt}F(z(t)) = \langle \nabla F(z(t)), \dot{z}(t)\rangle = -\langle \nabla F(z(t)), J\nabla F(z(t))\rangle = 0 \ ,$$

where the final equality is due to the skew-symmetry of $J = -J^\top$. Geometrically, this means that the skew-gradient vector $J\nabla F(z)$ is *orthogonal* to the gradient $\nabla F(z)$. Hence, the skew-gradient flow follows directions that are tangent to the energy levelset at the current point, thus conserving energy.

Skew-gradient descent. The first-order forward (Euler) discretization of (Skew-gradient flow) is the forward skew-gradient descent. This yields the dual iterates corresponding to the standard MWU algorithm from (MWU) with stepsize $\eta > 0$:

$$z_{t+1} = z_t - \eta J\nabla F(z_t) \ . \qquad \text{(MWU Dual)}$$

By convexity of F, the one-step change in energy under (MWU Dual) is always nonnegative, since

$$F(z_{t+1}) - F(z_t) \geq \langle \nabla F(z_t), z_{t+1} - z_t\rangle = -\eta\langle \nabla F(z_t), J\nabla F(z_t)\rangle = 0 \ ,$$

where again the final equality is due to $J = -J^\top$.

In contrast to the forward discretization, we show for (OMWU Dual) that the additional optimistic correction term (which can be viewed as an approximation of a backward discretization of skew-gradient flow), leads to strictly dissipating energy (see Lemma 4.1).

See also Wibisono et al. (2022) and Katona et al. (2026) for additional background on skew-gradient flows and their discretizations.

D Details on Equivalence of KL Minimization and Energy Dissipation

The energy function F and the regularizer R are closely related via duality. Here, we specify several key aspects of this relationship in more detail. For the OMWU iterates, this ultimately allows for establishing an equivalence between minimizing KL divergence in the primal space, and energy dissipation in the effective dual space.

Organization of Section. This section is organized as follows:

- **Section D.1** gives more details on the dual relationship of the log-sum-exp energy function F and the negative-entropy regularizer R, and their interplay over the effective dual space $\mathcal{Z}$. This includes establishing a certain modified Fenchel-Young identity (Proposition D.2) and proving that $\mathcal{Z}$ is orthogonal to the interior NE of A (Proposition D.6).
- **Section D.2** gives the proof of Proposition 3.2, which establishes an equivalence between differences in KL divergence and differences in energy F over the effective dual space $\mathcal{Z}$. The proof of the proposition relies on the properties introduced in Section D.1.
- **Section D.3** establishes additional properties of the energy gradient map ∇F over the effective dual space $\mathcal{Z}$. This culminates in establishing, under the assumption of a unique and interior NE, that the KL divergence from Nash is exactly the energy suboptimality gap in the effective dual space (Proposition D.11).

D.1 Duality of Regularizer and Energy Functions

The negative entropy regularizer R is the *convex (Fenchel) conjugate* F^* of the log-sum-exp energy function F. Here, recall by definition that $F^* : \mathbb{R}^{m+n} \to \mathbb{R} \cup \{+\infty\}$ is the convex function given by

$$F^*(w) \ := \ \sup_{z\in\mathbb{R}^{m+n}} \ \big\{\langle w, z\rangle - F(z)\big\} \ . \qquad (23)$$

for all $w \in \mathbb{R}^{m+n}$. As F and R are both separable, the fact that $F^* = R$ follows from standard relationships between the log-sum-exp and negative entropy functions. Moreover, it is also well

known that the Bregman divergence $D_R(w', w)$ induced by the negative entropy regularizer R is exactly the KL divergence $\mathrm{KL}(w', w)$. These relationships are stated in the following proposition:

Proposition D.1. *Let $F\colon \mathbb{R}^{m+n} \to \mathbb{R}$ and $R\colon \mathcal{W} \to \mathbb{R}$ be the energy and regularizer functions defined in (6) and (19), respectively. Let F^* be the convex conjugate of F as in (23). Then the following two properties hold:*

(i) $F^* = R$, meaning $F^*(w) = R(w)$ for $w \in \mathcal{W}$, and $F^*(w) = \infty$ otherwise.

(ii) For $w, w' \in \mathrm{ri}(\mathcal{W})$: $\mathrm{KL}(w', w) = D_R(w', w) = R(w') - R(w) - \langle \nabla R(w), w' - w \rangle$.

The proof is standard and thus we omit it here (see, e.g., Boyd and Vandenberghe (2004), Examples 3.19 and 3.25). Moreover, by standard properties of convex conjugates, note also that the energy $F = R^*$ is the conjugate function of R.

D.1.1 Modified Fenchel-Young Identity

Due to the dual relationship of F and R, we additionally establish the following useful properties that relate KL divergence to energy. Specifically, we prove the following proposition:

Proposition D.2 (Modified Fenchel-Young Identity)**.** *Let $F\colon \mathbb{R}^{m+n} \to \mathbb{R}$ and $R\colon \mathcal{W} \to \mathbb{R}$ be the energy and regularizer functions from (6) and (19), respectively. Fix $z \in \mathbb{R}^{m+n}$, and let $w = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Then the following properties hold:*

(i) $R(w) + F(z) = \langle w, z \rangle$.

(ii) $\mathrm{KL}(w', w) = D_R(w', w) = R(w') + F(z) - \langle z, w' \rangle$ for any $w' \in \mathrm{ri}(\mathcal{W})$.

The first statement is due to the standard Fenchel-Young inequality (see, e.g., Shalev-Shwartz et al. (2012), Section 2.7). The second statement requires first introducing several additional and important properties of the energy function, the negative entropy regularizer, and their gradients. We first introduce these components and then proceed with the proof of Proposition D.2 further below.

Energy function lacks global strict convexity. Recall from expression (7) in Section 3 that the linear subspace $\mathcal{S}$ of constant shift directions is defined as

$$\mathcal{S} = \mathrm{Span}\left(\begin{pmatrix} \mathbf{1}_m \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ \mathbf{1}_n \end{pmatrix} \right) .$$

While the regularizer R is strictly convex over the primal space $\mathcal{W} = \Delta_m \times \Delta_n$, the energy function F is *not* globally strictly convex over $\mathbb{R}^{m+n}$. In particular, the gradients of F are invariant to constant shifts (i.e., in directions $s \in \mathcal{S}$), meaning that F is affine in those directions. Specifically, the following properties hold:

Proposition D.3 (Energy function is not strictly convex)**.** *Let $F : \mathbb{R}^{m+n} \to \mathbb{R}$ be the energy function from (6). Fix any $z = (x, y) \in \mathbb{R}^{m+n}$. Then for any $s = (d\mathbf{1}_m, d'\mathbf{1}_n) \in \mathcal{S}$ for some $d, d' \in \mathbb{R}$:*

$$\nabla F(z + s) = \nabla F(z) \quad \textit{and} \quad F(z + s) = F(z) + (d + d') .$$

Proof. Recall from (21) that $\nabla F(z) = (\mathrm{softmax}_m(x), \mathrm{softmax}_n(y))$. Then using the definition of $\mathrm{softmax}_k(\cdot)$, we have

$$\begin{aligned} \nabla F(z + s) &= (\mathrm{softmax}_m(x + d\mathbf{1}_m), \mathrm{softmax}_n(y + d'\mathbf{1}_n)) \\ &= (\mathrm{softmax}_m(x), \mathrm{softmax}_n(y)) = \nabla F(z) . \end{aligned}$$

Similarly, using the definition of $F = \mathrm{LSE}_m + \mathrm{LSE}_n$, we have

$$F(z + s) = \mathrm{LSE}_m(x + d\mathbf{1}_m) + \mathrm{LSE}_n(y + d'\mathbf{1}_n) = F(z) + (d + d') ,$$

which follows from the definition of log-sum-exp. ■

Gradient of negative entropy regularizer. Let $\nabla R\colon \mathrm{ri}(\mathcal{W}) \to \mathbb{R}^{m+n}$ denote the gradient map $\nabla R(w) = (\nabla_p R(w), \nabla_q R(w))$ for $w = (p, q) \in \mathrm{ri}(\mathcal{W})$. By definition of R and the negative entropy functions from (19), it follows that

$$\nabla R(w) = \begin{pmatrix} \log(p) + \mathbf{1}_m \\ \log(q) + \mathbf{1}_n \end{pmatrix} , \tag{24}$$

where $\log(p) = (\log(p(1)), \ldots, \log(p(m))) \in \mathbb{R}^m$ and $\log(q) = (\log(q(1)), \ldots, \log(q(n))) \in \mathbb{R}^n$.

Although $R = F^*$ is the convex conjugate of F, the energy function's lack of global strict convexity shown in Proposition D.3 means that F is not a *Legendre function* (see, e.g., Cesa-Bianchi and Lugosi (2006), Section 11.2 for a definition). The main consequence of this fact is that the regularizer gradient $\nabla R = \nabla F^*$ *is not* the inverse map of ∇F (since, by Proposition D.3, $\nabla F \colon \mathbb{R}^{m+n} \to \text{ri}(\mathcal{W})$ is not injective). However, the map $\nabla R \colon \text{ri}(\mathcal{W}) \to \mathbb{R}^{m+n}$ does serve as an inverse to ∇F over the quotient space $\mathbb{R}^{m+n}/\mathcal{S}$, as formalized in the following proposition:

Proposition D.4 (Gradient map of regularizer)**.** *Let $R \colon \mathcal{W} \to \mathbb{R}$ be the joint regularizer from (19). Fix $z = (x, y) \in \mathbb{R}^{m+n}$, and let $w = (p, q) = \nabla F(z)$. Then there exists $s \in \mathcal{S}$ such that*

$$\nabla R(w) = z + s \ .$$

Specifically, $s = (c_x \mathbf{1}_m, c_y \mathbf{1}_n)$, where $c_x = 1 - \text{LSE}_m(x) \in \mathbb{R}$ and $c_y = 1 - \text{LSE}_n(y) \in \mathbb{R}$.

Proof. By definition of F and the assumption that $w = (p, q) = \nabla F(z)$, we have $p = \text{softmax}_m(x)$ and $q = \text{softmax}_n(y)$. Using the definition of softmax_m and softmax_n, it follows that

$$\log(p) + \mathbf{1}_m = \begin{pmatrix} \log\left(\frac{\exp(x(1))}{\sum_{j=1}^m \exp(x(j))}\right) + 1 \\ \vdots \\ \log\left(\frac{\exp(x(m))}{\sum_{j=1}^m \exp(x(j))}\right) + 1 \end{pmatrix} = x + (1 - \text{LSE}_m(x)) \cdot \mathbf{1}_m \ . \tag{25}$$

By a similar calculation, we also have $\log(q) + \mathbf{1}_n = y + (1 - \text{LSE}_n(y)) \cdot \mathbf{1}_n$. Substituting these simplifications into the definition of $\nabla R(w)$ from (24) yields the stated claim. ■

In contrast to Proposition D.4, due to the strict convexity of the joint negative entropy regularizer R, the energy gradient ∇F *does* serve as an inverse to ∇R. We state and prove this in the following proposition, which is used in later subsections.

Proposition D.5 (Energy gradient is inverse of entropy gradient)**.** *Let F and R be the functions from (6) and (19), respectively. Fix any $w = (p, q) \in \text{ri}(\mathcal{W})$. Then $\nabla F(\nabla R(w)) = w$.*

Proof. Let $v = (v_x, v_y) = \nabla R(w) \in \mathbb{R}^{m+n}$. By definition of ∇R from (24), this means that $v_x = (\log(p) + \mathbf{1}_m)$ and $v_y = (\log(q) + \mathbf{1}_n)$. By a straightforward calculation, observe that

$$\nabla \text{LSE}_m(v_x) = \text{softmax}_m(v_x) = \text{softmax}_m(\log(p) + \mathbf{1}_m) = p$$
$$\text{and } \nabla \text{LSE}_n(v_y) = \text{softmax}_n(v_y) = \text{softmax}_n(\log(q) + \mathbf{1}_n) = q \ .$$

As $\nabla F(v) = (\text{softmax}_m(v_x), \text{softmax}_n(v_y))$, we conclude that $\nabla F(\nabla R(w)) = \nabla F(v) = w$. ■

We are now equipped to give the proof of the "modified" Fenchel-Young identity of Proposition D.2:

Proof of Proposition D.2. For part (i), note that as F is convex and differentiable, the standard Fenchel-Young inequality (see, e.g., Boyd and Vandenberghe (2004), Section 3.3.2) says that, for any $z \in \mathbb{R}^{m+n}$ and $w' \in \mathcal{W}$:

$$F(z) + F^*(w') \geq \langle w', z \rangle \ ,$$

where equality holds if and only if $w' = \nabla F(z)$. Thus part (i) follows from the assumption that $w = \nabla F(z)$ and that $R = F^*$ (as established in the first claim of Proposition D.1).

For part (ii), recall from the second claim of Proposition D.1 that $\text{KL}(w', w) = D_R(w', w)$. Using the definition of $D_R(w', w)$, we can then write

$$\begin{aligned} \text{KL}(w', w) = D_R(w', w) &= R(w') - R(w) - \langle \nabla R(w), w' - w \rangle \\ &= R(w') + F(z) - \langle w, z \rangle - \langle \nabla R(w), w' - w \rangle \ , \end{aligned} \tag{26}$$

where the second equality follows from rewriting $R(w)$ using part (i) of the proposition. Moreover, by the assumption that $w = \nabla F(z)$, it follows from Proposition D.4 that $\nabla R(w) = z + s$ for some $s \in \mathcal{S}$. Thus we can further simplify the final term of (26) and write

$$\langle \nabla R(w), w' - w \rangle = \langle z + s, w' - w \rangle = \langle z, w' - w \rangle \ . \tag{27}$$

Here, observe that the term $\langle s, w' - w \rangle = 0$ by definition of $s \in \mathcal{S}$ and $w, w' \in \mathcal{W} = \Delta_m \times \Delta_n$. Then substituting (27) into (26), we find

$$\text{KL}(w', w) = R(w') + F(z) - \langle w, z \rangle - \langle z, w' - w \rangle = R(w') + F(z) - \langle z, w' \rangle \ ,$$

which concludes the proof. ■

D.1.2 Effective Dual Space

Recall from Section 3.2 that we define the linear subspace $\mathcal{Z} \subseteq \mathbb{R}^{m+n}$ by

$$\mathcal{Z} := \text{Span}(J\mathcal{W}) = \left\{ \sum_{i=1}^{k} \tau_i \cdot Jw_i \text{ for } k \in \mathbb{N}, \tau_i \in \mathbb{R}, w_i \in \mathcal{W} \right\}. \tag{28}$$

We refer to $\mathcal{Z}$ as the *effective dual space*, as by (4), all dual OMWU iterates $\{z_t\}$ lie in $\mathcal{Z}$ when $z_0 \in \mathcal{Z}$.

Effective Dual Space is Orthogonal to Interior NE. The defining structural property of $\mathcal{Z}$ is that, when A has an interior Nash equilibrium $w^\star$, every vector $z \in \mathcal{Z}$ is orthogonal to $w^\star$. Formally:

Proposition D.6 (Effective dual space is orthogonal to interior Nash)**.** *Let $A \in \mathbb{R}^{m \times n}$ have an interior Nash equilibrium $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$. Then $\langle z, w^\star \rangle = 0$ for every $z \in \mathcal{Z}$.*

Proof. First, let $c = \langle p^\star, Aq^\star \rangle$. By the definition of J from (3) and using part (i) of Proposition B.1, observe that

$$Jw^\star = \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix} \begin{pmatrix} p^\star \\ q^\star \end{pmatrix} = \begin{pmatrix} Aq^\star \\ -A^\top p^\star \end{pmatrix} = \begin{pmatrix} c\mathbf{1}_m \\ -c\mathbf{1}_n \end{pmatrix} . \tag{29}$$

Now fix $z \in \mathcal{Z}$, and observe by definition that we can write $z = J\big(\sum_{i=1}^{k} \tau_i w_i\big)$ for some $k \geq 1$, where all $\tau_i \in \mathbb{R}$ and $w_i = (p_i, q_i) \in \mathcal{W} = \Delta_m \times \Delta_n$. Together with the skew-symmetry of $J = -J^\top$, it then follows that

$$\langle z, w^\star \rangle = \left\langle J\big(\textstyle\sum_{i=1}^{k} \tau_i w_i\big), w^\star \right\rangle = -\left\langle \big(\textstyle\sum_{i=1}^{k} \tau_i w_i\big), Jw^\star \right\rangle = -\textstyle\sum_{i=1}^{k} \tau_i \langle w_i, Jw^\star \rangle . \tag{30}$$

Using (29), observe for each $i \in [k]$ that

$$\langle w_i, Jw^\star \rangle = \left\langle \begin{pmatrix} p_i \\ q_i \end{pmatrix}, \begin{pmatrix} c\mathbf{1}_m \\ -c\mathbf{1}_n \end{pmatrix} \right\rangle = c \cdot \big(\langle p_i, \mathbf{1}_m \rangle - \langle q_i, \mathbf{1}_n \rangle \big) = 0 ,$$

where the final equality comes from the fact that $p_i \in \Delta_m$ and $q_i \in \Delta_n$ are probability distributions. Thus each term of (30) vanishes, and it follows that $\langle z, w^\star \rangle = 0$. ■

D.2 Proof of Proposition 3.2 – Change in KL is Change in Energy

Using the modified Fenchel-Young identity of Proposition D.2 and the orthogonality of the effective dual space $\mathcal{Z}$ to an interior Nash equilibrium from Proposition D.6, we can now prove the key structural result of Proposition 3.2. Restated here:

Proposition 3.2 (Equivalence between energy and KL differences)**.** *Let $A \in \mathbb{R}^{m \times n}$ have an interior NE $w^\star \in \text{ri}(\mathcal{W})$. Fix $z, z' \in \mathcal{Z}$, and let $w = \nabla F(z) \in \text{ri}(\mathcal{W})$ and $w' = \nabla F(z') \in \text{ri}(\mathcal{W})$. Then it holds that*

$$\text{KL}(w^\star, w') - \text{KL}(w^\star, w) = F(z') - F(z) .$$

Proof. As $w = \nabla F(z)$ and $w' = \nabla F(z')$ by assumption, we have by Proposition D.2 that

$$\text{KL}(w^\star, w') = R(w^\star) + F(z') - \langle z', w^\star \rangle \quad \text{and} \quad \text{KL}(w^\star, w) = R(w^\star) + F(z) - \langle z, w^\star \rangle .$$

Taking their difference, we find

$$\text{KL}(w^\star, w') - \text{KL}(w^\star, w) = F(z') - F(z) - \langle z' - z, w^\star \rangle . \tag{31}$$

As both $z, z' \in \mathcal{Z}$ and $w^\star \in \text{ri}(\mathcal{W})$, we have from Proposition D.6 that both $\langle z', w^\star \rangle = 0$ and $\langle z, w^\star \rangle = 0$. Thus the final term in (31) vanishes, which yields the desired statement. ■

Change in KL is Change in Energy for OMWU iterates. By (4), the OMWU dual iterates satisfy $z_t - z_0 \in \mathcal{Z}$ for all $t \geq 0$. Hence $z_t \in \mathcal{Z}$ whenever $z_0 \in \mathcal{Z}$. Thus as an immediate corollary of Proposition 3.2, we have the following:

Corollary D.7 (Equivalence between energy and KL differences for OMWU)**.** *Let $w^\star$ be the unique and interior NE of A, and let $\{w_t\}$ and $\{z_t\}$ be the primal and dual iterates of OMWU on A with stepsize $\eta > 0$, and initialized from $w_0 \in \text{ri}(\mathcal{W})$ and $z_0 \in \mathcal{Z}$ such that $\nabla F(z_0) = w_0$. Then*

$$\text{KL}(w^\star, w_{t+1}) - \text{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t) \;\; \textit{for all } t \geq 0.$$

Note that Corollary D.7 assumes that the dual initialization z_0 satisfying $\nabla F(z_0) = w_0$ also lies in $\mathcal{Z}$. In the following subsection, we establish additional properties of the energy gradient map ∇F and the effective dual space $\mathcal{Z}$ that imply such an initialization always exists when A has a unique and interior NE. Thus, this assumption on $z_0 \in \mathcal{Z}$ holds without loss of generality.

D.3 KL Divergence is Energy Suboptimality Gap Over Effective Dual Space

In this section, we further establish the following properties about the energy gradient map when its domain is restricted to the effective dual space $\mathcal{Z}$. Specifically, when A has a unique and interior NE $w^\star$, we establish in this section the following properties:

- **In Section D.3.1**, we prove that $\nabla F : \mathcal{Z} \to \text{ri}(\mathcal{W})$ is *surjective* (Proposition D.9) and also *injective* when the value of the game is centered at zero (Proposition D.10).
- **In Section D.3.2**, we prove for $z \in \mathcal{Z}$ and $w = \nabla F(z) \in \text{ri}(\mathcal{W})$ that $\text{KL}(w^\star, w)$ is exactly the energy suboptimality gap $F(z) - \min_{z' \in \mathcal{Z}} F(z')$ over the effective dual space (Proposition D.11).

D.3.1 Surjectivity (and Injectivity) of Energy Gradient Map Over Effective Dual Space

Helper lemma under unique and interior Nash. We start with the following helper lemma, which establishes that the only vector $v \in \mathcal{S}^\perp$ satisfying $Jv \in \mathcal{S}$ is $v = 0$, under the assumption of a unique and interior NE:

Lemma D.8. *Let $A \in \mathbb{R}^{m \times n}$ have a unique and interior NE $w^\star \in \text{ri}(\mathcal{W})$. Fix any $v \in \mathbb{R}^{m+n}$. Suppose that $Jv \in \mathcal{S}$ and $v \in \mathcal{S}^\perp$. Then $v = 0$.*

Proof. Since $w^\star$ is a unique and interior Nash equilibrium, we have by part (i) of Proposition B.1 that $Jw^\star \in \mathcal{S}$. Now for $\tau \in \mathbb{R}$, define $w_\tau = w^\star + \tau v$. In other words, w_τ is on the line through $w^\star$ in the direction of v. It follows by linearity that

$$Jw_\tau = J(w^\star + \tau v) = Jw^\star + \tau Jv \in \mathcal{S} \, .$$

Here, the inclusion in $\mathcal{S}$ follows from the assumption that $Jv \in \mathcal{S}$.

Now, as also $v = (v_x, v_y) \in \mathcal{S}^\perp$, we have by definition that $\langle \mathbf{1}_m, v_x \rangle = 0$ and $\langle \mathbf{1}_n, v_y \rangle = 0$. Thus, since $w^\star$ is interior, there exists a sufficiently small $\epsilon > 0$ such that, for all $\tau \in (-\epsilon, \epsilon)$, the point $w_\tau \in \text{ri}(\mathcal{W})$ is also an interior (component-wise) probability distribution. These two zero-sum conditions ensure that each component of w_τ still has coordinates that sum to 1.

However, since $w^\star$ is unique, and since $Jw_\tau \in \mathcal{S}$, we have by part (ii) of Proposition B.1 that $w_\tau = w^\star + \tau v = w^\star$ for all $\tau \in (-\epsilon, \epsilon)$. Thus by the uniqueness of $w^\star$, we must have $v = 0$. ■

Surjectivity of energy gradient over effective dual space. We now show that, when A has a unique and interior NE, then over the effective dual space $\mathcal{Z}$, the gradient map $\nabla F : \mathcal{Z} \to \text{ri}(\mathcal{W})$ is surjective. Formally, we prove the following:

Proposition D.9 (Surjectivity of ∇F over $\mathcal{Z}$)**.** *Let $A \in \mathbb{R}^{m \times n}$ have a unique and interior NE $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$. Then for all $w \in \text{ri}(\mathcal{W})$, there exists $z \in \mathcal{Z}$ such that $\nabla F(z) = w$.*

Proof. We prove the claim via the following three steps:

1. Energy gradient is globally surjective:

First, we show that $\nabla F : \mathbb{R}^{m+n} \to \text{ri}(\mathcal{W})$ is surjective over the full domain $\mathbb{R}^{m+n}$. For this, fix $w \in \text{ri}(\mathcal{W})$. Recalling that R is the joint negative entropy regularizer, observe that Proposition D.4 shows that $\nabla R(w) = v$ for some $v \in \mathbb{R}^{m+n}$. Using the definition of ∇R, it follows by Proposition D.5 that $\nabla F(v) = \nabla F(\nabla R(w)) = w$. Thus ∇F is globally surjective.

2. Energy gradient is invariant to constant shifts:

Now, again fix $w \in \text{ri}(\mathcal{W})$, and let $v = \nabla R(w) \in \mathbb{R}^{m+n}$. Further recall from Proposition D.3 that ∇F is invariant to constant shifts, meaning for any $s \in \mathcal{S}$ that $\nabla F(v + s) = \nabla F(v) = w$. Thus to prove ∇F is surjective over $\mathcal{Z}$, it is sufficient to establish $v + s \in \mathcal{Z}$ for some $s \in \mathcal{S}$. As $v \in \mathbb{R}^{m+n}$ can lie in the full space, it suffices to equivalently show that $\mathcal{Z} + \mathcal{S} = \mathbb{R}^{m+n}$.

3. Sufficient property holds under a unique and interior NE:

We will prove that $\mathcal{Z} + \mathcal{S} = \mathbb{R}^{m+n}$ holds under the assumption of a unique and interior NE $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$. For this, first observe that the property $\mathcal{Z} + \mathcal{S} = \mathbb{R}^{m+n}$ is equivalent to establishing equality between the orthogonal complements of the two sets. Since $\mathcal{Z}$, $\mathcal{S}$ are linear

subspaces, we have $(\mathcal{Z}+\mathcal{S})^\perp = \mathcal{Z}^\perp \cap \mathcal{S}^\perp$. Moreover, $(\mathbb{R}^{m+n})^\perp = \{0\}$. Thus our goal is to show $\mathcal{Z}^\perp \cap \mathcal{S}^\perp = \{0\}$.

For this, pick any $v = (v_x, v_y) \in \mathcal{Z}^\perp$. By definition, this means $\langle v, z\rangle = 0$ for every $z \in \mathcal{Z}$. We will show this implies $Jv \in \mathcal{S}$. In particular, for any $w = (p, q) \in \mathcal{W}$, let $z = Jw$. Thus clearly $z \in \mathcal{Z}$, and thus we must have $\langle v, z\rangle = 0$. As $J = -J^\top$, we can further express this condition as

$$0 = \langle v, z\rangle = \langle v, Jw\rangle = -\langle w, Jv\rangle = \big(\langle q, A^\top v_x\rangle - \langle p, Av_y\rangle\big)$$

Since this constraint must be satisfied for all $w = (p, q) \in \mathcal{W}$ simultaneously, it follows that we must have $A^\top v_x = c_1 \mathbf{1}_n$ and $Av_y = c_2 \mathbf{1}_m$ for constants $c_1, c_2 \in \mathbb{R}$. Thus $Jv \in \mathcal{S}$ by definition of $\mathcal{S}$.

We have established that $v \in \mathcal{Z}^\perp$ implies $Jv \in \mathcal{S}$. Now suppose further that $v \in \mathcal{S}^\perp$, meaning $v \in \mathcal{Z}^\perp \cap \mathcal{S}^\perp$. However, since $w^\star$ is the unique and interior NE, we have by Lemma D.8 that $v = 0$. Thus $\mathcal{Z}^\perp \cap \mathcal{S}^\perp = \{0\}$, which establishes $\mathcal{Z} + \mathcal{S} = \mathbb{R}^{m+n}$ under the assumptions on $w^\star$. ■

Conditions for injectivity of energy gradient over effective dual space. We further establish that ∇F is injective over $\mathcal{Z}$ when the value of the game A is zero. Formally, we prove the following:

Proposition D.10 (∇F injective when game value is zero)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$. If $d = \langle p^\star, Aq^\star\rangle = 0$, then $\nabla F(z) \neq \nabla F(z')$ for all $z \neq z' \in \mathcal{Z}$.*

Proof. First, recall from Proposition D.3 that ∇F is invariant under constant shifts. Thus it holds for $z, z' \in \mathcal{Z}$ that $\nabla F(z) = \nabla F(z')$ if and only if $z - z' \in \mathcal{S}$. This means that $\nabla F : \mathcal{Z} \to \mathbb{R}^{m+n}$ is injective over $\mathcal{Z}$ if and only if $\mathcal{Z} \cap \mathcal{S} = \{0\}$. We will show this latter condition holds when $d = 0$.

For this, pick any $z \in \mathcal{Z} \cap \mathcal{S}$. As $z \in \mathcal{Z}$, there exists some $v = (v_x, v_y) \in \mathbb{R}^{m+n}$ with $z = Jv$, where $v = \sum_{i=1}^k \tau_i w_i$ for $k \geq 1$ and $\tau_i \in \mathbb{R}$, $w_i = (p_i, q_i) \in \mathcal{W}$. Since $z \in \mathcal{S}$, it then further holds that $z = Jv \in \mathcal{S}$. Using the fact that $w^\star$ is unique and interior and applying a similar argument as in the proof of Lemma D.8, $Jv \in \mathcal{S}$ then further implies that $v = \alpha w^\star$ for some constant $\alpha \in \mathbb{R}$. Together, this means that if $z \in \mathcal{Z} \cap \mathcal{S}$, then $z \in \text{Span}(Jw^\star)$. Moreover, as clearly $Jw^\star \in \mathcal{Z}$, and also since $Jw^\star \in \mathcal{S}$ by Proposition B.1, we have that

$$\mathcal{Z} \cap \mathcal{S} = \text{Span}(Jw^\star)\ . \tag{32}$$

Finally, note that part (i) of Proposition B.1 more specifically gives $Jw^\star = (d\mathbf{1}_m, -d\mathbf{1}_n) \in \mathcal{S}$, where $d = \langle p^\star, Aq^\star\rangle \in \mathbb{R}$ is the value of the game. Thus if $d = 0$, then $Jw^\star = 0$, and $\text{Span}(Jw^\star) = \{0\}$. Thus $d = 0$ means $\mathcal{Z} \cap \mathcal{S} = \{0\}$, which implies $\nabla F : \mathcal{Z} \to \mathbb{R}^{m+n}$ is injective. ■

D.3.2 KL from Nash is Energy Suboptimality Gap

Finally, we show that when A has a unique and interior NE $w^\star$, the KL divergence from Nash in the primal space is equivalent to the energy suboptimality gap in the effective dual space. Formally:

Proposition D.11 (KL as Energy Suboptimality Gap)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star = (p^\star, q^\star) \in \text{ri}(\mathcal{W})$. Let $z \in \mathcal{Z}$ such that $w = \nabla F(z)$. Then there exists $z^\star \in \mathcal{Z}$ such that $\nabla F(z^\star) = w^\star$, and moreover* $\text{KL}(w^\star, w) = F(z) - \min_{z'\in\mathcal{Z}} F(z') = F(z) - F(z^\star)$.

Proof. We proceed via the following steps:

1. Existence of minimizer in effective dual space:

Recall that F is convex and continuously differentiable. Thus over the linear subspace $\mathcal{Z} \subset \mathbb{R}^{m+n}$, the first-order optimality conditions for F give (see, e.g., Boyd and Vandenberghe (2004), Sec. 4.2.3):

$$z \in \text{argmin}_{z'\in\mathcal{Z}} F(z') \iff \langle \nabla F(z), z''\rangle = 0 \text{ for all } z'' \in \mathcal{Z}\ .$$

Now due to Proposition D.9, $\nabla F : \mathcal{Z} \to \text{ri}(\mathcal{W})$ is surjective over $\mathcal{Z}$. Thus, as $w^\star \in \text{ri}(\mathcal{W})$, there exists $z^\star \in \mathcal{Z}$ such that $\nabla F(z^\star) = w^\star$. Due to Proposition D.6, this further implies that

$$\langle \nabla F(z^\star), z\rangle = \langle w^\star, z\rangle = 0 \text{ for all } z \in \mathcal{Z} \implies z^\star \in \text{argmin}_{z\in\mathcal{Z}} F(z)\ .$$

2. Minimum energy function value attained:

Recall from expression (32) in the proof of Proposition D.10 that $\mathcal{Z} \cap \mathcal{S} = \text{Span}(Jw^\star)$. Thus fixing $z^\star \in \text{argmin}_{z\in\mathcal{Z}} F(z)$ from Step (1), then also $z_\tau = z^\star + \tau Jw^\star \in \text{argmin}_{z\in\mathcal{Z}} F(z)$ for any $\tau \in \mathbb{R}$.

This is because $\nabla F(z^\star + s) = \nabla F(z^\star)$ for any $s \in \mathcal{S}$, and since $Jw^\star \in \mathcal{S}$. Thus $z' \in \mathcal{Z}$ is in the argmin set if and only if $z' = z_\tau$ for some $\tau \in \mathbb{R}$. However, since $F(z_\tau) = F(z^\star)$ for every $\tau \in \mathbb{R}$, $\min_{z\in\mathcal{Z}} F(z)$ is well-defined, and thus $F(z^\star) = \min_{z\in\mathcal{Z}} F(z)$.

3. Apply Fenchel-Young identity:

Now by Part (ii) of the Fenchel-Young identity from Proposition D.2, we have for all $z \in \mathcal{Z}$ and $w = \nabla F(z)$ that

$$\mathrm{KL}(w^\star, w) = R(w^\star) + F(z) - \langle w^\star, z\rangle = R(w^\star) + F(z)\ , \tag{33}$$

where the second equality is due to Proposition D.6. Now let $z^\star \in \mathcal{Z}$ such that $\nabla F(z^\star) = w^\star$. Then it follows from (33) that $0 = \mathrm{KL}(w^\star, w^\star) = R(w^\star) + F(z^\star) \implies R(w^\star) = -F(z^\star)$. By Step (1), we have $z^\star \in \operatorname{argmin}_{z\in\mathcal{Z}} F(z)$, and by Step (2), we further have $F(z^\star) = \min_{z\in\mathcal{Z}} F(z)$. Then substituting $R(w^\star) = -F(z^\star)$ into (33), we find for all $z \in \mathcal{Z}$ and $w = \nabla F(z)$ that

$$\mathrm{KL}(w^\star, w) = F(z) - \min_{z'\in\mathcal{Z}} F(z')\ ,$$

which concludes the proof. ■

E Details on Energy Function Hessian Matrix

This section gives additional preliminaries on the Hessian of the energy function that are used in the proof of energy dissipation under OMWU.

Organization of Section. The section is organized as follows:

- **Section E.1** gives more details on the energy function's Hessian that are used throughout the sequel.
- **Section E.2** gives more details on the *local hessian stability* (LHS) property introduced in Section 4 and provides the proof of Proposition 4.3.
- **Section E.3** proves several bounds using the local seminorms $\|\cdot\|_z$ that hold under the LHS property.

E.1 Properties of Energy Hessian

Our analysis of OMWU relies on several properties of the energy function's Hessian $\nabla^2 F$. Formally, for any $z = (x, y) \in \mathbb{R}^{m+n}$, the Hessian $\nabla^2 F(z) \in \mathbb{R}^{(m+n)\times(m+n)}$ is the block diagonal symmetric matrix

$$\nabla^2 F(z) = \begin{pmatrix} \nabla^2 \mathrm{LSE}_m(x) & 0 \\ 0 & \nabla^2 \mathrm{LSE}_n(y) \end{pmatrix}\ . \tag{34}$$

Note by convexity of F that $\nabla^2 F(z) \succeq 0$ for all $z \in \mathbb{R}^{m+n}$.

In the next proposition, we state several basic properties about the energy Hessian. In particular, we note that $\nabla^2 F$ has a natural probabilistic characterization: for $z = (x, y) \in \mathbb{R}^{m+n}$ and $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$, the Hessian blocks $\nabla^2 \mathrm{LSE}_m(x)$ and $\nabla^2 \mathrm{LSE}_n(y)$ are the *covariance matrices* of the probability distributions p and q, respectively. Thus, quadratic forms of these Hessian blocks can be interpreted as variances under the distributions p and q. Before stating and proving these properties formally, we first introduce the following notation for discrete variances:

Notation for variances. Fix $(v_x, v_y) \in \mathbb{R}^{m+n}$ and $(p, q) \in \mathcal{W}$. We write $\mathrm{Var}_p(v_x) \in \mathbb{R}$ to denote the variance of $v_x \in \mathbb{R}^m$ under the distribution $p \in \Delta_m$, and we similarly write $\mathrm{Var}_q(v_y) \in \mathbb{R}$ to denote the variance of $v_y \in \mathbb{R}^n$ under the distribution $q \in \Delta_n$. Formally, we define:

$$\mathrm{Var}_p(v_x) := \textstyle\sum_{i=1}^m p(i)\big(v_x(i) - \langle p, v_x\rangle\big)^2 \ \text{ and } \ \mathrm{Var}_q(v_y) := \sum_{j=1}^n q(j)\big(v_y(j) - \langle q, v_y\rangle\big)^2\ . \tag{35}$$

Then for any $z \in \mathbb{R}^{m+n}$, the matrix $\nabla^2 F(z)$ has the following properties:

Proposition E.1 (Properties of Energy Hessian)**.** *Let F be the energy function from (6). For any $z \in \mathbb{R}^{m+n}$, let $\nabla^2 F(z)$ be the matrix in (34). Let $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Then:*

(i) The Hessian $\nabla^2 F(z)$ is given by

$$\nabla^2 F(z) = \begin{pmatrix} \big(\mathrm{Diag}(p) - pp^\top\big) & 0 \\ 0 & \big(\mathrm{Diag}(q) - qq^\top\big) \end{pmatrix}\ .$$

(ii) *For any* $v = (v_x, v_y) \in \mathbb{R}^{m+n}$*, it holds that* $\langle v, \nabla^2 F(z) v \rangle = \mathrm{Var}_p(v_x) + \mathrm{Var}_q(v_y)$.

(iii) $\mathrm{Null}\big(\nabla^2 F(z)\big) = \mathcal{S} = \mathrm{Span}\big(\binom{\mathbf{1}_m}{0}, \binom{0}{\mathbf{1}_n}\big)$.

(iv) $0 \preceq \nabla^2 F(z) \preceq I$.

Proof. We prove each of the statements separately:

Proof of (i). This result is standard, and a proof can be found in Boyd and Vandenberghe (2004), Section 3.1.5, or Gao and Pavel (2017), Proposition 2.

Proof of (ii). Using the structure of $\nabla^2 F(z)$ from claim (i), for $v = (v_x, v_y) \in \mathbb{R}^{m+n}$ we have

$$\langle v, \nabla^2 F(z) v \rangle = \langle v_x, (\mathrm{Diag}(p) - pp^\top) v_x \rangle + \langle v_y, (\mathrm{Diag}(q) - qq^\top) v_y \rangle \, .$$

Observe that the matrices $(\mathrm{Diag}(p) - pp^\top)$ and $(\mathrm{Diag}(q) - qq^\top)$ are exactly the covariance matrices of the distributions $p \in \Delta_m$ and $q \in \Delta_n$. Indeed, recalling the definitions of $\mathrm{Var}_p(v_x)$ and $\mathrm{Var}_q(v_y)$ from (35), it is straightforward to check that

$$\begin{aligned}
\langle v_x, (\mathrm{Diag}(p) - pp^\top) v_x \rangle &= \langle v_x, \mathrm{Diag}(p) v_x \rangle - \langle v_x, pp^\top v_x \rangle \\
&= \langle v_x, \mathrm{Diag}(p) v_x \rangle - \big(\langle p, v_x \rangle\big)^2 \\
&= \textstyle\sum_{i=1}^m p(i) \cdot (v_x(i))^2 - \big(\sum_{i=1}^m p(i) \cdot v_x(i)\big)^2 = \mathrm{Var}_p(v_x) \, .
\end{aligned}$$

Using identical calculations, it follows that also $\langle v_y, (\mathrm{Diag}(q) - qq^\top) v_y \rangle = \mathrm{Var}_q(v_y)$, which completes the proof of claim (ii).

Proof of (iii). By (ii), the quadratic form vanishes exactly when both blocks of v are constant, since p and q have full support. This gives the claimed nullspace; see also Proposition 2 of Gao and Pavel (2017).

Proof of (iv). Observe that $(\mathrm{Diag}(p) - pp^\top) \preceq I$, which follows from the fact that all $0 < p(i) \leq 1$ and $pp^\top \succeq 0$. Similarly, $(\mathrm{Diag}(q) - qq^\top) \preceq I$, and thus due to the block structure of $\nabla^2 F(z)$ from part (i) of the proposition, it follows that $\nabla^2 F(z) \preceq I$. By convexity of F, also $0 \preceq \nabla^2 F(z)$. ■

E.1.1 Hessian of Regularizer and Inverse Energy Hessian

We also consider the Hessian matrix of the joint negative entropy regularizer $R : \mathcal{W} \to \mathbb{R}$ from (19).

Negative Entropy Hessian. Fix any $w = (p, q) \in \mathrm{ri}(\mathcal{W})$. Then let $\nabla^2 R(w) \in \mathbb{R}^{(m+n)\times(m+n)}$ be the Hessian matrix of R given by

$$\nabla^2 R(w) = \begin{pmatrix} \nabla^2 R_m(p) & 0 \\ 0 & \nabla^2 R_n(q) \end{pmatrix} = \begin{pmatrix} \mathrm{Diag}(1/p) & 0 \\ 0 & \mathrm{Diag}(1/q) \end{pmatrix} \, . \tag{36}$$

Here, we write $1/p = (1/p(1), \ldots, 1/p(m)) \in \mathbb{R}^m$ and $1/q = (1/q(1), \ldots, 1/q(n)) \in \mathbb{R}^n$. Note that the second equality in (36) follows from the definitions of $\nabla R_m(p)$ and $\nabla R_n(q)$ from (24).

Negative Entropy Hessian as Restricted Inverse of Energy Hessian. Due to the duality of the energy function and the negative entropy regularizer, the inverse of the energy Hessian on $\mathcal{S}^\perp$ agrees with $\nabla^2 R$ as a quadratic form. In particular, let $\nabla^2 F(z)|_{\mathcal{S}^\perp}$ denote the restriction of the Hessian to $\mathcal{S}^\perp$, viewed as a linear operator from this subspace to itself. Then we have the following relationship:

Proposition E.2 (Energy hessian restricted inverse)**.** *Fix* $z \in \mathbb{R}^{m+n}$*, and let* $w = (p, q) = \nabla F(z)$*. Write* $(\nabla^2 F(z))^{-1}$ *for the inverse of* $\nabla^2 F(z)|_{\mathcal{S}^\perp}$*. Then* $\nabla^2 F(z) \nabla^2 R(w) v = v$ *for all* $v \in \mathcal{S}^\perp$*, and*

$$\langle v, (\nabla^2 F(z))^{-1} v \rangle = \langle v, \nabla^2 R(w) v \rangle, \qquad \text{for } v \in \mathcal{S}^\perp.$$

Proof. Let $z = (x, y)$, and let $v = (v_x, v_y)$. We first prove the claim for the blocks $\nabla^2 \mathrm{LSE}_m(x)$ and $\nabla^2 R_m(p)$. For this, we have by definition of $\nabla^2 R_m$ and $\nabla^2 \mathrm{LSE}_m$ that

$$\begin{aligned}
\nabla^2 \mathrm{LSE}_m(x) \nabla^2 R_m(p) v_x &= (\mathrm{Diag}(p) - pp^\top)(\mathrm{Diag}(1/p)) v_x \\
&= \mathrm{Diag}(p) \mathrm{Diag}(1/p) v_x - p(p^\top \mathrm{Diag}(1/p)) v_x \\
&= I v_x - p \cdot \langle \mathbf{1}_m, v_x \rangle = v_x \, .
\end{aligned}$$

Here, the penultimate equality is due to $p^\top \mathrm{Diag}(1/p) = \mathbf{1}_m^\top$, and the final equality is due to the assumption that $v \in \mathcal{S}^\perp$, and thus $\langle v_x, \mathbf{1}_m \rangle = 0$. By an identical calculation, we have $\nabla^2 \mathrm{LSE}_n(y) \nabla^2 R_n(q) v_y = v_y$, which proves the first identity by definition of $\nabla^2 F$ and $\nabla^2 R$. By Proposition E.1(iii), $\nabla^2 F(z)$ annihilates $\mathcal{S}$ and is positive definite on $\mathcal{S}^\perp$. Thus the orthogonal projection of $\nabla^2 R(w) v$ onto $\mathcal{S}^\perp$ equals $(\nabla^2 F(z))^{-1} v$. Taking its inner product with $v \in \mathcal{S}^\perp$ gives the equality of quadratic forms. ■

E.1.2 Gradient Differences and Bregman Divergence in Integral Form

Throughout, we will also use the energy Hessian $\nabla^2 F$ to express certain remainder terms when taking Taylor expansions of F. First, we can express the Bregman divergence $D_F(\cdot, \cdot)$ as follows:

Proposition E.3 (Bregman divergence in integral form)**.** *Let F be the function in (6), and let $z, z' \in \mathbb{R}^{m+n}$. Let $z_s := z - s(z - z')$ for $s \in [0, 1]$. Then the Bregman divergence $D_F(z', z)$ is*

$$D_F(z', z) = \int_0^1 (1 - s) \big\langle z - z', \nabla^2 F(z_s)(z - z') \big\rangle \, ds \, .$$

For $z, z' \in \mathbb{R}^{m+n}$, we can also express the difference of energy gradients using the energy Hessian:

Proposition E.4 (Gradient difference and Hessian remainder)**.** *Let F be the function in (6), and let $z, z' \in \mathbb{R}^{m+n}$. Let $z_s := z - s(z - z')$ for $s \in [0, 1]$. Let $G_F(z, z') \in \mathbb{R}^{m+n}$ be the remainder term*

$$G_F(z, z') = \int_0^1 (\nabla^2 F(z_s) - \nabla^2 F(z))(z - z') \, ds \, . \tag{37}$$

Then $\nabla F(z) - \nabla F(z') = \nabla^2 F(z)(z - z') + G_F(z, z')$.

Propositions E.3 and E.4 are both standard consequences of Taylor's theorem (see, e.g., Nocedal and Wright (2006), Theorem 2.1), and thus we omit the proofs.

E.2 Local Hessian Stability Property

In the sequel, our analysis of the energy dissipation under OMWU relies on the following *local Hessian stability* property of the energy function, introduced originally in Section 4:

Proposition 4.3 (LHS Property)**.** *For any $z, z' \in \mathbb{R}^{m+n}$ and $\alpha > 0$, if $\|z - z'\|_\infty \leq \alpha$, then:*

$$\exp(-2\alpha) \cdot \nabla^2 F(z) \preceq \nabla^2 F(z') \preceq \exp(2\alpha) \cdot \nabla^2 F(z) \, . \tag{LHS}$$

Proof. Recalling the block structure of the energy Hessian $\nabla^2 F(z)$ from Proposition E.1, it suffices to show that the property holds for the block components $\nabla^2 \mathrm{LSE}_m(x)$ and $\nabla^2 \mathrm{LSE}_n(y)$. We will prove the property holds for the former, and for the latter it will follow by identical calculations.

For this, fix any vector $u \in \mathbb{R}^m$, and let $z = (x, y)$ and $z' = (x', y')$. Observe that the assumption $\|z - z'\|_\infty \leq \alpha$ implies $\|x - x'\|_\infty \leq \alpha$. Now for every $s \in [0, 1]$, define $x_s = x + s(x' - x)$, and let $h = x' - x$. Moreover, let $p_s = \nabla \mathrm{LSE}_m(x_s) = \mathrm{softmax}_m(x_s)$, meaning

$$p_s = \mathrm{softmax}_m(x_s) = \mathrm{softmax}_m(x + sh) \, .$$

Further let $f_u(s)$ be the scalar function given by

$$f_u(s) = \big\langle u, \nabla^2 \mathrm{LSE}_m(x_s) u \big\rangle = \big\langle u, (\mathrm{Diag}(p_s) - p_s p_s^\top) u \big\rangle = \textstyle\sum_{i=1}^m p_s(i) \cdot \big(u(i) - \langle p_s, u \rangle\big)^2 \, ,$$

where the equalities follow from the structure of $\nabla^2 \mathrm{LSE}_m$ from Proposition E.1.

Now, differentiating $f_u(s)$ with respect to s and simplifying, we find

$$f_u'(s) = \textstyle\sum_{i=1}^m p_s(i) \cdot \big(h(i) - \langle p_s, h \rangle\big)\big(u(i) - \langle p_s, u \rangle\big)^2 \, .$$

By the assumption $\|x - x'\|_\infty = \|h\|_\infty \leq \alpha$, it follows that $|h(i) - \langle p_s, h \rangle| \leq 2\alpha$. Thus we can further bound

$$|f_u'(s)| \leq 2\alpha \cdot \sum_{i=1}^m p_s(i) \cdot \big(u(i) - \langle p_s, u \rangle\big)^2 = 2\alpha \cdot f_u(s) \, .$$

Applying Grönwall's inequality then yields

$$\exp(-2\alpha) \cdot f_u(0) \le f_u(1) \le \exp(2\alpha) \cdot f_u(0) \ .$$

Since $x_0 = x$ and $x_1 = x'$ and u was arbitrary, it follows by definition of f_u that

$$\exp(-2\alpha) \cdot \nabla^2 \mathrm{LSE}_m(x) \preceq \nabla^2 \mathrm{LSE}_m(x') \preceq \exp(2\alpha) \cdot \nabla^2 \mathrm{LSE}_m(x) \ ,$$

which concludes the proof. ■

Local Hessian Stability and Generalized Self-Concordance. Similar versions of this LHS property have also been exploited in the optimization literature as a consequence of *(generalized) self-concordance*. See, e.g., Bach (2010), Tran-Dinh et al. (2015) (Theorem 4.b, eq. (20)), Sun and Tran-Dinh (2019) (Proposition 8, eq. (16) when $\nu = 2$), or more recently Freund et al. (2026) (Proposition 16, and Lemma 15), and the references therein. For the purposes of the present paper, our proof of energy dissipation in Lemma 4.1 uses the semidefinite ordering of (LHS) via the application of several general inequalities involving local seminorms under the energy Hessian. We present these tools in the next subsection.

E.3 Consequences of Local Hessian Stability Property

E.3.1 Review of Notation for Local Seminorms

We first recall and expand on the notation of the local seminorms introduced in Section 4. Fix $z \in \mathbb{R}^{m+n}$. Then for all $u, v \in \mathbb{R}^{m+n}$, we define

$$\langle u, v \rangle_z := \langle u, \nabla^2 F(z) v \rangle \ \text{ and } \ \|v\|_z := \sqrt{\langle v, \nabla^2 F(z) v \rangle} \ .$$

By Proposition E.1(iii), the bilinear form is positive semidefinite with nullspace $\mathcal{S}$, and hence defines a degenerate inner product. In particular, $\|v\|_z = 0$ if and only if $v \in \mathcal{S}$, so this seminorm induces a norm on $\mathbb{R}^{m+n}/\mathcal{S}$ and restricts to a norm on $\mathcal{S}^\perp$. The triangle and Cauchy–Schwarz inequalities still hold: for all $u, v \in \mathbb{R}^{m+n}$, the square-root representation gives

$$\|v\|_z = \|(\nabla^2 F(z))^{1/2} v\|_2, \qquad |\langle u, v \rangle_z| \le \|u\|_z \|v\|_z.$$

See Rudin (1991), Chapter 1, for background on seminorms and quotient spaces.

We further define a local *dual* norm $\|\cdot\|_{z,*}$ over the linear subspace $\mathcal{S}^\perp \subseteq \mathbb{R}^{m+n}$ (as defined in (7)). Here $(\nabla^2 F(z))^{-1}$ denotes the inverse of $\nabla^2 F(z)|_{\mathcal{S}^\perp}$, meaning all inverse Hessians and their square roots are restricted to acting on vectors $v \in \mathcal{S}^\perp$. By Proposition E.2, its quadratic form agrees with that of $\nabla^2 R(w)$, where $w = \nabla F(z)$ and $\nabla^2 R(w)$ is defined in (36). Then over $\mathcal{S}^\perp$, for any $z \in \mathbb{R}^{m+n}$, the local dual norm $\|\cdot\|_{z,*}$ is given by

$$\|u\|_{z,*} = \sqrt{\langle u, (\nabla^2 F(z))^{-1} u \rangle} \quad \text{for } u \in \mathcal{S}^\perp.$$

For $z \in \mathbb{R}^{m+n}$, the primal-dual pair $(\|\cdot\|_z, \|\cdot\|_{z,*})$ satisfies the generalized Cauchy-Schwarz inequality

$$\big|\langle v, u \rangle\big| \le \|v\|_z \cdot \|u\|_{z,*} \quad \text{for all } v \in \mathbb{R}^{m+n} \text{ and } u \in \mathcal{S}^\perp.$$

E.3.2 Properties of Local Seminorms under LHS

We state and prove several useful properties of the local seminorms under the (LHS) property. The local seminorm comparisons and Bregman divergence bounds below use positive semidefiniteness and the Loewner inequalities in (LHS), and do not require the Hessian to be positive definite on the ambient space.

Relationships between global norms and local seminorms. The following proposition relates the local seminorm and its dual norm to the Euclidean norm:

Proposition E.5. *For all $z \in \mathbb{R}^{m+n}$, $v \in \mathbb{R}^{m+n}$, and $u \in \mathcal{S}^\perp$: $\|v\|_z \le \|v\|_2$ and $\|u\|_2 \le \|u\|_{z,*}$.*

Proof. For the first claim, recall by part (iv) of Proposition E.1 that $\nabla^2 F(z) \preceq I$. Thus for any $v \in \mathbb{R}^{m+n}$, it holds that

$$\|v\|_z^2 = \langle v, \nabla^2 F(z) v \rangle \le \langle v, v \rangle = \|v\|_2^2 \ ,$$

which implies $\|v\|_z \le \|v\|_2$. For the second claim, note that when restricted to $\mathcal{S}^\perp$, the ordering $\nabla^2 F(z) \preceq I$ implies $I \preceq (\nabla^2 F(z))^{-1}$. Then for any $u \in \mathcal{S}^\perp$ and $z \in \mathbb{R}^{m+n}$, it holds that

$$\|u\|_2^2 = \langle u, u\rangle \le \langle u, (\nabla^2 F(z))^{-1} u\rangle = \|u\|_{z,*}^2 ,$$

which implies $\|u\|_2 \le \|u\|_{z,*}$. ■

Local seminorm transfer. For vectors $z, z', v \in \mathbb{R}^{m+n}$, a key consequence of the (LHS) property is the ability to relate the local seminorm $\|v\|_z$ to the local seminorm $\|v\|_{z'}$ up to multiplicative factors depending on $\|z - z'\|_\infty$. Formally, we have the following relationships:

Lemma E.6. *For $\alpha > 0$, let $z, z' \in \mathbb{R}^{m+n}$ such that $\|z - z'\|_\infty \le \alpha$. Then for all $v \in \mathbb{R}^{m+n}$:*

$$\exp(-\alpha) \cdot \|v\|_z \le \|v\|_{z'} \le \exp(\alpha)\|v\|_z .$$

Proof. Due to the (LHS) property, for any $z, z' \in \mathbb{R}^{m+n}$ such that $\|z - z'\|_\infty \le \alpha$, we have

$$\exp(-2\alpha)\nabla^2 F(z) \preceq \nabla^2 F(z') \preceq \exp(2\alpha)\nabla^2 F(z) .$$

Using the definition of local seminorms, this implies that for any $v \in \mathbb{R}^{m+n}$,

$$\exp(-2\alpha) \cdot \|v\|_z^2 \le \|v\|_{z'}^2 = \langle v, \nabla^2 F(z')v\rangle \le \exp(2\alpha) \cdot \|v\|_z^2 . \tag{38}$$

Taking the square root in the above inequality concludes the proof. ■

Non-expansiveness of local seminorm. The next lemma establishes that the energy Hessian satisfies a certain non-expansiveness property under its own induced local seminorm:

Lemma E.7. *Let $M \in \mathbb{R}^{(m+n)\times(m+n)}$ be a matrix, and fix $z \in \mathbb{R}^{m+n}$. Then for any $v \in \mathbb{R}^{m+n}$:*

$$\|M\nabla^2 F(z)v\|_z \le \|M\|_2 \cdot \|v\|_z .$$

Proof. Applying Proposition E.5, we can write

$$\|M\nabla^2 F(z)v\|_z \le \|M\nabla^2 F(z)v\|_2 \le \|M\|_2 \cdot \|\nabla^2 F(z)v\|_2 , \tag{39}$$

where the final inequality comes from the definition of the spectral norm $\|M\|_2$. To further bound the term $\|\nabla^2 F(z)v\|_2$, observe by definition that

$$\|\nabla^2 F(z)v\|_2^2 = \langle v, (\nabla^2 F(z))^2 v\rangle \le \langle v, \nabla^2 F(z)v\rangle = \|v\|_z^2 . \tag{40}$$

Here, the inequality comes from the fact that $(\nabla^2 F(z))^2 \preceq \nabla^2 F(z)$ since, by Proposition E.1, $0 \preceq \nabla^2 F(z) \preceq I$. Then taking square roots and substituting (40) into (39) concludes the proof. ■

E.3.3 Bounds on Bregman Divergence and Hessian Remainder Terms under LHS

Additional notation. To start, we define the following two scalar functions which will be useful in our proofs for controlling certain constant multiplicative factors. Specifically, let $\mu : \mathbb{R} \to \mathbb{R}$ and $\nu : \mathbb{R} \to \mathbb{R}$ be the functions such that, for any $x \neq 0$:

$$\mu(x) = \frac{\exp(x) - x - 1}{x^2} , \quad \nu(x) = \frac{\exp(x) - x - 1}{x} . \tag{41}$$

We define $\mu(0) = 1/2$ and $\nu(0) = 0$ by continuity. In particular we will be using the following straightforward bounds:

Lemma E.8. *For any $\delta > 0$ and $|x| \le \delta$, it holds that*

$$\left|\mu(x) - \tfrac{1}{2}\right| \le \frac{\delta \exp(\delta)}{6} \quad \text{and} \quad \left|\nu(x) - \tfrac{x}{2}\right| \le \frac{\delta^2 \exp(\delta)}{6} .$$

Proof. For $x \neq 0$, taking a second-order Taylor expansion with Lagrange remainder, we have:

$$\exp(x) = 1 + x + \frac{x^2}{2} + \frac{\exp(\xi)x^3}{6} \tag{42}$$

for some ξ between 0 and x. Rearranging, we find $\mu(x) - \frac{1}{2} = \frac{\exp(\xi)x}{6}$. Noting that $|\exp(\xi)x| \le \delta\exp(\delta)$ for $|x| \le \delta$ then gives $|\mu(x) - \frac{1}{2}| \le \frac{\delta\exp(\delta)}{6}$, which yields the first claim of the lemma.

For the second claim, we can again rearrange (42) to find $\nu(x) - \frac{x}{2} = \frac{\exp(\xi)x^2}{6}$. Taking absolute values and again using $|\exp(\xi)x^2| \le \delta^2\exp(\delta)$ for $|x| \le \delta$ then gives $|\nu(x) - \frac{x}{2}| \le \frac{\delta^2\exp(\delta)}{6}$. The case $x = 0$ follows from the definitions above. ■

Bounds on Bregman divergence.

Lemma E.9. *Fix $z, z' \in \mathbb{R}^{m+n}$ such that $\|z - z'\|_\infty \le \alpha$ for $\alpha > 0$. Then the Bregman divergence $D_F(z', z)$ is bounded above and below by*

$$\mu(-2\alpha) \cdot \|z' - z\|_z^2 \;\le\; D_F(z', z) \;\le\; \mu(2\alpha) \cdot \|z' - z\|_z^2 \,.$$

Proof. Recall from Proposition E.3 that the Bregman divergence $D_F(z', z)$ can be written as

$$D_F(z', z) = \int_0^1 (1-s)\big\langle z - z', \nabla^2 F(z_s)(z - z')\big\rangle \, ds \,, \tag{43}$$

where $z_s := z - s(z - z') \in \mathbb{R}^{m+n}$ for $s \in [0, 1]$. Fixing s, observe that $\|z - z_s\|_\infty = s\|z - z'\|_\infty \le s\alpha$, which follows from the assumption that $\|z - z'\|_\infty \le \alpha$. Thus we have by (LHS) that

$$\exp(-2s\alpha)\nabla^2 F(z) \preceq \nabla^2 F(z_s) \preceq \exp(2s\alpha)\nabla^2 F(z) \,. \tag{44}$$

For the upper bound on $D_F(z', z)$, note that the relationship in (44) implies

$$\begin{aligned}\big\langle z - z', \nabla^2 F(z_s)(z - z')\big\rangle &\le \exp(2s\alpha) \cdot \langle z - z', \nabla^2 F(z)(z - z')\rangle \\ &= \exp(2s\alpha) \cdot \|z - z'\|_z^2 \,.\end{aligned}$$

Substituting back into (43) and integrating, we conclude (using the function μ from (41)):

$$\begin{aligned}D_F(z', z) &\le \int_0^1 (1-s)\exp(2s\alpha)\|z - z'\|_z^2 \, ds \\ &= \Big(\int_0^1 (1-s)\exp(2s\alpha)\, ds\Big) \cdot \|z - z'\|_z^2 = \mu(2\alpha) \cdot \|z' - z\|_z^2 \,.\end{aligned}$$

For the lower bound on $D_F(z', z)$, we similarly have from (44) that

$$\begin{aligned}\big\langle z - z', \nabla^2 F(z_s)(z - z')\big\rangle &\ge \exp(-2s\alpha) \cdot \langle z - z', \nabla^2 F(z)(z - z')\rangle \\ &= \exp(-2s\alpha) \cdot \|z - z'\|_z^2 \,.\end{aligned}$$

Thus we can again substitute back into (43) and integrate to find

$$D_F(z', z) \ge \Big(\int_0^1 (1-s)\exp(-2s\alpha)\, ds\Big) \cdot \|z - z'\|_z^2 = \mu(-2\alpha) \cdot \|z' - z\|_z^2 \,,$$

which concludes the proof. ∎

Bound on Hessian remainder term.

Lemma E.10. *Fix $z, z' \in \mathbb{R}^{m+n}$. Then $G_F(z, z') \in \mathcal{S}^\perp$. Moreover, if $\|z - z'\|_\infty \le \alpha$ for some $\alpha > 0$, then*

$$\|G_F(z, z')\|_{z,*} \le \nu(2\alpha) \cdot \|z - z'\|_z \,.$$

Proof. Recall from (37) that we define the remainder term $G_F(z, z')$ in integral form by

$$G_F(z, z') = \int_0^1 (\nabla^2 F(z - sv) - \nabla^2 F(z))v \, ds \,, \tag{45}$$

where $v = z - z'$. Now for the first claim of the lemma, recall from part (iii) of Proposition E.1 that $\text{Null}(\nabla^2 F(z'')) = \mathcal{S}$, and thus $\nabla^2 F(z'')b = 0$ for every $z'' \in \mathbb{R}^{m+n}$ and $b \in \mathcal{S}$. Then for all $u \in \mathbb{R}^{m+n}$, by symmetry of $\nabla^2 F$, it follows for every $b \in \mathcal{S}$ that

$$\langle b, \nabla^2 F(z'')u\rangle = \langle u, \nabla^2 F(z'')b\rangle = 0 \,.$$

By definition, this means $\nabla^2 F(z'')u \in \mathcal{S}^\perp$ for every $z'', u \in \mathbb{R}^{m+n}$, and thus also $G_F(z, z') \in \mathcal{S}^\perp$.

For the second claim of the lemma, taking local (dual) norms in (45) and using the triangle inequality:

$$\|G_F(z, z')\|_{z,*} \le \int_0^1 \|(\nabla^2 F(z - sv) - \nabla^2 F(z))v\|_{z,*} \, ds \,. \tag{46}$$

Now repeating the previous argument, observe that each term within the norm $\|\cdot\|_{z,*}$ in the right-hand-side of (46) belongs to $\mathcal{S}^\perp$. Similar to the proof of Lemma E.9, let $s \in [0,1]$ and write $z_s := z - sv$. For readability, we further write $P_z = \nabla^2 F(z)$, $P_{z_s} := \nabla^2 F(z_s)$, and $B_s := P_{z_s} - P_z$ for $s \in [0,1]$. Moreover, as we are working over $\mathcal{S}^\perp$, we write $P_z^{-1} = (\nabla^2 F(z))^{-1}$ to denote the energy Hessian inverse over $\mathcal{S}^\perp$. With this notation, our goal is thus to control the terms $\|B_s v\|_{z,*}$ in the integrand of (46). For this, define the matrix $M := P_z^{-1/2} P_{z_s} P_z^{-1/2}$, for which it holds that

$$P_z^{-1/2} B_s = (M - I) P_z^{1/2} \; . \tag{47}$$

Then taking local (dual) norms, we have

$$\begin{aligned} \|B_s v\|_{z,*}^2 = \langle B_s v, P_z^{-1} B_s v \rangle &= \langle B_s v, P_z^{-1/2} P_z^{-1/2} B_s v \rangle \\ &= \|P_z^{-1/2} B_s v\|_2^2 \\ &= \|(M - I) P_z^{1/2} v\|_2^2 \qquad \text{(using (47))} \\ &\le \|M - I\|_2^2 \cdot \|P_z^{1/2} v\|_2^2 \; . \end{aligned} \tag{48}$$

Now to bound the two terms of (48), observe first by definition of the local seminorm that

$$\|P_z^{1/2} v\|_2^2 = \langle P_z^{1/2} v, P_z^{1/2} v \rangle = \langle v, P_z^{1/2} P_z^{1/2} v \rangle = \langle v, P_z v \rangle = \|v\|_z^2 \; . \tag{49}$$

To bound $\|M - I\|_2^2$, recall from the (LHS) property that

$$\exp(-2s\alpha) P_z \preceq P_{z_s} \preceq \exp(2s\alpha) P_z \; , \tag{50}$$

which follows from the fact that $\|z - z_s\|_\infty = s\|z - z'\|_\infty \le s\alpha$ by assumption. Now due to the orderings of (50) and using the definition of M, it follows by congruence that

$$M = P_z^{-1/2} P_{z_s} P_z^{-1/2} \preceq \exp(2s\alpha) \cdot P_z^{-1/2} P_z P_z^{-1/2} = \exp(2s\alpha) I \tag{51}$$

and similarly

$$M \succeq \exp(-2s\alpha) I \; . \tag{52}$$

Thus rearranging (51) and (52) and using $\exp(2s\alpha) - 1 \ge 1 - \exp(-2s\alpha)$, it follows that

$$-(\exp(2s\alpha) - 1) I \preceq M - I \preceq (\exp(2s\alpha) - 1) I \; .$$

Thus, we have $\|M - I\|_2 \le (\exp(2s\alpha) - 1)$ by definition of spectral norm. Combining this with (49) and substituting back into (48) then yields

$$\|B_s v\|_{z,*} \le (\exp(2s\alpha) - 1) \cdot \|v\|_z \; .$$

Finally, putting this bound back into (46) and integrating, we find

$$\|G_F(z, z')\|_{z,*} \le \int_0^1 \|B_s v\|_{z,*} \, ds \le \Big(\int_0^1 (\exp(2s\alpha) - 1) \, ds \Big) \cdot \|v\|_z = \nu(2\alpha) \cdot \|v\|_z \; ,$$

which concludes the proof. ∎

F Details on Energy Dissipation Under OMWU

In this section, we develop the proof of Lemma 4.1, which gives matching upper and lower bounds on the one-step change in energy under (OMWU Dual). We first restate the lemma, before giving a more detailed overview of the proof:

Lemma 4.1 (Energy Dissipation). *For* (OMWU Dual) *with η satisfying Assumption 1, for all $t \ge 1$:*

$$-\tfrac{5}{4} \cdot \eta^2 \|J \nabla F(z_t)\|_{z_t}^2 \;\le\; F(z_{t+1}) - F(z_t) \;\le\; -\tfrac{1}{20} \cdot \eta^2 \|J \nabla F(z_t)\|_{z_t}^2 \; .$$

Remark F.1 (Constant Stepsize). Note that, for the sake of presentation, we do not attempt to optimize the constraint $0 < \eta \le \frac{1}{4(54\sigma_{\max}+9)}$ from Assumption 1. Up to the leading constants, the conclusion of Lemma 4.1 should hold under even larger absolute constant stepsizes.

Simplifying Notation. For readability, we write $\sigma = \sigma_{\max} = \|J\|_2$ to denote the spectral norm of J. By Proposition B.8, recall that $\|J\|_2 = \|A\|_2$.

F.1 Proof of Lemma 4.1

The proof follows two steps, both using several intermediate results. Here, we will state the conclusions of these results and derive their consequences for proving Lemma 4.1. We defer their full statements and proofs to the subsequent subsections.

1. Exact expansion of change in energy:

The starting point is to take a first-order Taylor expansion of $F(z_{t+1})$ around z_t. Using the structure of the (OMWU Dual) update, as well as the Hessian-based approximation of gradient differences from Proposition E.4, we derive an exact expansion of $\Delta F(z_t) = F(z_{t+1}) - F(z_t)$, the one-step change in energy over the dual OMWU iterates. Specifically, we prove the following:

Proposition F.3 (One-step expansion). *For every $t \geq 1$, there exists $\mathcal{E}_t \in \mathbb{R}$ such that*

$$\Delta F(z_t) = -\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 + D_F(z_{t+1}, z_t) + \eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle + \eta^2\mathcal{E}_t \ . \tag{53}$$

In particular, the error term $\mathcal{E}_t$ is comprised of four subterms that are defined explicitly in (58). The proof of Proposition F.3 is in Section F.2.

2. Controlling the error terms:

Notice that the first term of (53) is exactly the desired dissipation term $-\eta^2\|J\nabla F(z_t)\|_{z_t}^2$ (up to a constant factor) appearing in the statement of the lemma. Thus, the remaining technical challenge is to ensure the final three terms of (53) also scale at most like $O(\eta^2\|J\nabla F(z_t)\|_{z_t}^2)$, with a constant factor that can be made less than 1 via a sufficiently small stepsize η.

For this task, we distinguish between the middle two terms of (53), which we refer to as the *quadratic error terms*, and the final term of (53), which we refer to as the *cubic error terms*. We derive the requisite bounds on these two sets of terms using the local seminorm machinery developed in Section E.3, which hold due to the (LHS) property of the energy Hessian. In more detail:

Bounds on quadratic error terms. Under a sufficiently small stepsize, we establish:

Proposition F.7 (Bound on Bregman divergence term). *Fix $L \geq 21$. For $\eta \leq \frac{1}{L\sigma}$ and all $t \geq 1$:*

$$D_F(z_{t+1}, z_t) \leq \left(\tfrac{18}{25} + \phi(L)\right) \cdot \eta^2\|J\nabla F(z_t)\|_{z_t}^2 \ , \tag{54}$$

where $\phi(L) = \frac{36\exp(6/L)}{25L}$ is a strictly decreasing function of L that can be made arbitrarily small. In particular, it holds that $\frac{18}{25} + \phi(L) \leq \frac{149}{200}$, when $L \geq 72$.

We obtain a similar bound for the inner product involving the Hessian remainder term:

Proposition F.8 (Bound on inner product term). *Fix $L \geq 21$. For $\eta \leq \frac{1}{L\sigma}$ and all $t \geq 1$:*

$$\left|\eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle\right| \leq \psi(L) \cdot \eta^2\|J\nabla F(z_t)\|_{z_t}^2 \ , \tag{55}$$

where $\psi(L) = \frac{54}{5L} + \frac{108\exp(6/L)}{5L^2}$ is a strictly decreasing function of L that can be made arbitrarily small. In particular, it holds that $\psi(L) \leq \frac{31}{200}$ for all $L \geq 72$.

The proofs of Propositions F.7 and F.8 are given in Section F.3.2.

Bounds on cubic error terms. For the cubic error terms $\eta^2\mathcal{E}_t$, with $\mathcal{E}_t$ as defined in (58), we prove:

Proposition F.9 (Bound on cubic error terms). *Let $B = 16\sigma$. Then for $\eta \leq \frac{1}{21\sigma}$ and $t \geq 1$:*

$$\eta^2|\mathcal{E}_t| \leq B \cdot \eta^3\|J\nabla F(z_t)\|_{z_t}^2 \ . \tag{56}$$

The proof of Proposition F.9 is in Section F.3.3.

Combining the pieces for the upper bound.

For the upper bound of Lemma 4.1, first note that if $\sigma = 0$, the iterates are constant and the claim is immediate. Otherwise, Assumption 1 implies $\eta\sigma \leq 1/216$. Applying the quadratic bounds with $L = 216$ gives $18/25 + \phi(216) \leq 3/4$ and $\psi(216) \leq 1/16$, using $e^{1/36} \leq 36/35$; also $\eta B \leq 16/216 < 1/12$. Then applying Propositions F.7, F.8 and F.9, the error terms of (53) can be collectively bounded as

$$D_F(z_{t+1}, z_t) + \eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle + \eta^2\mathcal{E}_t \leq \left(\tfrac{3}{4} + \tfrac{1}{16} + \tfrac{1}{12}\right) \cdot \eta^2\|J\nabla F(z_t)\|_{z_t}^2 \ . \tag{57}$$

Then together with the expansion of (53), we conclude for $\Delta F(z_t) = F(z_{t+1}) - F(z_t)$ that

$$\begin{aligned}\Delta F(z_t) &\leq -\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 + \left(\tfrac{3}{4} + \tfrac{1}{16} + \tfrac{1}{12}\right) \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \\ &\leq -\tfrac{1}{20}\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \ ,\end{aligned}$$

where the final inequality uses $3/4 + 1/16 + 1/12 = 43/48 < 19/20$. This yields the upper bound statement of the lemma.

Combining the pieces for the lower bound.

For the lower bound on $\Delta F(z_t)$, note that $D_F(z_{t+1}, z_t) \geq 0$ by convexity of F. Thus again under Assumption 1, we apply Propositions F.8 and F.9 to bound

$$-\eta |\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle| - \eta^2 |\mathcal{E}_t| \geq -\left(\tfrac{31}{200} + \tfrac{1}{12}\right) \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \ .$$

Then again substituting this bound into (53), we obtain the matching lower bound

$$\begin{aligned}\Delta F(z_t) &\geq -\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 - \eta |\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle| - \eta^2 |\mathcal{E}_t| \\ &\geq -\left(1 + \tfrac{31}{200} + \tfrac{1}{12}\right) \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \\ &\geq -\left(1 + \tfrac{1}{4}\right) \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 = -\tfrac{5}{4} \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \ ,\end{aligned}$$

which completes the proof of the lemma. ■

Remark F.2 (Generality of Lemma 4.1)**.** We note that the proof of Lemma 4.1 only requires that (1) $J = -J^\top$ is a skew-symmetric matrix and (2) the function F satisfies the local Hessian stability property (LHS), together with the local seminorm bounds and uniform increment bound used below. While this latter property is established specifically for the log-sum-exp energy F in Proposition 4.3, Lemma 4.1 holds for any other function satisfying these same estimates.

Organization of remaining subsections. The remainder of this section is organized as follows:

- **Section F.2** gives the full proof of the exact one-step expansion of $F(z_{t+1}) - F(z_t)$.
- **Section F.3** gives the proofs of Propositions F.7, F.8, and F.9 for bounding the error terms. This also involves first introducing several technical lemmas related to the stability of the OMWU iterates (see Section F.3.1)
- **Section F.4** gives the proof of the core stability property of Proposition F.6.
- **Section F.5** gives a useful bound on the initial change in energy $F(z_1) - F(z_0)$ during the first step of the algorithm (which contains no optimistic correction term).

F.2 Expanding the One-Step Change in Energy

Proposition F.3. *Let $\{z_t\}$ denote the iterates of (OMWU Dual), with $z_{-1} = z_0$. Then for any $t \geq 1$:*

$$F(z_{t+1}) - F(z_t) = -\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 + D_F(z_{t+1}, z_t) + \eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle + \eta^2 \mathcal{E}_t \ ,$$

where $\mathcal{E}_t := \mathcal{E}_{t,1} + \mathcal{E}_{t,2} + \mathcal{E}_{t,3} + \mathcal{E}_{t,4}$ for

$$\begin{cases} \mathcal{E}_{t,1} := \langle J\nabla F(z_t), J\nabla^2 F(z_t)(z_t - z_{t-1})\rangle_{z_t} \\ \mathcal{E}_{t,2} := \langle J\nabla F(z_t), JG_F(z_t, z_{t-1})\rangle_{z_t} \\ \mathcal{E}_{t,3} := -\langle J\nabla F(z_t), J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2})\rangle_{z_t} \\ \mathcal{E}_{t,4} := -\langle J\nabla F(z_t), JG_F(z_{t-1}, z_{t-2})\rangle_{z_t} \ . \end{cases} \tag{58}$$

Proof. Let $\Delta F(z_t) := F(z_{t+1}) - F(z_t)$. By definition of the Bregman divergence D_F, recall that:

$$\Delta F(z_t) = \langle \nabla F(z_t), z_{t+1} - z_t\rangle + D_F(z_{t+1}, z_t) \ . \tag{59}$$

Thus, to obtain the statement of the proposition, our task is to expand the first-order term $\langle \nabla F(z_t), z_{t+1} - z_t\rangle$ in (59). For this, by the (OMWU Dual) update rule, observe that we can write

$$z_{t+1} - z_t = -\eta J\nabla F(z_t) - \eta J\big(\nabla F(z_t) - \nabla F(z_{t-1})\big) \ . \tag{60}$$

Now for $k \geq 0$, let us write $\mathcal{M}_k := \nabla F(z_k) - \nabla F(z_{k-1})$; in particular, $\mathcal{M}_0 = 0$. Then observe by substituting (60) into the first-order term of (59) that

$$\begin{aligned}\big\langle \nabla F(z_t), z_{t+1} - z_t \big\rangle &= -\eta \big\langle \nabla F(z_t), J\nabla F(z_t) \big\rangle - \eta \big\langle \nabla F(z_t), J\mathcal{M}_t \big\rangle \\ &= -\eta \big\langle \nabla F(z_t), J\mathcal{M}_t \big\rangle \\ &= \eta \big\langle J\nabla F(z_t), \mathcal{M}_t \big\rangle \,, \end{aligned} \tag{61}$$

where the second and third equalities are due to the skew-symmetry of $J = -J^\top$. We can then further expand the term $\mathcal{M}_t$ using the Hessian-based approximation of gradient differences from Proposition E.4. Specifically, applying the proposition yields

$$\mathcal{M}_t = \nabla F(z_t) - \nabla F(z_{t-1}) = \nabla^2 F(z_t)(z_t - z_{t-1}) + G_F(z_t, z_{t-1}) \,. \tag{62}$$

We can further use the update rule of (OMWU Dual) at time t to write

$$\begin{aligned} z_t - z_{t-1} &= -2\eta J\nabla F(z_{t-1}) + \eta J\nabla F(z_{t-2}) \\ &= -\eta J\nabla F(z_t) + \eta J(\nabla F(z_t) - \nabla F(z_{t-1})) - \eta J(\nabla F(z_{t-1}) - \nabla F(z_{t-2})) \\ &= -\eta J\nabla F(z_t) + \eta J\mathcal{M}_t - \eta J\mathcal{M}_{t-1} \,. \end{aligned} \tag{63}$$

Similar to $\mathcal{M}_t$, we can apply Proposition E.4 at time $t-1$ to obtain

$$\mathcal{M}_{t-1} = \nabla F(z_{t-1}) - \nabla F(z_{t-2}) = \nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2}) + G_F(z_{t-1}, z_{t-2}) \,. \tag{64}$$

Then combining (62), (63), and (64), we can write

$$\begin{aligned} z_t - z_{t-1} = -\eta J\nabla F(z_t) &+ \eta J\nabla^2 F(z_t)(z_t - z_{t-1}) \\ &+ \eta J G_F(z_t, z_{t-1}) \\ &- \eta J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2}) \\ &- \eta J G_F(z_{t-1}, z_{t-2}) \,, \end{aligned}$$

from which it follows by (62) that

$$\begin{aligned} \mathcal{M}_t = -\eta\nabla^2 F(z_t)J\nabla F(z_t) &+ \eta\nabla^2 F(z_t)J\nabla^2 F(z_t)(z_t - z_{t-1}) \\ &+ \eta\nabla^2 F(z_t)J G_F(z_t, z_{t-1}) \\ &- \eta\nabla^2 F(z_t)J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2}) \\ &- \eta\nabla^2 F(z_t)J G_F(z_{t-1}, z_{t-2}) \\ &+ G_F(z_t, z_{t-1}) \,. \end{aligned} \tag{65}$$

Now recall from (59) and (61) that $\Delta F(z_t) = D_F(z_{t+1}, z_t) + \eta\langle J\nabla F(z_t), \mathcal{M}_t\rangle$. Then plugging in the expansion of $\mathcal{M}_t$ from (65) and using the local seminorm and degenerate inner product notation $\|\cdot\|_{z_t}$ and $\langle\cdot,\cdot\rangle_{z_t}$, we conclude:

$$\begin{aligned} \Delta F(z_t) = -\eta^2\|J\nabla F(z_t)\|_{z_t}^2 + D_F(z_{t+1}, z_t) &+ \eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle \\ &+ \eta^2\langle J\nabla F(z_t), J\nabla^2 F(z_t)(z_t - z_{t-1})\rangle_{z_t} \\ &+ \eta^2\langle J\nabla F(z_t), J G_F(z_t, z_{t-1})\rangle_{z_t} \\ &- \eta^2\langle J\nabla F(z_t), J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2})\rangle_{z_t} \\ &- \eta^2\langle J\nabla F(z_t), J G_F(z_{t-1}, z_{t-2})\rangle_{z_t} \,. \end{aligned} \tag{66}$$

Here, up to the η^2 factor, the final four terms of (66) are the quantities $\mathcal{E}_{t,1}$, $\mathcal{E}_{t,2}$, $\mathcal{E}_{t,3}$, and $\mathcal{E}_{t,4}$ defined in the statement of the proposition, which concludes the proof. At $t = 1$, the last two terms vanish because $z_{-1} = z_0$, so the same expansion applies. ■

F.3 Controlling the Error Terms

F.3.1 Helper Propositions on Stability of OMWU Iterates

Here, we establish several key properties related to the stability of the OMWU dual iterates: First, the following proposition gives a bound in ℓ_∞ norm on the distance between consecutive iterates:

Proposition F.4. *Let $\{z_t\}$ be iterates of* (OMWU Dual)*. Then for all $t \geq 0$:*

$$\|z_{t+1} - z_t\|_\infty \leq 3\eta\sigma_{\max} \;\; \textit{and} \;\; \|z_{t+2} - z_t\|_\infty \leq 6\eta\sigma_{\max} \; .$$

Proof. Recall from the definition of (OMWU Dual) that for $t \geq 0$:

$$z_{t+1} - z_t = -2\eta J\nabla F(z_t) + \eta J\nabla F(z_{t-1}) \; .$$

As $\nabla F(z) \in \mathrm{ri}(\mathcal{W})$ for all $z \in \mathbb{R}^{m+n}$, we have by definition of J that $\|J\nabla F(z)\|_\infty \leq a_{\max}$, where $a_{\max} = \max_{(i,j)\in[m]\times[n]} |A(i,j)|$. By Proposition B.7, $a_{\max} \leq \sigma = \|J\|_2$, and thus by the triangle inequality $\|z_{t+1} - z_t\|_\infty \leq 3\eta\sigma$. The second claim for $\|z_{t+2} - z_t\|_\infty$ then follows similarly. ■

The next proposition relates the local seminorm of the dual increments $z_{t+1} - z_t$ to the magnitude of the payoff vector $J\nabla F(z_t)$ in local seminorm:

Proposition F.5. *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with $\eta \leq \frac{1}{21\sigma}$. Then for $t \geq 1$:*

$$\|z_{t+1} - z_t\|_{z_t} \leq \tfrac{6}{5} \cdot \eta \|J\nabla F(z_t)\|_{z_t} \; .$$

The proof of Proposition F.5, which is central to establishing the bounds on the error terms, relies further on the following stability property of the OMWU iterates:

Proposition F.6. *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with $\eta \leq \frac{1}{21\sigma}$. Then for all $t \geq 1$, the following inequalities hold:*

$$\textit{(i)} \;\; \|z_{t-1} - z_{t-2}\|_{z_t} \leq 2 \cdot \|z_t - z_{t-1}\|_{z_t} \;\; \textit{and} \;\; \textit{(ii)} \;\; \|z_t - z_{t-1}\|_{z_t} \leq 3 \cdot \|z_{t+1} - z_t\|_{z_t} \; .$$

We defer the proof of Proposition F.6 to Section F.4, and we assume it is true for now. We then proceed with the proof of Proposition F.5:

Proof of Proposition F.5 First, using the OMWU update rule at time $t+1$ and taking local seminorms, we have

$$\begin{aligned} \|z_{t+1} - z_t\|_{z_t} &= \eta \|J\nabla F(z_t) + J(\nabla F(z_t) - \nabla F(z_{t-1}))\|_{z_t} \\ &\leq \eta \|J\nabla F(z_t)\|_{z_t} + \eta \|J(\nabla F(z_t) - \nabla F(z_{t-1}))\|_{z_t} \; . \end{aligned} \tag{67}$$

Thus to prove the proposition, it suffices to control the second term of (67) in terms of $\|J\nabla F(z_t)\|_{z_t}$.

Controlling the second term of (67). Recall from Proposition E.4 that

$$\nabla F(z_t) - \nabla F(z_{t-1}) = \nabla^2 F(z_t)(z_t - z_{t-1}) + G_F(z_t, z_{t-1}) \; .$$

Thus it follows by taking norms and applying the triangle inequality that

$$\begin{aligned} \|J(\nabla F(z_t) - \nabla F(z_{t-1}))\|_{z_t} &= \|J\nabla^2 F(z_t)(z_t - z_{t-1}) + JG_F(z_t, z_{t-1})\|_{z_t} \\ &\leq \|J\nabla^2 F(z_t)(z_t - z_{t-1})\|_{z_t} + \|JG_F(z_t, z_{t-1})\|_{z_t} \; . \end{aligned} \tag{68}$$

To control the two terms of (68), the key is to upper bound both quantities in terms of $\|z_{t+1} - z_t\|_{z_t}$. For this, we rely on the stability property established in Proposition F.6. Specifically, for the first term of (68) we have

$$\begin{aligned} \|J\nabla^2 F(z_t)(z_t - z_{t-1})\|_{z_t} &\leq \|J\|_2 \cdot \|z_t - z_{t-1}\|_{z_t} && \text{(by Lemma E.7)} \\ &\leq 3\sigma \cdot \|z_{t+1} - z_t\|_{z_t} && \text{(by Proposition F.6) .} \end{aligned} \tag{69}$$

For the second term of (68), recall by Proposition F.4 that $\|z_t - z_{t-1}\|_\infty \leq 3\eta\sigma$. Then using similar calculations as in the proof of Part (c) of Lemma F.11, and also applying Proposition F.6, we have

$$\begin{aligned} \|JG_F(z_t, z_{t-1})\|_{z_t} &\leq \|JG_F(z_t, z_{t-1})\|_2 && \text{(by Proposition E.5)} \\ &\leq \sigma \cdot \|G_F(z_t, z_{t-1})\|_2 && \text{(by definition of } \|J\|_2\text{)} \\ &\leq \sigma \cdot \|G_F(z_t, z_{t-1})\|_{z_t,*} && \text{(by Proposition E.5)} \\ &\leq \sigma r_1 \cdot \|z_t - z_{t-1}\|_{z_t} && \text{(by Lemma E.10)} \\ &\leq 3\sigma r_1 \cdot \|z_{t+1} - z_t\|_{z_t} && \text{(by Proposition F.6) ,} \end{aligned} \tag{70}$$

where for readability we write $r_1 := \nu(6\eta\sigma)$ (for the function ν as defined in (41)). Then combining (69) and (70) and substituting back into (68) yields

$$\|J(\nabla F(z_t) - \nabla F(z_{t-1}))\|_{z_t} \leq 3\sigma(1 + r_1) \cdot \|z_{t+1} - z_t\|_{z_t} \; . \tag{71}$$

Combining the pieces. Combining (71) with (67), we can thus write

$$\|z_{t+1} - z_t\|_{z_t} \le \eta \|J\nabla F(z_t)\|_{z_t} + 3\eta\sigma(1 + r_1) \cdot \|z_{t+1} - z_t\|_{z_t} \,. \tag{72}$$

Using the definition of $r_1 := \nu(6\eta\sigma)$ and the upper bound on $\nu(\cdot)$ from Lemma E.8, it is then straightforward to check that $3\eta\sigma(1 + r_1) \le \frac{1}{6}$ when $\eta \le \frac{1}{21\sigma}$. Thus rearranging (72), we find

$$\|z_{t+1} - z_t\|_{z_t} \le \Big(\frac{1}{1 - 3\eta(1 + r_1)\sigma}\Big) \cdot \eta \|J\nabla F(z_t)\|_{z_t} \le \frac{6}{5} \cdot \eta \|J\nabla F(z_t)\|_{z_t} \,, \tag{73}$$

which concludes the proof. ■

F.3.2 Bounds on Quadratic Error Terms

We bound the Bregman divergence term in local seminorm as follows:

Proposition F.7. *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with $\eta \le \frac{1}{21\sigma}$. Then for $t \ge 1$:*

$$D_F(z_{t+1}, z_t) \le \tfrac{36}{25} \cdot \mu(6\eta\sigma) \cdot \eta^2 \|J\nabla F(z_t)\|^2_{z_t} \,.$$

Moreover, if $\eta \le \frac{1}{L\sigma}$ for $L \ge 72$, then $\frac{36}{25} \cdot \mu(6\eta\sigma) \le \frac{18}{25} + \frac{36\exp(6/L)}{25L} \le \frac{149}{200}$.

Proof. Recall from Proposition F.4 that $\|z_{t+1} - z_t\|_\infty \le 3\eta\sigma$. Then applying the upper bound on $D_F(z_{t+1}, z_t)$ from Lemma E.9, we have

$$D_F(z_{t+1}, z_t) \le \mu(6\eta\sigma) \cdot \|z_{t+1} - z_t\|^2_{z_t} \,. \tag{74}$$

Further using the bound $\|z_{t+1} - z_t\|_{z_t} \le \frac{6}{5} \cdot \eta \|J\nabla F(z_t)\|_{z_t}$ from Proposition F.5 (which holds under the constraint $\eta \le \frac{1}{21\sigma}$), we find

$$D_F(z_{t+1}, z_t) \le \tfrac{36}{25} \cdot \mu(6\eta\sigma) \cdot \eta^2 \|J\nabla F(z_t)\|^2_{z_t} \,.$$

Now using the definition of $\mu(\cdot)$ from (41) and the bounds of Lemma E.8, observe that if $\eta \le 1/(L\sigma)$ for some positive $L > 0$, then

$$\tfrac{36}{25} \cdot \mu(6\eta\sigma) \le \tfrac{18}{25} + \tfrac{36\exp(6/L)}{25L} \,.$$

The right hand side of this term is decreasing with L. It is then straightforward to check that, for $L \ge 72$, then $\frac{36\exp(6/L)}{25L} \le \frac{1}{40}$. Thus for all such $L \ge 72$, it holds that $\frac{18}{25} + \frac{36\exp(6/L)}{25L} \le \frac{149}{200}$. ■

For the inner product involving $G_F(z_t, z_{t-1})$, we derive a similar bound:

Proposition F.8. *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with $\eta \le \frac{1}{21\sigma}$. Then for $t \ge 1$:*

$$\big|\eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle\big| \le \tfrac{18}{5} \cdot \nu(6\eta\sigma) \cdot \eta^2 \|J\nabla F(z_t)\|^2_{z_t} \,.$$

Moreover, if $\eta \le \frac{1}{L\sigma}$ for $L \ge 72$, then $\frac{18}{5} \cdot \nu(6\eta\sigma) \le \frac{54}{5L} + \frac{108\exp(6/L)}{5L^2} \le \frac{31}{200}$.

Proof. Applying the generalized Cauchy-Schwarz inequality with the pair $(\|\cdot\|_{z_t}, \|\cdot\|_{z_t,*})$, we have

$$\big|\eta\langle J\nabla F(z_t), G_F(z_t, z_{t-1})\rangle\big| \le \eta \|J\nabla F(z_t)\|_{z_t} \cdot \|G_F(z_t, z_{t-1})\|_{z_t,*} \,. \tag{75}$$

Since by Proposition F.4 $\|z_t - z_{t-1}\|_\infty \le 3\eta\sigma$, we apply the dual norm bound of Lemma E.10 on $G_F(z_t, z_{t-1})$ to further write

$$\begin{aligned}\|G_F(z_t, z_{t-1})\|_{z_t,*} \le \nu(6\eta\sigma) \cdot \|z_t - z_{t-1}\|_{z_t} &\le 3 \cdot \nu(6\eta\sigma) \cdot \|z_{t+1} - z_t\|_{z_t} \\ &\le \tfrac{18}{5} \cdot \nu(6\eta\sigma) \cdot \eta \|J\nabla F(z_t)\|_{z_t} \,. \end{aligned} \tag{76}$$

Here, the second inequality comes from applying the stability bound of Proposition F.6, and the final inequality comes from Proposition F.5 (which holds under the constraint $\eta \le 1/(21\sigma)$). Combining expressions (76) and (75) then yields the first claim.

For the second claim, we have from the definition of $\nu(\cdot)$ from (41) and the bounds of Lemma E.8, that, if $\eta \le 1/(L\sigma)$ for some positive $L > 0$, then

$$\tfrac{18}{5} \cdot \nu(6\eta\sigma) \le \tfrac{54}{5L} + \tfrac{108\exp(6/L)}{5L^2} \,.$$

The right hand side of this term is decreasing with L, and it is straightforward to check that this term is at most $\frac{31}{200}$ for all $L \ge 72$. ■

F.3.3 Bounds on Cubic Error Terms

Proposition F.9. *Let $\{z_t\}$ be the iterates of (OMWU Dual) with $\eta \le \frac{1}{21\sigma}$, where $\sigma = \|J\|_2$, and set $\kappa = \exp(3\eta\sigma)$. Recall the terms $\mathcal{E}_{t,1}$, $\mathcal{E}_{t,2}$, $\mathcal{E}_{t,3}$, and $\mathcal{E}_{t,4}$ from (58) in Proposition F.3 are given by:*

$$\begin{cases} \mathcal{E}_{t,1} := \langle J\nabla F(z_t), J\nabla^2 F(z_t)(z_t - z_{t-1})\rangle_{z_t} \\ \mathcal{E}_{t,2} := \langle J\nabla F(z_t), JG_F(z_t, z_{t-1})\rangle_{z_t} \\ \mathcal{E}_{t,3} := -\langle J\nabla F(z_t), J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2})\rangle_{z_t} \\ \mathcal{E}_{t,4} := -\langle J\nabla F(z_t), JG_F(z_{t-1}, z_{t-2})\rangle_{z_t} \ . \end{cases}$$

Then for every $t \ge 1$, the following bounds hold:

$$|\mathcal{E}_{t,1}| \le \frac{18}{5}\sigma\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 \ , \quad |\mathcal{E}_{t,2}| \le \frac{3}{5}\sigma\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 \ ,$$
$$|\mathcal{E}_{t,3}| \le \frac{36}{5}\sigma\kappa^2\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 \ , \quad |\mathcal{E}_{t,4}| \le \frac{6}{5}\sigma\kappa\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 \ .$$

Moreover, recalling that $\mathcal{E}_t := \mathcal{E}_{t,1} + \mathcal{E}_{t,2} + \mathcal{E}_{t,3} + \mathcal{E}_{t,4}$, then $\eta^2|\mathcal{E}_t| \le 16\sigma\cdot\eta^3\|J\nabla F(z_t)\|_{z_t}^2$.

Proof. For each term, we use a combination of (i) (generalized) Cauchy-Schwarz (ii) the bounds from Lemma E.7 and Lemma E.10 stemming from the (LHS) property and (iii) the local stability property of OMWU from Proposition F.6 to show the magnitude of the term scales at most like a constant multiple of $\|J\nabla F(z_t)\|_{z_t}\cdot\|z_{t+1} - z_t\|_{z_t}$. From there, we apply to each case Proposition F.5, which establishes $\|z_{t+1} - z_t\|_{z_t} = O(\eta\|J\nabla F(z_t)\|_{z_t})$. In more detail:

- **Term $\mathcal{E}_{t,1}$.** Using Cauchy-Schwarz, we have

$$\begin{aligned} |\mathcal{E}_{t,1}| &= \left|\left\langle J\nabla F(z_t), J\nabla^2 F(z_t)(z_t - z_{t-1})\right\rangle_{z_t}\right| \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\|J\nabla^2 F(z_t)(z_t - z_{t-1})\|_{z_t} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\|J\|_2\cdot\|z_t - z_{t-1}\|_{z_t} && \text{(by Lemma E.7)} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma\cdot 3\cdot\|z_{t+1} - z_t\|_{z_t} && \text{(by Proposition F.6)} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma\cdot 3\cdot\tfrac{6}{5}\eta\cdot\|J\nabla F(z_t)\|_{z_t} && \text{(by Proposition F.5)} \\ &= \tfrac{18}{5}\sigma\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 \ . \end{aligned}$$

- **Term $\mathcal{E}_{t,2}$.** Recall by Proposition F.4 that $\|z_t - z_{t-1}\|_\infty \le 3\eta\sigma$. Let $r_1 := \nu(6\eta\sigma)$ for $\nu(\cdot)$ as defined in (41). Using the upper bound on $\nu(\cdot)$ from Lemma E.8, it is straightforward to check that $r_1 \le \frac{1}{6}$ when $\eta \le \frac{1}{21\sigma}$. By Proposition E.5, $\|JG_F(z_t, z_{t-1})\|_{z_t} \le \sigma\|G_F(z_t, z_{t-1})\|_{z_t,*}$. Then using Cauchy-Schwarz, we can write

$$\begin{aligned} |\mathcal{E}_{t,2}| &= \left|\left\langle J\nabla F(z_t), JG_F(z_t, z_{t-1})\right\rangle_{z_t}\right| \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma\|G_F(z_t, z_{t-1})\|_{z_t,*} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma r_1\cdot\|z_t - z_{t-1}\|_{z_t} && \text{(by Lemma E.10)} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma r_1\cdot 3\|z_{t+1} - z_t\|_{z_t} && \text{(by Proposition F.6)} \\ &\le \|J\nabla F(z_t)\|_{z_t}\cdot\sigma r_1\cdot 3\cdot\tfrac{6}{5}\cdot\eta\|J\nabla F(z_t)\|_{z_t} && \text{(by Proposition F.5)} \\ &\le \tfrac{3}{5}\sigma\cdot\eta\|J\nabla F(z_t)\|_{z_t}^2 && \text{(using } r_1 \le 1/6) \ . \end{aligned}$$

- **Term $\mathcal{E}_{t,3}$.** We repeat the steps for $\mathcal{E}_{t,1}$, transferring norms between z_t and z_{t-1}. Lemmas E.6 and E.7 give

$$\begin{aligned} &\|J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2})\|_{z_t} \\ &\quad\le \kappa\|J\nabla^2 F(z_{t-1})(z_{t-1} - z_{t-2})\|_{z_{t-1}} \\ &\quad\le \sigma\kappa\|z_{t-1} - z_{t-2}\|_{z_{t-1}} \le \sigma\kappa^2\|z_{t-1} - z_{t-2}\|_{z_t}. \end{aligned}$$

  Moreover, Proposition F.6 gives $\|z_{t-1} - z_{t-2}\|_{z_t} \le 6\|z_{t+1} - z_t\|_{z_t}$. Cauchy-Schwarz and Proposition F.5 therefore yield

$$\begin{aligned} |\mathcal{E}_{t,3}| &\le \sigma\kappa^2\|J\nabla F(z_t)\|_{z_t}\|z_{t-1} - z_{t-2}\|_{z_t} \\ &\le \tfrac{36}{5}\sigma\kappa^2\eta\|J\nabla F(z_t)\|_{z_t}^2. \end{aligned}$$

- **Term $\mathcal{E}_{t,4}$.** We repeat the steps for $\mathcal{E}_{t,2}$, now using the remainder bound at z_{t-1} and one norm transfer:

$$\begin{aligned}\|JG_F(z_{t-1}, z_{t-2})\|_{z_t} &\le \sigma\|G_F(z_{t-1}, z_{t-2})\|_{z_{t-1},*} \\ &\le \sigma r_1\|z_{t-1} - z_{t-2}\|_{z_{t-1}} \\ &\le \sigma r_1\kappa\|z_{t-1} - z_{t-2}\|_{z_t}.\end{aligned}$$

  Thus Cauchy-Schwarz and the same increment bound give

$$|\mathcal{E}_{t,4}| \le \tfrac{6}{5}\sigma\kappa\eta\|J\nabla F(z_t)\|_{z_t}^2.$$

  At $t = 1$, both the third and fourth terms vanish because $z_{-1} = z_0$, so their bounds also hold at this endpoint.

Since $\kappa \le e^{1/7} \le 7/6$, summing the coefficients gives $18/5+3/5+(36/5)\kappa^2+(6/5)\kappa \le 77/5 < 16$. Thus

$$|\mathcal{E}_t| \le |\mathcal{E}_{t,1}| + |\mathcal{E}_{t,2}| + |\mathcal{E}_{t,3}| + |\mathcal{E}_{t,4}| \le 16\sigma \cdot \eta\|J\nabla F(z_t)\|_{z_t}^2 ,$$

which concludes the proof. ■

F.4 Stability of Iterates Under Local Seminorm

In this section, we give the proof of Proposition F.6, which establishes the stability of the dual OMWU iterates in local seminorm. Restated here:

Proposition F.6. *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with $\eta \le \frac{1}{21\sigma}$. Then for all $t \ge 1$, the following inequalities hold:*

$$\textit{(i)}\ \ \|z_{t-1} - z_{t-2}\|_{z_t} \le 2 \cdot \|z_t - z_{t-1}\|_{z_t}\ \ \textit{and}\ \ \textit{(ii)}\ \ \|z_t - z_{t-1}\|_{z_t} \le 3 \cdot \|z_{t+1} - z_t\|_{z_t} .$$

F.4.1 Proof of Proposition F.6

We start by proving inequality (i), whose $t = 1$ case is immediate from $z_{-1} = z_0$. Inequality (ii) will follow from (i).

Shorthand notation. First, for readability, we define several pieces of shorthand notation that will be used throughout:

$$\begin{cases}\Delta_t := z_t - z_{t-1} \\ g_t := \nabla F(z_t) - \nabla F(z_{t-1})\end{cases} \qquad \text{for all } t \ge 0. \tag{77}$$

Here $z_{-1} = z_0$, so $\Delta_0 = g_0 = 0$. Moreover, for readability, we write $\sigma = \sigma_{\max} := \|J\|_2$ throughout to denote the spectral norm of J.

Definition of constants. Recall by Proposition F.4 that the dual iterates of OMWU satisfy the ℓ_∞ bounds of $\|z_t - z_{t-1}\|_\infty \le 3\eta\sigma$ and $\|z_t - z_{t-2}\|_\infty \le 6\eta\sigma$ for all $t \ge 2$. Thus, to streamline the proof of the proposition, we also define and use the following constants. These capture the multiplicative dependencies in applying the norm transfer inequality of Lemma E.6 and the Hessian remainder bound of Lemma E.10 with either $\alpha_1 := 3\eta\sigma$ or $\alpha_2 := 6\eta\sigma$:

$$\begin{aligned}\kappa_1 &:= \exp(3\eta\sigma) && \text{(for applying Lemma E.6 with } \alpha_1) \\ \kappa_2 &:= \exp(6\eta\sigma) && \text{(for applying Lemma E.6 with } \alpha_2) \\ r_1 &:= \nu(6\eta\sigma) && \text{(for applying Lemma E.10 with } \alpha_1) \\ r_2 &:= \nu(12\eta\sigma) && \text{(for applying Lemma E.10 with } \alpha_2) \\ \sigma &:= \|J\|_2 && \text{(spectral norm of } J)\end{aligned} \tag{78}$$

Here, recall that $\nu(x) = \frac{\exp(x)-x-1}{x}$ is the scalar function defined in (41). We then further define the constants B_1, B_2, and B by:

$$B_1 := \sigma\kappa_1^2 + \sigma r_1\kappa_1 \qquad B_2 := \sigma\kappa_2^2 + \sigma r_2\kappa_2 \qquad B := 2(B_1 + B_2) . \tag{79}$$

We will make use of the following properties which are straightforward to verify:

Proposition F.10. *Consider the constants defined in expressions* (78) *and* (79)*. Then the following inequalities hold for all $0 < \eta \le \frac{1}{21\sigma}$:*

$$\textit{(a)}\ \ 0 \le B_1 \le B_2 \qquad \textit{(b)}\ \ \eta \le \tfrac{1}{4\kappa_1^2(B_1+B_2)} \qquad \textit{(c)}\ \ \tfrac{1}{2\kappa_1^2} \le \tfrac{1}{2} \qquad \textit{(d)}\ \ \kappa_1^2 \le \tfrac{3}{2} .$$

Proof of Proposition F.10. We prove each of the relationships separately:

Part (a). Observe first that each term in the definition of B_1 and B_2, is non-negative, and thus $B_1, B_2 \geq 0$. Moreover, as $r_1 := \nu(6\eta\sigma)$ and $r_2 = \nu(12\eta\sigma)$, it follows by definition of ν (from (41)) that $r_1 \leq r_2$. Similarly, $\kappa_1 = \exp(3\eta\sigma) \leq \exp(6\eta\sigma) = \kappa_2$. Thus $B_1 \leq B_2$.

Part (b). Suppose $\eta \leq \frac{1}{L\sigma}$ for some $L \geq 1$. Then to verify $\eta \leq 1/(4\kappa_1^2(B_1 + B_2))$ under such a constraint on η, it is sufficient to check that $4\kappa_1^2(B_1 + B_2) \leq L\sigma$. For this, we start by simplifying $4\kappa_1^2(B_1 + B_2)$. Using the definition of B_1 and B_2 we have:

$$4\kappa_1^2(B_1 + B_2) = 4\sigma\kappa_1^2(\kappa_1^2 + r_1\kappa_1 + \kappa_2^2 + r_2\kappa_2) = 4\sigma(\kappa_1^4 + \kappa_1^3 r_1 + \kappa_1^2\kappa_2^2 + r_2\kappa_1^2\kappa_2) \; . \tag{80}$$

Note further that for $\eta \leq \frac{1}{L\sigma}$, we have $6\eta\sigma \leq \frac{6}{L}$ and $12\eta\sigma \leq \frac{12}{L}$. Thus by definition of r_1 and r_2, and using the bounds on $\nu(\cdot)$ from Lemma E.8, we can further write:

$$\begin{aligned} r_1 = \nu(6\eta\sigma) &\leq \tfrac{6\eta\sigma}{2} + \tfrac{1}{6} \cdot (\tfrac{6}{L})^2 \exp(\tfrac{6}{L}) \leq \tfrac{3}{L} + \tfrac{6}{L^2}\exp(\tfrac{6}{L}) \\ \text{and} \quad r_2 = \nu(12\eta\sigma) &\leq \tfrac{12\eta\sigma}{2} + \tfrac{1}{6} \cdot (\tfrac{12}{L})^2 \exp(\frac{12}{L}) \leq \tfrac{6}{L} + \tfrac{24}{L^2}\exp(\tfrac{12}{L}) \; . \end{aligned} \tag{81}$$

Combining expressions (81) and (80) and also using the definitions from (78) of $\kappa_1 = \exp(3\eta\sigma)$ and $\kappa_2 = \exp(6\eta\sigma)$, we can then write:

$$4\kappa_1^2(B_1+B_2) \leq 4\sigma \cdot \big(\exp(\tfrac{12}{L}) + \tfrac{3}{L}\exp(\tfrac{9}{L}) + \tfrac{6}{L^2}\exp(\tfrac{15}{L}) + \exp(\tfrac{18}{L}) + \tfrac{6}{L}\exp(\tfrac{12}{L}) + \tfrac{24}{L^2}\exp(\tfrac{24}{L})\big) \; . \tag{82}$$

Note that each term in the right hand side of (82) is decreasing in L. It is then straightforward to verify that the right hand side is at most $L\sigma$ for all $L \geq 21$, which corresponds to the constraint on L in the assumptions of the proposition and concludes the proof of part (b).

Part (c). Observe that $\frac{1}{2\kappa_1^2} \leq \frac{1}{2} \iff \kappa_1^2 \geq 1$. By definition, we have $\kappa_1^2 = \exp(6\eta\sigma) \geq 1$ for all $\eta > 0$, which proves the claim.

Part (d). By definition $\kappa_1^2 = \exp(6\eta\sigma)$. Thus for $\eta \leq \frac{1}{21\sigma}$, it follows that $\kappa_1^2 \leq \exp(\frac{6}{21}) \leq \frac{3}{2}$. $\square$

Using this new notation and relationships between constants, we now give the overall template for proving claim (i) of the proposition:

General template for proving (i). Our goal in proving (i) is to establish that

$$\|\Delta_{t-1}\|_{z_t} \leq C \cdot \|\Delta_t\|_{z_t} \tag{83}$$

for $C = 2$. For this, observe by the triangle inequality that

$$\|\Delta_{t-1}\|_{z_t} = \|\Delta_t - (\Delta_t - \Delta_{t-1})\|_{z_t} \leq \|\Delta_t\|_{z_t} + \|\Delta_t - \Delta_{t-1}\|_{z_t} \; . \tag{84}$$

Now suppose there exists a constant $B < 1/\eta$ such that

$$\|\Delta_t - \Delta_{t-1}\|_{z_t} \leq \eta B \|\Delta_{t-1}\|_{z_t} \; . \tag{♣}$$

Then it will follow from (84) that

$$\|\Delta_{t-1}\|_{z_t} \leq \Big(\frac{1}{1-\eta B}\Big)\|\Delta_t\|_{z_t} \; , \tag{85}$$

which establishes the inequality (i) so long as $\eta B \leq 1/2$. Thus, our goal is to prove the inequality (♣).

Helper inequalities. To prove the inequality (♣), we first establish the following helper inequalities that hold at general time indices. These follow primarily via application of the norm transfer inequality of Lemma E.6 and the dual norm bound on the Hessian remainder terms of Lemma E.10. We state and prove these inequalities in the following lemma:

Lemma F.11. *Let $\{z_t\}$ be the dual iterates of OMWU. Then the following inequalities hold:*

(a) For all $k \geq 2$: $\|J\nabla^2 F(z_{k-1})\Delta_{k-1}\|_{z_k} \leq \sigma\kappa_1^2 \cdot \|\Delta_{k-1}\|_{z_k}$.

(b) For all $k \geq 3$: $\|J\nabla^2 F(z_{k-2})\Delta_{k-2}\|_{z_k} \leq \sigma\kappa_2^2 \cdot \|\Delta_{k-2}\|_{z_k}$.

(c) For all $k \geq 2$: $\|JG_F(z_{k-1}, z_{k-2})\|_{z_k} \leq \sigma r_1\kappa_1 \cdot \|\Delta_{k-1}\|_{z_k}$.

(d) For all $k \geq 3$: $\|JG_F(z_{k-2}, z_{k-3})\|_{z_k} \leq \sigma r_2\kappa_2 \cdot \|\Delta_{k-2}\|_{z_k}$.

Proof of Lemma F.11 We prove each of the inequalities separately:

Part (a). By Proposition F.4, we have $\|z_k - z_{k-1}\|_\infty \le 3\eta\sigma$, and thus we can write

$$\begin{aligned}\|J\nabla^2 F(z_{k-1})\Delta_{k-1}\|_{z_k} &\le \kappa_1 \cdot \|J\nabla^2 F(z_{k-1})\Delta_{k-1}\|_{z_{k-1}} && \text{(by Lemma E.6)}\\ &\le \sigma\kappa_1 \cdot \|\Delta_{k-1}\|_{z_{k-1}} && \text{(by Lemma E.7)}\\ &\le \sigma\kappa_1^2 \cdot \|\Delta_{k-1}\|_{z_k} && \text{(by Lemma E.6)}\,.\end{aligned}$$

Part (b). Follows identically to that of (a), but now using $\|z_k - z_{k-2}\|_\infty \le 6\eta\sigma$.

Part (c). Using Proposition E.5, we have

$$\begin{aligned}\|JG_F(z_{k-1}, z_{k-2})\|_{z_k} &\le \|JG_F(z_{k-1}, z_{k-2})\|_2\\ &\le \|J\|_2 \cdot \|G_F(z_{k-1}, z_{k-2})\|_2 \le \sigma \cdot \|G_F(z_{k-1}, z_{k-2})\|_{z_{k-1},*}\,.\end{aligned}$$

Moreover, as $\|z_{k-1} - z_{k-2}\|_\infty \le 3\eta\sigma$, applying Lemma E.10 yields

$$\begin{aligned}\|G_F(z_{k-1}, z_{k-2})\|_{z_{k-1},*} &\le r_1\|z_{k-1} - z_{k-2}\|_{z_{k-1}} && \text{(by Lemma E.10)}\\ &\le r_1\kappa_1\|\Delta_{k-1}\|_{z_k} && \text{(by Lemma E.6)}\,.\end{aligned}$$

Part (d). Follows identically to that of (c), using $\|z_{k-2} - z_{k-3}\|_\infty \le 3\eta\sigma \le 6\eta\sigma$ for the remainder bound and $\|z_k - z_{k-2}\|_\infty \le 6\eta\sigma$ for the norm transfer. □

Proof of Inequality (♣). We prove this inequality by induction on t and with $B := 2(B_1 + B_2)$.

Base case: At time $t = 2$, recall first by the initialization of the dual OMWU iterates that we have $z_1 = z_0 - \eta J\nabla F(z_0)$. It follows that

$$\Delta_2 - \Delta_1 = -\eta J(2\nabla F(z_1) - \nabla F(z_0)) + \eta J\nabla F(z_0) = -2\eta J g_1\,,$$

which means that $\|\Delta_2 - \Delta_1\|_{z_2} \le 2\eta\|Jg_1\|_{z_2}$. Now by the definitions of g_1 and Δ_1 and using the gradient difference expression of Proposition E.4, we have

$$g_1 = \nabla F(z_1) - \nabla F(z_0) = \nabla^2 F(z_1)\Delta_1 + G_F(z_1, z_0)\,. \tag{86}$$

Thus we can write

$$\begin{aligned}\|Jg_1\|_{z_2} &\le \|J\nabla^2 F(z_1)\Delta_1\|_{z_2} + \|JG_F(z_1, z_0)\|_{z_2} &(87)\\ &\le \sigma\kappa_1^2\|\Delta_1\|_{z_2} + \sigma r_1\kappa_1\|\Delta_1\|_{z_2}\,, &(88)\end{aligned}$$

where the final inequality comes from applying Part (a) and Part (c) of Lemma F.11 to the two left-hand terms of (87), respectively. Recalling the definition of B_1 from (79), we thus conclude

$$\|\Delta_2 - \Delta_1\|_{z_2} \le \eta \cdot 2\|Jg_1\|_{z_2} \le \eta \cdot 2B_1\|\Delta_1\|_{z_2} \le \eta \cdot 2(B_1 + B_2)\|\Delta_1\|_{z_2}\,.$$

Here, the inequality $2B_1 \le 2(B_1 + B_2) = B$ comes from Part (a) of Proposition F.10. Thus, (♣) holds for the base case at time $t = 2$.

Inductive step: Fix $t \ge 3$, and suppose that (♣) holds for all $2 \le k \le t - 1$. To show the inequality holds also at time t, we start by repeating similar calculations as in proof of the base case. First, we decompose the difference $\Delta_t - \Delta_{t-1}$ by

$$\begin{aligned}\Delta_t - \Delta_{t-1} &= -\eta J\big(2\nabla F(z_{t-1}) - 2\nabla F(z_{t-2}) - (\nabla F(z_{t-2}) - \nabla F(z_{t-3}))\big)\\ &= -\eta J(2g_{t-1} - g_{t-2})\,.\end{aligned}$$

Taking local seminorms and using the triangle inequality, we can then further write

$$\|\Delta_t - \Delta_{t-1}\|_{z_t} \le 2\eta\|Jg_{t-1}\|_{z_t} + \eta\|Jg_{t-2}\|_{z_t}\,. \tag{89}$$

To bound $\|Jg_{t-1}\|_{z_t}$, using the definitions of g_{t-1} and Δ_{t-1}, and using Proposition E.4 to express the difference of gradients, we have

$$g_{t-1} = \nabla F(z_{t-1}) - \nabla F(z_{t-2}) = \nabla^2 F(z_{t-1})\Delta_{t-1} + G_F(z_{t-1}, z_{t-2})\,.$$

Thus it follows that

$$\begin{aligned}\|Jg_{t-1}\|_{z_t} &\le \|J\nabla^2 F(z_{t-1})\Delta_{t-1}\|_{z_t} + \|JG_F(z_{t-1}, z_{t-2})\|_{z_t} &(90)\\ &\le (\sigma\kappa_1^2 + \sigma r_1\kappa_1) \cdot \|\Delta_{t-1}\|_{z_t} &(91)\\ &= B_1 \cdot \|\Delta_{t-1}\|_{z_t}\,, &(92)\end{aligned}$$

where the final inequality comes from applying Part (a) and Part (c) of Lemma F.11 to the two terms in (90), respectively. Now using an analogous setup for $\|Jg_{t-2}\|_{z_t}$ and applying Part (b) and Part (d) of Lemma F.11, we similarly find

$$\|Jg_{t-2}\|_{z_t} \leq (\sigma\kappa_2^2 + \sigma r_2\kappa_2)\|\Delta_{t-2}\|_{z_t} \,. \tag{93}$$

Our task is now to further derive an upper bound on $\|Jg_{t-2}\|_{z_t}$ in terms of the $\|\Delta_{t-1}\|_{z_t}$. For this, as we assume the inductive hypothesis holds at time $t-1$, we know

$$\|\Delta_{t-1} - \Delta_{t-2}\|_{z_{t-1}} \leq \eta B\|\Delta_{t-2}\|_{z_{t-1}} \,. \tag{94}$$

Thus by two applications of the local-seminorm transfer of Lemma E.6 and using the inductive hypothesis from (94), we can write

$$\|\Delta_{t-1} - \Delta_{t-2}\|_{z_t} \leq \kappa_1 \cdot \|\Delta_{t-1} - \Delta_{t-2}\|_{z_{t-1}} \leq \kappa_1\eta B \cdot \|\Delta_{t-2}\|_{z_{t-1}} \leq \kappa_1^2\eta B \cdot \|\Delta_{t-2}\|_{z_t} \,.$$

Then using the reverse triangle inequality, we have

$$\|\Delta_{t-1}\|_{z_t} \geq \|\Delta_{t-2}\|_{z_t} - \|\Delta_{t-1} - \Delta_{t-2}\|_{z_t} \geq \big(1 - \kappa_1^2\eta B\big) \cdot \|\Delta_{t-2}\|_{z_t} \,. \tag{95}$$

Now recalling $B := 2(B_1 + B_2)$, it follows from Parts (b) and (c) in Proposition F.10 that $\eta\kappa_1^2 B \leq 1/2 < 1$ under the constraint that $\eta \leq 1/(21\sigma)$. Thus it follows that

$$\|\Delta_{t-2}\|_{z_t} \leq \Big(\frac{1}{1 - \kappa_1^2\eta B}\Big) \cdot \|\Delta_{t-1}\|_{z_t} \,. \tag{96}$$

Ultimately, substituting this bound on $\|\Delta_{t-2}\|_{z_t}$ back into (93) and recalling $B_2 := \sigma\kappa_2^2 + \sigma r_2\kappa_2$, we find

$$\|Jg_{t-2}\|_{z_t} \leq \Big(\frac{B_2}{1 - \kappa_1^2\eta B}\Big) \cdot \|\Delta_{t-1}\|_{z_t} \,. \tag{97}$$

Together with the bound on $\|Jg_{t-1}\|_{z_t}$ from (92), we conclude from (89) that

$$\begin{aligned}\|\Delta_t - \Delta_{t-1}\|_{z_t} &\leq 2\eta\|Jg_{t-1}\|_{z_t} + \eta\|Jg_{t-2}\|_{z_t} \\ &\leq \eta\Big(2B_1 + \frac{B_2}{1 - \kappa_1^2\eta B}\Big) \cdot \|\Delta_{t-1}\|_{z_t} \leq \eta \cdot 2(B_1 + B_2) \cdot \|\Delta_{t-1}\|_{z_t} \,.\end{aligned}$$

Here, the final inequality comes from the assumption that $\eta \leq 1/(21\sigma)$ and applying Part (b) of Proposition F.10. Thus we conclude that the inductive step of (♣) also holds for $B := 2(B_1 + B_2)$, which proves the inequality for all $t \geq 2$. Finally, substituting this setting of B into (85) yields

$$\|\Delta_{t-1}\|_{z_t} \leq \big(\tfrac{1}{1-\eta\cdot 2(B_1+B_2)}\big)\|\Delta_t\|_{z_t} \leq 2\|\Delta_t\|_{z_t} \,,$$

where the final inequality is due to $\eta B \leq 1/2$, which holds by Parts (b) and (c) of Proposition F.10. This completes the proof of part (i) of the proposition.

Proof of claim (ii) of main proposition. To prove the second inequality of the proposition, we apply the previous inequality (i) at time $t+1$. This yields

$$\|z_t - z_{t-1}\|_{z_{t+1}} \leq 2\|z_{t+1} - z_t\|_{z_{t+1}} \,. \tag{98}$$

As $\|z_{t+1} - z_t\|_\infty \leq 3\eta a_{\max}$, applying two norm transfers via Lemma E.6 to both sides of (98) yields:

$$\|z_{t+1} - z_t\|_{z_{t+1}} \leq \kappa_1\|z_{t+1} - z_t\|_{z_t} \;\text{ and }\; \|z_t - z_{t-1}\|_{z_{t+1}} \geq \frac{1}{\kappa_1}\|z_t - z_{t-1}\|_{z_t} \,. \tag{99}$$

Combining (98) and (99) then yields

$$\|z_t - z_{t-1}\|_{z_t} \leq 2(\kappa_1)^2 \cdot \|z_{t+1} - z_t\|_{z_t} \leq 3\|z_{t+1} - z_t\|_{z_t} \,,$$

where the final inequality is due to $\kappa_1^2 \leq 3/2$, which comes from Part (d) of Proposition F.10 under the assumption that $\eta \leq 1/(21\sigma)$. This completes the proof of part (ii) of the proposition. ■

F.5 Bound on Initial Change in Energy

By definition of (OMWU Dual), we have at time $t = 1$ that $z_1 = z_0 - \eta J \nabla F(z_0)$. In this initial step, the iterate z_1 is obtained only with a single skew-gradient step, with no optimistic correction term. Thus, the initial change in energy $F(z_1) - F(z_0)$ is nonnegative (although by Lemma 4.1, $F(z_t) - F(z_{t-1}) \leq 0$ for all subsequent $t \geq 2$). In the following lemma, we thus derive a worst-case upper bound on this initial change in energy at the first iteration:

Lemma F.12 (Initial Change in Energy)**.** *Let $\{z_t\}$ be the iterates of* (OMWU Dual) *with η satisfying Assumption 1. Then at time $t = 1$:*

$$F(z_1) - F(z_0) \leq \tfrac{5}{8} \cdot \eta^2 \|J \nabla F(z_0)\|_{z_0}^2 \, .$$

Proof. By a first-order Taylor expansion of $F(z_1)$ around z_0 and using the dual OMWU update rule $z_1 = z_0 - \eta J \nabla F(z_0)$, we have

$$\begin{aligned} F(z_1) - F(z_0) &= \langle \nabla F(z_0), z_1 - z_0 \rangle + D_F(z_1, z_0) \\ &= -\eta \langle \nabla F(z_0), J \nabla F(z_0) \rangle + D_F(z_1, z_0) \\ &= D_F(z_1, z_0) \, . \end{aligned}$$

Here, the final equality comes from the skew-symmetry of $J = -J^\top$. Now recall from Proposition F.4 that $\|z_1 - z_0\|_\infty \leq 3\eta\sigma$. Then applying the bound on D_F in local seminorm from Lemma E.9 and recalling the definition of $\mu(\cdot)$ from (41), we find

$$D_F(z_1, z_0) \leq \mu(6\eta\sigma) \cdot \|z_1 - z_0\|_{z_0}^2 = \mu(6\eta\sigma) \cdot \eta^2 \|J \nabla F(z_0)\|_{z_0}^2 \, .$$

Now if $\eta \leq \frac{1}{L\sigma}$ for any $L > 0$, the bound on $\mu(\cdot)$ from Lemma E.8 further gives

$$\mu(6\eta\sigma) \leq \tfrac{1}{2} + \tfrac{\exp(6/L)}{L} \, .$$

The second term in the above is strictly decreasing for all $L > 0$, and it is straightforward to check that $\frac{1}{2} + \frac{\exp(6/L)}{L} \leq \frac{1}{2} + \frac{1}{8} = \frac{5}{8}$ for all $L \geq 72$. Thus for $\eta \leq \frac{1}{72\sigma}$ (which holds under Assumption 1), we conclude that

$$F(z_1) - F(z_0) \leq \tfrac{5}{8} \cdot \eta^2 \|J \nabla F(z_0)\|_{z_0}^2 \, ,$$

which concludes the proof. ■

G Details on Non-Uniform Skew-Gradient Domination

In this section, we develop the proof of Proposition 4.4, which relates the *dissipation term* $\|J \nabla F(z)\|_z^2$ from Lemma 4.1 to the KL divergence $\mathrm{KL}(w^\star, w)$ via a *skew-gradient domination* inequality. To restate the proposition:

Proposition 4.4 (Non-uniform skew-gradient domination)**.** *Let $A \in \mathbb{R}^{m \times n}$ have a unique and interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$. Fix $z \in \mathcal{Z}$, and let $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Let $p_{\min} = \min_{i \in [m]} p(i)$ and $q_{\min} = \min_{j \in [n]} q(j)$, and define $w_{\min} := \min\{p_{\min}, q_{\min}\}$. Then $\sigma_{\min} > 0$, and moreover*

$$\|J \nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot w_{\min}^2 \cdot \mathrm{KL}(w^\star, w) \, .$$

Organization of Section. The proof of Proposition 4.4 first requires introducing several intermediate results, and thus this section is organized as follows:

- **Section G.1** introduces preliminaries on the spectrum of the energy function hessian $\nabla^2 F$ and on the spectrum of the payoff operator J when restricted to the subspace $\mathcal{S}^\perp$.
- **Section G.2** establishes a dual characterization of the KL $\leq \chi^2$ relationship from Proposition B.5 that bounds the energy suboptimality gap in the effective dual space $\mathcal{Z}$.
- **Section G.3** then gives the proof of Proposition 4.4 using the preliminaries of Sections G.1 and G.2.
- **Section G.4** presents a second proof of the skew-gradient domination inequality, but using a primal perspective (see Proposition G.13). This section also derives a general (but crude) upper bound on the dissipation term $\|J \nabla F(z)\|_z^2$.
- **Section G.5** gives examples that demonstrate the independence of the magnitude of $\sigma_{\min}$ with respect to the location of the interior NE $w^\star$.

G.1 Preliminaries on Restricted Spectral Parameters

G.1.1 Non-uniform local strong convexity in primal norm.

As mentioned in Section 3 (and explained further in Section D.1), the energy function F lacks global strict convexity over $\mathbb{R}^{m+n}$. In particular, Proposition D.3 implies that F is affine over the subspace of constant shifts $\mathcal{S}$, which reflects the fact that $\text{Null}(\nabla^2 F(z)) = \mathcal{S}$ for all $z \in \mathbb{R}^{m+n}$. However, when restricted to the orthogonal complement $\mathcal{S}^\perp$, the energy Hessian $\nabla^2 F(z)$ is positive definite, and F is strictly convex, although the degree of its curvature is highly dependent on the local point z. To quantify this dependency, we consider a local and non-uniform notion of restricted strong convexity that can be defined in terms of the minimum restricted eigenvalue $M(z)$ of $\nabla^2 F(z)$:

Definition G.1 (Minimum restricted eigenvalue)**.** *For $z \in \mathbb{R}^{m+n}$, let $M(z)$ be the scalar*

$$M(z) := \inf_{v \in \mathcal{S}^\perp \setminus \{0\}} \frac{\langle v, \nabla^2 F(z) v \rangle}{\|v\|_2^2} \, .$$

By definition of the local seminorm $\|\cdot\|_z$ and the subspaces $\mathcal{S}$ and $\mathcal{S}^\perp$, we then have the following:

Proposition G.2. *Fix $z \in \mathbb{R}^{m+n}$. Then for any $v \in \mathbb{R}^{m+n}$, it holds that*

$$\|v\|_z^2 \geq M(z) \cdot \|\Pi_{\mathcal{S}^\perp}(v)\|_2^2 \, .$$

Proof. Recall that every $v \in \mathbb{R}^{m+n}$ has a unique orthogonal decomposition with respect to $\mathcal{S}$ and $\mathcal{S}^\perp$, where $v = \Pi_{\mathcal{S}}(v) + \Pi_{\mathcal{S}^\perp}(v)$, where $\Pi_{\mathcal{S}}(\cdot)$ and $\Pi_{\mathcal{S}^\perp}(\cdot)$ denote the orthogonal projections onto $\mathcal{S}$ and $\mathcal{S}^\perp$, respectively. Then using this decomposition and also the definition of the local seminorm $\|\cdot\|_z$, we have

$$\begin{aligned} \|v\|_z^2 = \langle v, \nabla^2 F(z) v \rangle &= \big\langle (\Pi_{\mathcal{S}}(v) + \Pi_{\mathcal{S}^\perp}(v)), \nabla^2 F(z)(\Pi_{\mathcal{S}}(v) + \Pi_{\mathcal{S}^\perp}(v)) \big\rangle \\ &= \big\langle \Pi_{\mathcal{S}^\perp}(v), \nabla^2 F(z) \Pi_{\mathcal{S}^\perp}(v) \big\rangle \\ &\geq M(z) \cdot \|\Pi_{\mathcal{S}^\perp}(v)\|_2^2 \, . \end{aligned}$$

Here, the second equality comes from the fact that $\text{Null}(\nabla^2 F(z)) = \mathcal{S}$ (part (iii) of Proposition E.1), and the symmetry of the $\nabla^2 F(z)$, and the final inequality comes from the definition of $M(z)$ in Definition G.1, which yields the claim. ■

Moreover, using the characterization of the energy Hessian from Proposition E.1, the minimum restricted eigenvalue $M(z)$ is bounded below by the minimum coordinate of the primal variable $w = \nabla F(z)$ corresponding to the dual variable z:

Proposition G.3 (Lower bound on minimum restricted eigenvalue)**.** *Fix $z \in \mathbb{R}^{m+n}$, and let $w = (p, q) = \nabla F(z) \in \text{ri}(\mathcal{W})$. Let $p_{\min} = \min_{i \in [m]} p(i)$, $q_{\min} = \min_{j \in [n]} q(j)$, and $w_{\min} := \min\{p_{\min}, q_{\min}\}$. Then $M(z) \geq w_{\min}$.*

Proof. To prove the proposition, it suffices to show that, when restricted to $\mathcal{S}^\perp$, the minimum eigenvalue of $\nabla^2 F(z)$ is at least $w_{\min}$. Now recall by part (i) of Proposition E.1 that

$$\nabla^2 F(z) = \begin{pmatrix} \text{Diag}(p) - pp^\top & 0 \\ 0 & \text{Diag}(q) - qq^\top \end{pmatrix} \, .$$

Due to its separable block structure, we will prove the corresponding block-wise inequalities: e.g., when restricted to $\mathbf{1}_m^\perp$, the minimum eigenvalue of $\nabla^2 F_x(z) = \text{Diag}(p) - pp^\top$ is at least $p_{\min}$, and when restricted to $\mathbf{1}_n^\perp$, the minimum eigenvalue of $\nabla^2 F_y(z) = \text{Diag}(q) - qq^\top$ is at least $q_{\min}$.

We start with the claim for $\nabla^2 F_x(z)$. For this, let $\{\lambda_i^p\}$ denote the eigenvalues of $\nabla^2 F_x(z)$, and observe by part (iv) of Proposition E.1 that $0 \leq \lambda_i^p \leq 1$ for all $i \in [m]$. Moreover, $\nabla^2 F_x(z)$ has minimum eigenvalue 0 corresponding to eigenvector $\mathbf{1}_m$, since:

$$\nabla^2 F_x(z) \mathbf{1}_m = \text{Diag}(p) \mathbf{1}_m - p \langle p, \mathbf{1}_m \rangle = p - p = 0 \, .$$

Thus, because $\nabla^2 F_x(z)$ is symmetric, all other non-zero eigenvalues correspond to eigenvectors orthogonal to $\mathbf{1}_m$ (i.e., belonging to $\mathbf{1}_m^\perp$). Therefore, it suffices to prove that all positive eigenvalues

are bounded below by $p_{\min}$. For this, let $v \in \mathbf{1}_m^\perp \setminus \{0\}$ and $\lambda > 0$ be such an eigenvector-eigenvalue pair, which means $\mathrm{Diag}(p)v - p\langle p, v\rangle = \lambda v$. By rearranging, we must have at each coordinate

$$(p(i) - \lambda) \cdot v(i) = p(i) \cdot \langle p, v\rangle \quad \text{for all } i \in [m]. \tag{100}$$

Now assume by way of contradiction that $\lambda < p_{\min}$. This implies that $p(i) - \lambda > 0$ for all coordinates $i \in [m]$. Then by rearranging (100) and summing over all $i \in [m]$, we find

$$\langle \mathbf{1}_m, v\rangle = \textstyle\sum_{i=1}^m v(i) = \langle p, v\rangle \cdot \sum_{i=1}^m \frac{p(i)}{p(i)-\lambda} \ . \tag{101}$$

As $v \in \mathbf{1}_m^\perp$ by assumption, we then must have $\langle \mathbf{1}_m, v\rangle = 0$, meaning expression (101) evaluates to 0. However, the sum $\sum_{i=1}^m \frac{p(i)}{p(i)-\lambda} > 0$ by the assumption that $\lambda < p_{\min}$ (meaning $p(i) - \lambda > 0$ for all i) and the fact that $p \in \mathrm{ri}(\Delta_m)$. Moreover, if $\langle p, v\rangle = 0$, the equality in expression (100) further forces $v(i) = 0$ at every coordinate $i \in [m]$, which contradicts the assumption that $v \neq 0$. Thus the equality in (101) cannot be satisfied when $\lambda < p_{\min}$, and we conclude $\lambda \geq p_{\min}$.

Repeating an identical argument for $\nabla^2 F_y(z)$ similarly implies all positive eigenvalues corresponding to eigenvectors in $\mathbf{1}_n^\perp$ are bounded below by $q_{\min}$. By the block structure of $\nabla^2 F(z)$, it follows that all positive eigenvalues of $\nabla^2 F(z)$ are bounded below by $w_{\min} = \min(p_{\min}, q_{\min})$. Further using that $\mathcal{S}^\perp = \mathbf{1}_m^\perp \times \mathbf{1}_n^\perp$, it then follows by definition that $M(z) \geq w_{\min}$, which completes the proof. ■

G.1.2 Non-uniform local smoothness in dual norm.

We also establish a local notion of smoothness in the dual norm $\|\cdot\|_{z,*}$ when restricted to $\mathcal{S}^\perp$. For this, we define maximum restricted eigenvalue $L(z)$ of the inverse energy Hessian as follows:

Definition G.4 (Maximum restricted eigenvalue of inverse energy Hessian)**.** *For $z \in \mathbb{R}^{m+n}$, let $L(z)$ be the scalar*

$$L(z) := \sup_{v \in \mathcal{S}^\perp \setminus \{0\}} \frac{\langle v, (\nabla^2 F(z))^{-1} v\rangle}{\|v\|_2^2} \ .$$

Using the definition of the local dual norm $\|\cdot\|_{z,*}$, the following proposition is then immediate:

Proposition G.5. *Fix $z \in \mathbb{R}^{m+n}$. Then for any $v \in \mathcal{S}^\perp$, it holds that $\|v\|_{z,*}^2 \leq L(z) \cdot \|v\|_2^2$.*

Moreover, Proposition G.3 gives an upper bound on $L(z)$ in terms of the minimum coordinate of the primal variable w corresponding to z:

Proposition G.6. *Fix $z \in \mathbb{R}^{m+n}$, and let $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Let $p_{\min} = \min_{i\in[m]} p(i)$ and $q_{\min} = \min_{j\in[n]} q(j)$, and define $w_{\min} := \min(p_{\min}, q_{\min})$. Then $L(z) \leq \frac{1}{w_{\min}}$.*

Proof. Note in general that if $\nabla^2 F(z) \succeq cI$ (over the restricted subspace) for some $c > 0$, then $(\nabla^2 F(z))^{-1} \preceq \frac{1}{c} I$. Recall that in Proposition G.3, we established $\nabla^2 F(z) \succeq w_{\min} I$ over $\mathcal{S}^\perp$, and thus we have $(\nabla^2 F(z))^{-1} \preceq \frac{1}{w_{\min}} I$ over this subspace. The claim now follows from the definition of $L(z)$. ■

G.1.3 Minimum restricted singular value of J.

For the game's block skew-symmetric matrix payoff matrix $J \in \mathbb{R}^{(m+n)\times(m+n)}$ (as defined in (3)), recall that we define $\sigma_{\min} := \sigma_{\min}(J, \mathcal{S}^\perp)$ to be the minimum singular value of J when restricted to $\mathcal{S}^\perp$. This definition uses the projection operator $\Pi_{\mathcal{S}^\perp}$ onto $\mathcal{S}^\perp$, which we first explicitly define:

Definition G.7 (Projection onto $\mathcal{S}^\perp$)**.** *Let $\Pi_{\mathcal{S}^\perp} : \mathbb{R}^{m+n} \to \mathcal{S}^\perp$ be the orthogonal projection operator onto $\mathcal{S}^\perp$. Specifically, $\Pi_{\mathcal{S}^\perp}$ is given by the matrix*

$$\Pi_{\mathcal{S}^\perp} := \begin{pmatrix} \Pi_{\mathbf{1}_m^\perp} & 0 \\ 0 & \Pi_{\mathbf{1}_n^\perp} \end{pmatrix} := \begin{pmatrix} I_m - \frac{1}{m}\mathbf{1}_m\mathbf{1}_m^\top & 0 \\ 0 & I_n - \frac{1}{n}\mathbf{1}_n\mathbf{1}_n^\top \end{pmatrix} \ .$$

Here, $\Pi_{\mathbf{1}_m^\perp}$ and $\Pi_{\mathbf{1}_n^\perp}$ are the blockwise projection matrices onto $\mathbf{1}_m^\perp$ and $\mathbf{1}_n^\perp$, respectively. By slight abuse of notation, for $v \in \mathbb{R}^{m+n}$, we write $\Pi_{\mathcal{S}^\perp}(v)$ to denote the matrix-vector product $\Pi_{\mathcal{S}^\perp} v$.

Then recall from Section 2 that we define $\sigma_{\min}$ as follows:

Definition G.8 (Minimum restricted singular value of J)**.** *Let $J \in \mathbb{R}^{(m+n)\times(m+n)}$ be the matrix from (3). Then let $\sigma_{\min}$ denote the minimum singular value of J restricted to $\mathcal{S}^\perp$:*

$$\sigma_{\min} := \sigma_{\min}(J, \mathcal{S}^\perp) = \inf_{v \in \mathcal{S}^\perp \setminus \{0\}} \frac{\|\Pi_{\mathcal{S}^\perp}(Jv)\|_2}{\|v\|_2} \, .$$

When A has a unique and interior Nash equilibrium $w^\star$, then $\sigma_{\min} > 0$ is strictly positive:

Proposition G.9. *Let $A \in \mathbb{R}^{m\times n}$ have a unique interior Nash equilibrium $w^\star \in \mathrm{ri}(\mathcal{W})$, let J be the corresponding block skew-symmetric matrix from (3), and let $\sigma_{\min}$ be the minimum restricted singular value of J from Definition G.8. Then $\sigma_{\min} > 0$.*

Proof. First, note by definition of $\sigma_{\min}$ that we can equivalently write

$$\sigma_{\min} = \inf_{v \in \mathcal{S}^\perp \setminus \{0\}, \|v\|_2 = 1} \|\Pi_{\mathcal{S}^\perp}(Jv)\|_2 \, .$$

It is straightforward to see that $\sigma_{\min} = 0$ if and only if $\mathrm{Null}(\Pi_{\mathcal{S}^\perp} J)$ and $\mathcal{S}^\perp$ have a non-trivial intersection. Thus our goal is to establish that, when A has a unique interior Nash equilibrium, then

$$\mathrm{Null}(\Pi_{\mathcal{S}^\perp} J) \cap \mathcal{S}^\perp = \{0\} \, . \tag{102}$$

For this, suppose $v \in \mathrm{Null}(\Pi_{\mathcal{S}^\perp} J) \cap \mathcal{S}^\perp$. This implies that the following two properties hold:

(i) As $v \in \mathcal{S}^\perp$, then $\langle v, s \rangle = 0$ for all $s \in \mathcal{S}$.

(ii) As $v \in \mathrm{Null}(\Pi_{\mathcal{S}^\perp} J)$, then by linearity $0 = \Pi_{\mathcal{S}^\perp} Jv = \Pi_{\mathcal{S}^\perp}(Jv)$, which means $Jv \in \mathcal{S}$.

Recall that we assume A has a unique and interior NE. Then applying Lemma D.8, we have that $v \in \mathcal{S}^\perp$ and $Jv \in \mathcal{S}$ implies $v = 0$. Thus we conclude that if $v \in \mathrm{Null}(\Pi_{\mathcal{S}^\perp} J) \cap \mathcal{S}^\perp$, then $v = 0$. This establishes (102) and concludes the proof. ■

G.2 Properties of Energy Suboptimality Gap

Recall that Proposition D.11 established that, over the effective dual space $\mathcal{Z}$, and under the assumption of a unique and interior NE, then for $z \in \mathcal{Z}$ and $w = \nabla F(z)$:

$$\mathrm{KL}(w^\star, w) = F(z) - \min_{z' \in \mathcal{Z}} F(z') = F(z) - F(z^\star) \, , \tag{103}$$

where $z^\star \in \mathcal{Z}$ is a point such that $\nabla F(z^\star) = w^\star$ (and such a point $z^\star$ exists due to Proposition D.9). In the following proposition, using the primal relationship $\mathrm{KL}(w^\star, w) \le \chi^2(w^\star, w)$, we then further bound the energy suboptimality gap $F(z) - F(z^\star)$ in terms of the dual analogue of $\chi^2(w^\star, w)$. Formally, we have the following relationship:

Lemma G.10 (Dual KL vs. χ^2 relationship)**.** *Let $A \in \mathbb{R}^{m\times n}$ have unique interior Nash equilibrium $w^\star \in \mathrm{ri}(\mathcal{W})$, and let $z^\star \in \mathcal{Z}$ such that $\nabla F(z^\star) = w^\star$. Fix $z \in \mathcal{Z}$, and let $w = \nabla F(z)$. Then:*

$$F(z) - F(z^\star) \le \|\nabla F(z) - \nabla F(z^\star)\|_{z,*}^2 \, .$$

Proof. First, we have from (103) that $F(z) - F(z^\star) = \mathrm{KL}(w^\star, w)$. Then by Proposition B.5:

$$F(z) - F(z^\star) = \mathrm{KL}(w^\star, w) \le \chi^2(w^\star, w) = \sum\nolimits_{i=1}^{m+n} \frac{(w(i) - w^\star(i))^2}{w(i)} \, . \tag{104}$$

By definition of the dual local norm $\|\cdot\|_{z,*}^2$, we have for $v \in \mathcal{S}^\perp$ that

$$\|v\|_{z,*}^2 = \langle v, (\nabla^2 F(z))^{-1} v \rangle = \sum\nolimits_{i=1}^{m+n} \frac{(v(i))^2}{(\nabla F(z))(i)} \, . \tag{105}$$

In the final equality, we use the equality of quadratic forms from Proposition E.2, together with $\nabla^2 R(w) = \mathrm{Diag}(1/w)$. As $\nabla F(z) - \nabla F(z^\star) = w - w^\star \in \mathcal{S}^\perp$, substituting $v = \nabla F(z) - \nabla F(z^\star)$ in (105) then exactly yields the right hand side of (104). ■

G.3 Proof of Proposition 4.4

We now give the proof of Proposition 4.4 through the lens of non-uniform *skew-gradient domination*. Recall that we have $z \in \mathcal{Z}$ and $w = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. The proof uses the following four steps:

1. Non-uniform restricted strong convexity.

Using the definition of the minimum restricted eigenvalue $M(z)$ from Definition G.1, the inequality of Proposition G.2, and the lower bound on $M(z)$ from Proposition G.3, we have

$$\|J\nabla F(z)\|_z^2 \geq M(z) \cdot \|\Pi_{\mathcal{S}^\perp}\big(J\nabla F(z)\big)\|_2^2 \geq w_{\min} \cdot \|\Pi_{\mathcal{S}^\perp}\big(J\nabla F(z)\big)\|_2^2 \,.$$

2. Invariance to Nash and restricted spectrum of J.

Let $z^\star \in \mathcal{Z}$ such that $\nabla F(z^\star) = w^\star$ (recall that such a $z^\star$ exists due to Proposition D.9). Now by Proposition B.1, we have $J\nabla F(z^\star) = Jw^\star \in \mathcal{S}$. Thus by definition of the projection $\Pi_{\mathcal{S}^\perp}$, we have $\Pi_{\mathcal{S}^\perp}(J\nabla F(z^\star)) = 0$, and therefore by linearity

$$\Pi_{\mathcal{S}^\perp}(J\nabla F(z)) = \Pi_{\mathcal{S}^\perp}(J(\nabla F(z) - \nabla F(z^\star))) \,.$$

Moreover, as $w, w^\star \in \mathcal{W}$, it follows that $\langle w - w^\star, s\rangle = 0$ for all $s \in \mathcal{S}$, and thus $w - w^\star = \nabla F(z) - \nabla F(z^\star) \in \mathcal{S}^\perp$. Combining these pieces and using the definition of the minimum restricted singular value $\sigma_{\min}$ from Definition G.8, we can then further write:

$$\|\Pi_{\mathcal{S}^\perp}\big(J\nabla F(z)\big)\|_2^2 = \|\Pi_{\mathcal{S}^\perp}\big(J(\nabla F(z) - \nabla F(z^\star))\big)\|_2^2 \geq \sigma_{\min}^2 \cdot \|\nabla F(z) - \nabla F(z^\star)\|_2^2 \,.$$

Together with the inequality of Step 1, this means

$$\|J\nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot w_{\min} \cdot \|\nabla F(z) - \nabla F(z^\star)\|_2^2 \,. \tag{106}$$

3. Bounding the dual suboptimality gap.

We now derive the following lower bound on $\|\nabla F(z) - \nabla F(z^\star)\|_2^2$ in terms of the dual gap $F(z) - F(z^\star)$. We proceed in two steps. First, we have from Lemma G.10 the dual relationship between $\mathrm{KL}(w^\star, w)$ and $\chi^2(w^\star, w)$, which gives

$$F(z) - F(z^\star) \leq \|\nabla F(z) - \nabla F(z^\star)\|_{z,*}^2 \,.$$

Next, by the restricted local smoothness properties of Proposition G.5 and Proposition G.6, we further have

$$F(z) - F(z^\star) \leq \|\nabla F(z) - \nabla F(z^\star)\|_{z,*}^2 \leq \frac{1}{w_{\min}} \cdot \|\nabla F(z) - \nabla F(z^\star)\|_2^2 \,.$$

Rearranging then yields

$$\|\nabla F(z) - \nabla F(z^\star)\|_2^2 \geq w_{\min} \cdot \big(F(z) - F(z^\star)\big) \,. \tag{107}$$

4. Equivalence between primal and dual suboptimality gaps.

Combining (106) from Step 2 and (107) from Step 3 yields

$$\|J\nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot w_{\min}^2 \cdot (F(z) - F(z^\star)) \,. \tag{108}$$

This is exactly a non-uniform skew-gradient-domination property for the energy function F. To relate $\|J\nabla F(z)\|_z^2$ back to the primal space, we apply the equivalence of Proposition D.11 (also restated in (103)), which establishes for $z \in \mathcal{Z}$ and $w = \nabla F(z)$ that

$$F(z) - F(z^\star) = F(z) - \min_{z' \in \mathcal{Z}} F(z') = \mathrm{KL}(w^\star, w) \,. \tag{109}$$

Combining (108) with (109) yields the statement of the proposition and concludes the proof. ■

G.4 Primal Viewpoint of Non-Uniform Skew-Gradient Domination

The proof of Proposition 4.4 in Section G.3 is written purely in the dual perspective. In this section, we provide an alternative proof of the inequality from a primal perspective. This primal viewpoint stems from the structure of the energy Hessian $\nabla^2 F(z)$ from part (ii) of Proposition E.1. This implies that the dissipation term $\|J\nabla F(z)\|_z^2$ is the *variance* of the payoff vector $J\nabla F(z)$ under the component-wise pair of distributions $w = (p, q) = \nabla F(z)$.

The proof and slightly refined result via this second perspective is useful in later establishing the uniform best-iterate duality gap bound of Theorem 5.2. We state and prove the result further below in Proposition G.13. However, we first introduce several additional preliminaries used in the proof:

Component-wise restricted spectrum of A. First, we define the following component-wise analogues of the minimum restricted singular value $\sigma_{\min}(J, \mathcal{S}^\perp)$ from Definition G.8. Specifically, we define $\sigma_{\min,m}$ and $\sigma_{\min,n}$ as follows:

$$\begin{aligned} \sigma_{\min,n} &:= \sigma_{\min}(A, \mathbf{1}_n^\perp) = \inf_{v \in \mathbf{1}_n^\perp \setminus \{0\}} \frac{\|\Pi_{\mathbf{1}_m^\perp}(Av)\|_2}{\|v\|_2} \\ \text{and} \quad \sigma_{\min,m} &:= \sigma_{\min}(A^\top, \mathbf{1}_m^\perp) = \inf_{u \in \mathbf{1}_m^\perp \setminus \{0\}} \frac{\|\Pi_{\mathbf{1}_n^\perp}(A^\top u)\|_2}{\|u\|_2} \,. \end{aligned} \tag{110}$$

We have the following relationship between $\sigma_{\min}(J, \mathcal{S}^\perp)$ and the component-wise analogues in (110):

Proposition G.11. *Let $A \in \mathbb{R}^{m \times n}$, and consider the quantities $\sigma_{\min}$ from Definition G.8 and $\sigma_{\min,m}$ and $\sigma_{\min,n}$ from (110). Then $\sigma_{\min} = \min\{\sigma_{\min,m}, \sigma_{\min,n}\}$.*

Proof. Fix $x = (u, v) \in \mathcal{S}^\perp \setminus \{0\}$. By definition of $\mathcal{S}^\perp$, this means $u \in \mathbf{1}_m^\perp$ and $v \in \mathbf{1}_n^\perp$. Moreover, by the structure of $\Pi_{\mathcal{S}^\perp}$ from Definition G.7, we have:

$$\Pi_{\mathcal{S}^\perp}(Jx) = \begin{pmatrix} \Pi_{\mathbf{1}_m^\perp} & 0 \\ 0 & \Pi_{\mathbf{1}_n^\perp} \end{pmatrix} \begin{pmatrix} Av \\ -A^\top u \end{pmatrix} = \begin{pmatrix} \Pi_{\mathbf{1}_m^\perp}(Av) \\ \Pi_{\mathbf{1}_n^\perp}(-A^\top u) \end{pmatrix} \,.$$

Taking squared Euclidean norms, this means that both

$$\|x\|_2^2 = \|u\|_2^2 + \|v\|_2^2 \quad \text{and} \quad \|\Pi_{\mathcal{S}^\perp}(Jx)\|_2^2 = \|\Pi_{\mathbf{1}_m^\perp}(Av)\|_2^2 + \|\Pi_{\mathbf{1}_n^\perp}(A^\top u)\|_2^2 \,.$$

Moreover, by definition of $\sigma_{\min}$, we have

$$\sigma_{\min}^2 = \inf_{x=(u,v) \in \mathcal{S}^\perp \setminus \{0\}} \frac{\|\Pi_{\mathbf{1}_m^\perp}(Av)\|_2^2 + \|\Pi_{\mathbf{1}_n^\perp}(A^\top u)\|_2^2}{\|u\|_2^2 + \|v\|_2^2} \,. \tag{111}$$

Now let $\sigma_* = \min\{\sigma_{\min,m}, \sigma_{\min,n}\}$. To prove the claim, we show in two parts that both $\sigma_{\min} \geq \sigma_*$ and $\sigma_{\min} \leq \sigma_*$. For the first direction, observe by the definitions of $\sigma_{\min,m}$ and $\sigma_{\min,n}$ in (110) that both

$$\begin{aligned} \|\Pi_{\mathbf{1}_m^\perp}(Av)\|_2^2 &\geq \sigma_{\min,n}^2 \cdot \|v\|_2^2 \geq \sigma_*^2 \cdot \|v\|_2^2 \\ \text{and} \quad \|\Pi_{\mathbf{1}_n^\perp}(A^\top u)\|_2^2 &\geq \sigma_{\min,m}^2 \cdot \|u\|_2^2 \geq \sigma_*^2 \cdot \|u\|_2^2 \,, \end{aligned}$$

which implies that

$$\|\Pi_{\mathcal{S}^\perp}(Jx)\|_2^2 \geq \sigma_*^2 \cdot \big(\|u\|_2^2 + \|v\|_2^2\big) = \sigma_*^2 \cdot \|x\|_2^2 \,.$$

Dividing by $\|x\|^2$, taking an infimum over all $\mathcal{S}^\perp \setminus \{0\}$, and taking square roots then yield $\sigma_{\min} \geq \sigma_*$.

For the reverse direction, observe from expression (111) that by further restricting the domain, the infimum is non-decreasing, and thus

$$\sigma_{\min}^2 \leq \inf_{x=(u,v) \in (0, \mathbf{1}_n^\perp) \setminus \{0\}} \frac{\|\Pi_{\mathbf{1}_m^\perp}(Av)\|_2^2 + \|\Pi_{\mathbf{1}_n^\perp}(0)\|_2^2}{\|0\|_2^2 + \|v\|_2^2} = \sigma_{\min,n}^2 \,,$$

where the final equality follows by definition of $\sigma_{\min,n}$ in (110). Taking square roots, we thus have $\sigma_{\min} \leq \sigma_{\min,n}$. By an identical calculation, we also have $\sigma_{\min} \leq \sigma_{\min,m}$. Together, this yields $\sigma_{\min} \leq \min\{\sigma_{\min,n}, \sigma_{\min,m}\}$, which concludes the proof. ■

Variational characterization of variances. We also use the following variational characterization of variance terms $\mathrm{Var}_p(\cdot)$ and $\mathrm{Var}_q(\cdot)$ from (35), which we state and prove here:

Proposition G.12. *For any $v = (v_x, v_y) \in \mathbb{R}^{m+n}$, the following hold:*

$$\begin{aligned} \mathrm{Var}_p(v_x) &= \min_{c \in \mathbb{R}} \textstyle\sum_{i=1}^m p(i) \cdot \big(v_x(i) - c\big)^2 \\ \textit{and} \quad \mathrm{Var}_q(v_y) &= \min_{c \in \mathbb{R}} \textstyle\sum_{j=1}^n q(j) \cdot \big(v_y(j) - c\big)^2 \,. \end{aligned}$$

Proof. We prove the claim for $\mathrm{Var}_p(v_x)$, and the claim for $\mathrm{Var}_q(v_y)$ will follow by identical calculations. For this, define the scalar function $F(c) = \sum_{i=1}^m p(i) \cdot \big(v_x(i) - c\big)^2$. Using the fact that $\sum_{i=1}^m p(i) = 1$, expanding the definition of $F(c)$ yields

$$F(c) = \textstyle\sum_{i=1}^m p(i) \cdot \big((v_x(i))^2 - 2c v_x(i) + c^2\big) = c^2 - 2c\langle p, v_x \rangle + \sum_{i=1}^m p(i) \cdot (v_x(i))^2 \,.$$

Thus F is convex in c, and differentiating shows F is minimized at $c^* = \langle p, v_x \rangle$. It follows that

$$\min_{c\in\mathbb{R}} F(c) = F(c^*) = \sum_{i=1}^m p(i)\cdot(v_x(i))^2 - \big(\langle p, v_x\rangle\big)^2 = \mathrm{Var}_p(v_x)\ ,$$

which concludes the proof. ■

We now state the following proposition, which establishes the non-uniform skew-gradient domination result of Proposition 4.4 from a primal viewpoint:

Proposition G.13 (Primal view of non-uniform skew-gradient domination)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior Nash equilibrium $w^\star = (p^\star, q^\star) \in \mathrm{ri}(\mathcal{W})$. Fix $z \in \mathbb{R}^{m+n}$, and let $w = (p, q) = \nabla F(z)$. Let $p_{\min} = \min_{i\in[m]} p(i)$ and $q_{\min} = \min_{j\in[n]} q(j)$. Let $\sigma_{\min,m}$ and $\sigma_{\min,n}$ be defined as in (110), and let $\sigma_{\min}$ be the value from Definition G.8. Then the following hold:*

(i) $\|J\nabla F(z)\|_z^2 = \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p)$.

(ii) $\mathrm{Var}_p(Aq) \geq \sigma_{\min,n}^2 \cdot p_{\min} \cdot \|q - q^\star\|_2^2$ *and* $\mathrm{Var}_q(A^\top p) \geq \sigma_{\min,m}^2 \cdot q_{\min} \cdot \|p - p^\star\|_2^2$

(iii) $\|p - p^\star\|_2^2 \geq p_{\min} \cdot \mathrm{KL}(p^\star, p)$ *and* $\|q - q^\star\|_2^2 \geq q_{\min} \cdot \mathrm{KL}(q^\star, q)$.

Furthermore, it then holds that:

$$\|J\nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot p_{\min} \cdot q_{\min} \cdot \mathrm{KL}(w^\star, w)\ .$$

G.4.1 Proof of Proposition G.13

Granting parts (i), (ii), and (iii) of the proposition as true for now, the proof of the final statement follows by combining these inequalities. Specifically, we find

$$\begin{aligned}\|J\nabla F(z)\|_z^2 &\geq \min\{\sigma_{\min,m}^2, \sigma_{\min,n}^2\} \cdot p_{\min} \cdot q_{\min} \cdot \big(\mathrm{KL}(p^\star, p) + \mathrm{KL}(q^\star, q)\big) \\ &= \sigma_{\min}^2 \cdot p_{\min} \cdot q_{\min} \cdot \mathrm{KL}(w^\star, w)\ .\end{aligned}$$

Here, the final equality comes from the definition of $\mathrm{KL}(w^\star, w)$, and from Proposition G.11, which established that $\sigma_{\min} = \min\{\sigma_{\min,m}, \sigma_{\min,n}\}$. Thus it remains to prove the first three statements of the proposition, which we do in independent steps here:

Proof of part (i).

Using the definition of J, the relationship $w = (p, q) = \nabla F(z)$, the local seminorm $\|\cdot\|_z$, and the characterization of part (ii) in Proposition E.1, we have

$$\begin{aligned}\|J\nabla F(z)\|_z^2 &= \langle J\nabla F(z), \nabla^2 F(z) J\nabla F(z)\rangle \\ &= \mathrm{Var}_p(Aq) + \mathrm{Var}_q(-A^\top p) = \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p)\ .\end{aligned}$$

Here, the final equality uses the fact that $\mathrm{Var}_q(v) = \mathrm{Var}_q(-v)$ for any $v \in \mathbb{R}^n$.

Proof of part (ii).

Using the variational characterization of variance from Proposition G.12, we have

$$\begin{aligned}\mathrm{Var}_p(Aq) &= \min_{c\in\mathbb{R}} \sum_{i=1}^m p(i)\cdot\big((Aq - c\mathbf{1}_m)(i)\big)^2 \\ &\geq p_{\min}\cdot\min_{c\in\mathbb{R}}\sum_{i=1}^m \big((Aq - c\mathbf{1}_m)(i)\big)^2 = p_{\min}\cdot\|\Pi_{\mathbf{1}_m^\perp}(Aq)\|_2^2\ .\end{aligned} \tag{112}$$

Here, the final equality comes from the definition of the length of the orthogonal projection onto $\mathbf{1}_m^\perp$. Moreover, by Proposition B.1, we have that $Aq^\star = c\cdot\mathbf{1}_m$ for some constant $c \in \mathbb{R}$, and thus $\Pi_{\mathbf{1}_m^\perp}(Aq^\star) = 0$. Thus, it follows from the linearity of $\Pi_{\mathbf{1}_m^\perp}$ that we can write

$$\Pi_{\mathbf{1}_m^\perp}(Aq) = \Pi_{\mathbf{1}_m^\perp}(Aq) - \Pi_{\mathbf{1}_m^\perp}(Aq^\star) = \Pi_{\mathbf{1}_m^\perp}(A(q - q^\star))\ . \tag{113}$$

Observe also that since $q, q^\star \in \Delta_n$, we have $\langle q - q^\star, s\rangle = 0$ for all $s \in \mathrm{Span}(\mathbf{1}_n)$, which means by definition that $q - q^\star \in \mathbf{1}_n^\perp$. Using the definition of $\sigma_{\min,n}$ from (110), it then follows from (113) that

$$\|\Pi_{\mathbf{1}_m^\perp}(Aq)\|_2^2 \geq \sigma_{\min,n}^2 \cdot \|q - q^\star\|_2^2\ . \tag{114}$$

Combining (114) with (112) then yields $\mathrm{Var}_p(Aq) \geq \sigma_{\min,n}^2 \cdot p_{\min} \cdot \|q - q^\star\|_2^2$. Following identical calculations, we similarly conclude $\mathrm{Var}_q(A^\top p) \geq \sigma_{\min,m}^2 \cdot q_{\min} \cdot \|p - p^\star\|_2^2$.

Proof of part (iii).

Using the relationship $\mathrm{KL}(p^\star, p) \leq \chi^2(p^\star, p)$, we can write and further bound

$$\mathrm{KL}(p^\star, p) \leq \chi^2(p^\star, p) = \sum_{i\in[m]} \frac{(p(i) - p^\star(i))^2}{p(i)} \leq \frac{1}{p_{\min}} \cdot \sum_{i\in[m]} (p(i) - p^\star(i))^2 = \frac{1}{p_{\min}} \cdot \|p - p^\star\|_2^2 \,.$$

Rearranging yields $\|p - p^\star\|_2^2 \geq p_{\min} \cdot \mathrm{KL}(p^\star, p)$. Similarly, following identical calculations, we also conclude that $\|q - q^\star\|_2^2 \geq q_{\min} \cdot \mathrm{KL}(q^\star, q)$, which concludes the proof. ■

G.4.2 Uniform Upper Bound on Dissipation Term

Here, we additionally establish the following uniform *upper bound* on the dissipation term $\|J\nabla F(z)\|_z^2$ in terms of the KL divergence to Nash:

Proposition G.14 (Uniform Upper Bound on Dissipation Term)**.** *Fix $A \in \mathbb{R}^{m\times n}$ with unique and interior NE $w^\star = (p^\star, q^\star) \in \mathrm{ri}(\mathcal{W})$. Let $z \in \mathbb{R}^{m+n}$ and $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Recall that $\sigma_{\max} = \|A\|_2$. Then:*

$$\|J\nabla F(z)\|_z^2 \leq 2 \cdot \sigma_{\max}^2 \cdot \mathrm{KL}(w^\star, w) \,.$$

Proof. The proof of the upper bound uses the primal perspective established by Proposition G.13. For this, recall by Part (i) of that proposition that

$$\|J\nabla F(z)\|_z^2 = \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p) \,.$$

Now by Part (i) of Proposition B.1, recall that $Aq^\star = d\mathbf{1}_m$ and $A^\top p^\star = d\mathbf{1}_n$ for some constant $d \in \mathbb{R}$. Then by the variational characterization of variance from Proposition G.12, observe that

$$\begin{aligned}
\mathrm{Var}_p(Aq) &= \min_{c\in\mathbb{R}} \textstyle\sum_{i=1}^m p(i) \cdot \big((Aq - c\mathbf{1}_m)(i)\big)^2 \\
&\leq \textstyle\sum_{i=1}^m p(i) \cdot \big((A(q - q^\star))(i)\big)^2 \\
&\leq \|A(q - q^\star)\|_2^2 \leq \sigma_{\max}^2 \cdot \|q - q^\star\|_2^2 \qquad (115)
\end{aligned}$$

Here, the final inequality follows by definition of $\sigma_{\max} = \|A\|_2$. As $\|u\|_2^2 \leq \|u\|_1^2$ for any $u \in \mathbb{R}^m$, and recalling that $\mathrm{TV}(q^\star, q) = \frac{1}{2}\|q^\star - q\|_1$, we then further have by Pinsker's inequality

$$\|q - q^\star\|_2^2 \leq \|q - q^\star\|_1^2 = 4\mathrm{TV}(q^\star, q)^2 \leq 2\mathrm{KL}(q^\star, q) \,. \qquad (116)$$

Combining expressions (115) and (116) and repeating an identical calculation for $\mathrm{Var}_q(A^\top p)$, we thus conclude that

$$\|J\nabla F(z)\|_z^2 \leq 2\sigma_{\max}^2 \cdot \big(\mathrm{KL}(q^\star, q) + \mathrm{KL}(p^\star, p)\big) = 2 \cdot \sigma_{\max}^2 \cdot \mathrm{KL}(w^\star, w) \,,$$

which yields the desired claim. ■

Note that this upper bound holds uniformly over the dual and primal spaces and does not contain multiplicative factors depending on the local spectrum of $\nabla^2 F(z)$. However, we will later prove in Proposition J.1 a much tighter *non-uniform* upper bound on $\|J\nabla F(z)\|_z^2$ for the 2×2 setting. In particular, this latter bound *does* contain factors depending on the minimum coordinates $p_{\min}, q_{\min}$ that essentially matches the lower bound on $\|J\nabla F(z)\|_z^2$ from Proposition G.13. See Section J.

G.5 Independence Between Restricted Spectrum of J and Location of Nash Equilibrium

In this section, we show that the magnitude of the minimum restricted singular value $\sigma_{\min}$ of J is independent from the location of the game's Nash equilibrium. In particular, we show examples where (i) the interior NE is uniform but the $\sigma_{\min}$ can be arbitrarily small, and (ii) where the interior NE of the game is arbitrarily close to a vertex, but $\sigma_{\min}$ is constant.

In the context of Proposition 4.4, these examples also imply a certain independence between $\sigma_{\min}$ and the local strong convexity parameter $w_{\min}$ of the energy function: even if a primal iterate w is very close to the simplex boundary (which one would expect for the OMWU iterates if the NE $w^\star$ is near the boundary), the value of $\sigma_{\min}$ can still be an absolute constant. For this reason, the result of Proposition 4.4 indicates that the more salient geometric bottleneck in the skew-gradient domination inequality comes from the restricted strong convexity parameter of F, and not from $\sigma_{\min}$.

We proceed with the two examples:

G.5.1 Example: Uniform Nash, but Arbitrarily Small $\sigma_{\min}$.

Consider the scaled Matching Pennies game with payoff matrix:

$$A = \begin{pmatrix} \epsilon & -\epsilon \\ -\epsilon & \epsilon \end{pmatrix} \quad \text{for } 0 < \epsilon \leq 1. \qquad (\epsilon\text{-Scaled Matching Pennies})$$

Proposition G.15 (Minimum Restricted Singular Value of ϵ-Scaled MP)**.** *Fix any $0 < \epsilon \leq 1$, and let $A \in \mathbb{R}^{2\times 2}$ be the (ϵ-Scaled Matching Pennies) payoff matrix. Let $\sigma_{\min} = \min\{\sigma_{\min,m}, \sigma_{\min,n}\}$ be the minimum restricted singular value from Definition G.8 with $\sigma_{\min,m}$ and $\sigma_{\min,n}$ as in (110). Then the following hold:*

(i) The unique Nash equilibrium $w^\star = (p^\star, q^\star)$ of A is given by $p^\star = q^\star = (\frac{1}{2}, \frac{1}{2})$.

(ii) $\sigma_{\min} = 2\epsilon$.

Proof. We prove the two claims of the proposition separately:

Proof of part (i).

Observe by definition of A that

$$A\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = A^\top \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \epsilon & -\epsilon \\ -\epsilon & \epsilon \end{pmatrix}\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(\epsilon - \epsilon) \\ \frac{1}{2}(-\epsilon + \epsilon) \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} .$$

By Part (ii) of Proposition B.1, this means $w^\star = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is a Nash equilibrium of A. Since $A \in \mathbb{R}^{2\times 2}$ is nonconstant and $w^\star$ is interior, we further have by Part (ii) of Proposition I.3 that $w^\star$ is unique.

Proof of part (ii).

Let $\mathbf{1} = (1, 1)$. Observe that in this two-dimensional setting, the orthogonal subspace $\mathbf{1}^\perp$ is given by vectors $c \cdot (1, -1)$ for $c \in \mathbb{R}$. Thus, fix any $c \in \mathbb{R} \setminus \{0\}$, and let $v = c(1, -1)$. We can compute

$$Av = \begin{pmatrix} \epsilon & -\epsilon \\ -\epsilon & \epsilon \end{pmatrix}\begin{pmatrix} c \\ -c \end{pmatrix} = \begin{pmatrix} 2c\epsilon \\ -2c\epsilon \end{pmatrix} = (2c\epsilon) \cdot \begin{pmatrix} 1 \\ -1 \end{pmatrix} .$$

Observe by definition that $Av \in \mathbf{1}^\perp$. Thus taking norms, we have

$$\frac{\|\Pi_{\mathbf{1}^\perp}(Av)\|_2}{\|v\|_2} = \frac{\|(Av)\|_2}{\|v\|_2} = \frac{2|c|\epsilon \cdot \sqrt{2}}{|c|\sqrt{2}} = 2\epsilon . \qquad (117)$$

As (117) holds for any $c \in \mathbb{R} \setminus \{0\}$, and since $A = A^\top$, it follows that $\sigma_{\min,n} = \sigma_{\min,m} = 2\epsilon$. ■

G.5.2 Example: Vanishing Minimum Nash Coordinate, but Constant $\sigma_{\min}$.

Consider the payoff matrix $A_\delta \in \mathbb{R}^{2\times 2}$ given by

$$A_\delta = \begin{pmatrix} \delta^2 & -\delta(1-\delta) \\ -\delta(1-\delta) & (1-\delta)^2 \end{pmatrix} \quad \text{for } \delta \in (0, 0.5]. \qquad (118)$$

Proposition G.16 (Minimum Restricted Singular Value of A_δ)**.** *Fix $\delta \in (0, 0.5]$, and let $A_\delta \in \mathbb{R}^{2\times 2}$ be the payoff matrix from (118). Let $\sigma_{\min} = \min\{\sigma_{\min,m}, \sigma_{\min,n}\}$ be the minimum restricted singular value from Definition G.8 with $\sigma_{\min,m}$ and $\sigma_{\min,n}$ as in (110). Then the following hold:*

(i) The unique Nash equilibrium $w^\star = (p^\star, q^\star)$ of A_δ is given by $p^\star = q^\star = (1-\delta, \delta)$.

(ii) $\sigma_{\min} = \frac{1}{2}$.

Proof. We again prove the two claims of the proposition separately.

Proof of part (i).

Observe by definition of $A_\delta = A_\delta^\top$ that

$$\begin{aligned} A_\delta\begin{pmatrix} 1-\delta \\ \delta \end{pmatrix} = A_\delta^\top \begin{pmatrix} 1-\delta \\ \delta \end{pmatrix} &= \begin{pmatrix} \delta^2 & -\delta(1-\delta) \\ -\delta(1-\delta) & (1-\delta)^2 \end{pmatrix}\begin{pmatrix} 1-\delta \\ \delta \end{pmatrix} \\ &= \begin{pmatrix} \delta^2(1-\delta) - \delta^2(1-\delta) \\ -\delta(1-\delta)^2 + \delta(1-\delta)^2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} . \end{aligned}$$

By Part (ii) of Proposition B.1 and Part (ii) of Proposition I.3, this means $w^\star = ((1-\delta, \delta), (1-\delta, \delta))$ is the unique interior Nash equilibrium of A_δ.

Proof of part (ii).

Similarly to the proof of part (ii) of Proposition G.15, let $\mathbf{1} = (1,1)$, and thus $v \in \mathbf{1}^\perp$ when $v = c(1,-1)$ for some $c \in \mathbb{R}$. For such v, we can compute using the definition of A_δ that

$$A_\delta v = A_\delta^\top v = \begin{pmatrix} \delta^2 & -\delta(1-\delta) \\ -\delta(1-\delta) & (1-\delta)^2 \end{pmatrix} \begin{pmatrix} c \\ -c \end{pmatrix} = c \begin{pmatrix} \delta \\ \delta - 1 \end{pmatrix} .$$

Moreover, recalling from Definition G.7 that $\Pi_{\mathbf{1}^\perp} = I - \frac{1}{2}\mathbf{1}\mathbf{1}^\top$, it follows that

$$\Pi_{\mathbf{1}^\perp}(A_\delta v) = c \cdot \begin{pmatrix} \delta - \frac{1}{2}(2\delta - 1) \\ \delta - 1 - \frac{1}{2}(2\delta - 1) \end{pmatrix} = c \begin{pmatrix} \frac{1}{2} \\ -\frac{1}{2} \end{pmatrix} .$$

Taking norms, we find for any $c \in \mathbb{R} \setminus \{0\}$ that

$$\frac{\|\Pi_{\mathbf{1}^\perp}(A_\delta v)\|_2}{\|v\|_2} = \frac{|c|/\sqrt{2}}{|c|\sqrt{2}} = \frac{1}{2} .$$

It then follows by definition that $\sigma_{\min} = \frac{1}{2}$. ■

H Details on Universal Last-Iterate Convergence in KL

This section gives the proofs of Theorem 4.2 and Theorem 4.7, which establish asymptotic last-iterate convergence and a linear last-iterate convergence rate in KL divergence, respectively.

Organization of section. The section is organized as follows:

- **Section H.1** first proves a helper lemma that establishes an upper bound on the initial change in $\mathrm{KL}(w^\star, \cdot)$ over the first step of the algorithm.
- **Section H.2** gives the proof of Theorem 4.2 (asymptotic last-iterate convergence to NE).
- **Section H.3** gives the proof of the Theorem 4.7 (new, linear last-iterate convergence rate in KL).
- **Section H.4** gives a brief comparison with the prior result of Wei et al. (2021).
- **Section H.5** presents several numerical simulations.

H.1 Initial Change in KL Divergence

We start by proving the following lemma:

Lemma H.1 (Bound on Initial Change in KL)**.** *Fix $A \in \mathbb{R}^{m\times n}$ with unique and interior NE $w^\star \in \mathrm{ri}(\mathcal{W})$. Let $\{w_t\}$ denote the iterates of (OMWU) and let $\{z_t\}$ denote the iterates of (OMWU Dual) with η satisfying Assumption 1. Recall $\sigma_{\max} = \|A\|_2$. Then $\mathrm{KL}(w^\star, w_1) \le \frac{5}{4} \cdot \mathrm{KL}(w^\star, w_0)$.*

Proof. First, recall by Corollary D.7 the equivalence

$$\mathrm{KL}(w^\star, w_1) - \mathrm{KL}(w^\star, w_0) = F(z_1) - F(z_0) . \tag{119}$$

Then applying the upper bound on $F(z_1) - F(z_0)$ from Lemma F.12 and the uniform upper bound of Proposition G.14, we find

$$F(z_1) - F(z_0) \le \tfrac{5}{8} \cdot \eta^2 \|J\nabla F(z_0)\|_{z_0}^2 \le \tfrac{10}{8} \cdot \sigma_{\max}^2 \cdot \eta^2 \cdot \mathrm{KL}(w^\star, w_0) \le \tfrac{1}{4} \cdot \mathrm{KL}(w^\star, w_0) . \tag{120}$$

Here, the final inequality is due to $\eta \le \frac{1}{4(54\sigma_{\max}+9)} \le \frac{1}{4(54\sigma_{\max})}$ by Assumption 1, which implies crudely that $\frac{10}{8}\sigma_{\max}^2\eta^2 \le \frac{1}{4}$. Then substituting (120) into (119) and rearranging further gives

$$\mathrm{KL}(w^\star, w_1) \le \mathrm{KL}(w^\star, w_0) \cdot \left(1 + \tfrac{1}{4}\right) ,$$

which completes the proof. ■

H.2 Proof of Theorem 4.2 – Asymptotic Last-Iterate Convergence

We now give a proof of asymptotic last-iterate convergence to a unique and interior NE. We first restate the theorem:

Theorem 4.2 (Asymptotic Last-Iterate Convergence)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star$, and let $\{w_t\}$ be the iterates of* (OMWU) *on A with η satisfying Assumption 1. Then the sequence of iterates $\{w_t\}$ converges, and $\lim_{t\to\infty} w_t = w^\star$.*

Proof. Let $\{z_t\}$ denote the dual iterates of OMWU. The proof follows via the following steps:

1. Primal iterates lie in compact sublevelset:

Combining Corollary D.7 and the upper bound of Lemma 4.1, we have for all $t \geq 1$:

$$\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t) \leq -\tfrac{1}{20}\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \leq 0 \ .$$

Moreover, by Lemma H.1, we have $\mathrm{KL}(w^\star, w_1) \leq (5/4)\cdot \mathrm{KL}(w^\star, w_0)$. Together, this means that $\mathrm{KL}(w^\star, w_t) \leq (5/4)\cdot \mathrm{KL}(w^\star, w_0)$ for all $t \geq 0$. Letting $\kappa = (5/4)\cdot \mathrm{KL}(w^\star, w_0)$, this implies all iterates w_t lie in the compact sublevelset

$$\mathcal{U} = \{w \in \mathrm{ri}(\mathcal{W}) \ : \ \mathrm{KL}(w^\star, w) \leq \kappa\} \subset \mathrm{ri}(\mathcal{W}) \ .$$

2. Dissipation term is zero only at Nash:

For $z \in \mathbb{R}^{m+n}$, let $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. Observe by Part (i) of Proposition G.13 that $\|J\nabla F(z)\|_z^2 = \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p)$. By definition of variance (see expression (35)), note that $\mathrm{Var}_p(v) = 0$ and $\mathrm{Var}_q(u) = 0$ if and only if $v = c\mathbf{1}_m$ and $u = d\mathbf{1}_n$ for some $c, d \in \mathbb{R}$. Together, this implies that $\|J\nabla F(z)\|_z^2 = 0$ if and only if both Aq and $A^\top p$ are constant vectors. Due to the uniqueness of $w^\star$, this means by Part (ii) of Proposition B.1 that $\|J\nabla F(z)\|_z^2 = 0$ if and only if $\nabla F(z) = w = w^\star$. Now for $w = (p, q)$ let $D(w) = (\eta^2/20)\cdot(\mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p))$. Then for all $t \geq 1$:

$$\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) \leq -D(w_t) \quad \text{and} \quad \begin{cases} D(w_t) \geq 0 \text{ for all } w_t \\ D(w_t) = 0 \text{ iff } w_t = w^\star \ . \end{cases} \tag{121}$$

3. Subsequence and global convergence:

As all $w_t \in \mathcal{U}$ for the compact set $\mathcal{U}$, we have by the Bolzano-Weierstrass theorem that every infinite subsequence of $\{w_t\}$ has at least one limit point $w_\infty \in \mathcal{U}$. Moreover, by the continuity of $D(\cdot)$, if $w_{t_k} \to w_\infty$ for some subsequence $\{w_{t_k}\}$, then also $D(w_{t_k}) \to D(w_\infty)$. By (121), the sequence $\mathrm{KL}(w^\star, w_t)$ is nonincreasing for $t \geq 1$ and bounded below by zero, so it converges and its successive differences tend to zero. Thus $0 \leq D(w_t) \leq \mathrm{KL}(w^\star, w_t) - \mathrm{KL}(w^\star, w_{t+1})$ implies $D(w_t) \to 0$. By continuity, we obtain $D(w_\infty) = 0$, and hence $w_\infty = w^\star$. Every subsequential limit is therefore $w^\star$, and compactness of $\mathcal{U}$ then implies that the entire sequence $\{w_t\}$ converges to $w^\star$. ■

H.3 Proof of Theorem 4.7 – Universal Last-Iterate Convergence Rate

We first restate the theorem:

Theorem 4.7 (Last-Iterate Convergence Rate in KL)**.** *Let $A \in \mathbb{R}^{m\times n}$ have a unique and interior NE $w^\star = (p^\star, q^\star)$. Let $\{w_t\}$ denote the iterates of* (OMWU) *on A with η satisfying Assumption 1, initialized from $w_0 \in \mathrm{ri}(\mathcal{W})$. Let $\delta_p = \min_{i\in[m]} p^\star(i)$, $\delta_q = \min_{j\in[n]} q^\star(j)$, and $\delta := \min\{\delta_p, \delta_q\}$. For every t, let $p_{t,\min} = \min_{i\in[m]} p_t(i)$, $q_{t,\min} = \min_{j\in[n]} q_t(j)$, and $w_{t,\min} := \min\{p_{t,\min}, q_{t,\min}\}$. Let $\Lambda := \mathrm{CrossEnt}(w^\star, w_0) = \mathrm{KL}(w^\star, w_0) - R(w^\star)$. Then for all $t \geq 1$, the following hold:*

$$\textit{(1)} \quad \mathrm{KL}(w^\star, w_{t+1}) \leq \mathrm{KL}(w^\star, w_t)\cdot\left(1 - \tfrac{1}{20}\cdot\eta^2\sigma_{\min}^2\cdot w_{t,\min}^2\right) \ ,$$

$$\textit{(2)} \quad \mathrm{KL}(w^\star, w_{t+1}) \leq 2\mathrm{KL}(w^\star, w_0)\cdot\exp\left(-\tfrac{1}{20}\cdot\eta^2\sigma_{\min}^2\cdot\exp\left(\tfrac{-3\Lambda}{\delta}\right)\cdot t\right) \ .$$

Proof. The proof follows from the two main previously established properties: energy dissipation (Lemma 4.1) and skew-gradient domination (Proposition 4.4). We proceed to outline these steps.

1. Change in KL is change in energy:

First, let $\{z_t\}$ denote the dual OMWU iterates. Recall from Corollary D.7 that for all $t \geq 0$, we have

$$\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t) \ . \tag{122}$$

2. Upper bound on energy dissipation:

From Lemma 4.1, we have for all $t \geq 1$ under the constraint on η that

$$F(z_{t+1}) - F(z_t) \leq -\tfrac{1}{20} \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \ . \tag{123}$$

3. Non-uniform skew-gradient domination:

By Proposition 4.4, we further have the lower bound

$$\|J\nabla F(z_t)\|_{z_t}^2 \geq \sigma_{\min}^2 \cdot w_{t,\min}^2 \cdot \mathrm{KL}(w^\star, w_t) \ , \tag{124}$$

where $\sigma_{\min} > 0$ under the unique and interior Nash equilibrium assumption.

Combining expressions (122), (123), and (124), we find for all $t \geq 1$

$$\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) \leq -\tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 w_{t,\min}^2 \cdot \mathrm{KL}(w^\star, w_t) \ .$$

Rearranging then yields

$$\mathrm{KL}(w^\star, w_{t+1}) \leq \mathrm{KL}(w^\star, w_t) \cdot \left(1 - \tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 w_{t,\min}^2\right) \ , \tag{125}$$

which yields statement (1) of the theorem.

For the second statement, we require a lower bound on $w_{t,\min}^2$ that holds *uniformly* for all $t \geq 1$. We establish this bound in the following lemma, where $\Lambda = \mathrm{KL}(w^\star, w_0) - R(w^\star)$ is the cross-entropy between $w^\star$ and the initialization $w_0 \in \mathrm{ri}(\mathcal{W})$.

Lemma H.2 (Uniform lower bound on minimum coordinates)**.** *Assume the setting of Theorem 4.7. Then for all $t \geq 1$ it holds that $w_{t,\min}^2 \geq \exp(-3\Lambda/\delta)$.*

We note that the statement and proof of the lemma follow similarly to those of Lemma 19 in Wei et al. (2021), but generalized to hold under arbitrary initializations.

Proof of Lemma H.2 First, we have by Corollary D.7, Lemma 4.1, and Lemma H.1 that, for all $t \geq 1$:

$$\mathrm{KL}(w^\star, w_t) \leq \mathrm{KL}(w^\star, w_1) \leq \left(\tfrac{5}{4}\right) \cdot \mathrm{KL}(w^\star, w_0) \ . \tag{126}$$

Now for $w = (p, q)$, let $\mathrm{Ent}(w) = \mathrm{Ent}_m(p) + \mathrm{Ent}_n(q)$ be the sum of component-wise entropies. Moreover, recall that $R(w) = -\mathrm{Ent}(w)$, and let $\mathrm{CrossEnt}(w^\star, w)$ denote the cross entropy between $w^\star$ and w defined by

$$\mathrm{CrossEnt}(w^\star, w) = \mathrm{KL}(w^\star, w) - R(w^\star) \ .$$

Then using (126), we can write

$$\begin{aligned}
\mathrm{CrossEnt}(w^\star, w_t) &= \mathrm{KL}(w^\star, w_t) - R(w^\star) \\
&\leq \left(\tfrac{5}{4}\right) \cdot \mathrm{KL}(w^\star, w_0) - R(w^\star) &(127)\\
&\leq \left(\tfrac{5}{4}\right) \cdot \left(\mathrm{KL}(w^\star, w_0) - R(w^\star)\right) = \left(\tfrac{5}{4}\right) \cdot \mathrm{CrossEnt}(w^\star, w_0) \ . &(128)
\end{aligned}$$

Here, the second inequality comes from the fact that $-R(w^\star) \geq 0$. Now, let $i_{\min} \in \mathrm{argmin}_{i \in [m]}\, p_t(i)$, and $j_{\min} \in \mathrm{argmin}_{j \in [n]}\, q_t(j)$ denote the indices of p_t and q_t with smallest mass, respectively. Thus $p_{t,\min} = p_t(i_{\min})$ and $q_{t,\min} = q_t(j_{\min})$. Then by definition of $\mathrm{CrossEnt}(w^\star, w_t)$, we further have

$$\begin{aligned}
\mathrm{CrossEnt}(w^\star, w_t) &= \sum_{i=1}^{m} p^\star(i) \cdot \log\left(\frac{1}{p_t(i)}\right) + \sum_{j=1}^{n} q^\star(j) \cdot \log\left(\frac{1}{q_t(j)}\right) \\
&\geq p^\star(i_{\min}) \cdot \log\left(\frac{1}{p_{t,\min}}\right) + q^\star(j_{\min}) \cdot \log\left(\frac{1}{q_{t,\min}}\right) \ . &(129)
\end{aligned}$$

Observe that both terms in (129) are positive. Now let $\Lambda := \mathrm{CrossEnt}(w^\star, w_0)$ and $c = \frac{5}{4}$. Then combining (128) and (129), we have

$$p^\star(i_{\min}) \log\left(\frac{1}{p_{t,\min}}\right) \leq c\Lambda \implies \log\left(\frac{1}{p_{t,\min}}\right) \leq \frac{c\Lambda}{p^\star(i_{\min})} \leq \frac{c\Lambda}{\delta_p} \leq \frac{c\Lambda}{\delta} \ ,$$

where the final two inequalities are due to $p^\star(i_{\min}) \geq \delta_p \geq \delta$. Rearranging, we find

$$p_{t,\min} \geq \exp\left(\tfrac{-c\Lambda}{\delta}\right) \ . \tag{130}$$

By an identical calculation, we also have $q_{t,\min} \geq \exp(\frac{-c\Lambda}{\delta})$. As $w_{t,\min} = \min\{p_{t,\min}, q_{t,\min}\}$, it follows that $w_{t,\min} \geq \exp(\frac{-c\Lambda}{\delta})$, which holds for all $t \geq 1$. Squaring this inequality gives

$$w_{t,\min}^2 \geq \exp(-2c\Lambda/\delta) \geq \exp(-3\Lambda/\delta),$$

where the last inequality uses $2c = 5/2 \leq 3$. This completes the proof. □

We now return to the main proof of Theorem 4.7. Applying Lemma H.2 to expression (125) gives for all $t \geq 1$

$$\begin{aligned}
\mathrm{KL}(w^\star, w_{t+1}) &\leq \mathrm{KL}(w^\star, w_t) \cdot \big(1 - \tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 w_{t,\min}^2\big) \\
&\leq \mathrm{KL}(w^\star, w_t) \cdot \big(1 - \tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 \exp(-3\Lambda/\delta)\big) \\
&\leq \mathrm{KL}(w^\star, w_t) \cdot \exp\big(-\tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 \exp(-3\Lambda/\delta)\big)\ ,
\end{aligned}$$

where the final inequality is due to $1 - u \leq \exp(-u)$ for all $u \in \mathbb{R}$. Inductively, for $t \geq 1$ we have:

$$\mathrm{KL}(w^\star, w_{t+1}) \leq \mathrm{KL}(w^\star, w_1) \cdot \exp\big(-\tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 \exp(-3\Lambda/\delta) \cdot t\big)\ . \tag{131}$$

Finally, recall from Lemma H.1 that $\mathrm{KL}(w^\star, w_1) \leq \big(\frac{5}{4}\big) \cdot \mathrm{KL}(w^\star, w_0)$. Substituting this bound into (131) and using the crude bound $\frac{5}{4} \leq 2$ then establishes for all $t \geq 1$ that

$$\mathrm{KL}(w^\star, w_{t+1}) \leq 2\mathrm{KL}(w^\star, w_0) \cdot \exp\big(-\tfrac{1}{20}\eta^2 \cdot \sigma_{\min}^2 \exp(-3\Lambda/\delta) \cdot t\big)\ .$$

This yields statement (2) of the theorem and concludes the proof. ■

Improved dependence on local state in Part (1) of Theorem 4.7. We note that, as a corollary of the "primal" variant proof of non-uniform skew-gradient domination from Proposition G.13, we can obtain a slightly improved bound on the multiplicative one-step change in KL. In particular, assuming the setting of Theorem 4.7 and comparing the skew-gradient domination bounds of Proposition 4.4 (proven using primarily the "dual" perspective) and Proposition G.13 (proven using a slightly more involved "primal" perspective), we have:

$$\begin{aligned}
\text{Proposition 4.4}: &\quad \|J\nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot w_{\min}^2 \cdot \mathrm{KL}(w^\star, w) \\
\text{Proposition G.13}: &\quad \|J\nabla F(z)\|_z^2 \geq \sigma_{\min}^2 \cdot p_{\min} \cdot q_{\min} \cdot \mathrm{KL}(w^\star, w)\ .
\end{aligned}$$

Thus the bound of Proposition G.13 can be larger than that of Proposition 4.4 when only one of $p_{\min}, q_{\min}$ is small. Repeating an identical argument as in the first three steps of the proof of Theorem 4.7, we then have the following corollary:

Corollary H.3 (Improved one-step multiplicative decrease in KL)**.** *Assume the setting of Theorem 4.7. Then for every $t \geq 1$, it holds that* $\mathrm{KL}(w^\star, w_{t+1}) \leq \mathrm{KL}(w^\star, w_t) \cdot \exp\big(-\frac{1}{20}\eta^2\sigma_{\min}^2 \cdot p_{t,\min} \cdot q_{t,\min}\big)$.

H.4 Comparisons with Analysis of Wei et al. (2021)

In this section, we give a brief comparison of our result with that of Wei et al. (2021).

Overview of Rate in Wei et al. Theorem 3 of Wei et al. (2021) proves the following result: for a zero-sum game $A \in [-1, 1]^{m \times n}$ with a unique Nash equilibrium $w^\star \in \mathcal{W}$, it holds for the iterates $\{w_t\}$ of OMWU with stepsize $\eta \leq \frac{1}{8}$ and initialized from the uniform distribution that

$$\mathrm{KL}(w^\star, w_t) \leq C_3(1 + C_4)^{-t} \quad \text{for all } t \geq 1\ , \tag{132}$$

where $C_3, C_4 > 0$ are constants depending on the matrix A. Here, note that the constants C_3, C_4 depend on the minimum non-zero mass δ in the Nash equilibrium $w^\star$ of A, and one must explicitly track their dependence on δ throughout the proof to make a comparison.

Note also that Theorem 3 of Wei et al. does not necessarily assume that $w^\star \in \mathrm{ri}(\mathcal{W})$ is interior, as we require in Theorem 4.7. When $w^\star$ is interior, the unconditional intermediate estimate in Case 1 of their Lemma 2 proof, together with their Lemmas 16 and 19, gives a linear contraction without a startup phase. Tracking this dependence, including the square-root conversion from their potential to KL, yields positive game-dependent constants C_A, c_A, independent of t and η, for which one may take

$$C_3 = C_A \exp\big(\tfrac{\log mn}{\delta}\big) \quad \text{and} \quad C_4 = c_A\eta^2 \exp\big(\tfrac{-3\log mn}{\delta}\big)\ . \tag{133}$$

Here, we note that the $\log(mn)$ term is exactly the cross-entropy $\Lambda = \mathrm{CrossEnt}(w^\star, w_0)$ under the uniform initialization $p_0 = \frac{\mathbf{1}_m}{m}$ and $q_0 = \frac{\mathbf{1}_n}{n}$. Combining (132) and (133) then yields for Wei et al. a last-iterate convergence rate of

$$\mathrm{KL}(w^\star, w_t) \leq C_A \exp\big(\tfrac{\Lambda}{\delta}\big) \cdot \big(1 + c_A\eta^2 \exp(\tfrac{-3\Lambda}{\delta})\big)^{-t}\ . \tag{134}$$

Comparison with rate in Theorem 4.7. Comparing (134) to Part (2) of Theorem 4.7, both bounds have decay coefficients exponentially small in Λ/δ. The displayed rate of Wei et al. uses $c_A \eta^2 \exp(-3\Lambda/\delta)$, whereas Part (2) uses $\eta^2 \sigma_{\min}^2 \exp(-3\Lambda/\delta)/20$. Both proofs use *uniform* lower bounds on the minimum coordinate masses, established in our Lemma H.2 and their Lemma 19, respectively (note also this uniform bound on $w_{t,\min}$ is tight in general; see the proof of our lower bound of Theorem 4.10 in Section J). On the other hand, the rate of Wei et al. from (134) also contains the exponential prefactor $\exp(\Lambda/\delta)$, a term that grows unbounded as $\delta \to 0$. In comparison, the leading multiplicative factor in Part (2) of our Theorem 4.7 is $2\mathrm{KL}(w^\star, w_0)$. Thus our bound improves the leading prefactor, while both displayed decay coefficients have the same $\exp(-3\Lambda/\delta)$ dependence. The remaining game-dependent factors are displayed above. Moreover, the linear last-iterate lower bound in Theorem 4.10 shows that the exponential dependence on $1/\delta$ is unavoidable in general.

Comparison of proof techniques. Our technique for proving Theorem 4.7 is novel and distinct from the techniques of Wei et al. Note that the analysis of Wei et al. does not directly establish a bound on the one-step change $\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t)$ over the iterates of the algorithm (in contrast to our proof, which does). Instead, the proof of Wei et al. establishes a decreasing one step change of the potential function

$$\Theta_t = \mathrm{KL}(w^\star, \widehat{w}_t) + \tfrac{1}{16}\mathrm{KL}(\widehat{w}_t, w_{t-1}) \ ,$$

where $\{\widehat{w}_t\}$ is the sequence of half-step iterates when defining OMWU via its Optimistic Mirror Descent formulation (for the details of this formulation, see, e.g., Wei et al. (2021), expressions (1) and (2), Section 3.1.1 of Syrgkanis et al. (2015), or Section 2 of Cai et al. (2024)). Establishing a linear convergence rate for the sequence $\{\Theta_t\}$ then translates into the linear convergence rate for $\mathrm{KL}(w^\star, w_t)$.

In contrast, recall that our proof of Theorem 4.7 is based on the dual perspective of *energy dissipation* (via the one-step bounds of Lemma 4.1), which we then relate back to the primal space via the skew-gradient domination inequality of Proposition 4.4. While the calculations required to establish the tight energy dissipation bounds of Lemma 4.1 are still somewhat involved, the approach is conceptually simple, as outlined in the proof in Section F.1. Together, our analysis also gives a more precise and quantitative understanding of the geometric bottlenecks that cause the energy dissipation (and thus primal convergence in KL) to be slow over the iterates.

On the assumption of an interior Nash equilibrium. While our analysis does assume both a unique and interior Nash (while the results of Daskalakis and Panageas (2019) and Wei et al. (2021) only assume a unique Nash equilibrium), we believe our energy-based analysis can likely be adapted to also hold only under the uniqueness assumption. We leave this as a direction for future work.

H.5 Numerical Simulations

In this section, we present plots of several numerical simulations that highlight the dependence on δ in the linear rate of Theorem 4.7.

Setup. We run OMWU on two classes of payoff matrices: the 2×2 canonical matrix A_{δ_p, δ_q} from Definition I.1, and a diagonal payoff matrix $A \in [-1, 1]^{10 \times 10}$.

- **On canonical 2x2 matrix**: We instantiate $A_{\delta_p, \delta_q} \in [-1, 1]^{2 \times 2}$ with $\delta = \delta_p = \delta_q$, for a range of δ between 0 and 0.5. In this instantiation, $w^\star = (p^\star, q^\star)$ is given by $p^\star = q^\star = (1 - \delta, \delta)$. Thus, smaller δ corresponds to a Nash equilibrium closer to a vertex of $\mathcal{W}$.
- **On diagonal 10x10 matrix:** We instantiate a diagonal $A \in [-1, 1]^{10 \times 10}$ as follows: we construct a vector $v_\delta \in (0, 1)^{10}$ where the first 5 coordinates are $\delta/5$, and the last 5 coordinates are $(1-\delta)/5$. Then, we set $A = \mathrm{Diag}(d_\delta)$, where $d_\delta \propto 1/v_\delta$ (normalized so that the maximum entry of the vector d_δ is 1). It follows that the unique and interior NE $w^\star = (p^\star, q^\star)$ of A is given by $p^\star = q^\star = v_\delta$. Thus for smaller values of δ, the NE components $p^\star$ and $q^\star$ are roughly uniform among the latter 5 coordinates, and close to a boundary face of the simplex corresponding to the support of these coordinates.

Figure 2: Dependence of δ in linear rate. In Figure 2, we run OMWU on instantiations of the 2x2 and 10x10 diagonal matrices previously described with several values of δ. In each instance, we set $\eta = 0.2$, and we initialize the iterates at the uniform distributions. We plot the value of $\mathrm{KL}(w^\star, w_t)$ over the iterates in log scale. In both cases, the figure highlights the dependence on δ

seen in the last-iterate convergence rate of Part (2) of Theorem 4.7. For smaller δ, the slope of the linear convergence in the figure is much flatter, as suggested by the analytic bound.

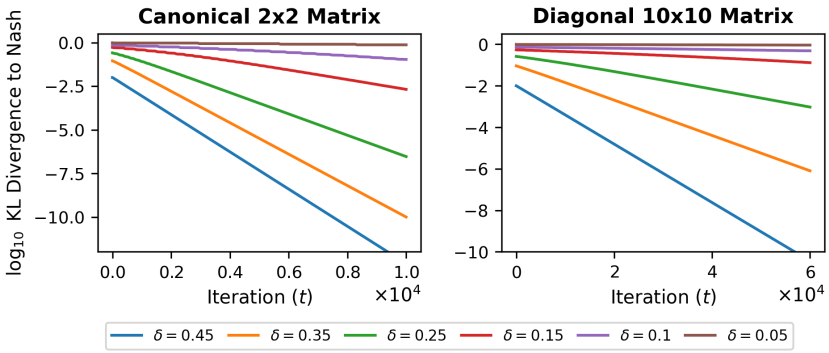

Figure 2: Plot of $\mathrm{KL}(w^\star, w_t)$ of OMWU in log-scale over time on instantiations of 2x2 and 10x10 payoff matrices with varying values of δ. The minimum NE coordinate is δ for the 2x2 instances and $\delta/5$ for the 10x10 instances.

Figure 3: Non-uniform one-step change. In Figure 3, we run OMWU on instantiations of the 2x2 and 10x10 diagonal matrices previously described with varying values of small δ. In each instance, we set $\eta = 0.2$, and we initialize the iterates at the uniform distributions. We plot the value of $\mathrm{KL}(w^\star, w_t)$ over the iterates in log scale. In both cases, the figure highlights the fact that the one-step change in $\mathrm{KL}(w^\star, w_t)$ can decay in a highly non-uniform manner, especially when δ is small, and when $\mathrm{KL}(w^\star, w_t)$ is still relatively large. This property is exactly indicated by the non-uniform dependence on $w_{t,\min}$ in Part (1) of Theorem 4.7. In this regime, the iterates will initially spend longer periods of time close to the simplex boundary (where $w_{t,\min}$ is smaller), which leads to smaller magnitudes of dual energy dissipation, and consequently a slower overall last-iterate convergence to NE.

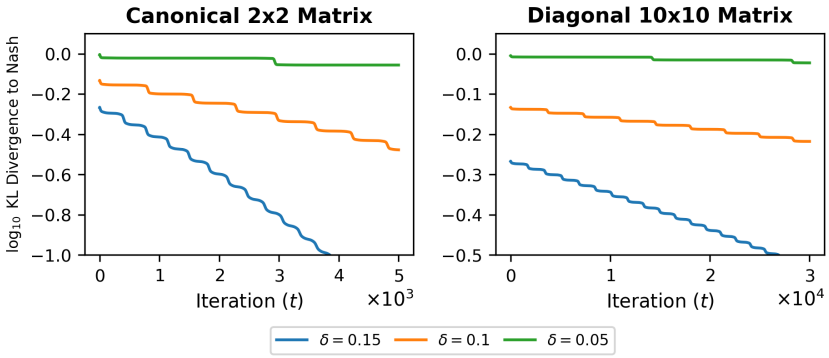

Figure 3: Plot of $\mathrm{KL}(w^\star, w_t)$ of OMWU in log-scale over time on instantiations of 2x2 and 10x10 payoff matrices with varying values of δ. The minimum NE coordinate is δ for the 2x2 instances and $\delta/5$ for the 10x10 instances. Observe that the one-step change in KL decays non-uniformly over time, as suggested by Part (1) of Theorem 4.7.

I Details on OMWU Dynamics in 2x2 Setting

This section gives additional preliminaries and details on OMWU in the 2×2 setting. These preliminaries are used in the proofs of the lower bounds of Theorem 4.10 and Theorem 5.1, and the upper bound of Theorem 5.2. In particular, we establish properties of the following canonical matrix:

Definition I.1. *Let $\delta_p, \delta_q \in (0,1)$. Then $A_{\delta_p,\delta_q} \in [-1,1]^{2\times 2}$ is the payoff matrix given by*

$$A_{\delta_p,\delta_q} = \begin{pmatrix} \delta_p\delta_q & -\delta_p(1-\delta_q) \\ -(1-\delta_p)\delta_q & (1-\delta_p)(1-\delta_q) \end{pmatrix}.$$

Significance of Canonical Matrix. For the 2×2 setting, the matrix A_{δ_p,δ_q} serves as a *canonical* payoff matrix. In particular, the unique and interior Nash equilibrium $w^\star = (p^\star, q^\star)$ of A_{δ_p,δ_q} is given by $p^\star = (1-\delta_p, \delta_p)$ and $q^\star = (1-\delta_q, \delta_q)$, and every other $A \in \mathbb{R}^{2\times 2}$ with $w^\star$ as an interior NE can be written as an affine transformation of A_{δ_p,δ_q}. We state and prove these structural properties in Proposition I.2 in the following subsection. Note also that running OMWU on the matrix A_{δ_p,δ_q} with $\delta_p = \delta \in (0, 1/2)$ and $\delta_q = 1/2$ captures the behavior identified in the hard 2×2 instance of Cai et al. (2024).

Organization of Section. The remainder of this section is organized as follows:

- **Sections I.2, I.4, and I.3** give details on the OMWU primal and dual iterates when run on A_{δ_p,δ_q}. These preliminaries are used in the lower bound proofs of Theorem 4.10 and Theorem 5.1, which are based on instantiations of A_{δ_p,δ_q} (see Sections J and K, respectively), and the upper bound proof of Theorem 5.2, which relies on the affine transformation of A_{δ_p,δ_q} (see Section L).
- **Section I.4** gives simplified expressions for duality gap, TV distance, and KL divergence in the 2×2 setting, as well as a visual comparison of their levelsets, in Section I.4.4.

Simplifications on notation. Throughout, we will drop the subscript and write $\mathbf{1} = (1, 1) \in \mathbb{R}^2$ to denote the two-dimensional all-ones vector, and softmax $=$ softmax$_2$ to denote the two-dimensional softmax map.

I.1 Structural Properties of Canonical Payoff Matrix

We prove the following properties of the matrix A_{δ_p,δ_q} from Definition I.1:

Proposition I.2. *Fix $\delta_p, \delta_q \in (0, 1)$. Let $A_{\delta_p,\delta_q} \in [-1, 1]^{2\times 2}$ be the matrix from Definition I.1. Let $w^\star = (p^\star, q^\star)$ with $p^\star = (1-\delta_p, \delta_p)$ and $q^\star = (1-\delta_q, \delta_q)$. Then the following properties hold:*

(i) The unique and interior NE of A_{δ_p,δ_q} is $w^\star$, and $\langle p^\star, A_{\delta_p,\delta_q} q^\star\rangle = 0$.

(ii) Fix $A \in [-1, 1]^{2\times 2}$, and suppose that $w^\star$ is an interior NE of A. Let $v = \langle p^\star, Aq^\star\rangle \in \mathbb{R}$, and let the entries of A be given by

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} .$$

Define $\gamma = \frac{a-v}{\delta_p\delta_q}$, and assume further that A is not a constant matrix of the form $A = u\mathbf{1}\mathbf{1}^\top$ for some $u \in \mathbb{R}$. Then $w^\star$ is the unique Nash equilibrium of A, $0 < |\gamma| \le 4$, and

$$A = \gamma \cdot A_{\delta_p,\delta_q} + v\mathbf{1}\mathbf{1}^\top .$$

Before giving the proof of Proposition I.2, we first state and prove the following useful facts about 2×2 zero-sum games with an interior equilibrium:

Proposition I.3 (Interior NE implies uniqueness in 2×2 games)**.** *Fix $A \in \mathbb{R}^{2\times 2}$ with entries*

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} .$$

Suppose A has an interior NE $w^\star = (p^\star, q^\star)$ where $p^\star = (1-\mu_p, \mu_p)$ and $q^\star = (1-\mu_q, \mu_q)$ for $\mu_p, \mu_q \in (0, 1)$. Let $v = \langle p^\star, Aq^\star\rangle \in \mathbb{R}$, and define $\phi = a - b - (c - d)$. Then the following properties hold:

(i) $\phi = 0$ if and only if $A = u\mathbf{1}\mathbf{1}^\top$ is a constant matrix for some $u \in \mathbb{R}$.

(ii) If $\phi \neq 0$, then $w^\star$ is the unique Nash equilibrium of A.

(iii) $\phi = \frac{a-v}{\mu_p\mu_q}$.

Proof. As $w^\star = (p^\star, q^\star)$ is an interior NE of A, by Part (i) of Proposition B.1, we must have $Aq^\star = v\mathbf{1}$ and $A^\top p^\star = v\mathbf{1}$ for a constant $v \in \mathbb{R}$. Using the definition of A and $p^\star = (1-\mu_p, \mu_p)$ and $q^\star = (1-\mu_q, \mu_q)$, this yields the constraints

$$\begin{cases} a - (a-b)\cdot\mu_q = v \\ c - (c-d)\cdot\mu_q = v \end{cases} \quad \text{and} \quad \begin{cases} a - (a-c)\cdot\mu_p = v \\ b - (b-d)\cdot\mu_p = v \end{cases} . \tag{135}$$

Rearranging further yields the constraints

$$\begin{cases} \mu_q \cdot (a - b - (c - d)) = a - c \\ \mu_p \cdot (a - b - (c - d)) = a - b \, . \end{cases} \tag{136}$$

For Part (i) of the proposition, recall that $\phi = a - b - (c - d)$. Now observe that if all $a = b = c = d$, meaning A is a constant matrix, then trivially $\phi = 0$. In the other direction, if $\phi = 0$, then (136) implies $a = c$ and $a = b$. By definition of ϕ, this further implies $a = d$. Thus if $\phi = 0$, then A is a constant matrix. In this degenerate case, every $(p, q) \in \mathrm{ri}(\Delta_2 \times \Delta_2)$ is an NE.

For Part (ii), suppose that $\phi \neq 0$. Then (136) implies that $\mu_q = (a - c)/\phi$ and $\mu_p = (a - b)/\phi$ have unique solutions that also uniquely define $p^\star = (1 - \mu_p, \mu_p)$ and $q^\star = (1 - \mu_q, \mu_q)$. This also further precludes A from simultaneously having a non-interior Nash equilibrium: if such an equilibrium w' exists, then every interior $\widetilde{w}$ on the line segment between w' and $w^\star = (p^\star, q^\star)$ would also be a Nash equilibrium (which follows by Part (ii) of Proposition B.1). However, this contradicts the uniqueness of the solution of (136). Thus, in the case that $\phi \neq 0$, an interior NE of A is the unique NE of A.

For Part (iii), observe from (135) that $a - v = (a - c) \cdot \mu_p$. Substituting the first equality of (136) further gives $a - v = \phi \mu_q \mu_p$. Rearranging for ϕ yields the desired statement. ■

We now proceed with the proof of Proposition I.2:

Proof of Proposition I.2. We prove the two parts of the proposition separately.

Proof of Part (i). Using the definition of A_{δ_p, δ_q} and $p^\star = (1 - \delta_p, \delta_p)$ and $q^\star = (1 - \delta_q, \delta_q)$, observe by a direct calculation that

$$A_{\delta_p, \delta_q} \begin{pmatrix} 1 - \delta_q \\ \delta_q \end{pmatrix} = \begin{pmatrix} \delta_p \delta_q (1 - \delta_q) - \delta_p (1 - \delta_q) \delta_q \\ -(1 - \delta_p) \delta_q (1 - \delta_q) + (1 - \delta_p)(1 - \delta_q) \delta_q \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \, .$$

By a similar calculation, we find $A_{\delta_p, \delta_q}^\top p^\star = 0$. Thus $A_{\delta_p, \delta_q} q^\star$ and $A_{\delta_p, \delta_q}^\top p^\star$ are both constant vectors. Thus by Part (ii) of Proposition B.1 it follows that $w^\star = (p^\star, q^\star)$ is an NE of A_{δ_p, δ_q}, and that $\langle p^\star, A_{\delta_p, \delta_q} q^\star \rangle = 0$. Now observe that since $\delta_p, \delta_q \in (0, 1)$, the equilibrium $w^\star$ is interior by definition of $p^\star$ and $q^\star$. Moreover, by Definition I.1, A_{δ_p, δ_q} is not a constant matrix. Thus, by Parts (i) and (ii) of Proposition I.3, it follows that $w^\star$ is also the unique NE of A_{δ_p, δ_q}.

Proof of Part (ii). Recalling the entries of $A \in \mathbb{R}^{2 \times 2}$ from the statement of the proposition, let $\phi = a - b - (c - d)$. Recall that we assume A has an interior Nash equilibrium $w^\star = (p^\star, q^\star)$, and that A is not a constant matrix. Thus, by Parts (i) and (ii) of Proposition I.3, $\phi \neq 0$, and $w^\star$ is the unique Nash equilibrium of A. Now, let $v = \langle p^\star, A q^\star \rangle$, and define

$$A' = A - v \mathbf{1}\mathbf{1}^\top \in \mathbb{R}^{2 \times 2} \, .$$

We will show that $A' = \gamma \cdot A_{\delta_p, \delta_q}$, for some $0 < |\gamma| \leq 4$, which will imply that $A = \gamma A_{\delta_p, \delta_q} + v \mathbf{1}\mathbf{1}^\top$.

For this, observe by Proposition B.1 that since $w^\star = (p^\star, q^\star)$ is an interior NE of A with $v = \langle p^\star, A q^\star \rangle$, then $A q^\star = v \mathbf{1}$ and $A^\top p^\star = v \mathbf{1}$. Thus by definition of $A' = A - v \mathbf{1}\mathbf{1}^\top$, it follows that both $A' q^\star = 0$ and $A'^\top p^\star = 0$. Using the definition of the entries of A' (which depend on the entries of A) it is straightforward to show that this further implies

$$b - v = \tfrac{-(a - v) \cdot (1 - \delta_q)}{\delta_q}, \; d - v = \tfrac{(a - v) \cdot (1 - \delta_p)(1 - \delta_q)}{\delta_p \delta_q}, \text{ and } c - v = \tfrac{-(a - v) \cdot (1 - \delta_p)}{\delta_p} \, . \tag{137}$$

By Part (iii) of Proposition I.3, recall that $\phi = (a - v)/(\delta_p \delta_q)$, and recall that $\phi \neq 0$. Then factoring out ϕ from each term of (137), we find exactly that $A' = \phi \cdot A_{\delta_p, \delta_q}$. By definition of $A' = A - v \mathbf{1}\mathbf{1}^\top$, we conclude $A = \phi \cdot A_{\delta_p, \delta_q} + v \mathbf{1}\mathbf{1}^\top$.

Set $\gamma := \phi$. By definition of $\phi = a - b - (c - d)$ and $A \in [-1, 1]^{2 \times 2}$, it follows that $|\phi| \leq 4$. Thus we conclude $A = \gamma A_{\delta_p, \delta_q} + v \mathbf{1}\mathbf{1}^\top$ for $0 < |\gamma| \leq 4$, which yields the statement of Part (ii) of the proposition. ■

Remark I.4. We note that Part (ii) of Proposition I.2 is similar to Lemma 5 of Cai et al. (2025). The key difference is that the "base" matrix in their lemma is not A_{δ_p, δ_q}, but a different 2×2 payoff matrix. The key property of A_{δ_p, δ_q} is that the value of the game is zero (i.e., $\langle p^\star, A_{\delta_p, \delta_q} q^\star \rangle = 0$). This allows for a more convenient primal-dual relationship under the OMWU iterates (c.f., Proposition D.10) that we introduce in the sequel.

Finally, we also show that A_{δ_p,δ_q} is well-conditioned in the sense that its minimum restricted singular value $\sigma_{\min}$ is always an absolute constant.

Proposition I.5 (Minimum restricted singular value of Canonical Matrix)**.** *Fix any $\delta_p, \delta_q \in (0,1)$. Let $A := A_{\delta_p,\delta_q} \in [-1,1]^{2\times 2}$ be the matrix from Definition I.1. Let $\sigma_{\min}$ be the minimum restricted singular value from Definition G.8. Then $\sigma_{\min} = \frac{1}{2}$.*

We note that the proof is the generalization of and follows similarly to that of Part (2) of Proposition G.16. There, we established for the special case $\delta = \delta_p = \delta_q$ that $\sigma_{\min} = \frac{1}{2}$. Proposition I.5 extends this result to hold for general $\delta_p, \delta_q \in (0,1)$.

Proof. First, let $\sigma_{\min,n}$ and $\sigma_{\min,m}$ be the component-wise restricted singular values from (110), and recall from Proposition G.11 that $\sigma_{\min} = \min\{\sigma_{\min}(A, \mathbf{1}^\perp), \sigma_{\min}(A^\top, \mathbf{1}^\perp)\}$. Using calculations similar to the proof of Part (2) of Proposition G.16, we will show that $\sigma_{\min}(A, \mathbf{1}^\perp) = \sigma_{\min}(A^\top, \mathbf{1}^\perp) = \frac{1}{2}$. For this, recall by definition that

$$\sigma_{\min}(A, \mathbf{1}^\perp) = \inf_{v \in \mathbf{1}^\perp \setminus \{0\}} \frac{\|\Pi_{\mathbf{1}^\perp}(Av)\|_2}{\|v\|_2} .$$

Moreover, for any $v \in \mathbf{1}^\perp \subset \mathbb{R}^2$, it follows that $v = c \cdot (1, -1)$ for some $c \in \mathbb{R}$. Then using the definition of $A = A_{\delta_p,\delta_q}$, and recalling that $\Pi_{\mathbf{1}^\perp} = I - \frac{1}{2}\mathbf{1}\mathbf{1}^\top$, it follows by a direct calculation that

$$\Pi_{\mathbf{1}^\perp}(Av) = c \cdot \begin{pmatrix} \delta_p - (\delta_p - (1-\delta_p))/2 \\ -(1-\delta_p) - (\delta_p - (1-\delta_p))/2 \end{pmatrix} = \frac{c}{2} \cdot \begin{pmatrix} 1 \\ -1 \end{pmatrix} .$$

It follows that $\|\Pi_{\mathbf{1}^\perp}(Av)\|_2 = |c|\sqrt{2}/2$. Moreover, we have by definition of $v = c \cdot (1, -1)$ that $\|v\|_2 = |c|\sqrt{2}$. Thus we conclude $\sigma_{\min}(A, \mathbf{1}^\perp) = 1/2$. By an identical calculation, we also find $\sigma_{\min}(A^\top, \mathbf{1}^\perp) = 1/2$, which completes the proof. ■

I.2 Primal and Dual OMWU Iterates

In this section, we give details on the OMWU iterates when run on the canonical 2×2 game A_{δ_p,δ_q} from Definition I.1. These derivations are used to prove the lower bounds of Theorems 4.10 and 5.1, and also the best-iterate upper bound of Theorem 5.2.

In this setting, note that while the dual iterates $\{z_t\}$ all belong to $\mathbb{R}^4$, the effective dual space $\mathcal{Z}$ (as introduced in Section 3.2) is a two-dimensional subspace. This allows for a convenient two-dimensional representation of both the primal and dual iterates that facilitates obtaining precise control of the OMWU trajectory as a function of the Nash equilibrium coordinates δ_p and δ_q.

Roadmap of subsection. Section I.2.1 presents structural properties of the lower-dimensional effective dual space. Section I.2.2 then gives simplified expressions for the energy function and primal-dual mapping over the effective dual space. Section I.2.3 states the simplified primal and dual OMWU updates when run on A_{δ_p,δ_q}. Throughout, we assume the following setting:

Setting I.1 (2×2 Setting Using Canonical Matrix)**.** *For fixed $\delta_p, \delta_q \in (0,1)$, assume that $A := A_{\delta_p,\delta_q}$ is the canonical 2×2 payoff matrix from Definition I.1.*

In this setting, we will also make use of the following simplified notation for readability:

Simplified notation and definitions. Under Setting I.1, let $\mathrm{LSE}\colon \mathbb{R}^2 \to \mathbb{R}$ denote the two-dimensional log-sum-exp function, and let $R : \Delta_2 \to \mathbb{R}$ denote the two-dimensional negative-entropy function (as defined in Section 3). Moreover, we write $\mathrm{sigmoid}\colon \mathbb{R} \to \mathbb{R}$ to denote the function given by $\mathrm{sigmoid}(u) = 1/(1 + \exp(-u))$ for $u \in \mathbb{R}$.

I.2.1 Effective Dual Space

Recall from Section 3 that $\mathcal{Z} = \mathrm{Span}(J\mathcal{W}) \subseteq \mathbb{R}^{2+2}$, where $J = ((0, A), (-A^\top, 0)) \in \mathbb{R}^{4\times 4}$ and $\mathcal{W} = \Delta_2 \times \Delta_2$. In order to introduce the lower-dimensional representation of $\mathcal{Z}$, it will be more convenient to further work with the individual components of the dual variables $z = (x, y) \in \mathbb{R}^{2+2}$. Specifically, we first introduce the component-wise effective dual spaces $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^2$:

Definition I.6 (Component-wise effective dual spaces)**.** *Fix $\delta_p, \delta_q \in (0,1)$ and assume Setting I.1. Let $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^2$ be the linear subspaces given by*

$$\mathcal{X} := \mathrm{Span}(A\Delta_2) \text{ and } \mathcal{Y} := \mathrm{Span}(A^\top \Delta_2) .$$

Given the product structure of $\mathcal{W} = \Delta_2 \times \Delta_2$, and using the definition of J, the following characterization of $\mathcal{Z} = \text{Span}(J\mathcal{W})$ is immediate:

Proposition I.7. *Fix $\delta_p, \delta_q \in (0,1)$, and assume Setting I.1. Then $\mathcal{Z} = \text{Span}(J\mathcal{W}) = \mathcal{X} \times \mathcal{Y}$.*

Now define $p^\star = (1-\delta_p, \delta_p)$ and $q^\star = (1-\delta_q, \delta_q)$, and recall from Part (i) of Proposition I.2 that $(p^\star, q^\star)$ is the unique and interior NE of $A := A_{\delta_p, \delta_q}$. The key property of $\mathcal{X}$ and $\mathcal{Y}$ is that these spaces are orthogonal to $p^\star$ and $q^\star$, respectively. Formally, we have the following refinement of Proposition D.6 for the present 2×2 setting:

Proposition I.8 (Component-wise orthogonality to Nash)**.** *Fix $\delta_p, \delta_q \in (0,1)$, and assume Setting I.1. Then for any $x = (x(1), x(2)) \in \mathcal{X}$ and $y = (y(1), y(2)) \in \mathcal{Y}$, it holds that:*

$$\begin{aligned} \langle x, p^\star \rangle = 0 \quad &\implies \quad x(2) = x(1) \cdot (1 - (1/\delta_p)) \\ \textit{and} \quad \langle y, q^\star \rangle = 0 \quad &\implies \quad y(2) = y(1) \cdot (1 - (1/\delta_q)) \ . \end{aligned}$$

Proof. First, fix $x \in \mathcal{X}$. By definition, this means that there exists $k \geq 1$ such that $x = \sum_{i=1}^k \tau_i A p_i$, where all $\tau_i \in \mathbb{R}$ and $p_i \in \Delta_2$. Then for each $i \in [k]$, observe that

$$\langle \tau_i A p_i, p^\star \rangle = \tau_i \langle p_i, A^\top p^\star \rangle = \tau_i \langle p_i, 0 \rangle = 0 \ .$$

Here, the final equality is due to the fact that $(p^\star, q^\star)$ is an interior Nash equilibrium of $A = A_{\delta_p, \delta_q}$ with $\langle p^\star, A q^\star \rangle = 0$ (Part (i) of Proposition I.2), which implies that $A^\top p^\star = 0$ is the zero vector (Part (i) of Proposition B.1). Thus $\langle x, p^\star \rangle = 0$ by definition of x. Now, as $p^\star = (1-\delta_p, \delta_p)$, a direct calculation and rearrangement then finds

$$\langle x, p^\star \rangle = x(1) \cdot (1-\delta_p) + x(2) \cdot \delta_p = 0 \quad \implies \quad x(2) = x(1) \cdot (1 - (1/\delta_p)) \ .$$

This proves the first statement of the proposition. The second statement follows using an identical argument for $y \in \mathcal{Y}$. ■

Before proceeding, we make the following two remarks:

Remark I.9 (Comparison with Proposition D.6)**.** For a matrix A with interior NE $w^\star = (p^\star, q^\star)$, Proposition D.6 establishes for the general case that $\langle z, w^\star \rangle = 0$ for all $z = (x, y) \in \mathcal{Z}$. In particular, this implies by definition of J that $\langle x, p^\star \rangle + \langle y, q^\star \rangle = 0$. Proposition I.8 establishes under Setting I.1 the slightly stronger *component-wise* property that both $\langle x, p^\star \rangle = 0$ and $\langle y, q^\star \rangle = 0$. As seen in the proof of Proposition I.8, this property is a consequence of the fact that $\langle p^\star, A_{\delta_p, \delta_q} q^\star \rangle = 0$, as proven in Part (i) of Proposition I.2. Indeed, it is then straightforward to see that this stronger property also holds for any other (general-dimension) A whose value of the game is zero.

Remark I.10 (Main consequence is lower-dimensional effective dual space)**.** The key consequence of Proposition I.8 is that, under Setting I.1, the dual variable components $x \in \mathcal{X} \subset \mathbb{R}^2$ and $y \in \mathcal{Y} \subset \mathbb{R}^2$ each lie on one-dimensional subspaces. Thus, a dual variable $z = (x, y) \in \mathcal{Z}$ can be fully specified by the first coordinates $(x(1), y(1)) \in \mathbb{R}^2$ of its components. This means that the effective dual space $\mathcal{Z} \subset \mathbb{R}^{2+2}$ is a *two-dimensional subspace*. This will allow for additional, convenient simplifications of the energy function F and the primal-dual relationship over $\mathcal{Z}$. We note that similar approaches were used in prior works of (Bailey and Piliouras, 2019a, Section 3.2) and (Lazarsfeld et al., 2025a, Proposition 4.2), albeit for different base matrices.

I.2.2 Primal-Dual Mapping Over Effective Dual Space

Under Setting I.1, working over the lower-dimensional structure of $\mathcal{Z}$ established by Proposition I.8 allows for simplified characterizations of the energy function F and its gradient map. Formally, we establish the following properties:

Proposition I.11 (Energy and primal-dual map over $\mathcal{Z}$)**.** *Fix $\delta_p, \delta_q \in (0,1)$ and assume Setting I.1. Let $z = (x, y) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, and let $w = (p, q) = \nabla F(z) \in \text{ri}(\mathcal{W}) = \text{ri}(\Delta_2) \times \text{ri}(\Delta_2)$. Then the following relationships hold:*

(i) $\text{LSE}(x) = x(1) + \log\left(1 + \exp\left(\frac{-x(1)}{\delta_p}\right)\right)$ *and* $\text{LSE}(y) = y(1) + \log\left(1 + \exp\left(\frac{-y(1)}{\delta_q}\right)\right)$.

(ii) $p(1) = \text{sigmoid}\left(\frac{x(1)}{\delta_p}\right)$ *and* $q(1) = \text{sigmoid}\left(\frac{y(1)}{\delta_q}\right)$.

(iii) $x(1) = \delta_p \cdot \log\left(\frac{p(1)}{1-p(1)}\right)$ *and* $y(1) = \delta_q \cdot \log\left(\frac{q(1)}{1-q(1)}\right)$.

Proof. Note that as $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, we have by Proposition I.8 that $x(2) = x(1) \cdot (1 - (1/\delta_p))$ and $y(2) = y(1) \cdot (1 - (1/\delta_q))$. We use this relationship to prove each of the three parts as follows:

Proof of Part (i). By definition of $\mathrm{LSE}(\cdot)$ and using Proposition I.8, we can compute

$$\begin{aligned}\mathrm{LSE}(x) &= \log\big(\exp(x(1)) + \exp(x(2))\big) \\ &= \log\big(\exp(x(1)) + \exp\big(x(1)\cdot(1-\tfrac{1}{\delta_p})\big)\big) \\ &= \log\big(\exp(x(1))\cdot(1+\exp\big(\tfrac{-x(1)}{\delta_p}\big)\big)\big) = x(1) + \log\big(1+\exp\big(\tfrac{-x(1)}{\delta_p}\big)\big)\,.\end{aligned}$$

The statement for $\mathrm{LSE}(y)$ follows by an identical calculation.

Proof of Part (ii). Recall by definition of F and the assumptions of the proposition that $\nabla F(z) = (p, q) = (\mathrm{softmax}(x), \mathrm{softmax}(y))$. Thus by definition of softmax and using Proposition I.8, we can write

$$\begin{aligned}p(1) = \frac{\exp\big(x(1)\big)}{\exp\big(x(1)\big) + \exp\big(x(2)\big)} &= \frac{\exp\big(x(1)\big)}{\exp\big(x(1)\big) + \exp\big(x(1)\cdot(1-\frac{1}{\delta_p})\big)} \\ &= \frac{\exp\big(x(1)\big)}{\exp\big(x(1)\big)\cdot\big(1+\exp\big(\frac{-x(1)}{\delta_p}\big)\big)} = \mathrm{sigmoid}\Big(\frac{x(1)}{\delta_p}\Big)\,.\end{aligned}$$

The statement for $q(1)$ follows by an identical calculation.

Proof of Part (iii). Recall for $u \in (0,1)$ that $\mathrm{sigmoid}^{-1}(u) = \log\left(\frac{u}{1-u}\right)$. Thus by Part (ii), we have

$$\frac{x(1)}{\delta_p} = \mathrm{sigmoid}^{-1}(p(1)) = \log\left(\frac{p(1)}{1-p(1)}\right),$$

Rearranging gives $x(1) = \delta_p \cdot \log\left(\frac{p(1)}{1-p(1)}\right)$. The claim for $y(1)$ follows identically. ■

Remark I.12 (Bijectivity of energy gradient map over $\mathcal{Z}$.)**.** Note that Parts (ii) and (iii) of Proposition I.11 establish that, under Setting I.1, $\nabla F : \mathcal{Z} \to \mathrm{ri}(\mathcal{W})$ is bijective. Indeed, note that as the zero-sum game value under $A := A_{\delta_p,\delta_q}$ is zero (Part (i) of Proposition I.2), the more general statements of Propositions D.9 and D.10 also together imply that $\nabla F : \mathcal{Z} \to \mathrm{ri}(\mathcal{W})$ is bijective. For the 2×2 setting, the closed-form expression of Part (iii) of Proposition I.11 thus allows for recovering the coordinates of the dual variable $z = (x, y) \in \mathcal{Z}$ solely using the coordinates of the corresponding primal variable $w \in \mathrm{ri}(\mathcal{W})$.

Part (iii) of Proposition I.11 also leads to several useful corollaries. First, it allows for defining the coordinates of the unique "dual" Nash variable $z^\star \in \mathcal{Z}$ that maps to the primal NE of $A := A_{\delta_p,\delta_q}$ under ∇F. More formally:

Corollary I.13 (Dual Nash coordinate in $\mathcal{Z}$)**.** *Fix* $\delta_p, \delta_q \in (0,1)$ *and assume Setting I.1. Let* $w^\star = (p^\star, q^\star) \in \mathrm{ri}(\mathcal{W})$. *Let* $z^\star = (x^\star, y^\star) \in \mathcal{Z}$ *be such that*

$$x^\star(1) = \delta_p \cdot \log\left(\tfrac{1-\delta_p}{\delta_p}\right) \;\; \textit{and} \;\; y^\star(1) = \delta_q \cdot \log\left(\tfrac{1-\delta_q}{\delta_q}\right).$$

Then $\nabla F(z^\star) = w^\star$.

As a second corollary, note that the expressions $\delta_p \log(p(1)/(1-p(1)))$ and $\delta_q \log(q(1)/(1-q(1)))$ are increasing in $p(1)$ and $q(1)$, respectively. Similarly, $\mathrm{sigmoid}(x(1)/\delta_p)$ and $\mathrm{sigmoid}(y(1)/\delta_q)$ are increasing in $x(1)$ and $y(1)$, respectively. Together this implies the following further relationships between $(p, q) \in \mathrm{ri}(\mathcal{W})$ and $(x, y) \in \mathcal{Z}$:

Corollary I.14. *Fix* $\delta_p, \delta_q \in (0,1)$ *and assume Setting I.1. Let* $z = (x, y) \in \mathcal{Z}$, *and let* $w = (p, q) = \nabla F(z) \in \mathrm{ri}(\mathcal{W})$. *Fix any* $\rho \in (0,1)$. *Then the following relationships hold:*

$$p(1) \geq \rho \iff x(1) \geq \delta_p \cdot \log\left(\frac{\rho}{1-\rho}\right) \;\; \textit{and} \;\; q(1) \geq \rho \iff y(1) \geq \delta_q \cdot \log\left(\frac{\rho}{1-\rho}\right).$$

I.2.3 OMWU Dynamics

Under Setting I.1, the primal-dual relationships over the effective dual space $\mathcal{Z}$ lead to simplified expressions for the primal and dual OMWU updates. Formally, we have the following:

Proposition I.15 (OMWU update rules on canonical 2×2 matrix)**.** *Fix $\delta_p, \delta_q \in (0,1)$ and assume Setting I.1. Let $\{w_t\}$ and $\{z_t\}$ denote the primal and dual iterates of OMWU on $A := A_{\delta_p,\delta_q}$ with stepsize $\eta > 0$ initialized from $w_0 = (p_0, q_0) \in \mathrm{ri}(\mathcal{W})$, where each $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$ and $z_t = (x_t, y_t) \in \mathcal{Z}$. Then for all $t \geq 1$:*

$$\begin{cases} x_{t+1}(1) = x_t(1) - \eta \cdot \delta_p \left(q_t(1) - (1-\delta_q) + \big(q_t(1) - q_{t-1}(1)\big)\right) \\ y_{t+1}(1) = y_t(1) + \eta \cdot \delta_q \left(p_t(1) - (1-\delta_p) + \big(p_t(1) - p_{t-1}(1)\big)\right) \end{cases},$$

and moreover

$$\begin{cases} p_{t+1}(1) = \mathrm{sigmoid}\left(\frac{x_{t+1}(1)}{\delta_p}\right) \\ q_{t+1}(1) = \mathrm{sigmoid}\left(\frac{y_{t+1}(1)}{\delta_q}\right) \end{cases}.$$

Proof. First, recalling the structure of $A = A_{\delta_p,\delta_q}$ from Definition I.1, a straightforward calculation shows that, for any $p, q \in \Delta_2$:

$$(Aq)(1) = \big(q(1) - (1-\delta_q)\big) \cdot \delta_p \quad \text{and} \quad (A^\top p)(1) = \big(p(1) - (1-\delta_p)\big) \cdot \delta_q \,. \tag{138}$$

Now recall by definition of the OMWU dual iterates $\{z_t\}$ from (4) that, for all $t \geq 1$:

$$z_{t+1} = (x_{t+1}, y_{t+1}) = z_t - \eta(Jw_t + J(w_t - w_{t-1})) \,,$$

which means component-wise that

$$\begin{aligned} x_{t+1} &= x_t - \eta(Aq_t + A(q_t - q_{t-1})) \\ \text{and} \quad y_{t+1} &= y_t + \eta(A^\top p_t + A^\top(p_t - p_{t-1})) \,. \end{aligned} \tag{139}$$

Considering only the first coordinates of the components in (139) and applying the simplifications from (138) yields the first statement of the proposition.

For the second statement, observe that all $z_t = (x_t, y_t) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ by definition. Moreover, by Proposition 3.1, we have $w_t = \nabla F(z_t)$ for all $t \geq 0$. Then applying Part (ii) of Proposition I.11 yields the second statement and concludes the proof. ■

Remark I.16 (Leading coordinates suffice for tracking OMWU iterates)**.** Due to Proposition I.8, for $z = (x, y) \in \mathcal{Z}$, observe that the leading coordinates $x(1) \in \mathbb{R}$ and $y(1) \in \mathbb{R}$ fully specify z. Similarly, for $w = (p, q) \in \mathcal{W}$, the leading coordinates $p(1) \in [0,1]$ and $q(1) \in [0,1]$ fully specify w. Thus, the simplified expressions in Proposition I.15 are sufficient for fully specifying the OMWU iterates $\{z_t\}$ and $\{w_t\}$ under Setting I.1.

I.3 Monotonicity and Cycling Properties of OMWU Iterates

Using the simplified OMWU update expressions from Proposition I.15, we prove in this section an additional set of monotonicity properties that the OMWU iterates satisfy under Setting I.1. In the statement of the lemma, we assume without loss of generality that $\delta_p, \delta_q < \frac{1}{2}$, and that the stepsize η is bounded by an absolute constant.

Lemma I.17. *Fix $0 < \delta_p \leq \delta_q < \frac{1}{2}$ and assume Setting I.1. Let $\{w_t\}$ and $\{z_t\}$ be the primal and dual iterates of OMWU with stepsize $0 < \eta < \frac{1}{8}$, initialized from some $w_0 \in \mathrm{ri}(\mathcal{W})$. For each $t \geq 1$, let $\widetilde{x}_t = x_t(1), \widetilde{y}_t = y_t(1)$ and $\widetilde{p}_t = p_t(1), \widetilde{q}_t = q_t(1)$. Fix any $k \geq 1$. Then the following hold:*

(i) Suppose $\widetilde{q}_{k-1} \leq 1 - 3\delta_q$. Then: $\widetilde{x}_{k+1} - \widetilde{x}_k \geq \frac{\eta}{2} \cdot \delta_p \delta_q$.

(ii) Suppose $\widetilde{q}_{k-1} \geq 1 - \frac{\delta_q}{3}$. Then: $\widetilde{x}_{k+1} - \widetilde{x}_k \leq -\frac{\eta}{3} \cdot \delta_p \delta_q$.

(iii) Suppose $\widetilde{p}_{k-1} \leq 1 - 3\delta_p$. Then: $\widetilde{y}_{k+1} - \widetilde{y}_k \leq -\frac{\eta}{2} \cdot \delta_p \delta_q$.

(iv) Suppose $\widetilde{p}_{k-1} \geq 1 - \frac{\delta_p}{3}$. Then: $\widetilde{y}_{k+1} - \widetilde{y}_k \geq \frac{\eta}{3} \delta_p \delta_q$.

Proof. We start by proving Part (i) of the lemma, and we note that Parts (ii), (iii), and (iv) will follow by nearly identical calculations. For this, using the update rule of Proposition I.15, it follows that

$$\widetilde{x}_{k+1} - \widetilde{x}_k \geq \tfrac{\eta}{2}\delta_p\delta_q \iff 1 - \delta_q - 2\widetilde{q}_k + \widetilde{q}_{k-1} \geq \tfrac{\delta_q}{2} \ . \tag{140}$$

For this, let $\widetilde{q}_{k-1} = 1 - \alpha_{k-1}$ for $\alpha_{k-1} \in (0,1)$. Our goal is to derive an upper bound on $\widetilde{q}_k$ in terms of $\widetilde{q}_{k-1}$ that will facilitate establishing the inequality in the right hand side of (140). For this, observe again by the dual update rule of Proposition I.15 that

$$\widetilde{y}_k - \widetilde{y}_{k-1} = \eta\delta_q\big(\widetilde{p}_{k-1} - (1-\delta_p) + \widetilde{p}_{k-1} - \widetilde{p}_{k-2}\big) \leq 2\eta\delta_q \tag{141}$$

where the inequality follows from the fact that $\widetilde{p}_{k-1} \leq 1$, that $\widetilde{p}_{k-1} - \widetilde{p}_{k-2} \leq 1$, and that $\delta_p > 0$, with $\widetilde{p}_{-1} = \widetilde{p}_0$ when $k = 1$. Then by definition of $\widetilde{q}_k = \text{sigmoid}(\widetilde{y}_k/\delta_q)$, and using the fact that $1 - \text{sigmoid}(u) = 1/(1+\exp(u))$ for any $u \in \mathbb{R}$, it follows that

$$\begin{aligned} 1 - \widetilde{q}_k = \frac{1}{1+\exp(\widetilde{y}_k/\delta_q)} &\geq \frac{1}{1+\exp(\widetilde{y}_{k-1}/\delta_q)\cdot\exp(2\eta)} \\ &\geq \frac{1}{(1+\exp(\widetilde{y}_{k-1}/\delta_q))\cdot\exp(2\eta)} = \exp(-2\eta)(1-\widetilde{q}_{k-1}) \ . \end{aligned}$$

Here, the first inequality comes from applying the bound $\widetilde{y}_k \leq \widetilde{y}_{k-1} + 2\eta\delta_q$ from (141), the second inequality comes from $\exp(2\eta) \geq 1$, and the final equality comes from $\widetilde{q}_{k-1} = \text{sigmoid}(\widetilde{y}_{k-1}/\delta_q)$. As $\exp(-u) \geq 1 - u$ for all $u \in \mathbb{R}$, we then further have $1 - \widetilde{q}_k \geq (1-2\eta)(1-\widetilde{q}_{k-1})$. Using the definition $1 - \widetilde{q}_{k-1} = \alpha_{k-1}$, this means that $\widetilde{q}_k \leq 1 - (1-2\eta)\alpha_{k-1}$, and thus we have

$$1 - \delta_q - 2\widetilde{q}_k + \widetilde{q}_{k-1} \geq 1 - \delta_q - 2\big(1 - (1-2\eta)\alpha_{k-1}\big) + 1 - \alpha_{k-1} = (1-4\eta)\alpha_{k-1} - \delta_q \ .$$

Now under the assumption that $\widetilde{q}_{k-1} = 1 - \alpha_{k-1} \leq 1 - 3\delta_q$, we have $\alpha_{k-1} \geq 3\delta_q$. Moreover as we assume $\eta \leq 1/8$, we have $1 - 4\eta \geq 1/2$. This allows us to further write

$$1 - \delta_q - 2\widetilde{q}_k + \widetilde{q}_{k-1} \geq \tfrac{3}{2}\delta_q - \delta_q = \tfrac{1}{2}\delta_q \ ,$$

which exactly establishes the right-hand side of (140) and thus the statement of Part (i) of the lemma.

The proof of Part (ii) follows via nearly identical calculations as the proof of Part (i). In particular, we further use the identity $\exp(u) = \text{sigmoid}(u)/(1 - \text{sigmoid}(u))$ for $u \in \mathbb{R}$ and the fact that $\exp(2\eta) \leq 1 + 4\eta$ when $\eta \leq 1/8$. The proofs of Part (iii) and Part (iv) then follow identically to Parts (i) and (ii), respectively. This is due to the symmetry of the dual update rules for $\widetilde{x}_{k+1}$ and $\widetilde{y}_{k+1}$ in Proposition I.15, and thus we omit these calculations. ■

I.3.1 Maximum Width of Dual Coordinates

In the following proposition, we establish a bound on the coordinates of a dual variable in terms of its energy value. This is used in the sequel to establish an upper bound on the maximum distance between any two dual (and also primal) OMWU iterates. Formally:

Proposition I.18 (Energy bound implies dual coordinate bound)**.** *Fix $\delta_p, \delta_q \in (0,1)$ and assume Setting I.1. Fix $z = (x,y) \in \mathcal{Z}$, and let $\alpha > 0$. If $F(z) \leq \alpha$, then the following hold:*

$$-\big(\tfrac{\delta_p}{1-\delta_p}\big)\cdot\alpha \leq x(1) \leq \alpha \quad \textit{and} \quad -\big(\tfrac{\delta_q}{1-\delta_q}\big)\cdot\alpha \leq y(1) \leq \alpha \ .$$

Proof. Fix $z = (x,y) \in \mathcal{Z}$, and let $\widetilde{x} = x(1) \in \mathbb{R}$ and $\widetilde{y} = y(1) \in \mathbb{R}$. Recall by definition of the energy function F and its simplification over $\mathcal{Z}$ from Proposition I.11 that

$$F(z) = \text{LSE}(x) + \text{LSE}(y) = \big(\widetilde{x} + \log(1+\exp(-\widetilde{x}/\delta_p))\big) + \big(\widetilde{y} + \log(1+\exp(-\widetilde{y}/\delta_q))\big) \ .$$

Note that over $\mathcal{Z}$, both terms $\widetilde{x} + \log(1+\exp(-\widetilde{x}/\delta_p))$ and $\widetilde{y} + \log(1+\exp(-\widetilde{y}/\delta_q))$ are convex (in particular, minimized at the coordinates $x^\star(1)$ and $y^\star(1)$ from Corollary I.13) and always positive for all $\delta_p, \delta_q \in (0,1)$. Thus, for any $z = (x,y) \in \mathcal{Z}$, both $\text{LSE}(x) \geq 0$ and $\text{LSE}(y) \geq 0$. Therefore, $F(z) \leq \alpha$ implies that both

$$\text{LSE}(x) \leq \widetilde{x} + \log\big(1+\exp(-\widetilde{x}/\delta_p)\big) \leq \alpha \ \text{ and } \ \text{LSE}(y) \leq \widetilde{y} + \log\big(1+\exp(-\widetilde{y}/\delta_q)\big) \leq \alpha \ .$$

For the $\widetilde{x}$ variable, observe further that $\log(1+\exp(-\widetilde{x}/\delta_p)) \geq 0$, and thus

$$\widetilde{x} \leq \alpha - \log(1+\exp(-\widetilde{x}/\delta_p)) \leq \alpha \ .$$

This proves the upper bound on $x(1) = \widetilde{x}$. For the lower bound, recall that $\log(1 + \exp(-u)) \geq -u$ for any $u \in \mathbb{R}$. Thus if $\mathrm{LSE}(x) \leq \alpha$, then

$$\alpha \geq \widetilde{x} + \log(1 + \exp(-\widetilde{x}/\delta_p)) \geq \widetilde{x} - (\widetilde{x}/\delta_p) = \widetilde{x}((\delta_p - 1)/\delta_p) = -\widetilde{x}((1 - \delta_p)/\delta_p) \; .$$

Then rearranging gives

$$\widetilde{x} \geq -(\delta_p/(1 - \delta_p)) \cdot \alpha \; ,$$

which proves the lower bound on $x(1)$. The bounds on $y(1) = \widetilde{y}$ follow by identical calculations. ■

I.4 Structural Properties of Distances to Nash in 2x2 Setting

In this section, we give simplified expressions for duality gap, TV distance, and KL divergence under Setting I.1 (i.e., when $A := A_{\delta_p, \delta_q}$ is the canonical matrix from Definition I.1). For this, we consider throughout the following subsections $w = (p, q) \in \mathcal{W} = \Delta_2 \times \Delta_2$, and we use the notation $\widetilde{p} = p(1) \in [0, 1]$ and $\widetilde{q} = q(1) \in [0, 1]$ for readability. Moreover, recall by Proposition I.2 that the unique and interior NE $w^\star = (p^\star, q^\star)$ of $A = A_{\delta_p, \delta_q}$ is given by $p^\star = (1 - \delta_p, \delta_p)$ and $q^\star = (1 - \delta_q, \delta_q)$.

I.4.1 Duality Gap under Setting I.1

For $w = (p, q) \in \mathcal{W}$, recall that duality gap is defined as

$$\mathrm{DG}(w) = \max_{q' \in \Delta_2} \langle q', A^\top p \rangle - \min_{p' \in \Delta_2} \langle p', Aq \rangle .$$

Then in terms of the leading coordinates $\widetilde{p} = p(1)$ and $\widetilde{q} = q(1)$, we have by definition of A_{δ_p, δ_q}:

$$\begin{aligned} \mathrm{DG}(w) &= \max\left\{(A^\top p)(1), (A^\top p)(2)\right\} + \max\left\{(-Aq)(1), (-Aq)(2)\right\} \\ &= \max\left\{(\widetilde{p} - (1 - \delta_p))\delta_q, (1 - \delta_p - \widetilde{p})(1 - \delta_q)\right\} \\ &\qquad + \max\left\{(1 - \delta_q - \widetilde{q})\delta_p, (\widetilde{q} - (1 - \delta_q))(1 - \delta_p)\right\} \; . \end{aligned}$$

Here, we also used the fact that $p(2) = 1 - \widetilde{p}$ and $q(2) = 1 - \widetilde{q}$. Thus $\mathrm{DG}(w)$ has a piecewise linear structure, which we can summarize depending on the location of $(\widetilde{p}, \widetilde{q})$ relative to $(1 - \delta_p, 1 - \delta_q)$ (i.e., the leading coordinates of the unique and interior NE of A_{δ_p, δ_q}). Then $\mathrm{DG}(w)$ is given by:

$$\mathrm{DG}(w) = \begin{cases} (\widetilde{p} - (1 - \delta_p))\delta_q + (\widetilde{q} - (1 - \delta_q))(1 - \delta_p) & \text{if } \widetilde{p} \geq 1 - \delta_p, \;\; \widetilde{q} \geq 1 - \delta_q \\ (1 - \delta_p - \widetilde{p})(1 - \delta_q) + (\widetilde{q} - (1 - \delta_q))(1 - \delta_p) & \text{if } \widetilde{p} \leq 1 - \delta_p, \;\; \widetilde{q} \geq 1 - \delta_q \\ (1 - \delta_p - \widetilde{p})(1 - \delta_q) + (1 - \delta_q - \widetilde{q})\delta_p & \text{if } \widetilde{p} \leq 1 - \delta_p, \;\; \widetilde{q} \leq 1 - \delta_q \\ (\widetilde{p} - (1 - \delta_p))\delta_q + (1 - \delta_q - \widetilde{q})\delta_p & \text{if } \widetilde{p} \geq 1 - \delta_p, \;\; \widetilde{q} \leq 1 - \delta_q \; . \end{cases} \tag{142}$$

I.4.2 TV Distance under Setting I.1

Recall by definition that the total variation distance $\mathrm{TV}(w^\star, w)$ is given by

$$\mathrm{TV}(w^\star, w) = \mathrm{TV}(p^\star, p) + \mathrm{TV}(q^\star, q) = \tfrac{1}{2}\|p^\star - p\|_1 + \tfrac{1}{2}\|q^\star - q\|_1 \; .$$

As $p = (\widetilde{p}, 1 - \widetilde{p})$ and $p^\star = (1 - \delta_p, \delta_p)$, observe that

$$\|p^\star - p\|_1 = |1 - \delta_p - \widetilde{p}| + |\delta_p - (1 - \widetilde{p})| = 2 \cdot |1 - \delta_p - \widetilde{p}| \; .$$

Similarly, $\|q^\star - q\|_1 = 2 \cdot |1 - \delta_q - \widetilde{q}|$. Thus we have that

$$\mathrm{TV}(w^\star, w) = \mathrm{TV}(p^\star, p) + \mathrm{TV}(q^\star, q) = |1 - \delta_p - \widetilde{p}| + |1 - \delta_q - \widetilde{q}| \; . \tag{143}$$

I.4.3 KL Divergence under Setting I.1

Recall by definition that the KL divergence $\mathrm{KL}(w^\star, w)$ is given by

$$\mathrm{KL}(w^\star, w) = \mathrm{KL}(p^\star, p) + \mathrm{KL}(q^\star, q) \; .$$

Then under this setting, the components $\mathrm{KL}(p^\star, p)$ and $\mathrm{KL}(q^\star, q)$ can be simplified as:

$$\begin{aligned} \mathrm{KL}(p^\star, p) &= \sum\nolimits_{i=1}^{2} p^\star(i) \log\left(\frac{p^\star(i)}{p(i)}\right) = (1 - \delta_p) \cdot \log\left(\frac{1 - \delta_p}{\widetilde{p}}\right) + \delta_p \cdot \log\left(\frac{\delta_p}{1 - \widetilde{p}}\right) \\ \mathrm{KL}(q^\star, q) &= \sum\nolimits_{j=1}^{2} q^\star(j) \log\left(\frac{q^\star(j)}{q(j)}\right) = (1 - \delta_q) \cdot \log\left(\frac{1 - \delta_q}{\widetilde{q}}\right) + \delta_q \cdot \log\left(\frac{\delta_q}{1 - \widetilde{q}}\right) \; . \end{aligned} \tag{144}$$

I.4.4 Comparison of Levelsets Between Distances

In the following set of figures, we plot the levelsets of $\mathrm{DG}(w)$, $\mathrm{TV}(w^\star, w)$, and $\mathrm{KL}(w^\star, w)$ for the 2×2 setting using the simplified expressions from Section I.4. We plot the levelsets over the leading coordinates $(p(1), q(1))$ for three different settings of $w^\star = (p^\star, q^\star)$. In each subplot, $w^\star$ is denoted by the yellow star. The figures highlight a fundamental distinction between the geometries of these distance functions over the simplex.

Figure 4: Uniform NE. Here, $p^\star = q^\star = (0.5, 0.5)$.

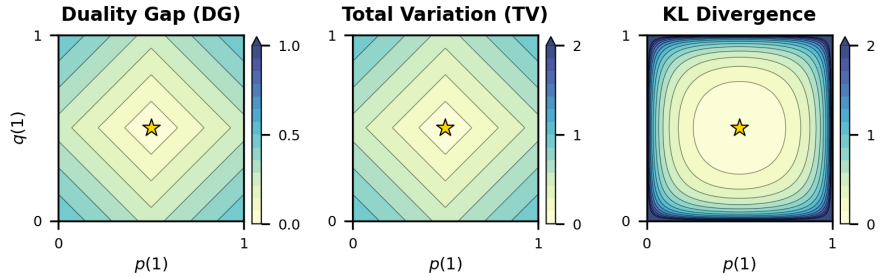

Figure 4: $\mathrm{DG}(w)$, $\mathrm{TV}(w^\star, w)$, and $\mathrm{KL}(w^\star, w)$ levelsets in the 2×2 setting for $p^\star = q^\star = (0.5, 0.5)$.

Figure 5: Near-boundary NE in one component. Here, $p^\star = (0.9, 0.1)$ and $q^\star = (0.5, 0.5)$.

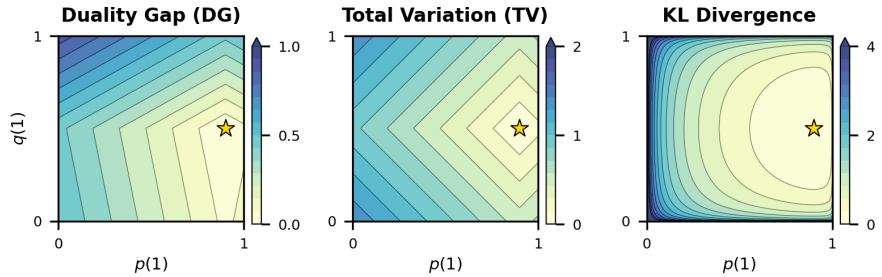

Figure 5: $\mathrm{DG}(w)$, $\mathrm{TV}(w^\star, w)$, and $\mathrm{KL}(w^\star, w)$ levelsets in the 2×2 setting for $p^\star = (0.9, 0.1)$, and $q^\star = (0.5, 0.5)$.

Figure 6: Near-boundary NE in both components. Here, $p^\star = q^\star = (0.9, 0.1)$.

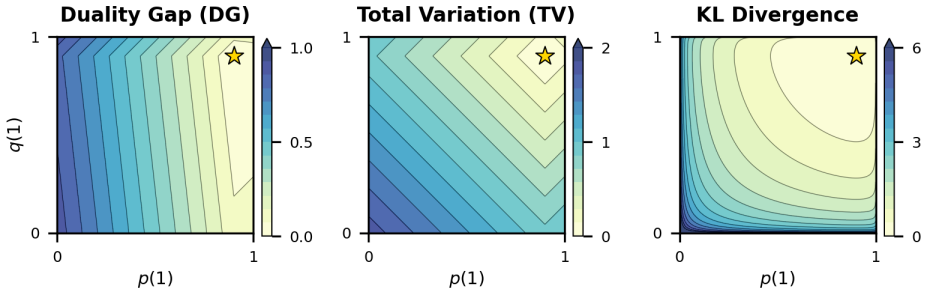

Figure 6: $\mathrm{DG}(w)$, $\mathrm{TV}(w^\star, w)$, and $\mathrm{KL}(w^\star, w)$ levelsets in the 2×2 setting for $p^\star = q^\star = (0.9, 0.1)$.

Discussion. Observe that the geometry of the levelsets of each distance function is highly dependent on the location of the Nash equilibrium $w^\star$. In particular, when $w^\star$ is closer to the simplex boundary (Figure 5 and Figure 6), the $\mathrm{KL}(w^\star, w)$ levelsets become "squished" near the boundary closest to $w^\star$. Note in these regions that a primal variable $w = (p, q)$ can be simultaneously close to $w^\star$ in TV distance (which is symmetric), yet still far away in $\mathrm{KL}(w^\star, w)$. This property is captured in our uniform best-iterate lower bound of Theorem K.1, Part (1). Additionally, observe in Figure 6 that for $w^\star$ close to a vertex, a primal variable $w = (p, q)$ with, for example, $p(1) = 0.95$ and $q(1) = 0.5$, can

simultaneously be close in $\mathrm{DG}(w)$, but far in $\mathrm{TV}(w^\star, w)$ and $\mathrm{KL}(w^\star, w)$. This property is captured by the separation in uniform best-iterate guarantees between TV and KL (via the constant lower bound of Theorem 5.1) and DG (via the $\widetilde{O}(T^{-1/2})$ best-iterate rate of Theorem 5.2). Finally, note that the location of the NE $w^\star$ in the example of Figure 5 roughly corresponds to that of the hard 2×2 instance used in the results of Cai et al. (2024).

J Details on Lower Bound for Universal Last-Iterate Convergence in KL

This section gives the proof of Theorem 4.10, which gives a universal lower bound on the convergence rate of OMWU in KL divergence. We first restate the theorem:

Theorem 4.10 (Lower bound in KL)**.** *Assume the definitions of Theorem 4.7. Fix any $T \geq 3$. Then there exists $A \in [-1, 1]^{2\times 2}$ with an interior NE $w^\star = (p^\star, q^\star)$ with $\delta := \delta_p = \delta_q$ such that, from a positive measure set of initializations w_0, each with $\mathrm{KL}(w^\star, w_0) \leq 6$, the OMWU iterates $\{w_t\}$ with η as in Assumption 1 satisfy*

$$\mathrm{KL}(w^\star, w_t) \geq \mathrm{KL}(w^\star, w_0) \cdot \exp(-40\eta^2 \exp(\tfrac{-1}{\delta}) \cdot t) \quad \textit{for all } t \in [T].$$

Observe that in Theorem 4.10, the dependence on δ in the lower bound matches that in the upper bound of Theorem 4.7, indicating that our analysis is tight. The proof of the theorem follows two steps that we outline here:

High-level overview of proof. Let $\{w_t\}$ and $\{z_t\}$ denote the primal and dual iterates of the OMWU algorithm. We have by Corollary D.7 and the lower bound on energy dissipation from Lemma 4.1 that, for all $t \geq 1$:

$$\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) = F(z_{t+1}) - F(z_t) \geq -\tfrac{5}{4}\eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \ .$$

Recall in Proposition 4.4 that we established a lower bound on the dissipation term $\|J\nabla F(z_t)\|_{z_t}^2$. This translates into the upper bound on last-iterate convergence rate in KL in Theorem 4.7. For the present lower bound of Theorem 4.10, we first establish in Proposition J.1 an analogous *non-uniform upper bound* on the dissipation term $\|J\nabla F(z)\|_z^2$.

In particular, this upper bound is a structural and general property that we prove for the 2×2 setting, and which holds over a region of the primal space $\mathcal{W}$ near the simplex boundary. This leads to a much tighter bound compared to the crude uniform upper bound of Proposition G.14. Similar to the lower bound on $\|J\nabla F(z)\|_z^2$, the upper bound of Proposition J.1 contains a non-uniform multiplicative factor that depends on the minimum coordinates $p_{\min}$ and $q_{\min}$ of the corresponding primal variable $w = (p, q) = \nabla F(z)$. This further indicates that the local state-dependence in the lower bound of Proposition 4.4 is necessary.

Then, the last-iterate lower bound of Theorem 4.10 follows by constructing a 2×2 payoff matrix and a set of initializations where we establish an upper bound on $p_{t,\min}$ and $q_{t,\min}$ that is exponentially small in the minimum Nash coordinate δ, and that holds for all iterates t over a fixed time horizon. This relies on the preliminaries for OMWU in the 2×2 setting that we developed in Section I.

Roadmap of the section. In Section J.1, we state and prove the intermediate result of Proposition J.1. The full proof of Theorem 4.10 then follows in Section J.2.

J.1 Non-Uniform Upper Bound on Dissipation Term

In this section, we establish a non-uniform *upper bound* on the dissipation term $\|J\nabla F(z)\|_z^2$. For simplicity, we focus on the low-dimensional setting where $A \in \mathbb{R}^{2\times 2}$.

Proposition J.1 (Non-uniform upper bound on dissipation term)**.** *Let $A \in \mathbb{R}^{2\times 2}$ have a unique and interior Nash equilibrium $w^\star = (p^\star, q^\star)$ such that $p^\star = q^\star = (1 - \delta, \delta)$ for $\delta \in (0, 0.5)$. Let $\sigma_{\max} = \|A\|_2$. Consider $z \in \mathbb{R}^{2+2}$ and $w = (p, q) = \nabla F(z)$ such that $p_{\min} = p(2) \leq \delta$ and $q_{\min} = q(2) \leq \delta$. Define $\widehat{w}_{\min} = \max(p_{\min}, q_{\min})$. Then:*

$$\|J\nabla F(z)\|_z^2 \leq 8 \cdot \sigma_{\max}^2 \cdot \widehat{w}_{\min} \cdot \delta \cdot \mathrm{KL}(w^\star, w) \ .$$

Proof. To prove the proposition, we leverage the primal perspective of the non-uniform gradient domination result from Proposition G.13. Specifically, using part (i) of Proposition G.13, we have

$$\|J\nabla F(z)\|_z^2 = \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p) \ . \tag{145}$$

We will thus derive upper bounds on the individual terms in (145) via the following steps:

1. Exact Characterization of Variance in 2×2 Setting.

Let $p = (p(1), p(2))$ and $q = (q(1), q(2))$. We first prove that

$$\mathrm{Var}_p(Aq) \le 2 \cdot p(1) \cdot p(2) \cdot \|A\|_2^2 \cdot \|q - q^\star\|_2^2 \tag{146}$$

$$\text{and} \quad \mathrm{Var}_q(A^\top p) \le 2 \cdot q(1) \cdot q(2) \cdot \|A\|_2^2 \cdot \|p - p^\star\|_2^2 \,. \tag{147}$$

To prove (146), let $v = Aq \in \mathbb{R}^2$. By definition of $\mathrm{Var}_p(v)$ from (35), it is straightforward to compute in the two-dimensional case that $\mathrm{Var}_p(v) = p(1) \cdot p(2) \cdot (v(1) - v(2))^2$. Now, let $\mathbf{1} \in \mathbb{R}^2$ denote the all-ones vector. By Part (i) of Proposition B.1, we have $Aq^\star = c\mathbf{1}$ for some $c \in \mathbb{R}$. Thus, we can write:

$$v = Aq = A(q^\star + (q - q^\star)) = c\mathbf{1} + A(q - q^\star) \,.$$

Further letting $d = (1, -1)$, it follows that we can write and simplify

$$\begin{aligned} |v(1) - v(2)| = |\langle d, v\rangle| &= |\langle d, c\mathbf{1} + A(q - q^\star)\rangle| \\ &= |\langle d, A(q - q^\star)\rangle| \\ &\le \|d\|_2 \cdot \|A(q - q^\star)\|_2 \le \sqrt{2} \cdot \|A\|_2 \cdot \|q - q^\star\|_2 \,. \end{aligned}$$

Substituting this bound on $|v(1) - v(2)|$ into the definition of $\mathrm{Var}_p(v)$ then exactly yields the desired inequality of (146). By an identical calculation, we also find $\mathrm{Var}_q(A^\top p) \le 2q(1)q(2)\|A\|_2^2\|p - p^\star\|_2^2$, where we additionally use the fact that $\|A\|_2 = \|A^\top\|_2$.

2. Relating KL to Euclidean Distance to Nash.

Next, we prove the following relationships between KL and Euclidean distance to Nash:

$$\|p - p^\star\|_2^2 \le 4 \cdot \max(p_{\min}, \delta) \cdot \mathrm{KL}(p^\star, p) \tag{148}$$

$$\text{and} \quad \|q - q^\star\|_2^2 \le 4 \cdot \max(q_{\min}, \delta) \cdot \mathrm{KL}(q^\star, q) \,. \tag{149}$$

We start by proving (148). For this, recall from Section 3 that for $w = (p, q) \in \mathcal{W}$, we write $R(w) = R_m(p) + R_n(q)$ to denote the (separable) negative entropy function, where $R_m = -\mathrm{Ent}_m$ and $R_n = -\mathrm{Ent}_n$. Moreover, $\mathrm{KL}(w^\star, w) = D_R(w^\star, w)$ is the Bregman divergence of R, and more specifically $\mathrm{KL}(p^\star, p) = D_{R_m}(p^\star, p)$ and $\mathrm{KL}(q^\star, q) = D_{R_n}(q^\star, q)$. Thus, by using the integral form of Bregman divergences (see, e.g., Nocedal and Wright (2006), Theorem 2.1), it follows that

$$\mathrm{KL}(p^\star, p) = D_{R_m}(p^\star, p) = \int_0^1 (1 - s)\big\langle (p - p^\star), \nabla^2 R_m(p_s)(p - p^\star)\big\rangle \, ds \,, \tag{150}$$

where we write $p_s = p - s(p - p^\star)$ for $s \in [0, 1]$.

We simplify the integrand in (150) as follows: first, by definition in (36), we have for all $s \in [0, 1]$

$$\nabla^2 R_m(p_s) = \mathrm{Diag}\left(\tfrac{1}{p_s(1)}, \ldots, \tfrac{1}{p_s(m)}\right) \,.$$

Thus for each s, using the facts that $m = 2$ and $p^\star = (1 - \delta, \delta)$ we can simplify

$$\begin{aligned} \big\langle (p - p^\star), \nabla^2 R_m(p_s)(p - p^\star)\big\rangle &= \textstyle\sum_{i=1}^2 \dfrac{(p(i) - p^\star(i))^2}{p_s(i)} \\ &= \frac{(p(1) - (1 - \delta))^2}{p_s(1)} + \frac{(1 - p(1) - \delta)^2}{p_s(2)} \\ &= \big(p(1) - (1 - \delta)\big)^2 \cdot \left(\frac{1}{p_s(1)} + \frac{1}{p_s(2)}\right) \\ &= \frac{1}{2} \cdot \|p - p^\star\|_2^2 \cdot \frac{1}{p_s(1) \cdot p_s(2)} \,. \end{aligned}$$

Here, the last line follows from the fact that $\|p - p^\star\|_2^2 = 2 \cdot (p(1) - (1 - \delta))^2$ in this two-dimensional setting. Thus, we can further rewrite (150) as

$$\mathrm{KL}(p^\star, p) = \frac{1}{2} \cdot \|p - p^\star\|_2^2 \cdot \int_0^1 \frac{(1 - s)}{p_s(1) \cdot p_s(2)} \, ds \,. \tag{151}$$

To further derive a lower bound on (151), we will show a uniform upper bound on $p_s(1) \cdot p_s(2) = p_s(2) \cdot (1 - p_s(2))$ for all $s \in [0,1]$. For this, we recall by the assumptions of the proposition that $p^\star(2) = \delta = \min(p^\star(1), p^\star(2)) \leq \frac{1}{2}$, and also $p(2) = p_{\min} = \min(p(1), p(2)) \leq \frac{1}{2}$. In particular, this also implies that $p_s(2) = \min(p_s(1), p_s(2)) \leq \frac{1}{2}$ for all $s \in [0,1]$. Thus, for all $s \in [0,1]$:

$$p_s(2) \cdot (1 - p_s(2)) \leq p_s(2) \leq \max(p(2), \delta) = \max(p_{\min}, \delta) \; .$$

Substituting this uniform bound into (151), and using the fact that $\int_0^1 (1-s) ds = \frac{1}{2}$, we find

$$\mathrm{KL}(p^\star, p) \geq \frac{1}{4} \cdot \|p - p^\star\|_2^2 \cdot \frac{1}{\max(p_{\min}, \delta)} \; .$$

Rearranging then yields the desired inequality from (148). Using identical calculations similarly yields the inequality from (149).

3. Combining the Pieces.

Combining the inequalities of (146) and (147) from Step 1 and expressions (148) and (149) from Step 2, we find

$$\begin{aligned} \mathrm{Var}_p(Aq) &\leq 8 \cdot \|A\|_2^2 \cdot p_{\min} \cdot \max(q_{\min}, \delta) \cdot \mathrm{KL}(q^\star, q) \\ \text{and} \quad \mathrm{Var}_q(A^\top p) &\leq 8 \cdot \|A\|_2^2 \cdot q_{\min} \cdot \max(p_{\min}, \delta) \cdot \mathrm{KL}(p^\star, p) \; . \end{aligned} \tag{152}$$

Here, we additionally used the fact that $p(1)p(2) = (1 - p(2))p(2) \leq p(2) = p_{\min}$ and similarly $q(1)q(2) \leq q_{\min}$. Then under the assumptions that $q_{\min} \leq \delta$ and $p_{\min} \leq \delta$, and using the notation $\widehat{w}_{\min} = \max(p_{\min}, q_{\min})$, substituting the bounds of (152) into (145) and simplifying yields

$$\begin{aligned} \|J \nabla F(z)\|_z^2 &= \mathrm{Var}_p(Aq) + \mathrm{Var}_q(A^\top p) \\ &\leq 8 \cdot \|A\|_2^2 \cdot \big(\widehat{w}_{\min} \cdot \delta \cdot \mathrm{KL}(p^\star, p) + \widehat{w}_{\min} \cdot \delta \cdot \mathrm{KL}(q^\star, q) \big) \\ &= 8 \cdot \sigma_{\max}^2 \cdot \widehat{w}_{\min} \cdot \delta \cdot \mathrm{KL}(w^\star, w) \; . \end{aligned}$$

This yields the statement of the proposition. ■

J.2 Proof of Theorem 4.10

We begin the proof by giving a more detailed overview of the main steps:

(1) We consider the payoff matrix $A \in [-1,1]^{2\times 2}$ that is the instantiation of the canonical 2×2 matrix A_{δ_p, δ_q} from Definition I.1, with $\delta_p = \delta_q = \delta \in (0, 1/2)$. Given the time horizon T, we set $\delta = 1/T$, which means the interior NE $w^\star = (p^\star, q^\star)$ of A will be near a vertex of $\mathcal{W}$. We also consider initializations $w_0 \in \mathcal{W}$ that are near the same vertex of $\mathcal{W}$.

(2) Using the control of the OMWU trajectory on the matrix A_{δ_p, δ_q} established in Section I.2.3, our setting of δ and the set of initializations allows for proving that the minimum coordinates $p_{t,\min}$ and $q_{t,\min}$ of the OMWU iterates have an *upper bound* of $\exp(-1/\delta)$ that holds uniformly over all iterates $t \in [T]$. Note that this matches (in its dependence on δ) the general uniform lower bound on $p_{t,\min}$ and $q_{t,\min}$ from Lemma H.2.

(3) Then, using the *lower bound* on energy dissipation from Lemma 4.1 and the *upper bound* on the dissipation term $\|J \nabla F(z)\|_z^2$ from Proposition J.1, we obtain the final rate.

We proceed to give the details of these steps below:

Setup of payoff matrix and set of initializations.

Given $T \geq 3$, let $\delta = 1/T < 1/2$. Let $A := A_{\delta,\delta} \in [-1,1]^{2\times 2}$ be the instantiation of the matrix A_{δ_p, δ_q} from Definition I.1 with $\delta_p = \delta_q = \delta$. Thus A is the matrix

$$A = A_{\delta,\delta} = \begin{pmatrix} \delta^2 & -\delta(1-\delta) \\ -\delta(1-\delta) & (1-\delta)^2 \end{pmatrix} \; .$$

Moreover, recall from Proposition I.2 that A has the unique and interior NE $w^\star = (p^\star, q^\star)$ with $p^\star = q^\star = (1-\delta, \delta) \in \mathrm{ri}(\Delta_2)$. By Proposition B.7, we further have $\sigma_{\max} := \|A\|_2 \leq 2$.

Now, for $w = (p, q) \in \mathcal{W}$ with $p = (p(1), p(2))$ and $q = (q(1), q(2))$, we define the following regions $\mathcal{W}_{\delta,0}, \mathcal{C}_\delta \subset \text{ri}(\mathcal{W})$:

$$\begin{aligned} \mathcal{W}_{\delta,0} &:= \big\{(p,q) \in \text{ri}(\mathcal{W}) \;:\; \exp(-3/\delta) \le p(2), q(2) \le \exp(-2/\delta)\big\} \\ \mathcal{C}_\delta &:= \big\{(p,q) \in \text{ri}(\mathcal{W}) \;:\; p(2), q(2) \le \exp(-1/\delta)\big\}\,. \end{aligned} \tag{153}$$

Observe by definition that $\mathcal{W}_{\delta,0} \subset \mathcal{C}_\delta$, and that $\mathcal{W}_{\delta,0}$ has positive Lebesgue measure for any $\delta > 0$. Now let $\{w_t\}$ be the primal iterates of OMWU initialized from $w_0 \in \mathcal{W}_{\delta,0}$, and let $\{z_t\}$ denote the corresponding dual OMWU iterates. For all $t \ge 0$, let $w_t = (p_t, q_t) \in \text{ri}(\mathcal{W})$ and $z_t = (x_t, y_t) \in \mathcal{Z}$.

Tracking the OMWU trajectory.

We show for all $t \in [T]$ that $w_t \in \mathcal{C}_\delta$. By definition of $\mathcal{C}_\delta$, this implies that the minimum coordinate of each primal iterate p_t and q_t remains bounded by $\exp(-1/\delta)$ for all times $t \in [T]$.

First, using the bounds on the trajectory of the dual iterates from Lemma I.17, observe that if ever $w_{t-1}, w_t \in \mathcal{C}_\delta$, then $x_{t+1}(1) - x_t(1) < 0$ and $y_{t+1}(1) - y_t(1) > 0$. By the primal-dual relationships of Proposition I.15, this further means that $p_{t+1}(1) - p_t(1) < 0$ and $q_{t+1}(1) - q_t(1) > 0$. By definition of $\mathcal{W}_{\delta,0}$, this latter bound implies that $q_t(1) \ge 1 - \exp(-2/\delta) \ge 1 - \exp(-1/\delta)$ for all $t \in [T]$. Thus, in order to show that $w_t \in \mathcal{C}_\delta$ for all $t \in [T]$, it suffices to establish for all iterates that $p_t(1) \ge 1 - \exp(-1/\delta)$.

For this, using the primal-dual relationship of Corollary I.14, observe for any $t \ge 0$ that

$$p_t(1) \ge 1 - \exp(-1/\delta) \iff x_t(1) \ge \delta \cdot \log\big(\tfrac{1-\exp(-1/\delta)}{\exp(-1/\delta)}\big)\,. \tag{154}$$

Moreover, as $w_0 \in \mathcal{W}_{\delta,0}$, it additionally holds by definition of $p_0(1) = 1 - p_0(2)$ and by Corollary I.14 that

$$p_0(2) \le \exp(-2/\delta) \implies p_0(1) \ge 1 - \exp(-2/\delta) \implies x_0(1) \ge \delta \cdot \log\big(\tfrac{1-\exp(-2/\delta)}{\exp(-2/\delta)}\big)\,. \tag{155}$$

Now, observe further from the dual update rule of Proposition I.15 that we can crudely bound

$$x_{t+1}(1) - x_t(1) = -\eta\delta \cdot \big(q_t(1) - (1-\delta)\big) - \eta\delta\big(q_t(1) - q_{t-1}(1)\big) \ge -2\eta\delta\,. \tag{156}$$

Here, the inequality is due to the fact that $\delta, q_t(1), q_{t-1}(1) \in (0,1)$. Thus, for all $t \in [T]$ it holds by (155) and (156) that $x_t(1) \ge x_0(1) - t \cdot 2\eta\delta \ge x_0(1) - t\delta$, where the last inequality is due to the assumption $\eta \le 1/2$. Now recalling that $\delta = 1/T$, this further gives for all $t \in [T]$ that

$$\begin{aligned} x_t(1) \ge x_0(1) - T\delta &\ge \delta \cdot \log\big(\tfrac{1-\exp(-2/\delta)}{\exp(-2/\delta)}\big) - 1 \\ &= \delta \cdot \log\big(\tfrac{1-\exp(-2/\delta)}{\exp(-2/\delta)}\big) - \delta \cdot \log\big(\exp(\tfrac{1}{\delta})\big) \\ &= \delta \cdot \log\big(\tfrac{1-\exp(-2/\delta)}{\exp(1/\delta)\cdot\exp(-2/\delta)}\big) \\ &= \delta \cdot \log\big(\tfrac{1-\exp(-2/\delta)}{\exp(-1/\delta)}\big) \ge \delta \cdot \log\big(\tfrac{1-\exp(-1/\delta)}{\exp(-1/\delta)}\big)\,. \end{aligned}$$

Thus $x_t(1)$ satisfies the bound in (154) for all $t \in [T]$, which implies that also $p_t(1) \ge 1 - \exp(-1/\delta)$ for all $t \in [T]$. Together with the previously established fact that all $q_t(1) \ge 1 - \exp(-1/\delta)$, we then have by definition that $w_t \in \mathcal{C}_\delta$ for all $t \in [T]$.

Bound on change in KL divergence.

Since $w_t \in \mathcal{C}_\delta$ for $0 \le t \le T$, it follows by definition of $\mathcal{C}_\delta$ on this horizon that

$$p_{t,\min} = \min(p_t(1), p_t(2)) \le \exp(-1/\delta) \;\text{ and }\; q_{t,\min} = \min(q_t(1), q_t(2)) \le \exp(-1/\delta)\,.$$

Letting $\widehat{w}_{t,\min} = \max(p_{t,\min}, q_{t,\min})$, we then also have $\widehat{w}_{t,\min} \le \exp(-1/\delta)$ for all $0 \le t \le T$. It follows that each iterate w_t satisfies the preconditions of Proposition J.1. This allows us to bound

$$\|J\nabla F(z_t)\|_{z_t}^2 \le 8 \cdot \sigma_{\max}^2 \cdot \widehat{w}_{t,\min} \cdot \delta \cdot \text{KL}(w^\star, w_t) \le 16 \cdot \exp(-1/\delta) \cdot \text{KL}(w^\star, w_t)\,. \tag{157}$$

Here, we use the fact that $\sigma_{\max} \le 2$ and $\delta \le 1/2$ by definition of A. Moreover, recall from Corollary D.7 that the primal and dual OMWU iterates satisfy $\text{KL}(w^\star, w_{t+1}) - \text{KL}(w^\star, w_t) =$

$F(z_{t+1}) - F(z_t)$ for all times $t \geq 0$. Then applying the *lower bound* on energy dissipation from Lemma 4.1, for $1 \leq t < T$ we can bound

$$\begin{aligned}\mathrm{KL}(w^\star, w_{t+1}) - \mathrm{KL}(w^\star, w_t) &= F(z_{t+1}) - F(z_t) \\ &\geq -\frac{5}{4} \cdot \eta^2 \|J\nabla F(z_t)\|_{z_t}^2 \\ &\geq -20 \cdot \eta^2 \exp(-1/\delta) \cdot \mathrm{KL}(w^\star, w_t) \ .\end{aligned}$$

Here, the final inequality comes from the upper bound on $\|J\nabla F(z_t)\|_{z_t}^2$ from expression (157). Rearranging, we find

$$\begin{aligned}\mathrm{KL}(w^\star, w_{t+1}) &\geq \mathrm{KL}(w^\star, w_t) \cdot (1 - 20\eta^2 \cdot \exp(-1/\delta)) \\ &\geq \mathrm{KL}(w^\star, w_t) \cdot \exp(-40\eta^2 \cdot \exp(-1/\delta)) \ . \end{aligned} \tag{158}$$

The final inequality follows from the fact that $20\eta^2 \cdot \exp(-1/\delta) \leq 1/2$ under the constraint on η in Assumption 1, and from the bound $1 - u \geq \exp(-2u)$ for all $u \in [0, 0.5]$. For the initial step, skew-symmetry gives $F(z_1) - F(z_0) = D_F(z_1, z_0) \geq 0$, as in the proof of Lemma F.12. Thus $\mathrm{KL}(w^\star, w_1) \geq \mathrm{KL}(w^\star, w_0)$. Iterating (158) for $1 \leq t < T$, we conclude for all $0 \leq t \leq T$ that

$$\mathrm{KL}(w^\star, w_t) \geq \mathrm{KL}(w^\star, w_0) \cdot \exp(-40\eta^2 \cdot \exp(-1/\delta) \cdot t) \ . \tag{159}$$

Initial bound on KL divergence.

It remains to show a constant upper bound on $\mathrm{KL}(w^\star, w_0)$ under initializations $w_0 = (p_0, q_0) \in \mathcal{W}_{\delta,0}$. For this, using the definition of $\mathrm{KL}(w^\star, w_0) = \mathrm{KL}(p^\star, p_0) + \mathrm{KL}(q^\star, q_0)$ from (144), we have

$$\mathrm{KL}(p^\star, p_0) = (1-\delta) \cdot \log\left(\tfrac{1-\delta}{p_0(1)}\right) + \delta \cdot \log\left(\tfrac{\delta}{p_0(2)}\right) \leq \delta \cdot \log\left(\tfrac{\delta}{p_0(2)}\right) \ .$$

Here, the inequality comes from the fact that $p_0(1) > 1 - \delta$ for all $(p_0, q_0) \in \mathcal{W}_{\delta,0}$. Additionally using the fact that $p_0(2) \geq \exp(-3/\delta)$ for all $(p_0, q_0) \in \mathcal{W}_{\delta,0}$, we can further derive the bound

$$\mathrm{KL}(p^\star, p_0) \leq \delta \cdot \log\left(\tfrac{\delta}{\exp(-3/\delta)}\right) = \delta \log \delta - \delta\left(\tfrac{-3}{\delta}\right) \leq 3 \ .$$

The final inequality is due to $\delta \log \delta \leq 0$ for $\delta \leq 1$.

By an identical calculation, we also find $\mathrm{KL}(q^\star, q_0) \leq 3$ for all $w_0 = (p_0, q_0) \in \mathcal{W}_{\delta,0}$. Thus it follows for all such w_0 that $\mathrm{KL}(w^\star, w_0) \leq 6$, which concludes the proof. ■

K Details on Uniform Convergence Rate Lower Bounds

In this section, we develop the proof of the uniform best-iterate convergence rate lower bounds of Theorem 5.1 that were introduced in Section 5. To restate the theorem:

Theorem 5.1 (Uniform Best-Iterate Lower Bounds)**.** *For every $T \geq 2$, there exists $A \in [-1, 1]^{2\times 2}$ with an interior Nash equilibrium $w^\star$ such that, for the iterates $\{w_t\}$ of (OMWU) on A with η satisfying Assumption 1, from a positive measure set of initializations that includes the joint uniform distributions, the following hold:*

$$\min_{t\in[T]} \mathrm{TV}(w^\star, w_t) \geq \tfrac{1}{3} \ \ \textit{and} \ \ \min_{t\in[T]} \mathrm{KL}(w^\star, w_t) \geq \tfrac{1}{9} \ .$$

Theorem 5.1 establishes a separation between TV distance and KL divergence and duality gap in terms of uniform best-iterate convergence guarantees for OMWU: while a quantitative uniform best-iterate convergence rate is attainable in duality gap under the uniform initialization (due to Cai et al. (2025), and also Theorem 5.2), no such rate is possible under the stronger notions of TV and KL divergence.

Using the intuition for the geometric bottlenecks developed in Section 4, the hard instances in the proof of the theorem have a (symmetric) NE close to a simplex vertex, and with δ shrinking with T. In this regime, we rely on the primal-dual relationship of the OMWU iterates to exactly track the OMWU dynamics in the 2×2 setting (see Section I for these preliminaries). Subsequently, this allows for establishing on these instances that *all T iterates* of OMWU remain a bounded distance away from equilibrium in TV and KL.

Additional Lower Bounds in KL and DG under Non-Uniform Initializations. In addition to Theorem 5.1, we also prove a second set of uniform best-iterate convergence rate lower bounds that further highlight the different geometries of the distance functions and their interplay with the OMWU iterates. Specifically, when initialized from some positive measure subsets of non-uniform distributions (that may be near the boundary of the simplex), we prove two additional lower bounds: (i) all iterates of OMWU can be a constant distance away from Nash in KL, but convergent in TV and DG, and (ii) all iterates of OMWU can also be a constant distance away from Nash in DG. Formally:

Theorem K.1 (Additional Uniform Best-Iterate Lower Bounds)**.** *For every $T \geq 3$, it holds that:*

(1) There exists $A \in [-1,1]^{2\times 2}$ with interior NE $w^\star$ and a positive measure set of initializations, not including the joint uniform distributions, such that, for the iterates $\{w_t\}$ of (OMWU) on A with stepsize η satisfying Assumption 1:

$$\text{for all } t \in [T], \text{ both } \begin{cases} \mathrm{TV}(w^\star, w_t) \leq O(T^{-1}) \\ \mathrm{DG}(w_t) \leq O(T^{-1}) \end{cases} \quad \text{while also } \min_{t\in[T]} \mathrm{KL}(w^\star, w_t) \geq \tfrac{1}{6} \,.$$

(2) There exists $A \in [-1,1]^{2\times 2}$ with interior NE $w^\star$ and a positive measure set of initializations, not including the joint uniform distributions, such that, for the iterates $\{w_t\}$ of (OMWU) on A with stepsize η satisfying Assumption 1:

$$\min_{t\in[T]} \mathrm{DG}(w_t) \geq \tfrac{1}{6} \,.$$

Part (1) of Theorem K.1 is based on the *universal* last-iterate lower bound in KL from Theorem 4.10. Note that this result gives a separation in uniform best-iterate convergence guarantees between TV and KL (which is not established by Theorem 5.1).

Moreover, Part (2) of Theorem K.1 further shows that, when initialized near the simplex boundary, no best-iterate convergence guarantee is possible in duality gap. Note that this does not contradict the uniform best-iterate convergence rate upper bounds of Theorem 5.2 and of Cai et al. (2025), both of which hold under the assumption of joint uniform initializations.

In the following subsections, we proceed to give the proofs of Theorem 5.1 and Theorem K.1, both of which use the preliminaries for OMWU on the canonical 2×2 instance A_{δ_p,δ_q} from Section I.

K.1 Proof of Theorem 5.1

To prove the theorem, we instantiate the payoff matrix A and the set of initializations as follows:

- Set $A := A_{\delta_p,\delta_q} \in [-1,1]^{2\times 2}$ from Definition I.1 with $\delta := \delta_p = \delta_q = \frac{1}{6\exp(T)}$.
- Initialize $w_0 = (p_0, q_0) \in \mathcal{P}_0 \times \mathcal{Q}_0$, where:

$$\mathcal{P}_0 = \{p \in \mathrm{ri}(\Delta_2) \,:\, p(1) \leq 1/2\} \text{ and } \mathcal{Q}_0 = \{q \in \mathrm{ri}(\Delta_2) \,:\, q(1) \leq 1/2\} \,.$$

Let $\mathcal{W}_0 = \mathcal{P}_0 \times \mathcal{Q}_0 \subset \mathrm{ri}(\mathcal{W})$, and observe that $\mathcal{W}_0$ has positive Lebesgue measure, and that $\mathcal{P}_0$ and $\mathcal{Q}_0$ both contain the uniform distributions $p_0 = (1/2, 1/2)$ and $q_0 = (1/2, 1/2)$. Moreover, by the setting of A, we have by Corollary I.14 that the corresponding dual initializations $(x_0, y_0) \in \mathcal{Z}$ under OMWU satisfy $x_0(1) \leq 0$ and $y_0(1) \leq 0$ for all $(p_0, q_0) \in \mathcal{W}_0$.

The proof then proceeds by (1) tracking the trajectory of the dual OMWU iterates over the time horizon T, which (2) leads to a lower bound on the TV distance and KL divergence from Nash that holds uniformly over all iterates. We proceed to give the details of these two steps:

Bounds on Trajectory. For every t, let $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$ and $z_t = (x_t, y_t) \in \mathcal{Z}$ denote the primal and dual iterate of OMWU. Moreover, let $\widetilde{p}_t = p_t(1) \in (0,1)$, $\widetilde{q}_t = q_t(1) \in (0,1)$ and $\widetilde{x}_t = x_t(1) \in \mathbb{R}$, $\widetilde{y}_t = y_t(1) \in \mathbb{R}$ denote the leading coordinate of each variable. Under $A := A_{\delta,\delta}$, Proposition I.15 then gives for $t+1 \geq 1$ the simplified OMWU dual update rule of

$$\widetilde{x}_{t+1} - \widetilde{x}_t = \eta\delta \cdot \left(1 - \delta - \widetilde{q}_t + \widetilde{q}_{t-1} - \widetilde{q}_t\right) < 2\eta\delta \leq \delta \,.$$

Here, the inequality comes from $0 < \widetilde{q}_t, \widetilde{q}_{t-1} < 1$ and $\delta > 0$, and the assumption that $\eta < 1/2$. It follows that, under the initialization $p_0 \in \mathcal{P}_0$:

$$\widetilde{x}_t \leq \widetilde{x}_0 + \delta \cdot t \leq \delta T \leq \delta \log\left(\tfrac{1}{6\delta}\right) \leq \delta \log\left(\tfrac{1-3\delta}{3\delta}\right) \text{ for all } t \in [T] \,.$$

Here, we use the increment bound above, $\widetilde{x}_0 \leq 0$, and $t \leq T$, together with $\delta = 1/(6\exp(T))$. The final inequality uses $\delta \leq 1/6$, which holds for all $T \geq 2$. By Corollary I.14, it then follows that

$$\widetilde{x}_t \leq \delta \log\left(\tfrac{1-3\delta}{3\delta}\right) \iff \widetilde{p}_t \leq 1 - 3\delta \text{ for all } t \in [T]. \tag{160}$$

By Part (iii) of Lemma I.17, this further implies that $\widetilde{y}_{t+1} - \widetilde{y}_t \leq 0$ for all $t \in [T]$. Under the initializations $q_0 \in \mathcal{Q}_0$ and again using Corollary I.14, it then holds that

$$\widetilde{y}_t \leq \widetilde{y}_0 \leq 0 \implies \widetilde{q}_t \leq \tfrac{1}{2} \text{ for all } t \in [T]. \tag{161}$$

Bounds on Distance to Nash. Recall for $A = A_{\delta,\delta}$ that the unique NE $w^\star = (p^\star, q^\star)$ of A is given by $p^\star = q^\star = (1-\delta, \delta)$. Then by the definition of $\mathrm{TV}(w^\star, w_t)$ from (143), and using the bound $\widetilde{q}_t \leq 1/2$ from (161), we have for all $t \in [T]$ that

$$\mathrm{TV}(w^\star, w_t) = |1 - \delta - \widetilde{p}_t| + |1 - \delta - \widetilde{q}_t| \geq |1 - \delta - \widetilde{q}_t| \geq |\tfrac{1}{2} - \delta| \,. \tag{162}$$

As $\delta \leq 1/6$ for all $T \geq 2$, it then follows that

$$\mathrm{TV}(w^\star, w_t) \geq \tfrac{1}{3} \text{ for all } t \in [T].$$

Finally, by the relationship $\mathrm{KL}(w^\star, w_t) \geq \mathrm{TV}(w^\star, w_t)^2$ from Corollary B.6, we also have

$$\mathrm{KL}(w^\star, w_t) \geq \mathrm{TV}(w^\star, w_t)^2 \geq \tfrac{1}{9} \text{ for all } t \in [T],$$

which concludes the proof. ■

K.2 Proof of Theorem K.1

In this section we give the proof of Theorem K.1. We prove the two parts separately:

K.2.1 Proof of Part 1

The proof follows as a corollary of the *universal* last-iterate lower bound in KL from Theorem 4.10. There, recall from the proof that for all $T \geq 3$, the game instance is $A = A_{\delta,\delta}$ with $\delta = 1/T$. We start by recalling the properties established in the proof of Theorem 4.10:

Conclusions from Proof of Theorem 4.10. Recall from (153) that we define the set of initializations $\mathcal{W}_{\delta,0} \subset \mathrm{ri}(\mathcal{W})$ and the set $\mathcal{C}_\delta \subset \mathrm{ri}(\mathcal{W})$ as:

$$\mathcal{W}_{\delta,0} := \left\{(p,q) \in \mathrm{ri}(\mathcal{W}) \;:\; \exp(-3/\delta) \leq p(2), q(2) \leq \exp(-2/\delta)\right\}$$
$$\mathcal{C}_\delta := \left\{(p,q) \in \mathrm{ri}(\mathcal{W}) \;:\; p(2), q(2) \leq \exp(-1/\delta)\right\}.$$

The proof of Theorem 4.10 established the following two conclusions: for all initial $w_0 \in \mathcal{W}_{\delta,0}$ and for all times $t \in [T]$:

$$w_t \in \mathcal{C}_\delta \text{ and } \mathrm{KL}(w^\star, w_t) \geq \mathrm{KL}(w^\star, w_0) \cdot \exp(-40\eta^2 \cdot \exp(-1/\delta) \cdot t) \,. \tag{163}$$

We now show further implications of these conclusions:

All iterates remain close in TV and duality gap. By definition of the set $\mathcal{C}_\delta$, the fact that $w_t \in \mathcal{C}_\delta$ for all $t \in [T]$ implies that

$$1 - \exp(-1/\delta) \leq p_t(1), q_t(1) \leq 1 \text{ for all } t \in [T]. \tag{164}$$

By the definition of $\mathrm{TV}(w^\star, w)$ from (143), it then follows for all $t \in [T]$ that:

$$\mathrm{TV}(w^\star, w_t) = |p_t(1) - (1-\delta)| + |q_t(1) - (1-\delta)| \leq 2\delta = \tfrac{2}{T} \,.$$

Here, the final equality is due to the setting $\delta = 1/T$. Moreover, using the relationship between $\mathrm{DG}(w_t)$ and $\mathrm{TV}(w^\star, w_t)$ from Corollary B.6, and using the fact that $A_{\delta,\delta} \in [-1,1]^{2\times 2}$, we further conclude for all $t \in [T]$ that $\mathrm{DG}(w_t) \leq 2\mathrm{TV}(w^\star, w_t) \leq 4/T$. Thus it holds on this instance that

$$\mathrm{DG}(w_t) \leq O\big(T^{-1}\big) \text{ and } \mathrm{TV}(w^\star, w_t) \leq O\big(T^{-1}\big) \text{ for all } t \in [T].$$

All iterates remain far in KL. We now show that the conclusions from (159) imply that all iterates $\mathrm{KL}(w^\star, w_t)$ remain bounded below by a constant. First, we show a constant lower bound on $\mathrm{KL}(w^\star, w_0)$ for all initializations $w_0 \in \mathcal{W}_{\delta,0}$. For this, observe that $\mathcal{W}_{\delta,0} \subset \mathcal{C}_\delta$, and thus the bounds in (164) also apply for any initialization $w_0 = (p_0, q_0) \in \mathcal{W}_{\delta,0}$. Then, using the definition of $\mathrm{KL}(w^\star, w_0) = \mathrm{KL}(p^\star, p_0) + \mathrm{KL}(q^\star, q_0)$ from (144), we have

$$\begin{aligned}\mathrm{KL}(p^\star, p_0) &= (1-\delta) \cdot \log\left(\tfrac{1-\delta}{p_0(1)}\right) + \delta \cdot \log\left(\tfrac{\delta}{p_0(2)}\right) \\ &\geq 2 + (1-\delta)\log(1-\delta) + \delta \log \delta \geq 2 - \log 2 > \tfrac{2}{3} \,.\end{aligned}$$

Here, the first inequality uses $p_0(1) \leq 1$ and $p_0(2) \leq \exp(-2/\delta)$, and the second uses the binary entropy bound $-(1-\delta)\log(1-\delta) - \delta \log \delta \leq \log 2$. An identical calculation also gives $\mathrm{KL}(q^\star, q_0) \geq 2/3$, and thus by definition $\mathrm{KL}(w^\star, w_0) \geq 4/3$ for all initializations $w_0 \in \mathcal{W}_{\delta,0}$. Further using that $T = 1/\delta$, it follows from (159) that for all $t \in [T]$:

$$\begin{aligned}\mathrm{KL}(w^\star, w_t) &\geq \mathrm{KL}(w^\star, w_0) \cdot \exp(-40\eta^2 \cdot \exp(-1/\delta) \cdot T) \\ &\geq \tfrac{4}{3} \cdot \exp(-40\eta^2 T/\exp(T)) \\ &\geq \tfrac{4}{3} \cdot \tfrac{1}{\exp(2)} \geq \tfrac{1}{6} \,.\end{aligned}$$

Here, the penultimate inequality comes from the assumption that $\eta \leq \frac{1}{2}$, and since $T/\exp(T) \leq \frac{1}{5}$ for all $T \geq 3$. Thus we find on this instance that $\mathrm{KL}(w^\star, w_t) \geq 1/6$. ■

K.2.2 Proof of Part 2

The structure of the proof is similar to that of Theorem 5.1. First, we instantiate the payoff matrix A and the set of initializations as follows:

- Set $A := A_{\delta_p, \delta_q} \in [-1,1]^{2\times 2}$ from Definition I.1 with $\delta := \delta_p = \delta_q = \frac{1}{\exp(T)}$.
- Initialize $w_0 = (p_0, q_0) \in \mathcal{P}_0 \times \mathcal{Q}_0$, where:

$$\mathcal{P}_0 = \{p \in \mathrm{ri}(\Delta_2) \,:\, \tfrac{1}{2} \leq p(1) \leq \tfrac{3}{4}\} \;\text{ and }\; \mathcal{Q}_0 = \{q \in \mathrm{ri}(\Delta_2) \,:\, 1 - \tfrac{\delta^2}{3} \leq q(1) \leq 1 - \tfrac{\delta^2}{6}\} \,.$$

Let $\mathcal{W}_0 = \mathcal{P}_0 \times \mathcal{Q}_0 \subset \mathrm{ri}(\mathcal{W})$, and observe that $\mathcal{W}_0$ has positive Lebesgue measure. Moreover, for every $T \geq 3$, we have $\delta < 1/8$, and thus all initializations $p_0 \in \mathcal{P}_0$ satisfy $p_0(1) < 1 - \delta$. On the other hand, note that $\mathcal{W}_0$ does not contain the uniform initialization $p_0 = q_0 = (1/2, 1/2)$, which is in contrast to the proof of Theorem 5.1.

We now proceed to (1) track the trajectory of the dual OMWU iterates over the time horizon T, which (2) will lead to a constant lower bound on the duality gap that holds over all iterates. In more details:

Bounds on Trajectory. For every t, let $w_t = (p_t, q_t) \in \mathrm{ri}(\mathcal{W})$ and $z_t = (x_t, y_t) \in \mathcal{Z}$ denote the primal and dual iterate of OMWU, and let $\widetilde{p}_t = p_t(1) \in (0,1)$, $\widetilde{q}_t = q_t(1) \in (0,1)$ and $\widetilde{x}_t = x_t(1) \in \mathbb{R}$, $\widetilde{y}_t = y_t(1) \in \mathbb{R}$. Now under $A := A_{\delta,\delta}$, we can use Proposition I.15 to crudely bound for any $t \geq 0$

$$\widetilde{y}_{t+1} - \widetilde{y}_t = -\eta\delta(1 - \delta - \widetilde{p}_t - (\widetilde{p}_t - \widetilde{p}_{t-1})) \geq -2\eta\delta \geq -\delta \,. \tag{165}$$

Here, the first inequality is due to $\widetilde{p}_t, \widetilde{p}_{t-1}, \delta \in (0,1)$, and the second is by the assumption that $\eta < 1/2$. Now under the initialization $q_0 \in \mathcal{Q}_0$, we have $\widetilde{q}_0 \geq 1 - \delta^2/3$. By Corollary I.14, this implies that

$$\widetilde{q}_0 \geq 1 - \tfrac{\delta^2}{3} \implies \widetilde{y}_0 \geq \delta \log\left(\tfrac{1-\delta^2/3}{\delta^2/3}\right) \,. \tag{166}$$

Then using (165), it follows for all $t \in [T]$ that we can bound

$$\begin{aligned}\widetilde{y}_t \geq \widetilde{y}_0 - \delta \cdot t \geq \widetilde{y}_0 - \delta \cdot T &\geq \delta \log\left(\tfrac{1-\delta^2/3}{\delta^2/3}\right) - \delta \log\left(\tfrac{1}{\delta}\right) \\ &= \delta \log\left(\tfrac{1-\delta^2/3}{\delta/3}\right) \geq \delta \log\left(\tfrac{1-\delta/3}{\delta/3}\right) \,.\end{aligned} \tag{167}$$

Here, the second inequality comes from applying the lower bound on $\widetilde{y}_t$ from (166) and using the setting $\delta = 1/\exp(T) \implies T = \log(1/\delta)$. Again by Corollary I.14, it follows for all $t \in [T]$ that

$$\widetilde{y}_t \geq \delta \log\left(\tfrac{1-\delta/3}{\delta/3}\right) \implies \widetilde{q}_t \geq 1 - \delta/3 \,. \tag{168}$$

Moreover, using the fact that $\widetilde{p}_0 \le 3/4 < 1-\delta$ for all $p_0 \in \mathcal{P}_0$, this further implies via Part (ii) of Lemma I.17 that all $\widetilde{x}_{t+1} - \widetilde{x}_t < 0 \implies \widetilde{p}_{t+1} - \widetilde{p}_t < 0$. Thus for all $t \in [T]$ it holds that

$$\widetilde{p}_t \le \widetilde{p}_0 \le 3/4 < 1-\delta \,. \tag{169}$$

We now derive the implications of the uniform coordinate bounds of (168) and (169) on the duality gap of the iterates.

Bounds on Distance to Nash. For $A = A_{\delta,\delta}$, the unique NE $w^\star = (p^\star, q^\star)$ of A is given by $p^\star = q^\star = (1-\delta, \delta)$. Then by the definition of $\mathrm{DG}(w_t)$ from (142), and using the bounds $\widetilde{q}_t \ge 1 - \delta/3$ and $\widetilde{p}_t < 1-\delta$, we have for all $t \in [T]$ that

$$\begin{aligned}
\mathrm{DG}(w_t) &= (1-\delta-\widetilde{p}_t)(1-\delta) + (\widetilde{q}_t - (1-\delta))(1-\delta) && (170)\\
&\ge (1/4 - \delta)\cdot(1-\delta) && (171)\\
&\ge (1/4 - 1/20)\cdot(19/20) \ge 1/6 \,. && (172)
\end{aligned}$$

The first inequality comes from the fact that the second term in (170) is non-negative, and also by applying the upper bound $\widetilde{p}_t \le 3/4$. The second inequality comes from the fact that $\delta = 1/\exp(T)$, and thus $\delta \le 1/20$ for all $T \ge 3$. As the bound of (172) holds uniformly for all $t \in [T]$, we conclude on this instance that $\min_{t\in[T]} \mathrm{DG}(w_t) \ge 1/6$, which concludes the proof. ■

L Details on Best-Iterate Convergence Rate in Duality Gap for 2x2 Setting

In this section, we develop the proof of Theorem 5.2, which establishes a fast uniform best-iterate convergence rate upper bound in duality gap for the 2×2 setting. To restate the theorem:

Theorem 5.2 (Uniform Best-Iterate Convergence Rate in DG)**.** *Let $A \in [-1,1]^{2\times 2}$ be a zero-sum game with an interior Nash equilibrium $w^\star$, and let $\{w_t\}$ denote the iterates of running* (OMWU) *on A with a fixed absolute constant stepsize η satisfying Assumption 1, initialized at the uniform distributions. Then for any $T \ge 2$, it holds that*

$$\min_{t\in[T]} \mathrm{DG}(w_t) \le O(T^{-1/2}\cdot\sqrt{\log T}) \,.$$

Note that this result improves over the prior best-known $O(T^{-1/6})$ rate of Cai et al. (2025) for the same setting. This improvement is due to a refined proof technique that leverages the new analysis of OMWU in KL divergence from Theorem 4.7. We give a brief comparison with the analysis of Cai et al. (2025) and a high-level overview of our proof here:

Comparison with Cai et al. (2025). We recall that the proof of Cai et al. (2025) uses a two-regime approach depending on the magnitude of the time horizon T with respect to the minimum Nash coordinate $\delta = \min(\delta_p, \delta_q)$. After an initial phase, they show a best-iterate bound in duality gap scaling like $\min(\delta, T^{-1/4}\delta^{-1/2})$. Depending on whether $T \ge \delta^{-6}$ or $T < \delta^{-6}$, this term always is at most $T^{-1/6}$. During the initial phase, their bound is $O(\log^2(T)/T)$, which is also $O(T^{-1/6})$. For large T, their analysis relies on a universal upper bound on the random-iterate convergence in duality gap. For small T, their analysis relies on tracking the trajectory of the OMWU iterates in the primal space.

High-level overview of proof. In contrast to the proof of Cai et al. (2025), our proof uses a three-regime approach that allows for much tighter bounds when T is large with respect to δ. In the largest regime, we use our new universal last-iterate analysis of OMWU in KL divergence (Theorem 4.7) to establish a fast *last-iterate* bound in DG. For the moderate and small T regimes, we directly track the OMWU trajectory and obtain tighter best-iterate bounds of $O(\delta_p\delta_q)$ and $O(\max\{1/(\eta T), \delta_p\})$, respectively. Together, using the definition of the three regimes, this leads to the overall $\widetilde{O}(T^{-1/2})$ best-iterate bound in duality gap.

We give more details on these regimes and their consequences in the proof overview in the following section. However, we first make two additional remarks regarding the assumption of uniform initialization, and on the possibility of generalizing our techniques to higher-dimensional settings.

Remark L.1 (On Uniform Initialization)**.** The proof of Theorem 5.2 assumes a uniform initialization (similar to the results of Daskalakis and Panageas (2019), Wei et al. (2021), Cai et al. (2024), Cai et al. (2025)). As we showed in part (ii) of Theorem K.1, there is in general no quantitative uniform convergence rate when initialized near the boundary of the simplex.

Remark L.2 (On Generalizing to Higher Dimensions)**.** Even in the 2×2 setting, the proof for obtaining the $\widetilde{O}(T^{-1/2})$ uniform best-iterate convergence rate in duality gap is intricate and involves understanding the interplay between (a) the trajectory of the OMWU iterates and (b) the non-uniform rate of energy dissipation over the iterates. In the 2×2 setting, the low-dimensional property of the effective dual space (as detailed in Section I.2.1) allows for obtaining fine-grained control of the iterate trajectory. This helps for establishing a phase transition in the rate of energy dissipation when T is sufficiently large (see Lemma L.5 below). While such a transition empirically holds for higher-dimensional settings, establishing an analogous bound remains a challenging and open technical question. Moreover, controlling the trajectory of the iterates in the 2×2 setting also allows for proving that the primal iterates always enter a region of the simplex with a sufficiently small duality gap, even when far from Nash under a metric like TV distance (see Sections I.4 and I.4.4). Thus, extending the result to higher-dimensions likely requires some additional control of the iterates (e.g., establishing the number of steps spent in various regions of the primal and dual spaces, as a function of the location of the interior NE). We conjecture that OMWU still obtains a $\widetilde{O}(T^{-1/2})$ uniform best-iterate convergence rate in the general-dimension setting, and we leave this for future work.

L.1 Proof of Theorem 5.2

In this section, we give the proof of Theorem 5.2. The proof uses several intermediate components and results that we outline here, and which we prove in the later subsections.

Assumptions on NE. If A is constant, its duality gap is zero and the theorem is immediate. Henceforth, assume that A is not a constant matrix. Then by Part (ii) of Proposition I.3, since $w^\star = (p^\star, q^\star)$ is an interior NE of A, then $w^\star$ is the unique NE of A. Let $p^\star = (1 - \delta_p, \delta_p) \in \mathrm{ri}(\Delta_2)$ and $q^\star = (1 - \delta_q, \delta_q) \in \mathrm{ri}(\Delta_2)$ for $\delta_p, \delta_q \in (0, 1)$. By relabeling actions, we may take $0 < \delta_p, \delta_q \leq 1/2$. If $\delta_p > \delta_q$, replace A by $-A^\top$ and exchange p_t and q_t. This preserves the OMWU updates, duality gap, uniform initialization, and $\|A\|_2$, and orders the masses. Henceforth, assume $0 < \delta_p \leq \delta_q \leq 1/2$. We retain both possible signs of the canonical coefficient.

Assumptions on Initialization. We assume that the initialization $w_0 = (p_0, q_0) \in \mathrm{ri}(\mathcal{W})$ is the uniform distribution $p_0 = (1/2, 1/2)$ and $q_0 = (1/2, 1/2)$, with $w_{-1} = w_0$.

Considering Canonical 2×2 Matrix Suffices. The proof relies on the properties of the 2×2 setting introduced in Section I. In particular, we recall from Proposition I.2 that, for $v \in \mathbb{R}$, $0 < \gamma \leq 4$, and $\varepsilon \in \{-1, 1\}$, the matrix A can be decomposed as

$$A = \gamma \varepsilon A_{\delta_p, \delta_q} + v \mathbf{1}\mathbf{1}^\top .$$

Here, $\gamma = |a - b - c + d|$ and $\varepsilon = \mathrm{sign}(a - b - c + d)$ for $A = ((a, b), (c, d))$. The matrix $A_{\delta_p, \delta_q} \in [-1, 1]^{2 \times 2}$ is the canonical 2×2 matrix from Definition I.1 with a unique and interior NE $w^\star = (p^\star, q^\star)$ (i.e., the same NE as A). The key consequence of this decomposition is that, to establish bounds on duality gap for the iterates of OMWU on A, it suffices to analyze the iterates of OMWU on $\varepsilon A_{\delta_p, \delta_q}$ with an appropriately scaled stepsize. In particular, we have the following equivalence:

Proposition L.3 (OMWU on Canonical Matrix)**.** *Fix $A \in [-1, 1]^{2 \times 2}$ with an interior Nash equilibrium $w^\star = (p^\star, q^\star)$, where $p^\star = (1 - \delta_p, \delta_p)$ and $q^\star = (1 - \delta_q, \delta_q)$ for $\delta_p, \delta_q \in (0, 1)$. Assume $A = ((a, b), (c, d))$ is not constant, and set $\varepsilon = \mathrm{sign}(a - b - c + d)$. Let $A_{\delta_p, \delta_q} \in [-1, 1]^{2 \times 2}$ be the payoff matrix from Definition I.1. Fix the initialization $w_0 = (p_0, q_0)$ for $p_0 = q_0 = (0.5, 0.5)$ and $w_{-1} = w_0$. Then there exists $0 < \gamma \leq 4$ such that the following holds:*

- *Let $\{w_t\}$ be the iterates of (OMWU) on A with stepsize $\eta > 0$, initialized from w_0.*
- *Let $\{w'_t\}$ be the iterates of (OMWU) on $\varepsilon A_{\delta_p, \delta_q}$ with stepsize $\eta' = \eta\gamma > 0$, initialized from w_0.*

Then $w_t = w'_t$ for all $t \geq 0$.

Moreover, we also have the following relationship between the duality gaps under A and $\varepsilon A_{\delta_p, \delta_q}$:

Proposition L.4 (Relationship between Duality Gaps)**.** *Let $\mathrm{DG}_A(\cdot)$ denote the duality gap under matrix A, and let $\mathrm{DG}_{A'}(\cdot)$ denote the duality gap under matrix A', where $A = \gamma \cdot A' + v\mathbf{1}\mathbf{1}^\top$ for constants $\gamma, v \in \mathbb{R}$ with $\gamma > 0$. Then for any $w = (p, q) \in \mathcal{W}$ and value $B > 0$:*

$$\mathrm{DG}_{A'}(w) \leq B \implies \mathrm{DG}_A(w) \leq \gamma B .$$

The proofs of Proposition L.3 and Proposition L.4 are in Section L.2 and we grant them as true for now. As a result, the remainder of the proof focuses on establishing bounds on the best-iterate in duality gap for the OMWU iterates when run on the matrix $\varepsilon A_{\delta_p,\delta_q}$ with scaled stepsize. For this, we record the stepsize convention used below:

Assumptions on Stepsize for OMWU on Canonical Matrix. For OMWU run on $\varepsilon A_{\delta_p,\delta_q}$, with $\varepsilon \in \{-1, 1\}$, we use the following convention:

Assumption 2 (Stepsize on Canonical Matrix)**.** The canonical run is obtained by Proposition L.3 from a run on the original matrix whose stepsize satisfies Assumption 1.

In particular, writing $\widehat{\eta}$ for the original stepsize, the canonical stepsize is $\eta = \gamma\widehat{\eta}$. For $e = (1, -1)/\sqrt{2}$, the decomposition gives $\gamma = 2|e^\top Ae| \le 2\|A\|_2$, and hence

$$\eta \le \frac{2\|A\|_2}{4(54\|A\|_2 + 9)} \le \frac{1}{108}. \tag{173}$$

The restriction of A to $\mathbf{1}^\perp$, with orthogonal projection onto $\mathbf{1}^\perp$, is multiplication by $\varepsilon\gamma/2$. Thus its minimum restricted singular value is $\gamma/2$. Applying Lemma H.1 and Corollary H.3 to the original run, and using equality of the primal iterates, gives

$$\begin{aligned} \mathrm{KL}(w^\star, w_1) &\le \frac{5}{4}\mathrm{KL}(w^\star, w_0), \\ \mathrm{KL}(w^\star, w_{t+1}) &\le \mathrm{KL}(w^\star, w_t)\exp\left(-\frac{\eta^2}{80}p_{t,\min}q_{t,\min}\right) \quad (t \ge 1). \end{aligned} \tag{174}$$

Note that Assumption 1 is imposed on the original matrix and its original stepsize. However, the resulting KL bounds above also hold for the canonical run because its primal iterates are identical, even if the rescaled stepsize fails the numerical condition in Assumption 1 when evaluated using the canonical matrix.

Uniform Best-Iterate Duality Gap Bounds on Canonical 2×2 Matrix. Given the equivalence of Proposition L.3, the core technical component of the proof is to establish uniform best-iterate bounds in duality gap for OMWU run on $\varepsilon A_{\delta_p,\delta_q}$. For this, we prove three separate bounds that depend on the magnitude of the time horizon T with respect to the minimum Nash coordinates δ_p and δ_q. We refer to these regimes as *large* T, *moderate* T, and *small* T, and we give details below.

First, we make use of the values $V, L > 0$, which we define as

$$V = \frac{\min\{\delta_p, \delta_q\}}{20}, \qquad L = 80, \qquad \tau_{\mathrm{ep}} = \left\lfloor \frac{L}{\eta\delta_p\delta_q} \right\rfloor + 1\,. \tag{175}$$

We then define the three regimes in terms of V, L and η, where η is the stepsize used for the iterates of OMWU on $\varepsilon A_{\delta_p,\delta_q}$.

Large T regime. In the large T regime, we assume that

$$\frac{150000 \cdot L \cdot \log(2/V)}{\eta^2(\delta_p\delta_q)^2} \le T\,, \tag{Large T}$$

In this regime, we prove an upper bound on the *last-iterate* in KL divergence, which translates into an upper bound on the last-iterate in duality gap. In particular, we prove the following:

Lemma L.5 (Large T Last-Iterate Bound)**.** *Fix $0 < \delta_p, \delta_q \le \frac{1}{2}$, and let $A := A_{\delta_p,\delta_q}$. Let $\{w_t\}$ be the iterates of OMWU on A with η satisfying Assumption 2. Let T be as in* (Large T)*. Then:*

$$\mathrm{DG}(w_T) \le \sqrt{4\mathrm{KL}(w^\star, w_T)} \le O\Big(\exp\big(-\tfrac{1}{2}\eta\sqrt{T}\big)\Big)\,.$$

The proof of the lemma is in Section L.3.1.

Moderate T regime. In the moderate T regime, we assume that

$$\tau_{\mathrm{ep}} \le T < \frac{150000 \cdot L \cdot \log(2/V)}{\eta^2(\delta_p\delta_q)^2}\,. \tag{Moderate T}$$

For this regime, we prove the following bound on the best-iterate in duality gap in terms of the NE parameters δ_p and δ_q. Formally, we prove:

Lemma L.6 (Moderate T Best-Iterate Bound). *Fix $0 < \delta_p, \delta_q \le \frac{1}{2}$, and let $A := A_{\delta_p,\delta_q}$. Let $\{w_t\}$ be the iterates of OMWU on A with η satisfying Assumption 2. Let T be as in (Moderate T). Then there exists $t \in [T]$ such that*

$$\mathrm{DG}(w_t) \le O\big(\delta_p \delta_q\big) .$$

The proof of the lemma is in Section L.3.2.

For the negative canonical matrix, exchanging the players gives $-(-A_{\delta_p,\delta_q})^\top = A_{\delta_q,\delta_p}$. In particular,

$$\mathrm{DG}_{-A_{\delta_p,\delta_q}}(p_t, q_t) = \mathrm{DG}_{A_{\delta_q,\delta_p}}(q_t, p_t).$$

The KL divergence, stepsize, and regime thresholds are unchanged. The exchanged run is induced by $-A^\top$, so it also satisfies Assumption 2. Since Lemmas L.5 and L.6 do not require an ordering of the masses, their conclusions therefore hold for both canonical signs.

Small T regime. Finally, in the small T regime, we assume that

$$1 \le T < \tau_{\mathrm{ep}} . \tag{Small T}$$

In this regime, we prove the following bound:

Lemma L.7 (Small T Best-Iterate Bound). *Fix $0 < \delta_p \le \delta_q \le \frac{1}{2}$ and $\varepsilon \in \{-1, 1\}$, and let $A := \varepsilon A_{\delta_p,\delta_q}$. Let $\{w_t\}$ be the iterates of OMWU on A with η satisfying Assumption 2. Let T be as in (Small T). Then there exists a time $t \in [T]$ such that*

$$\mathrm{DG}(w_t) \le O\Big(\max\Big\{ \frac{1}{\eta T}, \delta_p \Big\}\Big) .$$

We note that the statement of the lemma is similar to that of Theorem 5 of Cai et al. (2025). We give the proof in Section L.3.3.

Moreover, we note that, in the proofs of each of Lemmas L.5, L.6, and L.7, the key technical component involves controlling the amount of time until the iterates enter certain regions of the primal space (defined relative to $w^\star$), and the number of iterations spent within these regions. Using the tools introduced in Section I.3, we establish such bounds in Lemma L.10, which we state and prove in Section L.2.3.

Combining the Pieces. Using the preceding lemmas, the final bound on the best-iterate in duality gap is straightforward. For this, let $\{w_t\}$ be the iterates of running OMWU on A with stepsize η, and let $\{w'_t\}$ be the iterates of running OMWU on $\varepsilon A_{\delta_p,\delta_q}$ with stepsize $\eta' = \gamma\eta$, as described in Proposition L.3. Fix an integer $T \ge 2$. In this paragraph, all regime thresholds and τ_{ep} are evaluated at the canonical stepsize η'.

Then for the iterates $\{w'_t\}$, we have by Lemmas L.5, L.6, and L.7, together with the sign symmetry above, the following:

- If T satisfies (Large T), then by Lemma L.5,

$$\mathrm{DG}_{\varepsilon A_{\delta_p,\delta_q}}(w'_T) \le O(\exp(-\tfrac{1}{2}\eta'\sqrt{T})) \le O\left(\frac{1}{\eta'\sqrt{T}}\right). \tag{176}$$

- If T satisfies (Moderate T), write $\delta = \min\{\delta_p, \delta_q\}$. By its lower endpoint, $T \ge \tau_{\mathrm{ep}} > L/(\eta'\delta_p\delta_q)$. Since $\delta \ge \delta_p\delta_q$ and $L > \eta'$, this implies $\delta > 1/T$. Consequently, $\log(2/V) = \log(40/\delta) \le \log(40T)$. The upper endpoint of the regime therefore gives

$$\delta_p\delta_q \le \frac{\sqrt{150000 L \log(2/V)}}{\eta'\sqrt{T}} \le \frac{\sqrt{150000 L \log(40T)}}{\eta'\sqrt{T}}.$$

  Since $T \ge 2$, we have $\log(40T) \le 7\log T$. Thus Lemma L.6 gives

$$\min_{t\in[T]} \mathrm{DG}_{\varepsilon A_{\delta_p,\delta_q}}(w'_t) \le O(\delta_p\delta_q) \le O\left(\frac{\sqrt{\log T}}{\eta'\sqrt{T}}\right). \tag{177}$$

- Finally, if T satisfies (Small T), integrality and $\delta_p \le \delta_q$ give

$$T \le \frac{L}{\eta'\delta_p\delta_q} \le \frac{L}{\eta'\delta_p^2}, \qquad \delta_p \le \sqrt{\frac{L}{\eta' T}} \le \frac{\sqrt{L}}{\eta'\sqrt{T}}.$$

Here the last inequality uses $\eta' \leq 1$. By Lemma L.7, we conclude

$$\min_{t\in[T]} \mathrm{DG}_{\varepsilon A_{\delta_p,\delta_q}}(w'_t) \leq O\left(\max\left\{\frac{1}{\eta' T}, \delta_p\right\}\right) \leq O\left(\frac{1}{\eta'\sqrt{T}}\right). \tag{178}$$

Thus using the fact that $\eta' = \gamma\eta$, for every $T \geq 2$, combining expressions (176), (177), and (178) gives

$$\min_{t\in[T]} \mathrm{DG}_{\varepsilon A_{\delta_p,\delta_q}}(w'_t) \leq O\left(\frac{\sqrt{\log T}}{\gamma\eta\sqrt{T}}\right).$$

Finally, by Proposition L.4, it follows that

$$\min_{t\in[T]} \mathrm{DG}_{A}(w_t) \leq O\left(\frac{\sqrt{\log T}}{\eta\sqrt{T}}\right).$$

Noting that the original stepsize η is an absolute constant concludes the proof. ■

L.2 Sufficiency of OMWU on Canonical 2×2 Matrix

Here, we give the proofs of Propositions L.3 and L.4, which are used in the proof of Theorem 5.2.

L.2.1 Proof of Proposition L.3

Proof. Let A have entries $A = ((a,b),(c,d)) \in \mathbb{R}^{2\times 2}$. Let $v = \langle p^\star, Aq^\star\rangle \in \mathbb{R}$. For readability, let $A' = \varepsilon A_{\delta_p,\delta_q}$. For each $t \geq 0$, let $w_t = (p_t, q_t)$ and $w'_t = (p'_t, q'_t)$. Now, directly from the definition of (OMWU), iterating the dual updates under uniform initialization and $w_{-1} = w_0$ gives, for $t \geq 1$:

$$\begin{cases} p_t = \mathrm{softmax}(-\eta(\sum_{k=1}^{t-1} Aq_k + Aq_{t-1})) \\ q_t = \mathrm{softmax}(\eta(\sum_{k=1}^{t-1} A^\top p_k + A^\top p_{t-1})) \end{cases}$$

$$\text{and} \quad \begin{cases} p'_t = \mathrm{softmax}(-\eta'(\sum_{k=1}^{t-1} A'q'_k + A'q'_{t-1})) \\ q'_t = \mathrm{softmax}(\eta'(\sum_{k=1}^{t-1} A'^\top p'_k + A'^\top p'_{t-1})). \end{cases}$$

We will show inductively that $p_t = p'_t$ and $q_t = q'_t$ for all $t \geq 0$. The base case at time $t = 0$ trivially holds by assumption in the proposition. Now suppose the claim also holds for all $0 \leq k \leq t-1$. Now recall from Part (ii) of Proposition I.2 that for $\gamma = |a - b - c + d| \in (0, 4]$, A can be decomposed as

$$A = \gamma A' + v\mathbf{1}\mathbf{1}^\top .$$

Thus at time t, observe that the cumulative payoff vector can be written as

$$\begin{aligned} -\eta\Big(\sum_{k=1}^{t-1} Aq_k + Aq_{t-1}\Big) &= -\eta\Big(\sum_{k=1}^{t-1}(\gamma A'q_k + v\mathbf{1}) + \gamma A'q_{t-1} + v\mathbf{1}\Big) \\ &= -\eta\gamma\Big(\sum_{k=1}^{t-1} A'q'_k + A'q'_{t-1}\Big) - \eta v t\mathbf{1} . \end{aligned} \tag{179}$$

Here, we applied in the second line the inductive hypothesis $q_k = q'_k$ for all $0 \leq k \leq t-1$. Then applying the softmax operator to (179), it follows that

$$\begin{aligned} p_t &= \mathrm{softmax}\Big(-\eta\Big(\sum_{k=1}^{t-1} Aq_k + Aq_{t-1}\Big)\Big) \\ &= \mathrm{softmax}\Big(-\eta\gamma\Big(\sum_{k=1}^{t-1} A'q'_k + A'q'_{t-1}\Big) - \eta v t\mathbf{1}\Big) \\ &= \mathrm{softmax}\Big(-\eta'\Big(\sum_{k=1}^{t-1} A'q'_k + A'q'_{t-1}\Big)\Big) = p'_t , \end{aligned}$$

where the penultimate equality comes from the invariance of softmax to constant shifts, and the final equality comes from the definition of p'_t. Repeating an identical calculation, we similarly find $q_t = q'_t$, which establishes $w_t = w'_t$ for all $t \geq 0$ and concludes the proof. ■

Remark L.8. We note that the statement of the result is similar to Proposition 2 of Cai et al. (2025), but using a different "base matrix" (see also Remark I.4). Also, observe that the proof of Proposition L.3 only relies on the invariance of the softmax function to constant shifts. Thus, the statement also holds under other MWU variants (e.g., standard MWU, c.f. Section C.2.2). Finally, we assume for simplicity that the initialization $w_0 = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is uniform, but the proof also extends in a straightforward way to non-uniform initializations.

L.2.2 Proof of Proposition L.4

Proof. Recall definition of duality gap (i.e., from Section 2) that for any $w = (p, q) \in \mathcal{W}$, we have

$$\mathrm{DG}_A(w) = \max_{q' \in \Delta_2} \langle q', A^\top p \rangle - \min_{p' \in \Delta_2} \langle p', Aq \rangle$$
$$\text{and } \mathrm{DG}_{A'}(w) = \max_{q' \in \Delta_2} \langle q', A'^\top p \rangle - \min_{p' \in \Delta_2} \langle p', A'q \rangle \ .$$

Now using the fact that $A = \gamma A' + v\mathbf{1}\mathbf{1}^\top$ for $\gamma, v \in \mathbb{R}$ and $\gamma > 0$, it follows that

$$\begin{aligned}
\mathrm{DG}_A(w) &= \gamma \cdot \max_{q' \in \Delta_2} \langle q', A'^\top p \rangle + v - \big(\gamma \cdot \min_{p' \in \Delta_2} \langle p', A'q \rangle + v\big) \\
&= \gamma \cdot \big(\max_{q' \in \Delta_2} \langle q', A'^\top p \rangle - \min_{p' \in \Delta_2} \langle p', A'q \rangle \big) \\
&= \gamma \cdot \mathrm{DG}_{A'}(w) \ .
\end{aligned}$$

Thus it follows that if $\mathrm{DG}_{A'}(w) \le B$ for any $B > 0$, then $\mathrm{DG}_A(w) \le \gamma B$. ■

L.2.3 Properties of Iterate Trajectory

In this section, we state and prove two key technical lemmas regarding the OMWU iterates on the canonical 2×2 matrix A_{δ_p, δ_q}. In the following lemma, we start by establishing (somewhat crude) uniform bounds on the coordinates of the primal and dual iterates when initialized from the uniform distribution.

Lemma L.9 (Energy and Coordinate Bounds on Canonical 2x2 Matrix)**.** *Fix $0 < \delta_p, \delta_q \le \frac{1}{2}$. Let $A := A_{\delta_p, \delta_q} \in [-1, 1]^{2\times 2}$ be the matrix from Definition I.1. Let $\{w_t\}$ and $\{z_t\}$ be the primal and dual iterates of OMWU for η satisfying Assumption 2 and initialized from the uniform distribution with $w_{-1} = w_0$. For each $t \ge 0$, let $z_t = (x_t, y_t)$ and let $w_t = (p_t, q_t)$. Then the following bounds hold, with (ii) and (iii) for all $t \ge 0$:*

(i) $F(z_t) \le F(z_1) \le (5/4)F(z_0) < 2$ *for* $t \ge 1$, *and* $F(z_0) < 2$.

(ii) $-\big(\frac{\delta_p}{1-\delta_p}\big) \cdot 2 \le x_t(1) \le 2$ *and* $-\big(\frac{\delta_q}{1-\delta_q}\big) \cdot 2 \le y_t(1) \le 2$.

(iii) $\min\{p_t(1), p_t(2)\} \ge \min\{\frac{1}{60}, 1 - p_t(1)\}$ *and* $\min\{q_t(1), q_t(2)\} \ge \min\{\frac{1}{60}, 1 - q_t(1)\}$.

Proof. First, under the uniform initialization $p_0 = (1/2, 1/2)$ and $q_0 = (1/2, 1/2)$, observe by the primal-dual relationship in Part (iii) of Proposition I.11 that $x_0 = (0, 0)$ and $y_0 = (0, 0)$. Thus by definition of $F(z_0)$, we have by Part (i) of Proposition I.11 that

$$F(z_0) = x_0(1) + \log(1 + \exp(\tfrac{-x_0(1)}{\delta_p})) + y_0(1) + \log(1 + \exp(\tfrac{-y_0(1)}{\delta_q})) = 2\log 2 \ .$$

Write $F_\star = \mathrm{Ent}(p^\star) + \mathrm{Ent}(q^\star) \ge 0$. By Proposition D.11, $F(z_t) = F_\star + \mathrm{KL}(w^\star, w_t)$. Equation (174) gives $\mathrm{KL}(w^\star, w_t) \le \mathrm{KL}(w^\star, w_1) \le (5/4)\mathrm{KL}(w^\star, w_0)$ for $t \ge 1$. Consequently, $F(z_t) \le F(z_1) \le (5/4)F(z_0)$ for $t \ge 1$. Thus for all $t \ge 0$, we have

$$F(z_t) \le \big(\tfrac{5}{4}\big) \cdot F(z_0) \le \tfrac{5}{2}\log 2 \le 2 \ .$$

This proves claim (i).

Claim (ii) then follows by a direct application of Proposition I.18 using the bound $F(z_t) \le 2$.

For claim (iii), observe from Proposition I.15 (and also by Part (ii) of Proposition I.11) that $p_t(1) = \mathrm{sigmoid}(x_t(1)/\delta_p)$ for all $t \ge 0$. Thus, using $x_t(1) \ge -2(\delta_p/(1-\delta_p))$, we can further bound

$$p_t(1) \ge \frac{1}{1 + \exp(2/(1-\delta_p))} \ge \frac{1}{60} \ ,$$

where the final inequality holds for $\delta_p \le 1/2$, since $\exp(4) < 59$. Thus

$$\min\{p_t(1), p_t(2)\} = \min\{p_t(1), 1 - p_t(1)\} \ge \min\{1/60, 1 - p_t(1)\} .$$

Using an identical calculation, we similarly find $\min\{q_t(1), q_t(2)\} \ge \min\{1/60, 1 - q_t(1)\}$, which yields claim (iii) of the lemma and concludes the proof. ■

The next lemma establishes a more fine-grained control of the trajectory of the OMWU iterates on A_{δ_p,δ_q}. In particular, the lemma proves two bounds in parallel: first, it gives bounds on the amount of time until the iterates enter a region of the primal space near the interior NE. Second, it establishes a lower bound on the number of steps taken in regions where the minimum coordinates of the primal iterates are uniformly lower bounded. This latter property is used in the proof of Theorem 5.2 to establish a faster decay in KL divergence (compared to the worst-case bound of Theorem 4.7) when T is in the large regime. Formally, we prove the following:

Lemma L.10 (Epoch invariants)**.** *Fix $0 < \delta_p, \delta_q \le \frac{1}{2}$, and let $A := A_{\delta_p,\delta_q} \in [-1, 1]^{2\times 2}$ be the matrix from Definition I.1. Let $\{w_t\}$ be the iterates of OMWU on A with stepsize $\eta > 0$ satisfying Assumption 2, initialized uniformly with $w_{-1} = w_0$, and let $w_t = (p_t, q_t)$. Let $\tau = 80/(\eta\delta_p\delta_q)$. Then for any integer $t' \ge 0$:*

(a) There exists at least one iteration $t' \le t \le t' + \tau$ such that

$$p_t(1) \ge 1 - \delta_p \ \text{ and } \ 1 - 3\delta_q \le q_t(1) \le 1 - \delta_q .$$

(b) There exist at least $\frac{1}{\eta}$ consecutive iterations $t' \le t \le t' + \tau$ such that

$$p_t(1) \le 1 - \delta_p \ \text{ and } \ q_t(1) \le 1 - \delta_q .$$

Proof. We use the dual update rule to follow the iterates through the four quadrants around equilibrium. Bounding the duration of these visits will give a fourth-to-first quadrant crossing for (a), together with a sufficiently long visit to the third quadrant for (b).

If $\delta_p = \delta_q = 1/2$, the uniform initialization is an equilibrium and all iterates are equal to it. Both properties then hold at every index, and $\lfloor\tau\rfloor + 1 \ge 1/\eta$ gives the required count. We may therefore assume that the equilibrium is not uniform. We write $d = \delta_p\delta_q$, $M = \mathrm{Diag}(\delta_p, \delta_q)$, and $\widetilde{p}_t = p_t(1)$, $\widetilde{q}_t = q_t(1)$. The equilibrium dual coordinates are

$$x^\star(1) = \delta_p \log\frac{1 - \delta_p}{\delta_p}, \qquad y^\star(1) = \delta_q \log\frac{1 - \delta_q}{\delta_q}.$$

We center the dual coordinates at equilibrium and write the primal deviations as follows:

$$\zeta_t = \begin{pmatrix} x_t(1) - x^\star(1) \\ y_t(1) - y^\star(1) \end{pmatrix}, \qquad h_t = \begin{pmatrix} \widetilde{p}_t - (1 - \delta_p) \\ \widetilde{q}_t - (1 - \delta_q) \end{pmatrix}.$$

The coordinates of h_t have the same signs as those of ζ_t. We measure $\zeta \in \mathbb{R}^2$ in the weighted norm

$$\|\zeta\|^2_{M^{-2}} := \frac{\zeta(1)^2}{\delta_p^2} + \frac{\zeta(2)^2}{\delta_q^2}.$$

With $R(a, b) = (-b, a)$, Proposition I.15 gives the following update, including the initial step with $\zeta_{-1} = \zeta_0$ and $h_{-1} = h_0$:

$$\zeta_{t+1} - \zeta_t = \eta M R(2h_t - h_{t-1}) \qquad (t \ge 0). \tag{180}$$

In particular, the first and second coordinates change by at most $2\eta\delta_p$ and $2\eta\delta_q$, respectively, per step.

Control of the optimistic correction. Recall that the sigmoid function is $1/4$-Lipschitz. We first use this property to show that, for every $t \ge 0$,

$$\|h_{t-1}\|_2 \le 2\|h_t\|_2, \qquad \|h_t - h_{t-1}\|_2 \le 2\eta\|h_t\|_2. \tag{181}$$

Both inequalities hold at $t = 0$, since $h_{-1} = h_0$. For the induction step, we assume the first holds at time t and obtain

$$\|h_{t+1} - h_t\|_2 \le \frac{1}{4}\|\zeta_{t+1} - \zeta_t\|_{M^{-2}} \le \frac{\eta}{4}(2\|h_t\|_2 + \|h_{t-1}\|_2) \le \eta\|h_t\|_2.$$

Thus $\|h_{t+1}\|_2 \geq (1-\eta)\|h_t\|_2$, which proves both inequalities at time $t+1$, since $\eta \leq 1/108$. Consequently, writing $r = \eta(1+2\eta)/4$, we obtain

$$\|\zeta_{t+1} - \zeta_t\|_{M^{-2}} \leq \eta(1+2\eta)\|h_t\|_2 \leq r\|\zeta_t\|_{M^{-2}}. \tag{182}$$

The last inequality again uses the Lipschitz property, now between the current and equilibrium dual coordinates. Since $r < 1$ and $\zeta_0 \neq 0$, no ζ_t is zero.

To compare the dual and primal deviations, we consider the pairing $E_t = \langle M^{-1}\zeta_t, h_t\rangle$. Each term pairs a scaled dual-coordinate deviation with the corresponding primal probability deviation. Their signs agree by monotonicity of the primal-dual map, so E_t is positive away from equilibrium. More precisely, it is the symmetrized KL divergence:

$$E_t = \mathrm{KL}(w^\star, w_t) + \mathrm{KL}(w_t, w^\star).$$

Indeed, the primal-dual map and $p_t(2) - p^\star(2) = -h_t(1)$ give

$$\frac{\zeta_t(1)h_t(1)}{\delta_p} = \sum_{i=1}^{2}(p_t(i) - p^\star(i))\log\frac{p_t(i)}{p^\star(i)} = \mathrm{KL}(p^\star, p_t) + \mathrm{KL}(p_t, p^\star).$$

Adding the analogous equality for q_t proves the identity. This positive pairing controls the amount of turning in the update. We first record two bounds that relate it to the primal and dual deviations:

$$E_t \geq 4\|h_t\|_2^2, \qquad E_t \geq \frac{1}{6}\|\zeta_t\|_2^2. \tag{183}$$

The first bound follows coordinatewise from the Lipschitz property and the matching signs. For the second bound, we fix $\delta \in (0, 1/2]$, write $c = \delta\log((1-\delta)/\delta)$, and consider $f(s) = \mathrm{sigmoid}((c+s)/\delta) - (1-\delta)$. The explicit sigmoid formula and $e^a \geq 1 + a$ give

$$|f(s)| \geq \frac{\delta(1-\delta)|s|}{\delta + |s|} \qquad (s \in \mathbb{R}).$$

Indeed, with $a = |s|/\delta$ and $b \in \{\delta, 1-\delta\}$, $|f(s)|$ equals $\delta(1-\delta)(e^a-1)/(1+b(e^a-1))$, which is at least $\delta(1-\delta)a/(1+a)$. By Lemma L.9(ii), each relevant centered dual coordinate satisfies $\delta + |s| \leq 3$. For its upper endpoint, $s \leq 2 - c \leq 2$ and $\delta + 2 \leq 3$. For its lower endpoint, $s \geq -2\delta/(1-\delta) - c$, and

$$\delta + c + \frac{2\delta}{1-\delta} \leq 1 + \frac{2\delta}{1-\delta} \leq 3,$$

where $c \leq \delta\log(1/\delta) \leq 1 - \delta$ was used. Therefore $(s/\delta)f(s) \geq s^2/6$, proving the second bound.

We now use these bounds to estimate the optimistic correction in (180). Since $|\frac{d}{ds}\log f'(s)| \leq 1/\delta$, integration between 0 and s gives

$$|f(s)| \geq \delta f'(s)(1 - e^{-|s|/\delta}) \geq \delta f'(s)\frac{|s|}{\delta + |s|}.$$

We apply the mean value theorem between consecutive dual coordinates. Since their distance is at most $2\eta\delta$, we obtain the following bound, where $\delta_1 = \delta_p$ and $\delta_2 = \delta_q$:

$$\begin{aligned}|\langle M^{-1}\zeta_t, h_t - h_{t-1}\rangle| &\leq e^{2\eta}\sum_{i=1}^{2}\left(1 + \frac{|\zeta_t(i)|}{\delta_i}\right)|h_t(i)|\frac{|\zeta_t(i) - \zeta_{t-1}(i)|}{\delta_i} \\ &\leq e^{2\eta}\big(3\eta\|h_t\|_2^2 + 2\eta E_t\big) \leq 3\eta e^{2\eta}E_t \leq \frac{E_t}{20}.\end{aligned}$$

For the second line, we used the coordinate step bounds and

$$\|\zeta_t - \zeta_{t-1}\|_{M^{-2}} \leq \eta(1+2\eta)\|h_{t-1}\|_2 \leq 3\eta\|h_t\|_2.$$

We bound the part of the sum coming from 1 by Cauchy–Schwarz and this norm estimate, and the part coming from $|\zeta_t(i)|/\delta_i$ by the coordinate step bound. For $t = 0$ the correction is zero, so the same estimate holds. The last line uses (183) and $\eta \leq 1/108$. Consequently, the determinant of the matrix with columns ζ_t and ζ_{t+1} satisfies

$$\begin{aligned}\det(\zeta_t, \zeta_{t+1}) &:= \det\begin{pmatrix}\zeta_t(1) & \zeta_{t+1}(1)\\ \zeta_t(2) & \zeta_{t+1}(2)\end{pmatrix} = \zeta_t(1)\zeta_{t+1}(2) - \zeta_t(2)\zeta_{t+1}(1) \\ &= \eta d\langle M^{-1}\zeta_t, 2h_t - h_{t-1}\rangle \geq \frac{19}{20}\eta d E_t > 0.\end{aligned} \tag{184}$$

Cyclic quadrant crossings and maximum visit duration. Here $(a, b) \in \mathbb{R}^2$ denotes a generic pair of centered dual coordinates. We assign the boundary rays consistently by taking the following disjoint quadrants:

$$Q_{\mathrm{I}} = \{(a,b) : a \geq 0,\ b > 0\}, \quad Q_{\mathrm{II}} = \{(a,b) : a < 0,\ b \geq 0\},$$
$$Q_{\mathrm{III}} = \{(a,b) : a \leq 0,\ b < 0\}, \quad Q_{\mathrm{IV}} = \{(a,b) : a > 0,\ b \leq 0\}.$$

Equations (182) and (184) imply that a step either stays in its current quadrant or enters the next quadrant in the order $Q_{\mathrm{I}}, Q_{\mathrm{II}}, Q_{\mathrm{III}}, Q_{\mathrm{IV}}, Q_{\mathrm{I}}$. Indeed, the positive determinant excludes a crossing in the reverse order. Crossing both axes would require $|\zeta_{t+1}(i) - \zeta_t(i)| \geq |\zeta_t(i)|$ for both coordinates, contradicting (182).

To bound the duration of each visit, we use the following fractions in the four quadrants, respectively:

$$\frac{|\zeta_t(2)|}{\|\zeta_t\|_1}, \quad \frac{|\zeta_t(1)|}{\|\zeta_t\|_1}, \quad \frac{|\zeta_t(2)|}{\|\zeta_t\|_1}, \quad \frac{|\zeta_t(1)|}{\|\zeta_t\|_1}.$$

Each fraction measures progress from the entry axis toward the exit axis of its quadrant. For successive iterates in the same quadrant, the increase of this fraction is exactly $\det(\zeta_t, \zeta_{t+1})/(\|\zeta_t\|_1\|\zeta_{t+1}\|_1)$. Moreover, (183) gives $\|\zeta_t\|_2 \leq \sqrt{6}\sqrt{E_t}$, and

$$\|\zeta_{t+1}\|_2 \leq \left(\sqrt{6} + \frac{\eta(1+2\eta)}{4}\right)\sqrt{E_t} \leq \frac{5}{2}\sqrt{E_t}.$$

Here we used $\|M\|_2 \leq 1/2$ and $\|h_t\|_2 \leq \sqrt{E_t}/2$. Thus each such increase is at least

$$\frac{(19/20)\eta d E_t}{(25/2)E_t} \geq \frac{\eta d}{14}.$$

Since the fraction lies in $[0, 1]$, a consecutive visit to any one quadrant contains at most $14/(\eta d) + 1$ iterates. In particular, every visit ends, and the four quadrants are crossed in the stated order indefinitely.

Verifying the two properties. For (a), let t be the last iterate of a visit to Q_{IV}. Then $\zeta_t(1) > 0$, $\zeta_t(2) \leq 0$, and $\zeta_{t+1}(2) > 0$. The coordinate step bound implies $-2\eta\delta_q \leq \zeta_t(2) \leq 0$. By the primal-dual relationship, it follows that

$$\widetilde{p}_t \geq 1 - \delta_p, \qquad \delta_q \leq 1 - \widetilde{q}_t \leq e^{2\eta}\delta_q < 3\delta_q,$$

which proves property (a) whenever this last iterate lies in the epoch.

For (b), we need a lower bound on the duration of a complete visit to Q_{III}. Let $s \geq 1$ be its first index after entering from Q_{II}, and let e be its last index before entering Q_{IV}. We write $k = e - s + 1$ for the number of iterates in this visit. At entry, the second coordinate crosses zero, so (182) gives

$$\frac{|\zeta_s(2)|}{\delta_q} \leq r\|\zeta_{s-1}\|_{M^{-2}} \leq \frac{r}{1-r}\|\zeta_s\|_{M^{-2}} \leq \frac{1}{2}\|\zeta_s\|_{M^{-2}}.$$

The last inequality uses $\eta \leq 1/108$. Thus $|\zeta_s(1)|/\delta_p \geq \|\zeta_s\|_{M^{-2}}/2$. At exit, $\zeta_{s+k}(1) > 0$, whereas $\zeta_s(1) \leq 0$. Therefore,

$$\|\zeta_{s+k} - \zeta_s\|_{M^{-2}} \geq \frac{|\zeta_s(1)|}{\delta_p} \geq \frac{1}{2}\|\zeta_s\|_{M^{-2}}.$$

On the other hand, iterating (182) gives, for every integer $j \geq 0$,

$$\|\zeta_{s+j}\|_{M^{-2}} \leq (1+r)^j\|\zeta_s\|_{M^{-2}},$$
$$\|\zeta_{s+j} - \zeta_s\|_{M^{-2}} \leq r\sum_{i=0}^{j-1}(1+r)^i\|\zeta_s\|_{M^{-2}} = \big((1+r)^j - 1\big)\|\zeta_s\|_{M^{-2}}.$$

If $k < 1/\eta$, then $r \leq \eta/3$ implies

$$(1+r)^k - 1 \leq e^{rk} - 1 < e^{1/3} - 1 < \frac{1}{2},$$

contradicting the displacement bound at exit. Consequently, $e - s + 1 = k \geq 1/\eta$. These k consecutive iterates all satisfy $\widetilde{p}_t \leq 1 - \delta_p$ and $\widetilde{q}_t \leq 1 - \delta_q$, which proves property (b).

It remains to place both events within the stated epoch. Starting at any integer t', at most five successive quadrant visits (the first possibly incomplete) contain both a last iterate in Q_{IV} and a complete visit to Q_{III}. The worst case is when t' falls partway through Q_{III}: we then count its remaining part, followed by visits to Q_{IV}, Q_{I}, Q_{II}, and Q_{III}. Their total length is at most

$$5\left(\frac{14}{\eta d}+1\right) \le \frac{80}{\eta d} = \tau.$$

All the indicated indices therefore lie between t' and $t'+\tau$, completing the proof. ■

L.3 Proof of Lemmas L.5, L.6, and L.7

L.3.1 Proof of Lemma L.5

To prove the lemma for the regime (Large T), we analyze the iterates of OMWU in two phases. In the first startup phase, the energy dissipation (and thus contraction in KL) is fast for at least $1/\eta$ steps in every epoch of $O(1/(\eta\delta_p\delta_q))$ iterations. This follows from establishing a tighter lower bound on $p_{t,\min}$ and $q_{t,\min}$ at those steps. After $O(\log(2/V)/(\eta^2(\delta_p\delta_q)^2))$ iterations, KL divergence will be sufficiently small to establish lower bounds on the minimum coordinates for all remaining iterations. Under (Large T), this leads to the required last-iterate bound, as we show below.

KL threshold for uniformly bounded coordinates. Let $V>0$ be the threshold from (175). For $0<b<a<1$, binary KL divergence satisfies

$$a\log\frac{a}{b}+(1-a)\log\frac{1-a}{1-b} = \int_b^a \frac{a-s}{s(1-s)}\,ds \ge \frac{(a-b)^2}{2a}.$$

Thus if a coordinate b is less than two thirds of its equilibrium value a, that player's KL divergence exceeds $a/18$. Every equilibrium coordinate is at least $\min\{\delta_p,\delta_q\}=20V$. Consequently, $\mathrm{KL}(w^\star,w_t)\le V$ implies

$$p_{t,\min}\ge 2\delta_p/3, \qquad q_{t,\min}\ge 2\delta_q/3.$$

By (174), $\mathrm{KL}(w^\star,w_t)$ is nonincreasing for $t\ge 1$. Therefore, once this KL threshold is reached at a time $t\ge 1$, the two coordinate bounds hold at all subsequent times. We will show that the threshold is reached by time $T/2$.

Energy decay in startup phase. By property (b) of Lemma L.10, every epoch contains at least $1/\eta$ consecutive iterations at which $p_t(1)\le 1-\delta_p$ and $q_t(1)\le 1-\delta_q$. Part (iii) of Lemma L.9 therefore gives

$$p_{t,\min}\ge \min\{1/60,\delta_p\}\ge \delta_p/30, \qquad q_{t,\min}\ge \min\{1/60,\delta_q\}\ge \delta_q/30.$$

Here we used $\delta_p,\delta_q\le 1/2$. Define the integers K and T_{startup} by

$$K=\left\lceil \frac{72000\log(2/V)}{\eta\delta_p\delta_q}\right\rceil, \qquad T_{\text{startup}} = 1+K\tau_{\mathrm{ep}}, \tag{185}$$

where τ_{ep} is the integer epoch length in (175). Apply property (b) at the starting times $1+j\tau_{\mathrm{ep}}$, for $0\le j<K$. Its indices lie in the disjoint blocks $\{1+j\tau_{\mathrm{ep}},\dots,(j+1)\tau_{\mathrm{ep}}\}$. At every good index, the transported one-step KL bound (174) gives

$$\begin{aligned}\mathrm{KL}(w^\star,w_{t+1}) &\le \mathrm{KL}(w^\star,w_t)\cdot\exp\Big(-\frac{\eta^2}{80}p_{t,\min}q_{t,\min}\Big)\\ &\le \mathrm{KL}(w^\star,w_t)\cdot\exp\Big(-\frac{\eta^2\delta_p\delta_q}{72000}\Big).\end{aligned}$$

At all other indices $t\ge 1$, $\mathrm{KL}(w^\star,w_t)$ does not increase. Since $\mathrm{KL}(w^\star,w_1)\le (5/4)\cdot\mathrm{KL}(w^\star,w_0)\le (5/2)\log 2<2$, we obtain

$$\mathrm{KL}(w^\star,w_{T_{\text{startup}}})\le \mathrm{KL}(w^\star,w_1)\exp\Big(-\frac{\eta^2\delta_p\delta_q}{72000}\frac{K}{\eta}\Big)\le 2\cdot\frac{V}{2}=V.$$

To account for the rounding in K and τ_{ep}, put $d=\delta_p\delta_q$ and $\ell=\log(2/V)$. Since $\eta d\le 1/432$, $\ell\ge\log 80>1$, and $L=80$,

$$\begin{aligned}T_{\text{startup}} &\le 1+\Big(\frac{72000\ell}{\eta d}+1\Big)\Big(\frac{L}{\eta d}+1\Big)\\ &= \frac{L\ell}{\eta^2 d^2}\Big(72000+\frac{72000\eta d}{L}+\frac{\eta d}{\ell}+\frac{2\eta^2 d^2}{L\ell}\Big)\le \frac{75000L\ell}{\eta^2 d^2}\le\frac{T}{2}.\end{aligned}$$

Here, the last inequality follows from the definition of (Large T).

Final last-iterate bound in KL Divergence and Duality Gap. By the preceding arguments, for all $t \geq T_{\text{startup}}$, $p_{t,\min} \geq 2\delta_p/3$ and $q_{t,\min} \geq 2\delta_q/3$. Again applying (174), we find

$$\mathrm{KL}(w^\star, w_T) \leq V \exp\left(-\frac{\eta^2 \delta_p \delta_q}{180}(T - T_{\text{startup}})\right) \leq V \exp\left(-\frac{\eta^2 \delta_p \delta_q T}{360}\right). \tag{186}$$

By definition of the regime (Large T),

$$\delta_p \delta_q \geq \frac{\sqrt{150000 L \log(2/V)}}{\eta\sqrt{T}} \geq \frac{360}{\eta\sqrt{T}}.$$

Here the final inequality uses $L = 80$ and $V \leq 1/40$. Finally, Corollary B.6 and $V \leq 1$ give

$$\mathrm{DG}(w_T) \leq \sqrt{4\mathrm{KL}(w^\star, w_T)} \leq 2\exp\left(-\frac{\eta^2 \delta_p \delta_q T}{720}\right) \leq 2\exp(-\eta\sqrt{T}/2),$$

which concludes the proof. ■

L.3.2 Proof of Lemma L.6

For the regime (Moderate T), recall by part (a) of Lemma L.10, applied at $t' = 1$, that, within the first τ_{ep} iterations (for τ_{ep} defined in (175)) there exists at least one iteration $t \in [\tau_{\text{ep}}]$ such that

$$p_t(1) \geq 1 - \delta_p \;\text{ and }\; 1 - 3\delta_q \leq q_t(1) \leq 1 - \delta_q\,. \tag{187}$$

For such an iterate $w_t = (p_t, q_t)$, observe then that

$$p_t(1) - (1 - \delta_p) \leq \delta_p \;\text{ and }\; (1 - \delta_q) - q_t(1) \leq 2\delta_q\,.$$

Then using the characterization of $\mathrm{DG}(\cdot)$ under A_{δ_p,δ_q} from (142) in Section I.4.1, this means that we can bound for this iterate

$$\mathrm{DG}(w_t) \leq (p_t(1) - (1 - \delta_p)) \cdot \delta_q + (1 - \delta_q - q_t(1)) \cdot \delta_p \leq \delta_p\delta_q + 2\delta_p\delta_q \leq 3\delta_p\delta_q\,. \tag{188}$$

For the regime (Moderate T), recall that we have $T \geq \tau_{\text{ep}}$. Thus by definition, we are guaranteed the iterate w_t satisfying (187) (and thus (188)) occurs within T total iterations. ■

L.3.3 Proof of Lemma L.7

We prove the stated bound for every integer $T \geq 1$. Write $A = \varepsilon A_{\delta_p,\delta_q}$, where $\varepsilon \in \{-1, 1\}$ and $0 < \delta_p \leq \delta_q \leq 1/2$. The initialization is $p_0 = q_0 = (1/2, 1/2)$ and $w_{-1} = w_0$. For readability, write $\widetilde{p}_t = p_t(1)$ and $\widetilde{q}_t = q_t(1)$. We use $\eta \leq 1/108$, as in (173). If $\delta_p > 1/10$, then $\mathrm{DG}(w_t) \leq 2 \leq 20\delta_p$ for every t, so the claim holds for either sign. Henceforth assume $\delta_p \leq 1/10$.

Positive canonical matrix with small equilibrium masses. Suppose $\varepsilon = 1$ and $0 < \delta_p \leq \delta_q \leq 1/10$. Now, observe that the update formulas of Proposition I.15 also hold at $t = 0$ because $w_{-1} = w_0$. In particular, $y_1(1) - y_0(1) = \eta\delta_q(\delta_p - 1/2) < 0$. Thus, using the fact that initially $p_0 = q_0 = (1/2, 1/2)$, we have by Part (iii) of Lemma I.17 and the primal-dual relationship of Proposition I.15 the following: while $p_{t-1}(1), p_t(1) < 1 - 3\delta_p$, then

$$y_{t+1}(1) - y_t(1) < 0 \implies q_{t+1}(1) - q_t(1) < 0\,. \tag{189}$$

The proof then distinguishes between the following two cases:

(a) There exists $1 \leq t \leq T$ such that $p_t(1) \geq 1 - 3\delta_p$.

(b) It holds for all $1 \leq t \leq T$ that $p_t(1) < 1 - 3\delta_p$.

For Case (a), let t be the first such time where the iterate $w_t = (p_t, q_t)$ satisfies $p_t(1) \geq 1 - 3\delta_p$. By (189), this further implies $q_t(1) \leq q_0(1) \leq 1/2 < 1 - \delta_q$, where the final inequality holds under the assumption $\delta_q \leq 1/10$. If $p_t(1) \leq 1 - \delta_p$, then by definition of $\mathrm{DG}(\cdot)$ under $A = A_{\delta_p,\delta_q}$ from (142) in Section I.4.1, it holds that

$$\mathrm{DG}(w_t) = (1 - \delta_p - p_t(1))(1 - \delta_q) + (1 - \delta_q - q_t(1))\delta_p \leq 2\delta_p + \delta_p = 3\delta_p\,. \tag{190}$$

If instead $p_t(1) > 1 - \delta_p$, the other applicable case of (142) gives

$$\mathrm{DG}(w_t) = \delta_p(1 - q_t(1)) - \delta_q(1 - p_t(1)) \leq \delta_p.$$

Thus the bound $\mathrm{DG}(w_t) \le 3\delta_p$ holds in either case.

For Case (b), expression (189) further implies that $q_t(1) \le q_{t-1}(1) \le q_0(1) = 1/2 \le 1 - 3\delta_q$ for all $t \in [T]$, where the final inequality holds for $\delta_q \le 1/10$. Now for each t, let $z_t = (x_t, y_t) \in \mathcal{Z}$ be the corresponding dual OMWU iterate. By the update rule of Proposition I.15, it follows that

$$\begin{aligned} x_{t+1}(1) - x_t(1) &= \eta\delta_p \cdot \big(1 - \delta_q - q_t(1) - (q_t(1) - q_{t-1}(1))\big) \\ &\ge \eta\delta_p \cdot \big(1 - \delta_q - q_t(1)\big) && \text{(since } q_t(1) - q_{t-1}(1) \le 0\text{)} \\ &\ge \eta\delta_p \cdot \big(\tfrac{1}{2} - \delta_q\big) && \text{(since } q_t(1) \le 1/2\text{)} \\ &\ge \tfrac{1}{4} \cdot \eta\delta_p && \text{(since } \delta_q \le 1/10 \le 1/4\text{)} \,. \end{aligned}$$

Thus for all $t \in [T]$ we find

$$x_t(1) \ge x_0(1) + t \cdot \tfrac{\eta}{4}\delta_p = t \cdot \tfrac{\eta}{4}\delta_p \quad \Longrightarrow \quad p_t(1) \ge \mathrm{sigmoid}\big(\tfrac{t\eta\delta_p}{4\delta_p}\big) = \mathrm{sigmoid}\big(\tfrac{t\eta}{4}\big) \,. \tag{191}$$

Here, the final equality in the lefthand side comes from $x_0(1) = 0$ under $p_0 = (1/2, 1/2)$, and the implication is due to the primal-dual relationship of Proposition I.15. Again using the definition of $\mathrm{DG}(\cdot)$ under $A = A_{\delta_p,\delta_q}$ from (142), we then find for all $t \in [T]$ that

$$\begin{aligned} \mathrm{DG}(w_t) &= (1 - \delta_p - p_t(1))(1 - \delta_q) + (1 - \delta_q - q_t(1))\delta_p \\ &\le 1 - \delta_p - p_t(1) + \delta_p \\ &= 1 - p_t(1) \le 1 - \mathrm{sigmoid}(\tfrac{t\eta}{4}) \le \tfrac{4}{\eta t} \,, \end{aligned}$$

where the final inequality holds due to $t\eta > 0$. Thus we conclude for Case (b) that

$$\min_{t\in[T]} \mathrm{DG}(w_t) \le \frac{4}{\eta T} \,. \tag{192}$$

The statement of the lemma in this case follows by combining the conclusions of expression (190) for Case (a) and expression (192) for Case (b).

Positive canonical matrix with mixed equilibrium masses. Suppose $\varepsilon = 1$ and $0 < \delta_p \le 1/10 < \delta_q < 1/2$. Let h be the first index with $\widetilde{p}_h \ge 1 - 3\delta_p$, and set $\tau_{\mathrm{init}} = \lceil 1/(10\eta) \rceil$. The preceding monotonicity argument shows that $\widetilde{q}_t$ decreases up to h, and the update gives $x_{t+1}(1) - x_t(1) \ge \eta\delta_p(1/2 - \widetilde{q}_t) \ge 0$ for $t < h$. Thus $\widetilde{p}_t \ge 1/2$ up to h. Using $x_0(1) = 0$, $|x_{t+1}(1) - x_t(1)| \le 2\eta\delta_p$, and $\eta\tau_{\mathrm{init}} < 11/100$, the primal-dual map gives $\widetilde{p}_j < \mathrm{sigmoid}(22/100) < 3/5$ for $j \le \tau_{\mathrm{init}}$. Hence $h > \tau_{\mathrm{init}}$. For $0 \le j < \tau_{\mathrm{init}}$, the update and its sum from $y_0(1) = 0$ give

$$\begin{aligned} \frac{y_{j+1}(1) - y_j(1)}{\delta_q} &\le \eta\left(\frac{6}{5} - \frac{1}{2} - \frac{9}{10}\right) = -\frac{\eta}{5}, \\ \widetilde{q}_{\tau_{\mathrm{init}}} &\le \mathrm{sigmoid}(-1/50) \le \frac{50}{101}. \end{aligned}$$

Therefore, for $\tau_{\mathrm{init}} \le t < h$, monotonicity strengthens our drift bound to $x_{t+1}(1) - x_t(1) \ge \eta\delta_p/202$. This ensures that the crossing occurs, where Case (a) gives $\mathrm{DG}(w_h) \le 3\delta_p$ whenever $h \le T$. For $\tau_{\mathrm{init}} \le T < h$, the gap estimate of Case (b), followed by summing this drift, gives

$$\mathrm{DG}(w_T) \le 1 - \widetilde{p}_T \le \exp\left(-\frac{\eta(T - \tau_{\mathrm{init}})}{202}\right) \le \frac{202}{\eta T}.$$

Here we used $x_{\tau_{\mathrm{init}}}(1) \ge 0$, $\eta\tau_{\mathrm{init}} < 11/100$, and $ze^{1-z} \le 1$ for $z > 0$. If $T < \tau_{\mathrm{init}}$, the bound follows from $\mathrm{DG}(w_T) \le 1 \le 1/(\eta T)$. This proves the required bound. The endpoint $\delta_q = 1/2$ follows by continuity of the finite OMWU trajectory and its best-iterate gap, since the constants above are independent of δ_q.

Negative canonical matrix. Now suppose $\varepsilon = -1$, and write $u_t = 1 - \widetilde{p}_t$ and $v_t = 1 - \widetilde{q}_t$. Negating the matrix leaves the effective dual space unchanged, so the same primal-dual formulas apply. As in the preceding increment estimates, the changes in $x_t(1)$ and $y_t(1)$ have magnitude at most $2\eta\delta_p$ and $2\eta\delta_q$, respectively. Thus consecutive deficit ratios lie between $e^{-2\eta}$ and $e^{2\eta}$, giving $97r_t/100 \le 2r_t - r_{t-1} \le 103r_t/100$ for $r = u, v$. Let $s = \inf\{t : v_t \le 9\delta_q/10\}$ and set $\tau_{\mathrm{init}} = \lceil 4\log(2/\delta_q)/\eta \rceil$. For $t \le \min\{s, \tau_{\mathrm{init}}\}$, the update rule gives

$$\eta\delta_q\tau_{\mathrm{init}} \le 4\delta_q \log(2/\delta_q) + \eta\delta_q \le 2\log 4 + 1/216 < 3,$$

$$\text{and} \quad \frac{x_t(1)}{\delta_p} = \eta\textstyle\sum_{j=0}^{t-1}(\delta_q - v_j) + \eta(v_0 - v_{t-1}) \le \dfrac{\eta\delta_q t}{10} + \dfrac{\eta}{2} < \dfrac{31}{100}.$$

Hence $u_t > 2/5$, and the other update gives $y_{t+1}(1) - y_t(1) \geq \eta\delta_q(97u_t/100 - \delta_p) \geq \eta\delta_q/4$ for $t < \min\{s, \tau_{\text{init}}\}$. If $s > \tau_{\text{init}}$, then $v_{\tau_{\text{init}}} \leq e^{-\eta\tau_{\text{init}}/4} \leq \delta_q/2$, a contradiction. Thus $s \leq \tau_{\text{init}}$, $\eta\delta_q s < 3$, and $v_t \leq e^{-\eta t/4}$ for $t \leq s$. Since v_t decreases up to s, the same telescoping identity gives

$$\frac{x_s(1)}{\delta_p} \geq -\eta \sum_{j=0}^{s-1} v_j \geq -\frac{\eta}{1 - e^{-\eta/4}} \geq -5.$$

Now let $h = \inf\{t > s : u_t \leq 3\delta_p\}$. We have $u_s > 2/5 > 3\delta_p$; while $s \leq t < h$, the ratio bounds give

$$y_{t+1}(1) - y_t(1) \geq \eta\delta_q(97u_t/100 - \delta_p) > 0,$$
$$x_{t+1}(1) - x_t(1) \geq \eta\delta_p(\delta_q - 103v_t/100) \geq \frac{\eta\delta_p\delta_q}{20}.$$

Since the first inequality makes $y_t(1)$ increase, v_t decreases and remains at most $v_s \leq 9\delta_q/10$, yielding the final bound in the second line. This positive drift also forces u_t below $3\delta_p$ in finite time. Applying (142) after exchanging players, for $T < s$ we obtain $\text{DG}(w_T) \leq \max\{v_T, \delta_q\} \leq 3v_T$; for $s \leq T < h$, they give $\text{DG}(w_T) = \delta_q u_T - \delta_p v_T$. Consequently, the preceding bounds imply for $T < s$ that

$$\text{DG}(w_T) \leq 3e^{-\eta T/4} \leq \frac{12}{\eta T},$$

and for $s \leq T < h$ that

$$\text{DG}(w_T) \leq \delta_q \exp\left(5 - \frac{\eta\delta_q(T-s)}{20}\right) \leq \delta_q \exp\left(6 - \frac{\eta\delta_q T}{20}\right) \leq \frac{20e^5}{\eta T}.$$

Here the last line uses $\eta\delta_q s < 3$ and $ze^{-z/20} \leq 20/e$ for $z > 0$. At h, the same gap formulas give $\text{DG}(w_h) \leq 3\delta_p\delta_q$ if $u_h \geq \delta_p$, and $\text{DG}(w_h) \leq \delta_p$ otherwise. Retaining this iterate for $T \geq h$ proves the required bound $O(\max\{1/(\eta T), \delta_p\})$ for the negative canonical matrix.

The cases above exhaust both signs and all ordered masses in $(0, 1/2]$, and prove Lemma L.7. ■

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Yes, the paper's main contribution is to give new theoretical insights on the convergence of OMWU, for which we prove several new properties.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer:

   [Yes]

   Justification:

   Yes, we clearly state our assumptions regarding the class of games (two-player zero-sum) and the Nash equilibrium (assumed to be interior and unique).

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer:

[Yes]

Justification:

Yes, we include the proofs of all theoretical results, mostly in the appendix sections.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   [N/A]

   Justification:

   Our paper does not include experimental results as a contribution.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer:

   [N/A]

   Justification:

   Our paper does not include experimental results as a contribution.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer:

   [N/A]

   Justification:

   Our paper does not include experimental results as a contribution.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer:

   [N/A]

   Justification:

   Our paper does not include experimental results as a contribution.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer:

   [N/A]

   Justification:

   Our paper does not include experimental results as a contribution.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer:

   [Yes]

   Justification:

   Yes, the research of this paper conforms in every respect with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer:

    [N/A]

Justification:

Our work is theoretical and foundational research. We acknowledge that there are many potential societal consequences of our theoretical results, however none of which we feel must be specifically higlighted.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification:

Our paper poses no such risks.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer:

[N/A]

Justification:

Our paper does not use existing assets.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: Our paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer:

    [N/A]

    Justification: Our paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer:

    [N/A]

    Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer:

    [N/A]

    Justification:

    The core method development in our research does not involve LLMs as any important, original, or non-standard component.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.