

Breaking the $1/\epsilon$ Barrier in Quantum Zero-Sum Games: Generalizing Metric Subregularity for Spectraplexes

MOTIVATION

Min-max optimization lies at the intersection of optimization, game theory, and machine learning. In its most general form, given an objective function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$, one seeks to solve $(\chi^*, \psi^*) = \arg \min_{\chi \in X} \arg \max_{\psi \in \Psi} f(\chi, \psi)$. Since von Neumann’s 1926 Göttingen lecture on *Gesellschaftsspiele* [Von Neumann, 1928], such *min-max points* have underpinned zero-sum game theory: over simplices (Δ^n, Δ^m) they are stable outcomes in which no player profits from unilateral deviation, later formalized as *Nash equilibria*. Nearly a century later, attention has shifted to settings where strategies are quantum states and payoffs arise from measurements, a framework capturing quantum-information exchange with implications for complexity and cryptography.

Quantum Game Theory (QGT) foundations trace back to the EPR paradox and the Bell–CHSH line of work [Bell, 1964, Clauser et al., 1969], which revealed correlations inexplicable by classical local theories and led naturally to *non-local games* [Aravind, 2002, Greenberger et al., 1989, Hardy, 1993, Mermin, 1990, Reichardt et al., 2013], culminating in the landmark result $\text{MIP}^* = \text{RE}$ [Ji et al., 2021]. In parallel, quantum zero-sum games have attracted renewed interest, with applications ranging from quantum robust machine learning to quantum generative adversarial networks (QGANs).

As in the classical setting, computing Nash equilibria in *general* quantum games is PPAD-complete even to approximate [Bostanci and Watrous, 2022]. Quantum *zero-sum* games admit far more efficient algorithms: building on Jain and Watrous [2009] and the refinements presented at QTML 2023 by Vasconcelos et al. [2025], one achieves ϵ -accuracy in pseudo-linear time, $O(1/\epsilon)$ iterations under full channel feedback [Lotidis et al., 2023]. Yet the geometry of the spectraplex has led to the conjecture [Vasconcelos et al., 2025, Wei et al., 2021] that no single-gradient method can surpass this $\Omega(1/\epsilon)$ barrier—in sharp contrast to the polyhedral setting, where logarithmic dependence $O(\log(1/\epsilon))$ is attainable¹.

CONJECTURE 1 (VASCONCELOS ET AL. [2025]). *There exists a quantum two-player zero-sum game for which any single-gradient method requires $\Omega(1/\epsilon)$ iterations to compute an ϵ -Nash equilibrium.*

We revisit this conjecture across a range of established methods and precision regimes:

Can gradient-based methods achieve linear convergence,
i.e., $O(\log(1/\epsilon))$ iteration complexity, in quantum zero-sum games? (*)

QUANTUM ZERO-SUM GAME SETUP

We study a restricted but fundamental class of non-interactive refereed quantum zero-sum games: two adversarial players, Alice and Bob, independently prepare quantum states and send them to a referee, who performs a joint measurement on the received states and, based on a fixed utility function, assigns a payoff to Alice. Alice prepares n -qubit states within her spectraplex $\mathcal{A} = \{\alpha \in \mathcal{H}_+^n : \text{Tr}[\alpha] = 1\}$, while Bob prepares m -qubit states within $\mathcal{B} = \{\beta \in \mathcal{H}_+^m : \text{Tr}[\beta] = 1\}$, where $\mathcal{H}_+^d$ is the set of $2^d \times 2^d$ -dimensional, positive Hermitian matrices.

Following the setting and notation of Vasconcelos et al. [2025], which builds upon the framework of Jain and Watrous [2009], we denote the joint state space as $\Psi = (\alpha, \beta) \in \mathcal{A} \oplus \mathcal{B}$. The prepared states are jointly measured through a positive operator-valued measurement (POVM) $\{P_\omega\}_{\omega \in \Omega}$ with finite outcome set Ω and $\sum_{\omega \in \Omega} P_\omega = \mathbb{I}$. Given a joint state Ψ , the probability of outcome $\omega \in \Omega$ is $p_\omega(\Psi) = \text{Tr}[P_\omega(\alpha \otimes \beta)]$, where $\otimes$ denotes the tensor product. Each outcome $\omega \in \Omega$ carries a normalized, finite real-valued utility $\mathcal{U}(\omega) \in [-1, 1]$, and Alice’s payoff observable is $U = \sum_{\omega \in \Omega} \mathcal{U}(\omega) P_\omega$. The expected payoff $f : \mathcal{A} \oplus \mathcal{B} \rightarrow \mathbb{R}$ to Alice is $f(\Psi) = \sum_{\omega \in \Omega} \mathcal{U}(\omega) p_\omega(\Psi)$; being zero-sum, Bob’s utility, observable, and payoff are $-\mathcal{U}$, $-U$, and $-f(\Psi)$. The payoff gradients,

¹We use computer-science iteration-complexity notation: error rates $O(e^{-\rho^T})$, $O(1/T)$, and $O(1/\sqrt{T})$ correspond respectively to $O(\log(1/\epsilon))$, $O(1/\epsilon)$, and $O(1/\epsilon^2)$ iterations.

serving as first-order feedback, are $F_\alpha(\beta) = \text{Tr}_{\mathcal{B}}[U^\dagger(\mathbb{I}_{\mathcal{A}} \otimes \beta)]$ and $F_\beta(\alpha) = -\text{Tr}_{\mathcal{A}}[U^\dagger(\alpha \otimes \mathbb{I}_{\mathcal{B}})]$. This gradient viewpoint of Vasconcelos et al. [2025] lets us transfer first-order techniques from classical game theory and optimization into the quantum setting.

PRIOR WORK

Our work builds directly on Jain and Watrous [2009] and Vasconcelos et al. [2025] for non-interactive quantum refereed zero-sum games. Jain and Watrous brought the *Matrix Multiplicative Weights Update* (MMWU) algorithm to this setting, a quantum analogue of the classical matrix exponentiated gradient method of Tsuda et al. [2005], itself a special case of Mirror Descent [Beck and Teboulle, 2003, Nemirovski and Yudin, 1983]. Vasconcelos et al. [2025] subsequently adapted a hierarchy of first-order methods to quantum zero-sum games; to our knowledge, the $O(1/\epsilon)$ average-iterate guarantee of Optimistic Matrix Multiplicative Weights Update (OMMWU) remains the state of the art among matrix multiplicative-weights methods for ϵ -approximate Nash equilibria.

From an optimization perspective, solving a zero-sum game $\min_{x \in \Delta^n} \max_{y \in \Delta^m} \langle x, Uy \rangle$ is equivalent to solving a variational inequality (VI) for a monotone operator over the polyhedral constraint set $\Delta^n := \{x \in \mathbb{R}_{\geq 0}^n : \mathbf{1}^\top x = 1\}$, $\Delta^m := \{y \in \mathbb{R}_{\geq 0}^m : \mathbf{1}^\top y = 1\}$: a pair $(x^*, y^*) \in \Delta^n \times \Delta^m$ is a Nash equilibrium—no player profits from a unilateral deviation—if and only if $\langle F(z^*), z - z^* \rangle \geq 0$ for all $z \in Z$, where $F(z) := (\nabla_x \langle x, Uy \rangle, -\nabla_y \langle x, Uy \rangle) = (Uy, -U^\top x)$, $Z := \Delta^n \times \Delta^m$, and $z := (x, y)$.

In a line of work pioneered by Tseng [1995] and later by Auslender and Teboulle [2005], Nemirovski [2004], the *mirror-prox* family was shown to achieve an iteration complexity of $O(1/\epsilon)$, which Ouyang and Xu [2021] recently proved optimal in the general convex-concave setting.

Sharper rates require additional structure. In strongly convex–concave problems—or strongly monotone VIs—gradient descent–ascent converges linearly in $O(\kappa \log(1/\epsilon))$ iterations, κ the condition number. Enforcing this via quadratic (Tikhonov) regularization of order ϵ inflates κ to $O(1/\epsilon)$, giving only $O((1/\epsilon) \log(1/\epsilon))$ complexity. This led to *Nesterov’s iterative smoothing* (IterSmooth) [Nesterov, 2005], which starts from a constant, well-conditioned regularization and gradually decays it toward the original game; Gilpin et al. [2012] showed that for polyhedral sets this yields $O(\log(1/\epsilon))$ complexity.

A well-known limitation of IterSmooth is its restarting, double-loop design, often viewed as impractical in applied settings. In contrast, polyhedrality has recently emerged as a particularly powerful structural property: advances on optimistic algorithms, including projected *Optimistic Gradient Descent–Ascent* (OGDA) and *Optimistic Multiplicative Weights Update* (OMMWU) [Applegate et al., 2023, Fercoq, 2023, Wei et al., 2020], demonstrate that, for bilinear games over simplices, polyhedral constraints act as an implicit strong convexifier, yielding exponentially fast last-iterate convergence to Nash equilibria. Conversely, when the feasible region is specified by higher-order polynomial constraints, such acceleration is no longer achievable, delineating a sharp boundary on the power of first-order methods [Wei et al., 2020].

OUR CONTRIBUTIONS & METHODOLOGY

We give a positive answer to question (*). To our knowledge, this is the first work showing that natural gradient-based last-iterate dynamics achieve *linear convergence* in quantum zero-sum games, refuting the conjecture of Vasconcelos et al. [2025]. Concretely, Euclidean methods such as IterSmooth and OGDA achieve $O(\log(1/\epsilon))$ last-iterate convergence to Nash equilibria, with constants controlled by a game-dependent spectraplex error-bound modulus; for the entropic method OMMWU, our analysis gives a non-asymptotic last-iterate guarantee of $\tilde{O}(1/\epsilon)$ via an entropy-regularized QRE approximation.

Let $\mathcal{G}(\Psi)$ denote the duality gap of a joint state $\Psi = (\alpha, \beta)$ and $\mathcal{Z}^*$ the set of Nash equilibria: $\mathcal{G}(\Psi) \leq \epsilon$ is the weak, exploitability-based notion of an ϵ -Nash equilibrium, while $\text{dist}(\Psi, \mathcal{Z}^*)$ measures last-iterate proximity. To the best of our knowledge, the following result is the first error-bound characterization for quantum zero-sum games over spectraplexes, and it is our geometric ingredient behind the logarithmic high-precision rates below.

MAIN RESULT 1 (SPECTRAPLEX SP–MS AND ERROR BOUND). *Quantum zero-sum games over spectraplexes satisfy the SP–MS condition with parameter $s = 0$: there exists $C > 0$ such that, for every $\Psi \in \mathcal{Z}$,*

$$\sup_{\Psi' \in \mathcal{Z}} \langle F(\Psi), \Psi - \Psi' \rangle_F / \|\Psi - \Psi'\|_F \geq C \operatorname{dist}(\Psi, \mathcal{Z}^*).$$

Consequently, there exists a game-dependent constant $\delta(\Xi) > 0$, independent of ϵ , such that $\operatorname{dist}(\Psi, \mathcal{Z}^*) \leq \mathcal{G}(\Psi)/\delta(\Xi)$: over the spectraplex, small duality gap controls Frobenius distance to the Nash set.

The above result makes the proofs of two classical gradient methods nearly immediate when those methods are formulated in matrix form.

MAIN RESULT 2 (ITERATIVE SMOOTHING). *The matrix adaptation of Nesterov’s iterative smoothing computes an ϵ -Nash equilibrium in $O(\kappa(\Xi) \log(\|F\|_{\text{op}}/\epsilon))$ first-order iterations, where $\kappa(\Xi) = \|F\|_{\text{op}}/\delta(\Xi)$ is the condition number induced by the spectraplex error bound.*

MAIN RESULT 3 (OPTIMISTIC GRADIENT DYNAMICS). *Let L be the Frobenius-norm Lipschitz constant of the saddle-point operator and take OGDA step size $\eta \leq 1/(8L)$. Since Main Result 1 gives SP–MS with $s = 0$, the standard last-iterate theory for projected optimistic dynamics applies in Frobenius geometry, so the matrix adaptation of OGDA converges geometrically to $\mathcal{Z}^*$, computing an ϵ -Nash equilibrium in $O_d(\log(1/\epsilon))$ iterations.*


The proof replaces finite active constraints by tight spectral constraints and semidefinite normal-cone geometry, giving a spectraplex analogue of Hoffman-type error bounds: for IterSmooth it keeps each outer stage close to $\mathcal{Z}^*$ despite the inner q-smoothing routine’s apparent $1/\epsilon$ dependence, and for OGDA it supplies the SP–MS residual bound for projected optimistic dynamics. These Euclidean results strengthen rather than contradict Vasconcelos et al. [2025]: their guarantee is ergodic $O(1/T)$, ours geometric last-iterate [Lin et al., 2023, Mertikopoulos et al., 2019, Vasconcelos et al., 2025, Wei et al., 2021].

Additionally, this raises a natural question for the von Neumann entropic regularizer, which offers the most favorable dimension dependence over the spectraplex in average-iterate:

Can entropic regularization yield a fast last-iterate theorem for the quantum zero-sum case?

Our answer is regime-dependent: in the high-precision regime, even with δ -explicit exploration, one must take $\delta \asymp \epsilon$, which turns geometric convergence to the quantum QRE into a $\tilde{O}_d(1/\epsilon)$ guarantee for approximating Nash equilibria. Trying to regularize further only exposes the bias–contraction tradeoff: smaller entropic bias comes with weaker contraction. In the low-precision, dimension-sensitive regime, entropy is too weak to control nearly flat spectral directions of the spectraplex; indeed, diagonal quantum games embed the classical hard instances of Cai et al. [2025]. Thus any direct $O(\log(1/\epsilon))$ last-iterate theorem for unregularized OMMWU must pay a conditioning cost, and our lower bound shows that the resulting performance can deteriorate arbitrarily badly; a promising direction is to further understand quantum game dynamics through the geometric intricacies of the spectraplex.

In conclusion, we resolve the open conjecture by showing that Euclidean dynamics may achieve fast $O(\log(1/\epsilon))$ last-iterate convergence in quantum zero-sum games, driven by a novel spectraplex Hoffman-type error bound. Conversely, we prove that non-Euclidean entropic methods (OMMWU) are fundamentally constrained by a sharp bias-contraction trade-off and flat spectral directions. Our lower bound establishes an explicit $\Omega(d^\alpha \log(1/\epsilon))$ barrier for unregularized entropic dynamics, proving that fast convergence is dictated by the game’s spectral conditioning. Ultimately, these results map the exact geometric boundaries where optimization, spectraplex geometry, and quantum game dynamics intersect.



REFERENCES

- David Applegate, Oliver Hinder, Haihao Lu, and Miles Lubin. 2023. Faster first-order primal-dual methods for linear programming using restarts and sharpness. *Mathematical Programming* 201, 1 (2023), 133–184.
- P. K. Aravind. 2002. Bell's Theorem Without Inequalities and Only Two Distant Observers. *Journal of Genetic Counseling* 15, 4 (Aug. 2002), 397–405. <https://doi.org/10.1023/A:1021272729475>
- Alfred Auslender and Marc Teboulle. 2005. Interior projection-like methods for monotone variational inequalities. *Mathematical Programming* 104 (2005), 39–68. <https://doi.org/10.1007/s10107-004-0568-x>
- Amir Beck and Marc Teboulle. 2003. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* 31, 3 (2003), 167–175. [https://doi.org/10.1016/S0167-6377\(02\)00231-6](https://doi.org/10.1016/S0167-6377(02)00231-6)
- J. S. Bell. 1964. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika* 1 (Nov 1964), 195–200. Issue 3. <https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195>
- John Bostanci and John Watrous. 2022. Quantum game theory and the complexity of approximating quantum Nash equilibria. *Quantum* 6 (Dec. 2022), 882. <https://doi.org/10.22331/q-2022-12-22-882>
- Yang Cai, Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Weiqiang Zheng. 2025. Fast Last-Iterate Convergence of Learning in Games Requires Forgetful Algorithms. arXiv:2406.10631 [cs.GT] <https://arxiv.org/abs/2406.10631>
- John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. 1969. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.* 23 (Oct 1969), 880–884. Issue 15. <https://doi.org/10.1103/PhysRevLett.23.880>
- Olivier Fercoq. 2023. Quadratic error bound of the smoothed gap and the restarted averaged primal-dual hybrid gradient. *Open Journal of Mathematical Optimization* 4 (2023), 1–34.
- Andrew Gilpin, Javier Peña, and Tuomas Sandholm. 2012. First-order algorithm with $O(\ln(1/\epsilon))$ convergence for ϵ -equilibrium in two-person zero-sum games. *Mathematical Programming* 133, 1–2 (June 2012), 279–298. <https://doi.org/10.1007/s10107-010-0430-2>
- Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. 1989. *Going Beyond Bell's Theorem*. Springer Netherlands, 69–72. https://doi.org/10.1007/978-94-017-0849-4_10
- Lucien Hardy. 1993. Nonlocality for two particles without inequalities for almost all entangled states. *Phys. Rev. Lett.* 71 (Sep 1993), 1665–1668. Issue 11. <https://doi.org/10.1103/PhysRevLett.71.1665>
- Rahul Jain and John Watrous. 2009. Parallel Approximation of Non-interactive Zero-sum Quantum Games. In *2009 24th Annual IEEE Conference on Computational Complexity*. 243–253. <https://doi.org/10.1109/CCC.2009.26>
- Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. 2021. $MIP^* = RE$. *Communications of the ACM* 64, 11 (oct 2021), 131–138. <https://doi.org/10.1145/3485628>
- Wayne Lin, Georgios Piliouras, Ryann Sim, and Antonios Varvitsiotis. 2023. Quantum Potential Games, Replicator Dynamics, and the Separability Problem. arXiv:2302.04789 [cs.GT]
- Kyriakos Lotidis, Panayotis Mertikopoulos, Nicholas Bambos, and Jose Blanchet. 2023. Payoff-based learning with matrix multiplicative weights in quantum games. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NIPS '23)*. Article 1652, 33 pages. <https://doi.org/10.5555/3666122.3667774>
- N. David Mermin. 1990. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.* 65 (Dec 1990), 3373–3376. Issue 27. <https://doi.org/10.1103/PhysRevLett.65.3373>
- Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. 2019. Optimistic mirror descent in saddle-point problems: Going the extra-(gradient) mile. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg8jjC9KQ>
- Arkadi Nemirovski. 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 (2004), 229–251. <https://doi.org/10.1137/S1052623403425629>
- Arkadi Semen Nemirovski and David Berkovich Yudin. 1983. *Problem Complexity and Method Efficiency in Optimization*. Wiley, New York, NY.
- Yu. Nesterov. 2005. Smooth minimization of non-smooth functions. *Mathematical Programming* 103, 1 (May 2005), 127–152. <https://doi.org/10.1007/s10107-004-0552-5>
- Yuyuan Ouyang and Yangyang Xu. 2021. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming* 185 (2021), 1–35. <https://doi.org/10.1007/s10107-019-01420-0>

- Ben W. Reichardt, Falk Unger, and Umesh Vazirani. 2013. Classical command of quantum systems. *Nature* 496, 7446 (April 2013), 456–460. <https://doi.org/10.1038/nature12035>
- Paul Tseng. 1995. On linear convergence of iterative methods for the variational inequality problem. *J. Comput. Appl. Math.* 60, 1-2 (1995), 237–252.
- Koji Tsuda, Gunnar Rätsch, and Manfred K. Warmuth. 2005. Matrix Exponentiated Gradient Updates for On-line Learning and Bregman Projection. *Journal of Machine Learning Research* 6 (Dec. 2005), 995–1018. <https://doi.org/10.5555/1046920.1088706>
- Francisca Vasconcelos, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Panayotis Mertikopoulos, Georgios Piliouras, and Michael I. Jordan. 2025. A Quadratic Speedup in Finding Nash Equilibria of Quantum Zero-Sum Games. *Quantum* 9 (May 2025), 1737. <https://doi.org/10.22331/q-2025-05-06-1737>
- John Von Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100, 1 (Dec. 1928), 295–320. <https://doi.org/10.1007/BF01448847>
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. 2020. Linear last-iterate convergence in constrained saddle-point optimization. *arXiv preprint arXiv:2006.09517* (2020).
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. 2021. Linear Last-iterate Convergence in Constrained Saddle-point Optimization. In *International Conference on Learning Representations*. https://openreview.net/forum?id=dx11_7vm5_r