

A Quadratic Speedup in Finding Nash Equilibria of Quantum Zero-Sum Games

Francisca Vasconcelos¹, Emmanouil-Vasileios Vlatakis-Gkaragkounis¹,
Panayotis Mertikopoulos², Georgios Piliouras³, and Michael I. Jordan¹

¹UC Berkeley

²Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France

³Singapore University of Technology and Design

Recent developments in domains such as non-local games, quantum interactive proofs, and quantum generative adversarial networks have renewed interest in quantum game theory and, specifically, quantum zero-sum games. Central to classical game theory is the efficient algorithmic computation of Nash equilibria, which represent optimal strategies for both players. In 2008, Jain and Watrous proposed the first classical algorithm for computing equilibria in quantum zero-sum games using the Matrix Multiplicative Weight Updates (MMWU) method to achieve a convergence rate of $\mathcal{O}(d/\epsilon^2)$ iterations to ϵ -Nash equilibria in the 4^d -dimensional spectraplex. In this work, we propose a hierarchy of quantum optimization algorithms that generalize MMWU via an extra-gradient mechanism. Notably, within this proposed hierarchy, we introduce the Optimistic Matrix Multiplicative Weights Update (OMMWU) algorithm and establish its average-iterate convergence complexity as $\mathcal{O}(d/\epsilon)$ iterations to ϵ -Nash equilibria. This quadratic speed-up relative to Jain and Watrous' original algorithm sets a new benchmark for computing ϵ -Nash equilibria in quantum zero-sum games.

Francisca Vasconcelos: francisca@berkeley.edu

Emmanouil-Vasileios Vlatakis-Gkaragkounis: emvlatakis@berkeley.edu

Contents

1	Introduction	3
1.1	Motivation	3
1.2	Prior Work	4
1.3	Our Contributions and Methodology	5
1.4	Paper Organization	7
2	Preliminaries	10
2.1	Quantum Zero-Sum Games	10
2.2	Quantum Nash Equilibria	11
2.3	Convexity, Smoothness, & Duality	14
2.4	Bregman Divergence & Mirror Map	16
2.5	Mirror & Proximal Steps	17
2.5.1	Mirror Step	18
2.5.2	Proximal Step	19
2.5.3	Regularizers	21
3	Algorithms for Quantum Zero-Sum Games	22
3.1	Matrix Dual Averaging (MDA) Method	22
3.2	Matrix Mirror Prox (MMP)	24
3.3	Optimistic Matrix Mirror-Prox (OMMP) Methods	25
4	Convergence Analysis	26
4.1	Monotonicity of \mathcal{F}	27
4.2	Error Decomposition	27
4.3	Optimality Analysis of (A_t)	27
4.4	Optimality Analysis of (C_t)	28
4.5	An Upper Bound for (B_t)	29
4.6	Average-Iterate Error Upper Bound	30
4.7	Regularizer-Induced Dimensionality Dependence	33
5	Conclusions	34
6	Acknowledgements	35
A	Complex Matrix Differentiation and Gradients	41
B	Mirror & Proximal Maps	42
C	Properties of Finite-Valued Zero-Sum Quantum Games	47
C.1	Properties of Spectraplexes	47
C.2	Properties of \mathcal{F}	48
D	Convergence Analysis Proofs	53

1 Introduction

1.1 Motivation

Ever since von Neumann’s groundbreaking work in the 1920’s [1], game theory has become a foundational pillar of modern mathematics, economics, and computer science. The central solution concept in non-cooperative game theory is that of a *Nash equilibrium*, i.e., a strategy of joint play in which no single player can benefit from a unilateral deviation [2]. Despite its mathematical significance, the computational perspective on general Nash equilibria is murkier, since even approximating such an equilibrium is a PPAD-complete problem for general classical games. On the other hand, Nash equilibria in two-player zero-sum games (i.e., where one player’s gain is the other’s loss) can be computed in polynomial time, and they have found practical applications in a wide variety of domains, from resource allocation problems [3], to political strategy [4], and the training of machine learning models [5, 6, 7].

When the players’ actions are the states of a quantum system, classical game theory no longer applies, and one must consider *quantum games*. These games constitute a natural framework for the study of the exchange and processing of quantum information, resulting in various useful applications to theoretical, computational, and cryptographic settings. The original interest in quantum games dates back to seminal work by Bell [8] and CSHS [9], explaining the Einstein-Podolsky-Rosen (EPR) Paradox. Since then, substantial work on quantum “non-local games” [10, 11, 12, 13, 14] has led to novel methods for testing whether two or more spatially separated, non-communicating quantum devices can generate correlations not reproducible by any pair of classical devices—resulting in entanglement-based tests of quantum advantage. Much of this work culminated in establishing $\text{MIP}^* = \text{RE}$ [15], which, among many notable results, demonstrated that there exists an efficient reduction from the Halting Problem to deciding whether a two-player non-local game has entangled value at most 1 or $\frac{1}{2}$. Furthermore, it was recently proven that non-local games can be compiled into single-prover interactive games [16]. In general, quantum interactive proofs with competing provers can be modeled as competitive quantum refereed games [17, 18, 19] and multi-prover quantum interactive proofs [20, 21, 22, 23] can be modeled as cooperative quantum games. Finally, there has been substantial work on quantum coin-flipping [24, 25, 26, 27], a game model where two players directly exchange quantum information, studying how two parties with competing interests can carry out a fair coin flip across a quantum communication channel.

In this work, we consider the specific class of quantum *zero-sum* games, in which two players are in direct competition and do not share entanglement. One of the earliest works in quantum game theory [28] studied a “matching pennies”-type zero-sum game to prove that quantum strategies are at least as good as and, in some cases, can outperform classical strategies. Further work on quantum zero-sum games led to a proof that the quantum complexity class $\text{QRG}(1)$, of problems having one-turned quantum refereed games, is contained in PSPACE [29]. Finally, in the context of quantum machine learning, there has been substantial recent interest in quantum generative adversarial networks (QGANs) [30, 31], for which the training of the competing generator and discriminator networks can be modeled as a quantum zero-sum game.

As in the classical setting, finding Nash equilibria of *general* quantum games is computationally prohibitive—PPAD-complete to be exact [32]. However, the situation is more favorable for quantum *zero-sum* games. As demonstrated by Jain and Watrous [29], Nash equilibria of d -qubit two-player quantum zero-sum games can be calculated to ϵ -accuracy in polynomial time and linear dependence on the number of qubits (logarithmic in the

spectraplex dimensionality), i.e., $\mathcal{O}(d/\epsilon^2)$, even with payoff-based information on the players’ side [33]. In view of this, equilibrium strategies in quantum zero-sum games can in principle be computed algorithmically, thus providing the required guarantees of implementability for proof schemes or cryptographic ciphers that rely on the computation of quantum Nash equilibria. Our goal is to pursue this objective further, bringing tools from online learning and computational learning theory to improve the design of algorithms for the computation of approximate Nash equilibria in quantum zero-sum games.

1.2 Prior Work

Our work builds upon the work of Jain and Watrous [29] for non-interactive quantum zero-sum games. To provide some context, [29] leveraged feedback in the form of quantum channels (governed by superoperators) to introduce the Matrix Multiplicative Weights Update (MMWU) algorithm, a two-player variant of the matrix exponentiated gradient algorithm of Tsuda et al. [34], which is itself a special case of Mirror Descent [35, 36]. Beyond the setting of quantum zero-sum games, variants of the MMWU algorithm have also been used in many important applications such as proving QIP=PSPACE [37], solving SDPs [38], finding balanced separators [39], enhancing spectral sparsification [40], covariance optimization tasks [41, 42, 43], and matrix learning [44]. For the quantum zero-sum games of interest in this work, despite recent work on no-regret learning dynamics for zero-sum and general quantum games [33, 45], to the best of our knowledge, the $\mathcal{O}(d/\epsilon^2)$ convergence rate of the MMWU algorithm is the state of the art for finding ϵ -approximate Nash equilibria. In this work, we ask:

Can we improve on the performance of MMWU in quantum zero-sum games?

More precisely, our paper’s goal and main contribution is to provide an accelerated version of the MMWU algorithm which achieves ϵ -accuracy in linear, $\mathcal{O}(d/\epsilon)$ time, instead of $\mathcal{O}(d/\epsilon^2)$.

To achieve this quadratic speedup, our point of departure is the literature on classical *finite* games. The classical version of the MMWU algorithm—known, among other names, as the Multiplicative Weights Update (MWU) method [46]—similarly achieves an $\mathcal{O}_d(1/\epsilon^2)$ iteration complexity.¹ However, Nemirovski [47] and Auslender & Teboulle [48] showed that, by a suitable modification of the MWU method—now commonly known as the *mirror-prox* family of algorithms—it is possible to accelerate this rate to $\mathcal{O}_d(1/\epsilon)$, a rate which was recently shown by Ouyang and Xu [49] to be order-optimal in the setting of min-max convex-concave problems.

The main algorithmic insight of the mirror-prox template is the combination of an extra-gradient step in the spirit of Korpelevich [50] with an iterate-averaging mechanism à la Polyak & Juditsky [51]. Coupled with the smoothness of the underlying objective, this combination yields a $\mathcal{O}(1/\epsilon)$ convergence rate, albeit with a possibly suboptimal dependence on the dimension of the problem. In the case of the simplex—which is the state space of classical, finite games—the dependence on the dimension can be improved dramatically by instantiating the mirror-prox with an entropic regularizer, which essentially boils down to a variant of MWU “with advice” [47]. The only notable drawback of the resulting algorithm is that it requires not one, but two oracle queries per iteration, which doubles the computation cost, and requires additional coordination / communication from the players’ side. Building on an original idea by Popov [52], this “cost doubling” issue

¹Hereon, we will use the $\mathcal{O}_d(\cdot)$ notation to hide the dimensionality dependence of the iteration complexity. There will be explicit discussion of dimensionality dependence in the quantum setting in Section 4.7.

was mitigated by the so-called *optimistic mirror descent* (OMD) proposal of Rakhlin and Sridharan [53], which optimistically reused past gradient information as a surrogate for the extra-gradient step of the mirror-prox algorithm. In this way, when applied to mixed extensions of classical, finite games, OMD ultimately achieves an order-optimal convergence rate in terms of both ϵ and d , all the while using a single oracle call per iteration. We summarize this hierarchy of optimization methods and the relevant trade-offs in Figure 1.

1.3 Our Contributions and Methodology

Our work focuses on developing an accelerated algorithm with similar, linear-time performance guarantees in the *quantum* setting. The main contributions of this work are thus twofold: *a*) we propose the *Optimistic Matrix Multiplicative Weights Update (OMMWU) algorithm* (Algorithm 1) for finding ϵ -Nash equilibria in quantum zero-sum games; and *b*) we show that OMMWU achieves the following iteration complexity guarantees.

Corollary 1 (OMMWU Iteration Complexity) *In a 4^d -dimensional spectraplex, OMMWU computes average-iterate ϵ -Nash equilibria in $\mathcal{O}(d/\epsilon)$ iterations.*

In order to achieve this result, we forge a link between quantum zero-sum games and established works in classical game theory and optimization. In addition to OMMWU—which surpasses the $\mathcal{O}(d/\epsilon^2)$ rate of MMWU with a single gradient call per iteration—we also propose a hierarchical family of algorithms and techniques that can be applied to various semidefinite programming problems (quantum or otherwise) and thus may be of independent interest.

Methodology. In terms of methodology, we move away from the *channel-based* view of quantum zero-sum games in [29] and instead adopt a *gradient-based* perspective, similar to prior work on learning in general quantum games [54] (see also [45] connecting the same online learning, regret minimization perspective to new classes of quantum correlated equilibria). In this framework, Alice and Bob’s feedback is characterized not by traditional quantum superoperators, but rather by mathematical gradient operators, which comprise the game’s gradient feedback operator \mathcal{F} . We prove several important properties of \mathcal{F} , namely that Lipschitz continuity and monotonicity. In doing so, we conclude that solving a quantum zero-sum game to a desired accuracy essentially boils down to solving a smooth, semidefinite convex-concave problem with first-order oracle information.

Our approach is based on endowing the problem’s feasible region—a product of spectraplexes—with an OMD template in the spirit of Rakhlin and Sridharan [53]. A difficulty that arises when trying to combine these elements is that the product structure of standard exponential weight algorithms does not carry over automatically to the non-commutative matrix variables that arise in the quantum setting. The key observation that allows us to retain the speed-up of a mirror prox method and the state-of-art dimensionality dependence is that the von Neumann entropy (which is the underlying regularizer of the MMWU algorithm) can be encoded in a matrix exponentiation step, as per the original MMWU algorithm (see also [55] for generalizations of this connection for optimization over symmetric cones). This allows us to implement an optimistic update structure in the problem’s dual space, where the primary arithmetic operation is ordinary addition—which is an abelian operation on the space of Hermitian matrices. Thus, drawing inspiration from a classical hierarchy of optimization methods, we propose a *design of quantum zero-sum game algorithms* (depicted in Figure 2) achieving an optimized convergence rate with a single gradient call per iteration.

Algorithm 1 Optimistic Matrix Multiplicative Weights Update (OMMWU)

Accuracy Parameter: ϵ

Regularization Function: $h : \{\mathcal{A}, \mathcal{B}\} \rightarrow \mathbb{R}$

Strong Convexity Parameter (of h): μ_h

Diameter (of \mathcal{BD}_h): $\mathcal{D}_h = \sup_{X, Y \in \mathcal{C}} \mathcal{BD}_h(X \| Y)$

Lipschitz Parameter (of \mathcal{F}): $\gamma_{\mathcal{F}}$

Proximal Map: $\text{Prox}\Pi_{\mathcal{A}, \mathcal{B}}^{h, \eta}(X, Y) := \arg \min_{C \in \mathcal{C}} \{\langle Y, C - X \rangle - \frac{1}{\eta} \mathcal{BD}_h(C \| X)\}$.

$\eta \leftarrow \mu_h / (2\gamma_{\mathcal{F}})$

▷ Step Size

$N \leftarrow \lceil \mathcal{D}_h / (\eta \cdot \epsilon) \rceil$

▷ Number of Rounds

$(\alpha_0, \beta_0) \leftarrow \left(\frac{1}{2^n} \mathbb{1}_{\mathcal{A}}, \frac{1}{2^m} \mathbb{1}_{\mathcal{B}} \right)$

▷ State Initialization

$(\hat{\alpha}_0, \hat{\beta}_0) \leftarrow (\alpha_0, \beta_0)$

▷ Intermediate State Initialization

for $t \in [1, N - 1]$ **do**

$\alpha_{t+1} \leftarrow \Lambda(\log \hat{\alpha}_t + \eta \mathcal{F}_{\alpha}(\beta_t))$

▷ State Updates

$\beta_{t+1} \leftarrow \Lambda(\log \hat{\beta}_t + \eta \mathcal{F}_{\beta}(\alpha_t))$

$\hat{\alpha}_{t+1} \leftarrow \Lambda(\log \hat{\alpha}_t + \eta \mathcal{F}_{\alpha}(\beta_{t+1}))$

▷ Momentum Updates

$\hat{\beta}_{t+1} \leftarrow \Lambda(\log \hat{\beta}_t + \eta \mathcal{F}_{\beta}(\alpha_{t+1}))$

end for

return $\left(\bar{\alpha} = \frac{1}{N} \sum_{t=0}^{N-1} \alpha_t, \bar{\beta} = \frac{1}{N} \sum_{t=0}^{N-1} \beta_t \right)$

Within this design, we demonstrate that the MMWU algorithm of [29] instantiates (via a von Neumann entropy regularizer) the *Matrix Dual Averaging (MDA)* method (Algorithm 4). Following a classical proof technique in the spirit of Ene and Nguyen [56], we also prove that another of the proposed algorithms in our hierarchy—the *Optimistic Matrix Mirror Prox (OMMP)* method (Algorithm 5)—obtains a quadratic speedup relative to the convergence rate of MDA, while still requiring only one gradient call per iteration:

Theorem 1 (Main Result) *The OMMP method computes ϵ -Nash equilibria in $\mathcal{O}_d(1/\epsilon)$ steps.*

Thus, OMMP establishes a new benchmark for computing ϵ -Nash finite-valued quantum zero-sum games. It is by instantiating OMMP with a von Neumann entropy regularizer, that the OMMWU algorithm and $\mathcal{O}(d/\epsilon)$ iteration complexity is achieved.

Proof Insights. The core idea of an *optimistic* algorithm is to leverage the gradient from the previous iteration as a forecast for the subsequent iteration’s gradient. This is predicated on the presumption that if the the game’s adversary employs a consistent (or stabilized) strategy, then the discrepancy between the gradients of successive iterations will be negligible. Therefore, the formulation of our algorithm is crafted by devising a parallel to such methods, circumventing the limitations imposed by the non-commutativity of the matrix operator, which is a common hurdle in quantum problems. Upon establishing the algorithm in this manner, we demonstrate that the (quantum) superoperators, corresponding to the method’s gradients, exhibit two significant characteristics: *monotonicity* and *Lipschitz continuity/smoothness*.

Monotonicity is a valuable property that allows us to: (i) correlate any Nash equilibrium with a more manageable variational inequality problem, and (ii), akin to convexity in single-agent optimization, it guarantees a consistent progression towards the equilibrium.

Lipschitz continuity, in contrast, is the essential quality that accelerates the convergence rate from $O(1/\epsilon^2)$ to $O(1/\epsilon)$. To comprehend the significance of this attribute, let’s view the situation through the lens of the minimization participant—named Bob, with his maximization counterpart being Alice. Bob finds himself involved in an online convex minimization game. In the absence of any predictive information about the forthcoming functions, which are influenced by Alice’s choices, folklore results suggests that the optimal rate is $\Theta(1/\sqrt{T})$, or equivalently $\Theta(1/\epsilon^2)$ iterations. However, should Alice’s strategy remain relatively static, Bob can anticipate that the upcoming minimization challenges will be similar, affording him a critical advantage for achieving a rate of $O(1/T)$, or $O(1/\epsilon)$.

This insight reveals that the algorithm’s performance is deeply tied to the Lipschitz parameter of the gradient operator. Here, the concept of mirror proxies plays a vital role in the proof. Utilizing mirror steps, distinct regularizers enable either more assertive or more refined movements within the strategy space of the spectraplex for both parties. Contrary to the exponential dependency implied by a Frobenius norm, we show that the Von Neumann entropy regularizer maintains a dimension-independent Lipschitz constant under a norm that is congruent with the specific needs of our problem.

Future Work. In addition to our novel methodology, we conclude with the following conjecture that aims to tighten our method’s performance guarantees and is likely of interest to both the quantum and optimization communities:

Conjecture 1 *The rates of [Theorem 1](#) are tight: specifically, there exists a quantum two-player, zero-sum game for which OMMP methods require $\Omega_d(1/\epsilon)$ iterations to compute ϵ -Nash equilibria.*

1.4 Paper Organization

The paper is divided as follows. [Section 2](#) reviews preliminaries, such as the quantum zero-sum games setup and mathematical concepts necessary for our proofs. [Section 3](#) presents the series of quantum optimization methods for computing approximate Nash equilibria of quantum zero-sum games, as depicted in [Figure 2](#). It begins with the MMWU method of [29] and shows how the algorithm is an instantiation of the more general Matrix Dual Averaging (MDA) method. It then introduces the Matrix Mirror Prox (MMP) method, which achieves the optimal rate, but requires two gradients per iteration. The section concludes with our proposed OMMP algorithm, which requires only a single gradient call per iteration. Finally, [Section 4](#) elaborates the convergence analysis of OMMP algorithm, proving that it achieves an $\mathcal{O}_d(1/\epsilon)$ rate. The paper concludes with a discussion of different implementations of the OMMP method, via different regularizers. We show that a von Neumann entropy regularizer, instantiating the OMMWU algorithm, achieves a better dimensionality dependence than the Frobenius regularizer. Thus, the OMMWU algorithm achieves the desired $\mathcal{O}(d/\epsilon)$ and a quadratic speedup relative to MMWU. For concision, we defer some details and proofs to the Appendix.

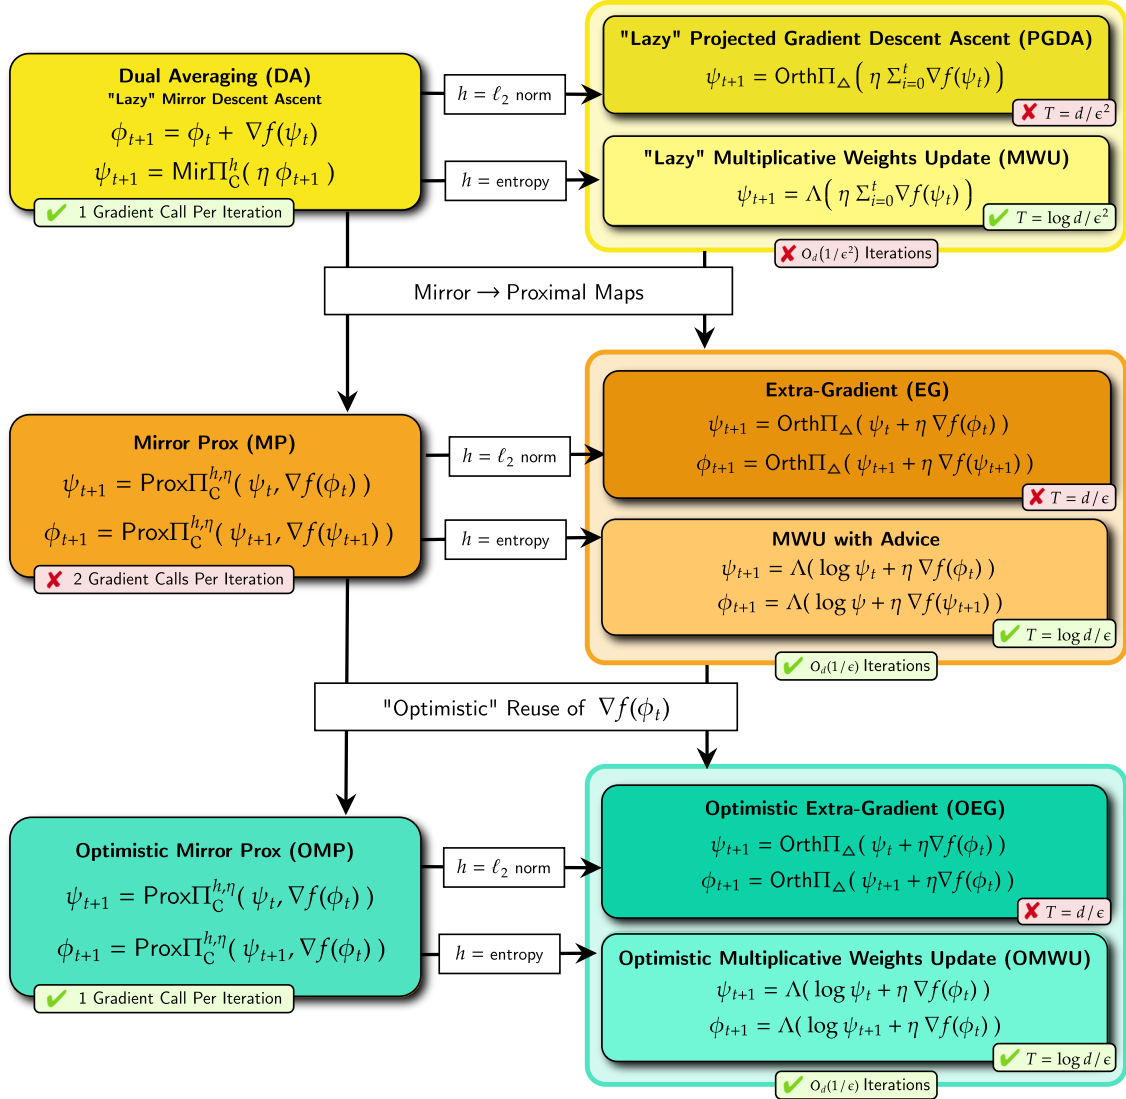


Figure 1: **Design of Classical Zero-Sum Game Algorithms.** This diagram provides the update rules for and relationships between different learning algorithms for classical zero-sum games. The left-hand side presents the Bregman generalized methods, parameterized by distance-generating/regularization function h . The right-hand side presents instantiations based on the ℓ_2 norm and entropy function, inducing an orthogonal projection $\text{Orth}\Pi$ and logit map Λ , respectively. Note that moving from $h = \ell_2$ to $h = \text{entropy}$ results in a logarithmic improvement in the total number of rounds T , by reducing the dependence on the simplex dimension d . Furthermore, moving from the mirror map $\text{Mir}\Pi$ of the Dual Averaging method (a "lazy" variant of the classic Mirror Descent Ascent algorithm) to the proximal maps $\text{Prox}\Pi$ of the Mirror Prox method results in a quadratic improvement in convergence, achieving the desired $O(1/\epsilon)$ rate. Furthermore, by reusing the "past gradient," the Single-Call Mirror Prox method retains the rate of Mirror Prox, but reduces the total number of gradient calls per iteration from two to one.

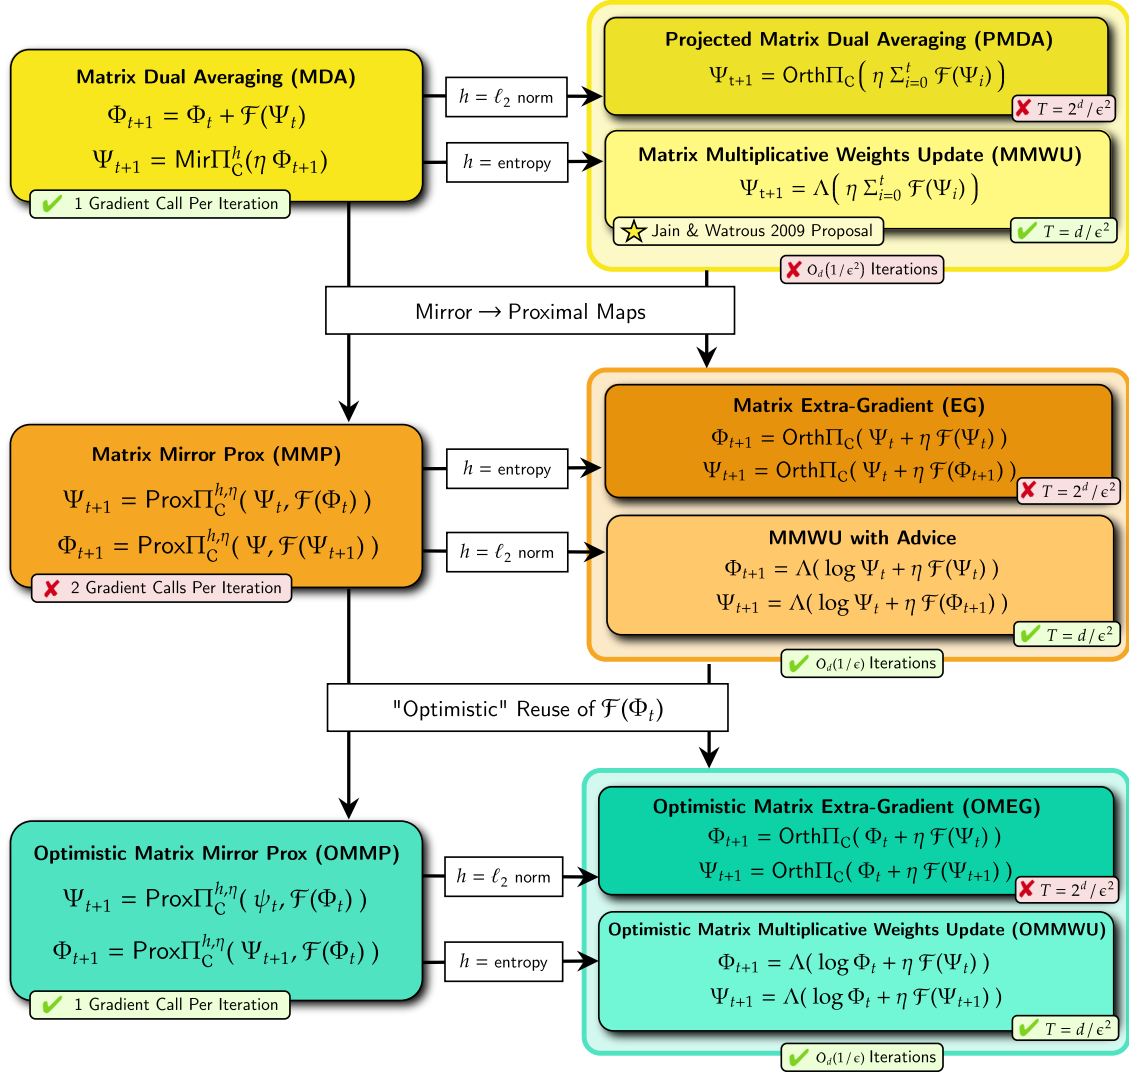


Figure 2: **Design of Quantum Zero-Sum Game Algorithms.** This diagram provides the update rules for and relationships between the learning algorithms for quantum zero-sum games as proposed in this work. The left-hand side presents the Bregman generalized methods, parameterized by distance-generating/regularization function h . The right-hand side presents instantiations based on the von Neumann entropy function and Frobenius (ℓ_2) norm, inducing an orthogonal projection $\text{Orth}\Pi$ and logit map Λ , respectively. Note that moving from ℓ_2 to an entropy function as the regularizer results in a logarithmic improvement in the total number of rounds T , by reducing the dependence on the spectraplex dimension $D = 4^d$. Furthermore, moving from the mirror map $\text{Mir}\Pi$ of the Matrix Dual Averaging method (Jain and Watrous' MMWU proposal [29]) to the proximal maps $\text{Prox}\Pi$ of the Matrix Mirror Prox method results in a quadratic improvement in convergence, achieving the desired $O(1/\epsilon)$ rate. Furthermore, by reusing the "past gradient," the Single-Call Matrix Mirror Prox method retains the rate of Matrix Mirror Prox, but reduces the total number of gradient calls per iteration from two to one.

2 Preliminaries

2.1 Quantum Zero-Sum Games

In this work, we consider the restricted class of non-cooperative zero-sum quantum games. Within this setting, two competing players, Alice and Bob, each transmit a mixed quantum state to a referee (henceforth referred to as Roger), who performs a joint measurement on both states. Roger also possesses a utility function that attributes payoffs to the players based on the measurement outcomes. Since this is a zero-sum game, Bob’s payoff will be the negative value of Alice’s—e.g., if Alice wins \$5, Bob will lose \$5.

Concretely, Alice and Bob, play the game by independently preparing multi-qubit quantum states α and β , respectively, to be sent to Roger. It is important to note that these are *mixed* states, corresponding to “mixed strategies,” and that Alice and Bob do not share entanglement. Thus, Alice and Bob’s set of possible moves is constrained to the set of density matrices lying in their respective *spectraplexes*. We will assume Alice prepares n -qubit states in her spectraplex,

$$\mathcal{A} = \{\alpha \in \mathcal{H}_+^n : \text{Tr}(\alpha) = 1\}, \quad (1)$$

while Bob prepares m -qubit states in his spectraplex,

$$\mathcal{B} = \{\beta \in \mathcal{H}_+^m : \text{Tr}(\beta) = 1\}, \quad (2)$$

where \mathcal{H}_+^d denotes the set of $2^d \times 2^d$ -dimensional, positive Hermitian matrices. In [Appendix C.1](#), we show that all spectraplexes are compact and convex sets. Denote the joint set of their states as the direct sum

$$\Psi = (\alpha, \beta) \in \mathcal{C} = \mathcal{A} \oplus \mathcal{B}, \quad (3)$$

with tensor product denoted by $\Psi^\otimes = \alpha \otimes \beta$. As a general notational note, we will use lower-case variables (e.g. α or β) to refer to states in the individual spectraplexes of either Alice (\mathcal{A}) or Bob (\mathcal{B}) and upper-case variables (e.g. Ψ or Φ) to denote joint states, which lie in the joint space (\mathcal{C}).

Alice and Bob send their states to Roger, who performs a joint measurement. This joint measurement has a finite number of possible measurement outcomes, Ω , defining an $(n + m)$ -qubit *positive operator-valued measurement (POVM)*,

$$\{P_\omega\}_{\omega \in \Omega}, \text{ where } \sum_{\omega \in \Omega} P_\omega = \mathbb{1}. \quad (4)$$

Thus, Roger’s probability of observing outcome $\omega \in \Omega$ is

$$p_\omega(\Psi) = \text{Tr}\left(P_\omega^\dagger \Psi^\otimes\right) \quad (5)$$

Dependent on the measurement outcome, Roger gives Alice and Bob a reward/payoff. Let

$$\mathcal{U} : \Omega \rightarrow \mathbb{R} \quad (6)$$

denote Alice’s *utility* function, which we will assume to be finite-valued and have image $[-1, 1]$ under normalization. We define Alice’s *payoff observable* as

$$U = \sum_{\omega \in \Omega} \mathcal{U}(\omega) P_\omega, \quad (7)$$

such that Alice's *expected payoff* is

$$u_{\text{Alice}} := u(\Psi_t) = \text{Tr}(U^\dagger \Psi^\otimes) = \sum_{\omega \in \Omega} \mathcal{U}(\omega) \text{Tr}(P_\omega^\dagger \Psi^\otimes) = \sum_{\omega \in \Omega} \mathcal{U}(\omega) p_\omega(\Psi). \quad (8)$$

Since the game is zero-sum, this implies that Bob's utility is $-\mathcal{U}$, with payoff observable $-U$ and expected payoff $u_{\text{Bob}} := -u(\Psi)$. In playing the game, Alice and Bob aim to maximize their expected payoff, written as $\max_\alpha u(\alpha, \beta)$ and $\max_\beta u(\alpha, \beta)$, respectively. Since the game is zero-sum, in maximizing his expected payoff, Bob can be seen as trying to minimize Alice's payoff: $\max_\beta -u(\alpha, \beta) = \min_\beta u(\alpha, \beta)$.

Finally, since $u(\alpha, \beta)$ is a bilinear function and the spectraplexes are compact and convex sets, von Neumann's Min-Max Theorem [1] states

$$\max_\alpha \min_\beta u(\alpha, \beta) = \min_\beta \max_\alpha u(\alpha, \beta). \quad (9)$$

2.2 Quantum Nash Equilibria

A *Nash equilibrium* of the quantum game is a pair of states (α^*, β^*) , such that each player has no incentive to change to a different state unilaterally:

$$u(\alpha^*, \beta^*) \geq u(\alpha, \beta^*), \quad \forall \alpha \in \mathcal{A} \quad (10)$$

$$u(\alpha^*, \beta^*) \leq u(\alpha^*, \beta), \quad \forall \beta \in \mathcal{B}. \quad (11)$$

Since \mathcal{C} is convex and u is linear in α and β , the existence of Nash equilibria follows from Debreu's equilibrium existence theorem [57].

In playing the game, Alice and Bob receive feedback from Roger. In the original work on quantum zero-sum games, Jain and Watrous [29] define Alice's feedback as a quantum channel, given by the superoperator $\Xi : \mathcal{B} \rightarrow \mathcal{A}$, applied solely to Bob's state. The superoperator is parameterized by the payoff observable as

$$\Xi(\beta) = \text{Tr}_{\mathcal{B}}[U(\mathbb{1}_{\mathcal{A}} \otimes \beta^\top)], \quad (12)$$

such that the expected utility can be expressed as

$$u(\alpha, \beta) = \text{Tr}[U(\alpha \otimes \beta)] = \text{Tr}[\alpha \text{Tr}_{\mathcal{B}}[U(\mathbb{1}_{\mathcal{A}} \otimes \beta)]] = \text{Tr}[\alpha \Xi(\beta^\top)]. \quad (13)$$

Similarly, Bob's feedback is defined by the quantum channel, corresponding to adjoint superoperator $\Xi^* : \mathcal{A} \rightarrow \mathcal{B}$, applied solely to Alice's state. This adjoint superoperator is defined as

$$\Xi^*(\alpha) = \text{Tr}_{\mathcal{A}}[U(\alpha^\top \otimes \mathbb{1}_{\mathcal{B}})]. \quad (14)$$

and is uniquely determined by the condition

$$u(\alpha, \beta) = \text{Tr}[\alpha \Xi(\beta^\top)] = \text{Tr}[\Xi^*(\alpha^\top) \beta]. \quad (15)$$

In this work, we will move away from thinking about Roger's feedback in terms of quantum channels towards a gradient-based view, as leveraged in recent work on learning general quantum games [54]. It is mathematically equivalent to think of the channel-based feedback in terms of gradients of the expected utility with respect to Alice and Bob's states. For a review of complex differentiation and gradients, refer to [Appendix A](#). Namely, we define Alice's feedback, or *payoff gradient*, as

$$\mathcal{F}_\alpha(\beta) = \nabla_{\alpha^\top} u(\alpha, \beta) = \nabla_{\alpha^\top} \text{Tr}[U^\dagger(\alpha \otimes \beta)] = \text{Tr}_{\mathcal{B}}[U^\dagger(\mathbb{1}_n \otimes \beta)], \quad (16)$$

which is equivalent to $\Xi(\beta^\top)$ of Equation (12). Similarly Bob's feedback is defined as the payoff gradient

$$\mathcal{F}_\beta(\alpha) = -\nabla_{\beta^\top} u(\alpha, \beta) = -\nabla_{\beta^\top} \text{Tr} \left[U^\dagger(\alpha \otimes \beta) \right] = -\text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha \otimes \mathbb{1}_m) \right], \quad (17)$$

which is equivalent to $-\Xi^*(\alpha^\top)$ of Equation (14). Similarly to Equation (15), we have that

$$u(\alpha, \beta) = \text{Tr} [\alpha \mathcal{F}_\alpha(\beta)] = -\text{Tr} [\beta \mathcal{F}_\beta(\alpha)]. \quad (18)$$

Under the gradient-based view, $\text{Tr}[\alpha \mathcal{F}_\alpha(\beta)] = \text{Tr}[\alpha \nabla_{\alpha^\top} u(\alpha, \beta)]$ can be interpreted as the directional derivative of u , as a function of Alice's state, in the direction of Alice's state. Similarly, $-\text{Tr}[\mathcal{F}_\beta(\alpha) \beta] = \text{Tr}[\nabla_{\beta^\top} u(\alpha, \beta) \beta]$ can be interpreted as the directional derivative of u , as a function of Bob's state, in the direction of Bob's state. As will be demonstrated throughout the work, moving away from a channel-based to this gradient-based view of the quantum zero-sum game feedback proves invaluable for mapping gradient-based insights from the classical games and optimization literature directly to the quantum setting.

Henceforth, we will denote the joint state as $\Psi = (\alpha, \beta)$, the joint feedback of both players, as follows:²

$$\mathcal{F}(\Psi) = (\mathcal{F}_\alpha(\beta), \mathcal{F}_\beta(\alpha)), \quad (19)$$

and the game's Nash equilibrium as $\Psi^* = (\alpha^*, \beta^*)$. Thus, as depicted in Figure 3, in each round t of the game, after playing the joint state Ψ_t , Alice will receive the individual payoff gradient $\mathcal{F}_{\alpha_t}(\beta_t)$ and Bob will receive $\mathcal{F}_{\beta_t}(\alpha_t)$, comprising joint feedback $\mathcal{F}(\Psi_t)$.




Figure 3: An illustration of a single round of the quantum zero-sum game (gradient-based view).

By using standard arguments [58], it can be shown that the solutions of the variational inequality

$$\text{Tr}[(\Psi - \Psi^*) \mathcal{F}(\Psi^*)] \leq 0, \quad \forall \Psi \in \mathcal{C}, \quad (20)$$

are the game's Nash equilibria, Ψ^* .³ Specifically, solutions satisfying Equation (20) are known as *strong solutions* to the variational inequality. Alternatively, *weak solutions* to

²In a slight abuse of notation, let $\mathcal{F}(\Psi) = \mathcal{F}(\alpha, \beta)$.

³Note that in the direct sum notation, for matrices A_1, A_2, B_1 , and B_2 , where $\dim(A_1) = \dim(A_2) = n \times n$ and $\dim(B_1) = \dim(B_2) = m \times m$, we have that $\text{Tr}[(A_1, B_1)(A_2, B_2)] = \text{Tr}[A_1 A_2] + \text{Tr}[B_1 B_2]$.




Figure 4: Intuitive illustrations of the strong and weak solutions Ψ^* of the variational inequalities given in Equation (20) and Equation (21), respectively. For strong solutions, the inner product between $\Psi - \Psi^*$ and $\mathcal{F}(\Psi^*)$ must be non-negative for all Ψ in the joint spectraplex \mathcal{C} . Meanwhile, for weak solutions, the inner product between $\Psi^* - \Psi$ and $\mathcal{F}(\Psi)$ must be non-positive for all $\Psi \in \mathcal{C}$.

the variational inequality are $\Psi^* \in \mathcal{C}$ satisfying

$$\text{Tr}[(\Psi^* - \Psi) \mathcal{F}(\Psi)] \geq 0, \quad \forall \Psi \in \mathcal{C}. \quad (21)$$

We offer intuitive visualizations of both the weak and strong solutions of variational inequalities as vector inner products in Figure 4.

As we will discuss in Section 4 and prove in Appendix C.2, \mathcal{F} is both monotone (Lemma 8) and Lipschitz-continuous (Lemma 9). Additionally, in Appendix C, we prove the following two important lemmas:

Lemma 1 *If \mathcal{F} is monotone, then a strong solution is a weak solution.*

Lemma 2 *If \mathcal{F} is Lipschitz continuous, then a weak solution is a strong solution.*

Therefore, by the monotonicity and Lipschitzness of \mathcal{F} , the strong and weak solutions (described, respectively, by Equation (20) and Equation (21)) are equally valid criterion for Nash equilibria of the finite-valued zero-sum quantum games under consideration. We will generally leverage the weak-solution criterion in our analysis.

Finally, in this work, we focus on the iteration complexity of our methods for computing an ϵ -approximate Nash equilibrium, $\tilde{\Psi}^* = (\tilde{\alpha}^*, \tilde{\beta}^*)$, which—as defined by [29]—must satisfy

$$\max_{\alpha} u(\alpha, \tilde{\beta}^*) - \epsilon \|U\|_{\infty} \leq u(\tilde{\alpha}^*, \tilde{\beta}^*) \leq \min_{\beta} u(\tilde{\alpha}^*, \beta) + \epsilon \|U\|_{\infty},$$

where $\|U\|_{\infty}$ denotes the infinity norm of the payoff observable U . From this definition and the maximal utility of Nash equilibria it follows that

$$\epsilon \|U\|_{\infty} \geq u(\tilde{\alpha}^*, \tilde{\beta}^*) - \min_{\beta} u(\tilde{\alpha}^*, \beta) \geq 0 \quad (22)$$

$$\epsilon \|U\|_{\infty} \geq \max_{\alpha} u(\alpha, \tilde{\beta}^*) - u(\tilde{\alpha}^*, \tilde{\beta}^*) \geq 0. \quad (23)$$

Summing these equations gives the *duality gap*,

$$\mathcal{DG}(\alpha', \beta') = \max_{\alpha} u(\alpha, \beta') - \min_{\beta} u(\alpha', \beta), \quad (24)$$

which effectively serves as a measure of proximity between any proposed states, (α', β') , and an exact Nash equilibrium, (α^*, β^*) . From Equation (22) it follows that for an ϵ -Nash equilibrium, the duality gap is bounded as

$$\epsilon' = 2\epsilon \|U\|_{\infty} \geq \mathcal{DG}(\alpha', \beta') \geq 0. \quad (25)$$

Therefore, to evaluate the error of a proposed solution $\tilde{\Psi}$, relative to any true Nash equilibrium, we leverage the notion of an *error/merit function*, which corresponds with the duality gap in the setting of min-max optimization:

$$\text{Error}(\tilde{\Psi}) = \sup_{\Psi \in \mathcal{C}} \text{Tr} \left[(\Psi - \tilde{\Psi}) \mathcal{F}(\Psi) \right]. \quad (26)$$

It is easy to check that $\text{Error}(\tilde{\Psi}) \geq 0$ for all $\tilde{\Psi}$ and $\text{Error}(\tilde{\Psi}) = 0$ if and only if $\tilde{\Psi} = \Psi^*$.

2.3 Convexity, Smoothness, & Duality

Crucial to our analysis will be functional notions of convexity, smoothness, and duality [59, 60]. We begin by defining convexity and strong convexity of a differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$.

Definition 1 (Convex Function) *A differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ is convex if*

$$f(X) \geq f(Y) + \langle \nabla f(Y), X - Y \rangle, \quad \forall X, Y \in \mathcal{C}. \quad (27)$$

Definition 2 (μ -Strongly Convex Function) *A convex differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ is μ -strongly convex with respect to norm $\|\cdot\|$ if*

$$f(X) \geq f(Y) + \langle \nabla f(Y), X - Y \rangle + \frac{\mu}{2} \|X - Y\|^2, \quad \forall X, Y \in \mathcal{C}, \quad (28)$$

or equivalently, following from the definition of the Bregman divergence in Equation (40),

$$\mathcal{BD}_f(X\|Y) \geq \frac{\mu}{2} \|X - Y\|^2, \quad \forall X, Y \in \mathcal{C}. \quad (29)$$

Since strong convexity implies standard convexity (with $\mu = 0$), the optimal points (minima/maxima) of any strongly convex function must satisfy first-order optimality.

Fact 1 (First-Order Optimality Condition) *For an optimization problem,*

$$\min f(X) \quad \text{subject to } X \in \mathcal{X},$$

with convex differentiable f , a feasible point X is optimal if and only if

$$\langle \nabla f(X), Y - X \rangle \geq 0, \quad \forall Y \in \mathcal{X}.$$

Another property of interest for the convex differentiable function f is smoothness, which implies Lipschitz continuity of ∇f .

Definition 3 (β -Smooth Function) *A convex differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ is β -smooth with respect to norm $\|\cdot\|$ if*

$$f(X) \leq f(Y) + \langle \nabla f(Y), X - Y \rangle + \frac{\beta}{2} \|X - Y\|^2, \quad \forall X, Y \in \mathcal{C}. \quad (30)$$

or equivalently, following from the definition of the Bregman divergence in Equation (40),

$$\mathcal{BD}_f(X\|Y) \leq \frac{\beta}{2} \|X - Y\|^2, \quad \forall X, Y \in \mathcal{C}. \quad (31)$$

Before establishing the Lipschitz condition of ∇f , we introduce the *dual norm*.

Definition 4 (Dual Norm) *If the generic norm of a matrix X is denoted $\|X\|$, the dual norm is defined as*

$$\|X\|_* = \sup\{\langle Y, X \rangle : \|Y\| \leq 1\}. \quad (32)$$

As shown in [61, 62], the Frobenius norm $\|\cdot\|_F$ is self-dual, meaning $\|X\|_* = \|X\|_F$.

Proposition 1 (β -Lipschitz) *A convex differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ is β -smooth with respect to norm $\|\cdot\|$ if its gradient ∇f is β -Lipschitz continuous,*

$$\|\nabla f(X) - \nabla f(Y)\|_* \leq \beta \|X - Y\| \quad \forall X, Y \in \mathcal{C}. \quad (33)$$

For the equivalence of Proposition 1 and Definition 3, see [63]. Geometrically, β -smoothness can be interpreted as a convex quadratic *upper* bound,

$$g_X^\dagger(Y) = f(Y) + \langle \nabla f(Y), X - Y \rangle + \frac{\beta}{2} \|X - Y\|^2, \quad (34)$$

of f such that $g_X^\dagger(X) = f(X)$ and $g_X^\dagger(Y) \geq f(Y)$, $\forall Y \in \mathcal{C}$. Meanwhile, μ -strong convexity can be interpreted as a convex quadratic *lower* bound,

$$g_X^-(Y) = f(Y) + \langle \nabla f(Y), X - Y \rangle + \frac{\mu}{2} \|X - Y\|^2, \quad (35)$$

of f such that $g_X^-(X) = f(X)$ and $g_X^-(Y) \leq f(Y)$, $\forall Y \in \mathcal{C}$. In this way, strong convexity is dual to smoothness.

Fenchel conjugacy relates two equivalent representations of a convex function. In the standard representation, a convex function f is represented by pairs $(X, f(X))$ consisting of coordinates X and corresponding function evaluations $f(X)$. Alternatively, as depicted in Figure 5, a convex function can be represented by its tangent at each point. These tangents are encoded as the pairs $(\theta, f^*(\theta))$, where $\theta = \nabla f(X)$ for some X is the gradient and $-f^*(\theta)$ is the intercept. The function f^* , which calculates the tangent's intercept $f^*(\theta)$ from a given gradient θ , is known as the Fenchel conjugate function, formally defined as follows.

Definition 5 (Fenchel Conjugate) *The Fenchel conjugate of a function $f : \mathcal{C} \rightarrow \mathbb{R}$ is the function $f^* : \mathcal{C} \rightarrow \mathbb{R}$ such that*

$$f^*(\theta) = \max_{X \in \mathcal{C}} \{\langle X, \theta \rangle - f(X)\}. \quad (36)$$

Relating back to Definition 4 of the dual norm, if it were the case that for all X , $f(X) = 0$, then $f^*(\theta) = \max_{X \in \mathcal{C}} \langle X, \theta \rangle = \|\theta\|_*$. Therefore, the Fenchel conjugate can be seen as maximizing the dual norm while also minimizing the penalty induced by function f . In this way, f can be seen as a “regularizer” function.

The definition of Fenchel conjugacy immediately implies the *Fenchel-Young Inequality*,

$$\forall X \in \mathcal{C}, \quad f^*(\theta) \geq \langle X, \theta \rangle - f(X). \quad (37)$$

Furthermore, taking the gradient of the Fenchel conjugate should result in the value of X which maximizes $\langle X, \theta \rangle - f(X)$, resulting in *Danskin's Theorem*, as follows. For more details, refer to Proposition 4.5.1 of [64].

Proposition 2 (Danskin's Theorem) *Let $f : \mathcal{C} \rightarrow \mathbb{R}$ be a strongly convex function. For all $X \in \mathcal{C}$, we have*

$$\nabla f^*(\theta) = \arg \min_{X \in \mathcal{C}} \{f(X) - \langle X, \theta \rangle\}.$$




Figure 5: Illustration of a function f and its Fenchel conjugate f^* . The function f is represented by a collection of pairs $(X, f(X))$, while its Fenchel conjugate f^* is visualized as the set of tangents to f , characterized by pairs $(\theta, f^*(\theta))$. [Adapted from Figure 2.3 in [60].]

Finally, as stated below and explained further in Corollary 3.5.11 of [65], there is a duality between the the convexity of f and the smoothness of f^* .

Lemma 3 (Strong-Smooth Duality) *A closed convex function $f : \mathcal{C} \rightarrow \mathbb{R}$ is α -strongly convex with respect to norm $\|\cdot\|$ if and only if f^* is $\frac{1}{\alpha}$ -smooth with respect to the corresponding dual norm $\|\cdot\|_*$.*

2.4 Bregman Divergence & Mirror Map

Critical to understanding our proposed algorithm and its improved convergence over Jain and Watrous’ MMWU proposal are the notions of “mirror maps” and Bregman divergence.

Historically, the first proposals for constrained Gradient Descent Ascent leveraged Euclidean (ℓ_2 norm) projections in order to ensure that the algorithm output would lie within the feasible set. However, it was eventually realized that for certain problems, as we will discuss in the following subsection, regularization based on this ℓ_2 norm results in suboptimal constants in the convergence rates and, by leveraging different distance metrics, faster rates can be achieved. In order to develop algorithms which could encompass any such desired distance metric, the Bregman divergence was introduced.

The Bregman divergence is parameterized by a “distance-generating” or “regularizer” function. Formally, a *distance-generating function* (DGF) on the d -dimensional spectraplex \mathcal{H}_+^d is defined as a convex, lower semi-continuous function $h : \mathbb{C}^d \rightarrow \mathbb{R} \cup \{\infty\}$, where the *effective* domain of h is \mathcal{X}^d and \mathbb{C}^d is the space of $2^d \times 2^d$ complex matrices. Denoting the space of $2^d \times 2^d$ real matrices as \mathbb{R}^d , let $\partial h : \mathcal{H}_+^d \rightarrow \mathbb{R}^d$ denote the *subdifferential* of h defined as⁴

$$\partial h(X) := \left\{ V \in \mathbb{R}^d : h(Y) \geq h(X) + \text{Tr}[(Y - X)^\dagger V], \forall Y \in \mathcal{H}_+^d \right\}. \quad (38)$$

The *domain of subdifferentiability* of h is defined as

$$\text{dom} \partial h := \{X \in \mathcal{X}^d : \partial h(X) \neq \emptyset\}. \quad (39)$$

⁴The subdifferential is the set-valued generalization of the derivative for convex functions which are not necessarily differentiable.




Figure 6: **Bregman Divergence.** The Bregman divergence $\mathcal{BD}_h(U||V)$ between two states in the primal space $U, V \in \mathcal{X}^d$, relative to regularizer h , is given by the difference in the dual space ($\mathbb{R}^{2^d \times 2^d}$) between $h(U)$ and its linear approximation via the first-order Taylor expansion of h around V , $\tilde{h}_V(U)$.

Having defined the DGF h and its subdifferential ∂h , we are ready to define the Bregman divergence. Given two states $U, V \in \mathcal{H}_+^d$, the Bregman divergence induced by h is the non-negative function $\mathcal{BD}_h : \text{dom}h \times \text{dom}\partial h \rightarrow \mathbb{R}$ defined as

$$\mathcal{BD}_h(U||V) := h(U) - h(V) - \text{Tr}[(U - V)^\dagger \partial h(V)], \quad (40)$$

where the subdifferential maps from states in the *feasible region* \mathcal{H}_+^d of the *primal space* \mathbb{C}^d , to the *dual space* of subgradients \mathbb{R}^d . The function ∂h , which performs a primal-to-dual mapping, is often referred to as a *mirror map*.

Intuitively, as illustrated in Figure 6, the Bregman divergence $\mathcal{BD}_h(U||V)$ between two states U and V in the primal space, relative to regularizer h , is given by the difference in the dual space between $h(U)$ and its linear approximation via the first-order Taylor expansion of h around V ,

$$\tilde{h}_V(U) = h(V) + \text{Tr}[(U - V)^\dagger \partial h(V)]. \quad (41)$$

Since h is convex, this difference is always non-negative. Note that $\mathcal{BD}_h(U||V)$ is not necessarily equivalent to $\mathcal{BD}_h(V||U)$.

Finally, our analysis will use the well-known *Three-Point Identity*,

$$\text{Tr}[(Z - X)^\dagger \nabla_{X^\top} \mathcal{BD}_h(X||Y)] = \mathcal{BD}_h(Z||Y) - \mathcal{BD}_h(X||Y) - \mathcal{BD}_h(Z||X), \quad (42)$$

which follows from the definition of the Bregman divergence [66].

2.5 Mirror & Proximal Steps

We will now introduce “mirror steps” and “proximal steps,” which incorporate the mirror map generalization of distance in constrained optimization updates.⁵ These steps implement the algorithmic state updates of the classical games algorithms discussed in the Prior

⁵For further background, we refer the interested reader to [66] and Chapter 4 of [67].




Figure 7: **Mirror & Proximal Step.** Visualizations of the mirror and proximal steps described in [Algorithm 2](#) and [Algorithm 3](#), respectively. [Adapted from Figures 4.1 and 4.2 in [67].]

Work ([Section 1.2](#)). Similar to a standard constrained optimization gradient step, the goal of these procedures is to produce an update that lies in the feasible set, closer to the optimum. However, by generalizing the Euclidean distance of traditional Gradient Descent Ascent steps, mirror and proximal steps enable incorporation of domain-specific knowledge about the geometry of the optimization task into the updates. This can lead to substantial improvements in convergence rate.

For intuition, let us first consider a classical single-agent case as an illustrative example. When optimizing a function f over a $(d - 1)$ -dimensional simplex, restricted by the ℓ_∞ -norm, standard Projected Subgradient Descent (PSD) reaches a convergence rate of $\mathcal{O}(\sqrt{d/T})$. In contrast, classical Mirror Descent Ascent using a negative Shannon entropy regularizer significantly improves the rate to $\mathcal{O}(\sqrt{\log d/T})$.

Our objective is to adapt this methodology to the more challenging, exponentially larger $(4^d - 1)$ -dimensional spectraplex \mathcal{H}_+^d , in a min-max framework. Specifically, we aim to incorporate domain-specific regularization techniques to achieve an $\mathcal{O}(d/T)$ convergence rate. This would offer an exponential dimensionality improvement from standard Projected Subgradient Ascent Descent, which has an $\mathcal{O}(2^d/\sqrt{T})$ convergence rate [68, 69], and a quadratic speed-up over the Jain-Watrous algorithm, which has an $\mathcal{O}(d/\sqrt{T})$ convergence rate [29].

Note that intuitions for the following definitions and proofs are provided in [Appendix B](#).

2.5.1 Mirror Step

Given a function $f : \mathbb{C}^d \rightarrow \mathbb{R}$ that we aim to optimize, convex feasible set $\mathcal{H}_+^d \subset \mathbb{C}^d$, DGF h , subgradients ∂h , and initial primal state $X_t \in \mathcal{H}_+^d$, each call to a *mirror step* oracle generates an updated state X_{t+1} via [Algorithm 2](#) (as depicted in [Figure 7](#)). In Step (3), the mirror map, ∂h^{-1} , can be implemented as follows.

Definition 6 (Mirror Map) For a state D in the dual space \mathcal{D} , the mirror map, or regularized best response,

$$\text{Mir}\Pi_{\mathcal{P}}^h(D) := \arg \max_{P \in \mathcal{P}} \{\text{Tr}[D^\dagger P] - h(P)\} \quad (43)$$

implements the inverse mirror map, $\partial h^{-1} : \mathcal{D} \rightarrow \mathcal{P}$, mapping $D \in \mathcal{D}$ to the corresponding state P in the feasible region \mathcal{P} of the primal space.

Algorithm 2 Mirror Step

1. Initial primal state $X_t \in \mathcal{H}_+^d$ is mapped to a dual state $\nabla h(X_t) \in \partial h(X_t) \subset \mathbb{R}^d$.
2. Gradient $\nabla f : \mathbb{C}^d \rightarrow \mathbb{R}^d$ is calculated with respect to X_t .
3. In the dual space, the state is updated according to the gradient with step size η , resulting in a new dual state

$$\nabla h(X_{t+1}) = \nabla h(X_t) + \eta \nabla f(X_t) \in \mathbb{R}^d.$$

4. Mapping back to the primal space, this corresponds to the updated feasible state

$$X_{t+1} \in \partial h^{-1}(\nabla h(X_{t+1}) + \eta \nabla f(X_t)) \in \mathcal{H}_+^d.$$

Algorithm 3 Proximal Step

1. From initial feasible primal state $X_t \in \mathcal{H}_+^d$, updated feasible primal state $Y_{t+1} \in \mathcal{H}_+^d$ is generated via the Mirror Step procedure of [Algorithm 2](#).
2. Initial feasible primal state X_t is mapped to a dual state $\nabla h(X_t) \in \partial h(X_t) \subset \mathbb{R}^d$.
3. Gradient $\nabla f : \mathbb{C}^d \rightarrow \mathbb{R}^d$ is calculated with respect to Y_{t+1} .
4. In the dual space, the state is updated according to the gradient with step size η , resulting in a new dual state

$$\nabla h(X_{t+1}) = \nabla h(X_t) + \eta \nabla f(Y_{t+1}) \in \mathbb{R}^d.$$

5. Mapping back to the primal space, this corresponds to the updated feasible state

$$X_{t+1} \in \partial h^{-1}(\nabla h(X_t) + \eta \nabla f(Y_{t+1})) \in \mathcal{H}_+^d.$$

Therefore, the full mirror step procedure, mapping from X_t to X_{t+1} , is achieved via the mirror map of $Y_t = \nabla h(X_t) + \eta \nabla f(X_t)$ on the feasible set \mathcal{H}_+^d , as

$$X_{t+1} = \text{Mir}\Pi_{\mathcal{H}_+^d}^h(\nabla h(X_t) + \eta \nabla f(X_t)). \quad (44)$$

Alternatively, the mirror step procedure can be expressed as a *mirror proximal map*.

Lemma 4 (Mirror Step) For state X_t in the feasible set \mathcal{H}_+^d of primal space \mathbb{C}^d ,

$$\text{MirProx}\Pi_{\mathcal{H}_+^d}^{h,\eta,f}(X_t) := \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(X_t)^\dagger X_{t+1}] + \frac{1}{\eta} \mathcal{BD}_h(X_{t+1} || X_t) \right\}, \quad (45)$$

implements the mirror step mapping to an updated feasible set primal state $X_{t+1} \in \mathcal{H}_+^d$.

2.5.2 Proximal Step

We refer to the mirror prox projection as a *proximal map*. Its generalization comprises the *proximal step* oracle, described in [Algorithm 3](#) and depicted in [Figure 7](#).

Definition 7 (Proximal Map) For a state P in the feasible set of the primal space \mathcal{P} and a state D in the dual space, the proximal map,

$$\begin{aligned} \text{Prox}\Pi_{\mathcal{P}}^h(P, D) &:= \arg \min_{P^* \in \mathcal{P}} \left\{ \text{Tr}[D^\dagger(P^* - P)] + \mathcal{B}\mathcal{D}_h(P^*||P) \right\} \\ &= \arg \min_{P^* \in \mathcal{P}} \left\{ \text{Tr}[D^\dagger P^*] + \mathcal{B}\mathcal{D}_h(P^*||P) \right\}, \end{aligned} \quad (46)$$

implements the proximal mapping to an updated feasible set primal state $P^* \in \mathcal{P}$.

Applying this definition to our problem of interest yields the following map.

Lemma 5 (Proximal Step) For state $X_t \in \mathcal{H}_+^d$ in the feasible set of the primal space,

$$\begin{aligned} \text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(Y_{t+1})) &:= \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(Y_{t+1})^\dagger(X_{t+1} - X_t)] + \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1}||X_t) \right\} \\ &= \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(Y_{t+1})^\dagger X_{t+1}] + \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1}||X_t) \right\} \end{aligned} \quad (47)$$

implements the proximal step mapping to an updated feasible set primal state $X_{t+1} \in \mathcal{H}_+^d$.

Note that $\text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(X_t)) = \text{MirProx}\Pi_{\mathcal{H}_+^d}^{h,\eta,f}(X_t)$, which, by [Lemma 4](#), implies that the proximal step is a generalization of the mirror step and, thus, can be used to implement a mirror step. Thus, the proximal step procedure of [Algorithm 3](#) can be achieved algorithmically via two proximal maps as:

$$Y_{t+1} = \text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(X_t)) \quad (48)$$

$$X_{t+1} = \text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(Y_{t+1})). \quad (49)$$

Note that the proximal step can be alternatively written as a single update:

$$X_{t+1} = \text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta} \left(X_t, \nabla f \left(\text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(X_t)) \right) \right). \quad (50)$$

Standard gradient-based optimization methods are generally based on a single step in the direction of the gradient. In contrast, proximal step methods are based on a two-step approach:

1. From initial point X_t , apply a “look-ahead” step, resulting in a “prediction” Y_{t+1} .
2. Use prediction Y_{t+1} in a “correction” step that achieves the update X_{t+1} .

As we will discuss in [Section 3.3](#), replacing mirror with proximal steps provides the desired acceleration from $\mathcal{O}_d(1/\epsilon^2)$ to $\mathcal{O}_d(1/\epsilon)$. Furthermore, the usage of the Bregman divergence enables a geometry-aware projection and, specifically, can be used to achieve a logarithmic improvement on dependence of the dimensionality of the spectraplex, as will be discussed in [Section 4.7](#).

2.5.3 Regularizers

In order to guarantee the well-posedness of the mirror and proximal projections, we will restrict h to the class of functions which are strongly convex. In the classical games setting, as discussed in the Prior Work (Section 1.2), the most widely used regularizers are the Euclidean (ℓ_2) norm ($h(x) = \frac{1}{2}\|x\|_2^2$) and the negative entropy ($h(x) = \sum_i x_i \log(x_i)$). In this work, we will leverage the quantum analogs of these classical regularizers: the Frobenius norm ($h(X) = \frac{1}{2}\|X\|_F^2$) and the von Neumann entropy ($h(X) = \text{Tr}[X \log X]$). As highlighted in the following lemmas (proved in Appendix B), the mirror and proximal maps take well-known forms when these regularizers are used.

Lemma 6 (Bregman Divergence of the Frobenius Norm) *If h is the squared Frobenius norm,*

$$h(X) = \frac{1}{2}\|X\|_F^2 = \frac{1}{2}\text{Tr}[X^\dagger X], \quad (51)$$

then the Bregman divergence is the squared Frobenius distance,

$$\mathcal{BD}_h(X||Y) = \frac{1}{2}\|X - Y\|_F^2, \quad (52)$$

the corresponding mirror map (regularized best-response) is the orthogonal projection,

$$\text{Mir}\Pi_{\mathcal{C}}^h(Y) = \arg \min_{C \in \mathcal{C}} \|Y - C\|_F^2 = \text{Orth}\Pi_{\mathcal{C}}(Y), \quad (53)$$

and the corresponding proximal map is

$$\text{Prox}\Pi_{\mathcal{C}}^h(X, Y) = \arg \min_{C \in \mathcal{C}} \|X + Y - C\|_F^2 = \text{Orth}\Pi_{\mathcal{C}}(X + Y). \quad (54)$$

Lemma 7 (Bregman Divergence of the von Neumann Entropy) *If h is the von Neumann entropy,*

$$h(X) = \text{Tr}[X \log X], \quad (55)$$

then the Bregman divergence is the quantum relative entropy

$$\mathcal{BD}_h(X||Y) = \text{Tr}[X(\log X - \log Y)], \quad (56)$$

the corresponding mirror map (regularized best-response) is the logit map,

$$\text{Mir}\Pi_{\mathcal{C}}^h(Y) = \frac{\exp(Y)}{\text{Tr}[\exp(Y)]} = \Lambda(Y), \quad (57)$$

and the corresponding proximal map is

$$\text{Prox}\Pi_{\mathcal{C}}^h(X, Y) = \frac{\exp(\log X + Y)}{\text{Tr}[\exp(\log X + Y)]} = \Lambda(\log X + Y). \quad (58)$$

In Section 3, we will leverage these mirror and proximal steps (viewed as oracles) to craft efficient iterative algorithms for finding Nash equilibria of zero-sum quantum games.

3 Algorithms for Quantum Zero-Sum Games

As previously discussed, our goal is to efficiently find Nash equilibria, $\Psi^* = (\alpha^*, \beta^*)$, satisfying the variational inequality of Equation (20). This section describes existing algorithms, as well as our proposed algorithm for finding these Nash equilibria. Section 4 will demonstrate that our proposed algorithm achieves a quadratic speedup relative to existing algorithms.

3.1 Matrix Dual Averaging (MDA) Method

In 2008, Jain and Watrous introduced a parallel algorithm to find approximate Nash-equilibria of quantum zero-sum games [29]. Their algorithm was based on the Matrix Multiplicative Weights Update (MMWU) method [46, 34, 70]. Here, we will instead present and assess a more general method, which we will refer to as Matrix Dual Averaging (MDA),⁶ which, with different DGFs, parameterizes a family of algorithms that include Jain and Watrous’ MMWU algorithm.

MDA algorithms are played over several rounds, $t \in [0, N]$, of the quantum zero-sum game. Suppose that, in round t , Alice and Bob play states $\Psi_t = (\alpha_t, \beta_t)$.⁷ After performing his joint measurement, Roger returns feedback to Alice and Bob in the form of their respective payoff gradients, $\mathcal{F}_\alpha(\Psi_t)$ and $\mathcal{F}_\beta(\Psi_t)$, as defined in Equation (16) and Equation (17). Alice and Bob use this feedback to individually update their states so as to iteratively maximize their individual payoff. The general MDA family of algorithms (Algorithm 4) can be expressed concisely as the updates:

$$\Phi_{t+1} = \Phi_t + \mathcal{F}(\Psi_t) \quad (59)$$

$$\Psi_{t+1} = \text{Mir}\Pi_{\mathcal{C}}^h(\eta \Phi_{t+1}), \quad (60)$$

where η is a chosen step-size and h is a regularizing, convex DGF. As we will see, it is the choice of DGF h which instantiates different MDA algorithms, including MMWU. Furthermore, relating back to the mirror step picture of Figure 7, the intermediate states Φ are averaged in the dual space, but are mapped onto feasible states Ψ of the primal space by the mirror projection $\text{Mir}\Pi_{\mathcal{C}}^h$. Overall, the algorithm only incorporates first-order derivatives in the update steps.

We will now specify how each player chooses their next state via the previously described MDA method. Note that both Alice and Bob perform the same procedure, so we will only specify the procedure from Alice’s perspective. Upon receiving feedback $\mathcal{F}_\alpha(\Psi_t)$ from Roger, Alice calculates her cumulative feedback,

$$\mathcal{C}\mathcal{F}_\alpha^{(t)} = \sum_{i=0}^{t-1} \mathcal{F}_\alpha(\Psi_i), \quad (61)$$

which is simply the sum of the payoff gradients across all prior rounds. To calculate the next state she will play, α_{t+1} , Alice re-weights the cumulative feedback by step-size η and performs a mirror projection onto her spectraplex,

$$\alpha_{t+1} = \text{Mir}\Pi_{\mathcal{A}}^h(\eta \mathcal{C}\mathcal{F}_\alpha^{(t)}) = \arg \max_{\alpha \in \mathcal{A}} \{\langle \eta \mathcal{C}\mathcal{F}_\alpha^{(t)}, \alpha \rangle - h(\alpha)\}. \quad (62)$$

⁶The name is a reference to the analogous algorithm for classical games, i.e., Nesterov’s Dual Averaging [68]. Dual Averaging is sometimes also known as “Lazy” Mirror Descent, since the key difference from standard Mirror Descent is that averaging happens in the dual space (rather than mapping from primal to dual space and back in each iteration).

⁷Note that, in the first round ($t = 0$), Alice and Bob play the maximally mixed state $\Psi_0 = (\frac{1}{2^n}, \frac{1}{2^m})$.

Algorithm 4 Matrix Dual Averaging

Accuracy Parameter: ϵ

Regularization Function: $h : \{\mathcal{A}, \mathcal{B}\} \rightarrow \mathbb{R}$

Diameter (of \mathcal{BD}_h): $\mathcal{D}_h = \sup_{X, Y \in \mathcal{C}} \mathcal{BD}_h(X \| Y)$

Regularized Best Response: $\text{Mir}\Pi_{\mathcal{A}, \mathcal{B}}^h(\cdot) = \arg \max_{x \in \{\mathcal{A}, \mathcal{B}\}} \{\langle \cdot, x \rangle - h(x)\}$

$\eta \leftarrow \mu_h / (2\gamma_{\mathcal{F}})$

▷ Step Size

$N \leftarrow \lceil \mathcal{D}_h / (\eta \cdot \epsilon^2) \rceil$

▷ Number of Rounds

$(\alpha_0, \beta_0) \leftarrow \left(\frac{1}{2^n} \mathbb{1}_{\mathcal{A}}, \frac{1}{2^m} \mathbb{1}_{\mathcal{B}} \right)$

▷ State Initialization

for $t \in [1, N - 1]$ **do**

$\mathcal{CF}_{\alpha}^{(t)} \leftarrow \sum_{i=0}^{t-1} \mathcal{F}_{\alpha}(\Psi_i)$

▷ Calculate Cumulative Feedback

$\mathcal{CF}_{\beta}^{(t)} \leftarrow \sum_{i=0}^{t-1} \mathcal{F}_{\beta}(\Psi_i)$

$\alpha_t \leftarrow \text{Mir}\Pi_{\mathcal{A}}^h \left(\eta \mathcal{CF}_{\alpha}^{(t)} \right)$

▷ State Updates

$\beta_t \leftarrow \text{Mir}\Pi_{\mathcal{B}}^h \left(\eta \mathcal{CF}_{\beta}^{(t)} \right)$

end for

return $\left(\bar{\alpha} = \frac{1}{N} \sum_{t=0}^{N-1} \alpha_t, \bar{\beta} = \frac{1}{N} \sum_{t=0}^{N-1} \beta_t \right)$

We will now specify how different choices of h lead to different instantiations of the MDA algorithm. For example, following from [Lemma 6](#), if the regularizer is the squared Frobenius norm,

$$h(\alpha) = \frac{1}{2} \|\alpha\|_F^2, \quad (63)$$

then the mirror projection maps Alice's next state onto the orthogonal projection

$$\alpha_{t+1} = \text{Orth}\Pi_{\mathcal{A}}(\eta \mathcal{CF}_{\alpha}^{(t)}). \quad (64)$$

Meanwhile, following from [Lemma 7](#), if the regularizer is the negative von Neumann entropy,

$$h(\alpha) = \text{Tr}[\alpha \log \alpha], \quad (65)$$

then the mirror projection maps Alice's next state onto the logit map

$$\alpha_{t+1} = \Lambda_{\mathcal{A}} \left(\eta \mathcal{CF}_{\alpha}^{(t)} \right) = \frac{\exp \left(\eta \mathcal{CF}_{\alpha}^{(t)} \right)}{\text{Tr} \left(\exp \left(\eta \mathcal{CF}_{\alpha}^{(t)} \right) \right)} \quad (66)$$

and, thus, the MMWU algorithm of [\[29\]](#). However, as proven in [\[29\]](#),

Theorem 2 (MDA Rate) *MDA methods compute an ϵ -Nash equilibrium in $\mathcal{O}_d(1/\epsilon^2)$ steps.*

Thus, MMWU for quantum zero-sum games obtains a worse convergence rate than the best known algorithms for classical zero-sum games, which obtain an $\mathcal{O}_d(1/\epsilon)$ rate, as described in [Section 1.2](#) and depicted in [Figure 1](#). In the remainder of this work, we will show how the classically efficient Mirror Prox algorithm and its variants can be lifted to the quantum zero-sum game setting in order to obtain the desired $\mathcal{O}_d(1/\epsilon)$ rate.

3.2 Matrix Mirror Prox (MMP)

The classical Mirror Prox (MP) algorithm is described in [Section 1.2](#), with updates given in [Figure 1](#). In each step of the iterative MP optimization procedure, the gradient of the players’ utilities is calculated followed by the execution of a mirror step ([Algorithm 2](#)), as illustrated in [Figure 7](#). The mirror step can be implemented via a mirror map of the dual gradient, as in [Equation \(44\)](#), or mirror map of the primal state, as in [Equation \(45\)](#). The addition of proximal steps—or extra-gradient steps—is crucial as they introduce a refinement to the basic gradient step. After the initial gradient is computed and the first mirror step is taken, the algorithm does not immediately proceed with this information. Instead, it calculates a second gradient at the point reached after the first mirror step. This second gradient offers a more accurate direction for the update because it accounts for the immediate effects of the initial step, thus incorporating more information about the curvature and constraints of the utility space. This two-pronged gradient approach is what distinguishes the MP algorithm from simpler gradient descent methods, offering a sophisticated tool for navigating the intricate landscapes of variational inequalities.

Drawing inspiration from classical MP, we propose an analog for the quantum zero-sum games of interest. This quantum analog, which we will refer to as the Matrix Mirror Prox (MMP) method, has updates:

$$\Phi_{t+1} = \text{Prox}\Pi_{\mathcal{C}}^{h,\eta}(\Psi_t, \mathcal{F}(\Psi_t)) \quad (67)$$

$$\Psi_{t+1} = \text{Prox}\Pi_{\mathcal{C}}^{h,\eta}(\Psi_t, \mathcal{F}(\Phi_{t+1})). \quad (68)$$

Notice that these update steps are of the same form as the proximal step updates given in [Equation \(48\)](#) and [Equation \(49\)](#). Alternatively, the updates can be expressed purely in terms of the joint states Ψ_t :

$$\Psi_{t+1} = \text{Prox}\Pi_{\mathcal{C}}^{h,\eta}\left(\Psi_t, \mathcal{F}\left(\text{Prox}\Pi_{\mathcal{C}}^{h,\eta}(\Psi_t, \mathcal{F}(\Psi_t))\right)\right). \quad (69)$$

As discussed in [Section 2.5](#), the incorporation of proximal steps in the Matrix Mirror Prox algorithm offers two significant advantages over the conventional Projected Subgradient Descent Ascent method:

1. **Geometry-Awareness:** MMP replaces the Frobenius norm of PGDA with a generalized Bregman divergence. Similar to MDA, utilizing a mirror map with this Bregman divergence enables MMP to adapt to the inherent geometry of the problem space. This makes the algorithm highly versatile, capable of effectively handling a wide array of problems, including those with intricate constraints.
2. **Intermediate Proximal Step:** Proximal steps in MMP introduce an intermediate “prediction” state (i.e. state Y_t in [Section 2.5.2](#)), which plays a crucial role. At a high level, this prediction state extrapolates the future trajectory of the descent/ascent, thereby guiding the subsequent state update. This leverages the intuition that the gradient at the next state offers a more informative update gradient than that at the current state.

Importantly, this combination of geometry-aware mirror maps and forward-looking intermediate proximal steps accelerate the convergence of MMP with von Neumann regularization to min-max points, thereby making MMP a potent tool for solving min-max problems efficiently.

Algorithm 5 Optimistic Matrix Mirror Prox (OMMP)

Accuracy Parameter: ϵ
Regularization Function: $h : \{\mathcal{A}, \mathcal{B}\} \rightarrow \mathbb{R}$
Strong Convexity Parameter (of h): μ_h
Diameter (of \mathcal{BD}_h): $\mathcal{D}_h = \sup_{X, Y \in \mathcal{C}} \mathcal{BD}_h(X \| Y)$
Lipschitz Parameter (of \mathcal{F}): $\gamma_{\mathcal{F}}$
Proximal Map: $\text{Prox}\Pi_{\mathcal{A}, \mathcal{B}}^{h, \eta}(X, Y) := \arg \min_{C \in \mathcal{C}} \{\langle Y, C - X \rangle - \frac{1}{\eta} \mathcal{BD}_h(C \| X)\}$.

 $\eta \leftarrow \mu_h / (2\gamma_{\mathcal{F}})$
 \triangleright Step Size

 $N \leftarrow \lceil \mathcal{D}_h / (\eta \cdot \epsilon) \rceil$
 \triangleright Number of Rounds

 $(\alpha_0, \beta_0) \leftarrow \left(\frac{1}{2^n} \mathbb{1}_{\mathcal{A}}, \frac{1}{2^m} \mathbb{1}_{\mathcal{B}} \right)$
 \triangleright State Initialization

 $(\hat{\alpha}_0, \hat{\beta}_0) \leftarrow (\alpha_0, \beta_0)$
 \triangleright Optimistic-Momentum Initialization

for $t \in [1, N - 1]$ **do**
 $\alpha_{t+1} \leftarrow \text{Prox}\Pi_{\mathcal{A}}^{h, \eta}(\hat{\alpha}_t, \mathcal{F}_{\alpha}(\Psi_t))$
 \triangleright State Updates

 $\beta_{t+1} \leftarrow \text{Prox}\Pi_{\mathcal{B}}^{h, \eta}(\hat{\beta}_t, \mathcal{F}_{\beta}(\Psi_t))$
 $\hat{\alpha}_{t+1} \leftarrow \text{Prox}\Pi_{\mathcal{A}}^{h, \eta}(\hat{\alpha}_t, \mathcal{F}_{\alpha}(\Psi_{t+1}))$
 \triangleright Momentum Updates

 $\hat{\beta}_{t+1} \leftarrow \text{Prox}\Pi_{\mathcal{B}}^{h, \eta}(\hat{\beta}_t, \mathcal{F}_{\beta}(\Psi_{t+1}))$
end for
return $\left(\bar{\alpha} = \frac{1}{N} \sum_{t=0}^{N-1} \alpha_t, \bar{\beta} = \frac{1}{N} \sum_{t=0}^{N-1} \beta_t \right)$

As is the case for classical MP (discussed in [Section 1.2](#)), a key downside of MMP is that it requires two gradient calls per iteration. Specifically, in each round t , to calculate the updated joint state Ψ_{t+1} , Alice and Bob must query Roger with two distinct sets of states, Ψ_t and Φ_{t+1} , so as to obtain the respective feedback gradients $\mathcal{F}(\Psi_t)$ and $\mathcal{F}(\Phi_{t+1})$. Classically, there exist variants of MP with the same convergence rate that require only a single gradient call per iteration. The remainder of this work will focus on a proposed classically-inspired variant of MMP which requires only a single gradient call per iteration, and a proof that it obtains the desired $\mathcal{O}_d(1/\epsilon)$ convergence rate.

3.3 Optimistic Matrix Mirror-Prox (OMMP) Methods

We will now discuss the main contribution of this work, a novel approach to finding Nash equilibria of quantum zero-sum games, which we refer to as the Optimistic Matrix Mirror Prox (OMMP) method. This algorithm draws inspiration from seminal classical game theory work on Optimistic Mirror Prox (OMP)⁸ methods, as described in the Prior Work ([Section 1.2](#)). The OMMP algorithm, detailed in [Algorithm 5](#), can be succinctly represented through the following update equations:

$$\Psi_{t+1} = \text{Prox}\Pi_{\mathcal{C}}^{h, \eta}(\Phi_t, \mathcal{F}(\Psi_t)) \quad (70)$$

$$\Phi_{t+1} = \text{Prox}\Pi_{\mathcal{C}}^{h, \eta}(\Phi_t, \mathcal{F}(\Psi_{t+1})) \quad (71)$$

⁸There are a variety of classical algorithms of this general form, known as Popov's modified Arrow-Hurwicz algorithm, Past Extra-Gradient, Optimistic Extra-Gradient, or the Optimistic Gradient Variant.

Here, η is a predetermined step size, and ProxII refers to the proximal map, as elaborated in [Algorithm 3](#) and [Definition 7](#).

Although the update equations for OMMP closely resemble those of the standard MMP algorithm, a critical difference exists. As alluded to in the previous section, the OMMP variant necessitates only a single gradient call per iteration. While it may appear that both Alice and Bob require two feedback gradients, $\mathcal{F}(\Psi_t)$ and $\mathcal{F}(\Psi_{t+1})$, for each round t , they will have already acquired $\mathcal{F}(\Psi_t)$ during round $t - 1$. This allows them to reuse $\mathcal{F}(\Psi_t)$ and only make a query for a new gradient with one state, Ψ_{t+1} .

To gain some intuition for this algorithm, we will revisit the original proposal of “optimism” in this context, due to Popov [52]. In conventional Gradient Descent Ascent, both players update their strategies concurrently, moving in directions opposite to their individual gradients—akin to a tug-of-war. Optimistic Gradient Descent Ascent (OGDA) enhances this by incorporating a “look-ahead” feature. Players consider not only the current gradient but also anticipate the opponent’s forthcoming strategic adjustments. This anticipatory or “optimistic” approach is formalized by adding an extra term to the update rule, essentially an estimate of the future gradient.

This anticipatory element is also how OMMP diverges from standard MMP. Rather than employing a new prediction, optimistic OMMP utilizes a past “inertia” point, Φ_t . Unlike MMP, which focuses solely on future states, OMMP integrates past gradient information to stabilize its learning trajectory.

4 Convergence Analysis

We will now prove one of the main results of the paper:

Theorem 1 (Main Result) *The OMMP method computes ϵ -Nash equilibria in $\mathcal{O}_d(1/\epsilon)$ steps.*

This establishes that our proposed OMMP method yields a quadratic speedup compared to previous MDA-based approaches. Concretely, we will leverage the proof structure of Ene and Nguyen [56] to demonstrate that, with an appropriately chosen step-size η and for a given accuracy parameter ϵ , the *average iterate*

$$\bar{\Psi}_T = \frac{1}{T} \sum_{t=1}^T \Psi_t \tag{72}$$

converges to an ϵ -Nash equilibrium in $\mathcal{O}_d(1/\epsilon)$ iterations.⁹

We will conclude with a discussion of the dimensionality-dependence hidden in the $\mathcal{O}_d(\cdot)$ notation. In particular in [Section 4.7](#), we will show that the von Neumann entropy regularizer reduces OMMP’s dimensionality-dependence logarithmically relative to the Frobenius norm regularizer. Since OMMP with the von Neumann entropy regularizer instantiates the OMMWU algorithm, we thus prove that OMMWU obtains the desired $\mathcal{O}(d/\epsilon)$ iteration complexity and, thus, achieves a quadratic speedup to MMWU’s $\mathcal{O}(d/\epsilon^2)$ iteration complexity.

⁹It’s important to note the difference in the notation used to describe algorithmic convergence rates between the optimization and computer science communities. In the realm of optimization, *convergence rates* such as $\mathcal{O}(e^{-\rho T})$, $\mathcal{O}(1/T)$, and $\mathcal{O}(1/\sqrt{T})$ are commonly employed to indicate how swiftly an algorithm approximates the optimal solution as the iteration count T increases. In contrast, the computer science literature tends to recast these convergence rates as *number of iterations* required to achieve an ϵ -accurate solution, thereby giving a more direct account of computational efficiency. Expressed in terms of accuracy parameter ϵ , this leads to expressions like $\mathcal{O}(\log(1/\epsilon))$, $\mathcal{O}(1/\epsilon)$, and $\mathcal{O}(1/\epsilon^2)$.

4.1 Monotonicity of \mathcal{F}

Critical to our analysis is the monotonicity of the feedback operator \mathcal{F} , which is a key property of the finite-valued quantum zero sum games of interest.

Lemma 8 (Monotonicity of \mathcal{F}) \mathcal{F} is monotone or, equivalently,

$$\text{Tr}[(\mathcal{F}(X) - \mathcal{F}(Y))(X - Y)] \geq 0, \quad \forall X, Y \in \mathcal{C}. \quad (73)$$

The proof of this result employs techniques involving the Pauli decomposition of the feedback operator, thereby deviating from the conventional proof for monotonicity of classical normal game operators. We defer the proof of \mathcal{F} 's monotonicity (in addition to many other properties) to [Appendix C.2](#).

4.2 Error Decomposition

Following from our discussion of Nash equilibria in [Section 2.2](#) and the natural error function, given in [Equation \(26\)](#), the error of the average iterate is defined as

$$\text{Error}(\bar{\Psi}_T) = \sup_{Z \in \mathcal{C}} \text{Tr}[\mathcal{F}(Z)(\bar{\Psi}_T - Z)] = \frac{1}{T} \sup_{Z \in \mathcal{C}} \sum_{t=1}^T \text{Tr}[\mathcal{F}(Z)(\Psi_t - Z)]. \quad (74)$$

From the monotonicity of \mathcal{F} , given in [Equation \(73\)](#), for Ψ_t and some state $Z \in \mathcal{C}$,

$$\text{Tr}[\mathcal{F}(Z)(\Psi_t - Z)] \leq \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - Z)], \quad (75)$$

which implies that the average-iterate error is upper bounded as

$$\text{Error}(\bar{\Psi}_T) \leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \sum_{t=1}^T \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - Z)]. \quad (76)$$

In order to produce an explicit bound, we begin by decomposing the terms of the summation, as follows, resulting in a sum of three distinct terms: (A_t) , (B_t) , and (C_t) .

$$\begin{aligned} \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - Z)] &= \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - \Phi_t + \Phi_t - Z)] \\ &= \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - \Phi_t) + \mathcal{F}(\Psi_t)(\Phi_t - Z)] = \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - \Phi_t)] + \text{Tr}[\mathcal{F}(\Psi_t)(\Phi_t - Z)] \\ &= \text{Tr}[(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1}) + \mathcal{F}(\Psi_{t-1}))(\Psi_t - \Phi_t)] + \text{Tr}[\mathcal{F}(\Psi_t)(\Phi_t - Z)] \\ &= \text{Tr}[(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1}))(\Psi_t - \Phi_t) + \mathcal{F}(\Psi_{t-1})(\Psi_t - \Phi_t)] + \text{Tr}[\mathcal{F}(\Psi_t)(\Phi_t - Z)] \\ &= \underbrace{\text{Tr}[\mathcal{F}(\Psi_t)(\Phi_t - Z)]}_{(A_t)} + \underbrace{\text{Tr}[(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1}))(\Psi_t - \Phi_t)]}_{(B_t)} + \underbrace{\text{Tr}[\mathcal{F}(\Psi_{t-1})(\Psi_t - \Phi_t)]}_{(C_t)}. \end{aligned} \quad (77)$$

In the following analysis, we will show that terms (A_t) and (C_t) can be upper bounded via the optimality of the proximal updates in the OMMP algorithm. Furthermore, the term (B_t) can be upper bounded via the strong convexity, smoothness, and duality of the OMMP proximal updates.

4.3 Optimality Analysis of (A_t)

We begin by upper bounding the term (A_t) . Recall from the definition of the OMMP algorithm in [Section 3.3](#) that Φ_t is the outcome of the proximal projection of [Equation \(71\)](#), for a regularization function h . Therefore, Φ_t can be expressed as

$$\Phi_t = \arg \min_{Z \in \mathcal{C}} \phi_t(Z), \quad (78)$$

where $\phi_t(Z)$ is the proximal optimization term

$$\phi_t(Z) = \text{Tr} [\mathcal{F}(\Psi_t) (Z - \Phi_{t-1})] + \frac{1}{\eta} \mathcal{BD}_h(Z \| \Phi_{t-1}). \quad (79)$$

Note that ϕ_t is differentiable, with gradient

$$\nabla_{Z^\top} \phi_t(Z) = \mathcal{F}(\Psi_t) + \frac{1}{\eta} \nabla_{Z^\top} \mathcal{BD}_h(Z \| \Phi_{t-1}). \quad (80)$$

Using the definitions of [Section 2.3](#), in the supplement ([Appendix D](#)) we further show that:

Proposition 3 *For μ -strongly convex regularization function h with respect to norm $\|\cdot\|$, ϕ_t is $\frac{\mu}{\eta}$ -strongly convex with respect to norm $\|\cdot\|$.*

Since ϕ_t is strongly convex and differentiable it satisfies the First-Order Optimality Condition ([Fact 1](#)) for feasible solution momentum Φ_t ,

$$\text{Tr} \left[\nabla_{\Phi_t^\top} \phi_t(\Phi_t) (Z - \Phi_t) \right] \geq 0, \quad \forall Z \in \mathcal{C}. \quad (81)$$

Plugging the gradient of [Equation \(80\)](#) into this expression gives

$$\begin{aligned} & \text{Tr} \left[\left(\mathcal{F}(\Psi_t) + \frac{1}{\eta} \nabla_{\Phi_t^\top} \mathcal{BD}_h(\Phi_t \| \Phi_{t-1}) \right) (\Phi_t - Z) \right] \leq 0 \\ & \text{Tr} [\mathcal{F}(\Psi_t) (\Phi_t - Z)] + \frac{1}{\eta} \text{Tr} \left[\nabla_{\Phi_t^\top} \mathcal{BD}_h(\Phi_t \| \Phi_{t-1}) (\Phi_t - Z) \right] \leq 0 \\ & \text{Tr} [\mathcal{F}(\Psi_t) (\Phi_t - Z)] \leq \frac{1}{\eta} \text{Tr} \left[-\nabla_{\Phi_t^\top} \mathcal{BD}_h(\Phi_t \| \Phi_{t-1}) (\Phi_t - Z) \right]. \end{aligned} \quad (82)$$

From our definition of (A_t) in [Equation \(77\)](#) and the three-point identity for the Bregman divergence given in [Equation \(42\)](#), it follows that, $\forall Z \in \mathcal{C}$,

$$(A_t) \leq \frac{1}{\eta} \left(\mathcal{BD}_h(Z \| \Phi_{t-1}) - \mathcal{BD}_h(\Phi_t \| \Phi_{t-1}) - \mathcal{BD}_h(Z \| \Phi_t) \right). \quad (83)$$

4.4 Optimality Analysis of (C_t)

We will now upper bound (C_t) , similarly to how we upper bounded (A_t) . Recall from the definition of the OMMP algorithm in [Section 3.3](#) that Ψ_t is the outcome of the proximal projection of [Equation \(70\)](#), for a regularization function h . Therefore, Ψ_t can be expressed as

$$\Psi_t = \arg \min_{Z \in \mathcal{C}} \psi_t(Z), \quad (84)$$

where $\psi_t(Z)$ is the proximal optimization term

$$\psi_t(Z) = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (Z - \Phi_{t-1})] + \frac{1}{\eta} \mathcal{BD}_h(Z \| \Phi_{t-1}). \quad (85)$$

Note that ψ_t is differentiable, with gradient

$$\nabla_{Z^\top} \psi_t(Z) = \mathcal{F}(\Psi_{t-1}) + \frac{1}{\eta} \nabla_{Z^\top} \mathcal{BD}_h(Z \| \Phi_{t-1}). \quad (86)$$

Similarly with the previous section, in the supplement ([Appendix D](#)) we further show that:

Proposition 4 For μ -strongly convex regularization function h with respect to norm $\|\cdot\|$, ψ_t is $\frac{\mu}{\eta}$ -strongly convex with respect to norm $\|\cdot\|$.

Since ψ_t is strongly convex and differentiable, it satisfies the First-Order Optimality Condition ([Fact 1](#)) for a feasible solution state Ψ_t ,

$$\text{Tr} \left[\nabla_{\Psi_t^\top} \psi_t(\Psi_t) (Z - \Psi_t) \right] \geq 0, \quad \forall Z \in \mathcal{C}. \quad (87)$$

Plugging the gradient of [Equation \(86\)](#) into this expression gives, $\forall Z \in \mathcal{C}$,

$$\text{Tr} \left[\left(\mathcal{F}(\Psi_{t-1}) + \frac{1}{\eta} \nabla_{\Psi_t^\top} \mathcal{BD}_h(\Psi_t \| \Phi_{t-1}) \right) (\Psi_t - Z) \right] \leq 0 \quad (88)$$

$$\text{Tr} [\mathcal{F}(\Psi_{t-1}) (\Psi_t - Z)] \leq \frac{1}{\eta} \text{Tr} \left[\nabla_{\Psi_t^\top} \mathcal{BD}_h(\Psi_t \| \Phi_{t-1}) (Z - \Psi_t) \right]. \quad (89)$$

Since the expression holds for all $Z \in \mathcal{C}$, we can choose to set $Z = \Phi_t$, which implies

$$\text{Tr} [\mathcal{F}(\Psi_{t-1}) (\Psi_t - \Phi_t)] \leq \frac{1}{\eta} \text{Tr} \left[\nabla_{\Psi_t^\top} \mathcal{BD}_h(\Psi_t \| \Phi_{t-1}) (\Phi_t - \Psi_t) \right]. \quad (90)$$

From our definition of (C_t) in [Equation \(77\)](#) and the Three-Point Identity for the Bregman divergence, given in [Equation \(42\)](#), it follows that

$$(C_t) \leq \frac{1}{\eta} \left(\mathcal{BD}_h(\Phi_t \| \Phi_{t-1}) - \mathcal{BD}_h(\Psi_t \| \Phi_{t-1}) - \mathcal{BD}_h(\Phi_t \| \Psi_t) \right). \quad (91)$$

4.5 An Upper Bound for (B_t)

By the Cauchy-Schwarz inequality for Frobenius norm, we can upper bound (B_t) as

$$(B_t) = \text{Tr} [(F(\Psi_t) - F(\Psi_{t-1})) (\Psi_t - \Phi_t)] \leq \|F(\Psi_t) - F(\Psi_{t-1})\|_F \|\Psi_t - \Phi_t\|_F. \quad (92)$$

or more generally, with respect to norms $(\|\cdot\|, \|\cdot\|_*)$

$$(B_t) = \text{Tr} [(F(\Psi_t) - F(\Psi_{t-1})) (\Psi_t - \Phi_t)] \leq \|F(\Psi_t) - F(\Psi_{t-1})\|_* \|\Psi_t - \Phi_t\|. \quad (93)$$

To further manipulate this upper bound for (B_t) , we leverage the ideas from Fenchel duality and convexity introduced in [Section 2.3](#).

Following from [Equation \(84\)](#) and [Equation \(79\)](#), respectively, Ψ_t is the minimizer of ψ_t and Φ_t is the minimizer of ϕ_t . Re-expressing ϕ_t in terms of ψ_t ,

$$\phi_t(U) = \psi_t(U) + \text{Tr} [(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1})) (U - \Phi_{t-1})] \quad (94)$$

implies that Φ_t can be expressed in terms of a minimization of ψ_t ,

$$\Phi_t = \arg \min_{U \in \mathcal{C}} \{ \psi_t(U) + \text{Tr} [(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1})) U] \}. \quad (95)$$

Since we assumed that h is μ -strongly convex, by [Proposition 4](#) ψ_t must be $\frac{\mu}{\eta}$ -strongly convex. Denoting the Fenchel conjugate of ψ_t as ψ_t^* , since ψ_t is strongly convex, Danskin's Theorem ([Proposition 2](#)) states that $\nabla \psi_t^*$ is well-defined, for all $V \in \mathcal{C}$, as

$$\nabla \psi_t^*(V) = \arg \min_{U \in \mathcal{C}} \{ \psi_t(U) - \text{Tr} [V U] \}. \quad (96)$$

In the case that $V = 0$, Equation (78) implies that

$$\nabla\psi_t^*(0) = \arg \min_{U \in \mathcal{C}} \psi_t(U) = \Psi_t. \quad (97)$$

Meanwhile, if $V = \mathcal{F}(\Psi_{t-1}) - \mathcal{F}(\Psi_t)$, Equation (95) implies that

$$\nabla\psi_t^*(\mathcal{F}(\Psi_{t-1}) - \mathcal{F}(\Psi_t)) = \arg \min_{U \in \mathcal{C}} \{\psi_t(U) + \text{Tr}[(\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1}))U]\} = \Phi_t. \quad (98)$$

Therefore,

$$\|\Psi_t - \Phi_t\|_F = \|\nabla\psi_t^*(0) - \nabla\psi_t^*(\mathcal{F}(\Psi_{t-1}) - \mathcal{F}(\Psi_t))\|_F. \quad (99)$$

Furthermore, since ψ_t is $\frac{\mu}{\eta}$ -strongly convex with respect to the norm $\|\cdot\|$, Strong-Smooth Duality (Lemma 3) implies that ψ_t^* is $\frac{\eta}{\mu}$ -smooth with respect to the norm $\|\cdot\|_*$. By Proposition 1, this implies that $\nabla\psi_t^*$ is $\frac{\eta}{\mu}$ -Lipschitz continuous with respect to $\|\cdot\|_*$,

$$\|\nabla\psi_t^*(X) - \nabla\psi_t^*(Y)\| = \|\nabla\psi_t^*(X) - \nabla\psi_t^*(Y)\|_{**} \leq \frac{\eta}{\mu} \|X - Y\|_*. \quad (100)$$

Equivalently for a pair of norms, we get Plugging this and Equation (99) into Equation (92) produces the final (B_t) upper bound,

$$(B_t) \leq \frac{\eta}{\mu} \|\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1})\|_*^2. \quad (101)$$

4.6 Average-Iterate Error Upper Bound

Combining the average-iterate error upper bound of Equation (76) and its decomposition, as expressed in Equation (77), we have that

$$\text{Error}(\bar{\Psi}_T) \leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \sum_{t=1}^T \text{Tr}[\mathcal{F}(\Psi_t)(\Psi_t - Z)] = \frac{1}{T} \sup_{Z \in \mathcal{C}} \sum_{t=1}^T (A_t) + (B_t) + (C_t).$$

Plugging in the upper bounds for (A_t) from Equation (83), (B_t) from Equation (101), and (C_t) from Equation (91) gives

$$\begin{aligned} \text{Error}(\bar{\Psi}_T) &\leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \sum_{t=1}^T \left(\frac{1}{\eta} \left(\mathcal{BD}_h(Z\|\Phi_{t-1}) - (\mathcal{BD}_h(Z\|\Phi_t) + \mathcal{BD}_h(\Psi_t\|\Phi_{t-1}) + \mathcal{BD}_h(\Phi_t\|\Psi_t)) \right) \right. \\ &\quad \left. + \frac{\eta}{\mu} \|\mathcal{F}(\Psi_t)^\top - \mathcal{F}(\Psi_{t-1})^\top\|_*^2 \right) \\ &\leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \underbrace{\frac{1}{\eta} \sum_{t=1}^T (\mathcal{BD}_h(Z\|\Phi_{t-1}) - \mathcal{BD}_h(Z\|\Phi_t))}_{(I)} - \frac{1}{\eta} \sum_{t=1}^T (\mathcal{BD}_h(\Psi_t\|\Phi_{t-1}) + \mathcal{BD}_h(\Phi_t\|\Psi_t)) \\ &\quad + \frac{\eta}{\mu} \sum_{t=1}^T \|\mathcal{F}(\Psi_t)^\top - \mathcal{F}(\Psi_{t-1})^\top\|_*^2. \end{aligned}$$

Via telescoping sums, the first term (I) can be reduced to

$$(I) = \frac{1}{\eta} \sum_{t=1}^T (\mathcal{BD}_h(Z\|\Phi_{t-1}) - \mathcal{BD}_h(Z\|\Phi_t)) = \frac{1}{\eta} (\mathcal{BD}_h(Z\|\Phi_0) - \mathcal{BD}_h(Z\|\Phi_T)). \quad (102)$$

Additionally, since h is μ -strongly convex with respect $\|\cdot\|$,

$$\mathcal{BD}_h(X\|Y) \geq \frac{\mu}{2}\|X - Y\|^2. \quad (103)$$

Therefore, we can re-express the error as

$$\begin{aligned} \text{Error}(\bar{\Psi}_T) &\leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \frac{1}{\eta} \mathcal{BD}_h(Z\|\Phi_0) - \underbrace{\frac{\mu}{2\eta}\|Z - \Phi_T\|^2 + \frac{\mu}{2\eta} \sum_{t=1}^T (\|\Phi_{t-1} - \Psi_t\|^2 + \|\Phi_t - \Psi_t\|^2)}_{(II)} \\ &\quad + \underbrace{\frac{\eta}{\mu} \sum_{t=1}^T \|\mathcal{F}(\Psi_t)^\top - \mathcal{F}(\Psi_{t-1})^\top\|_*^2}_{(III)}, \end{aligned}$$

and proceed to upper bound terms (II) and (III). We begin by expanding (II) as

$$\begin{aligned} (II) &= -\frac{\mu}{2\eta}\|\Phi_T - Z\|^2 - \frac{\mu}{2\eta} \sum_{t=1}^T (\|\Phi_{t-1} - \Psi_t\|^2 + \|\Phi_t - \Psi_t\|^2) \\ &= \frac{\mu}{2\eta} \left(\|\Phi_0 - \Psi_0\|^2 - \|\Phi_T - \Psi_T\|^2 - \|\Phi_T - Z\|^2 - \sum_{t=1}^T (\|\Phi_{t-1} - \Psi_t\|^2 + \|\Phi_{t-1} - \Psi_{t-1}\|^2) \right) \end{aligned}$$

Since we initialized $\Phi_0 = \Psi_0$,

$$\|\Phi_0 - \Psi_0\|^2 = 0. \quad (104)$$

Furthermore, since $\|\Phi_T - Z\|^2 \geq 0$ for all $Z \in \mathcal{C}$, including $Z = \Psi_T$,

$$(II) \leq -\frac{\mu}{2\eta} \sum_{t=1}^T (\|\Phi_{t-1} - \Psi_t\|^2 + \|\Phi_{t-1} - \Psi_{t-1}\|^2). \quad (105)$$

Defining

$$\Gamma := \sum_{t=1}^T (\|\Psi_t - \Phi_{t-1}\|^2 + \|\Psi_{t-1} - \Phi_{t-1}\|^2) \quad (106)$$

implies the following upper bound on (II)

$$(II) \leq -\frac{\mu}{2\eta}\Gamma. \quad (107)$$

In order to bound term (III), we will leverage an important property of finite-valued quantum zero-sum games. Namely, in [Appendix D](#), we prove several properties of the \mathcal{F} operator, including the following lemma regarding Lipschitz continuity.

Lemma 9 (Lipschitz Continuity of \mathcal{F}) *For finite-valued quantum zero-sum games, \mathcal{F} is a Lipschitz continuous operator, meaning there exists a Lipschitz constant $\gamma \in \mathbb{R}$ such that*

$$\|\mathcal{F}(Z) - \mathcal{F}(Z')\|_* \leq \gamma\|Z - Z'\|, \quad \forall Z, Z' \in \mathcal{C}. \quad (108)$$

For the case of $(\|\cdot\|_F, \|\cdot\|_F)$ and $(\|\cdot\|_\infty, \|\cdot\|_1)$, $\gamma_{F,F} = \mathcal{O}(2^d)$ while $\gamma_{\infty,1} = \mathcal{O}(1)$.

Leveraging the Lipschitz continuity of \mathcal{F} , as given in Equation (108), and the fact that for an arbitrary norm $\|\cdot\|$

$$\|X - Y\|^2 \leq (\|X\| + \|Y\|)^2 \leq \|X\|^2 + \|Y\|^2 + 2\|X\|\|Y\| \leq 2\|X\|^2 + 2\|Y\|^2, \quad (109)$$

(III) can be bounded as

$$\begin{aligned} (III) &= \frac{\eta}{\mu} \sum_{t=1}^T \|\mathcal{F}(\Psi_t) - \mathcal{F}(\Psi_{t-1})\|_*^2 \leq \frac{\eta}{\mu} \sum_{t=1}^T \gamma^2 \|\Psi_t - \Psi_{t-1}\|^2 \\ &= \frac{\gamma^2 \eta}{\mu} \sum_{t=1}^T \|\Psi_t - \Phi_{t-1} + \Phi_{t-1} - \Psi_{t-1}\|^2 \\ &= \frac{\gamma^2 \eta}{\mu} \sum_{t=1}^T \|(\Psi_t - \Phi_{t-1}) - (\Psi_{t-1} - \Phi_{t-1})\|^2 \\ &\leq \frac{2\gamma^2 \eta}{\mu} \left(\sum_{t=1}^T \|\Psi_t - \Phi_{t-1}\|^2 + \|\Psi_{t-1} - \Phi_{t-1}\|^2 \right) \\ &= \frac{2\gamma^2 \eta}{\mu} \Gamma. \end{aligned} \quad (110)$$

Therefore, combining the bound on (II), from Equation (107), with that of (III), from Equation (110), results in

$$(II) + (III) \leq \left(\frac{2\gamma^2 \eta}{\mu} - \frac{\mu}{2\eta} \right) \Gamma. \quad (111)$$

Note that there exist step-sizes η , such that $(II) + (III) \leq 0$. Namely,

$$\eta \leq \frac{\mu}{2\gamma} \implies (II) + (III) \leq 0. \quad (112)$$

Recall that our objective was to upper bound $\text{Error}(\bar{\Psi}_T)$, where

$$\text{Error}(\bar{\Psi}_T) \leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \frac{1}{\eta} \mathcal{BD}_h(\Phi_0 \| Z) + (II) + (III), \quad (113)$$

If we choose the step-size such that $\eta \leq \frac{\mu}{2\gamma}$, the upper bound simplifies to

$$\text{Error}(\bar{\Psi}_T) \leq \frac{1}{T} \sup_{Z \in \mathcal{C}} \frac{1}{\eta} \mathcal{BD}_h(\Phi_0 \| Z). \quad (114)$$

Since $\mathcal{BD}_h(\Phi_0 \| Z)$ is bounded above by the Bregman diameter of \mathcal{C} , i.e.

$$\mathcal{BD}_h(\Phi_0 \| Z) \leq \sup_{X, Y \in \mathcal{C}} \mathcal{BD}_h(X \| Y), \quad (115)$$

it has no time dependence, meaning

$$\text{Error}(\bar{\Psi}_T) = \mathcal{O}_d \left(\frac{1}{T} \right),$$

where \mathcal{O}_d is used to refer to the convergence dependence on the spectraplex dimension d , governed by the Bregman diameter of \mathcal{C} , as given in Equation (115). Therefore, setting $\text{Error}(\bar{\Psi}_T)$ to a desired accuracy parameter, i.e. $\text{Error}(\bar{\Psi}_T) = \epsilon$, we obtain that the number of steps necessary for average-iterate convergence to an ϵ -Nash equilibrium is

$$T = \mathcal{O}_d \left(\frac{1}{\epsilon} \right). \quad (116)$$

4.7 Regularizer-Induced Dimensionality Dependence

Following from this analysis, for a fixed dimension d , the family of algorithms encompassed by OMMP can compute an ϵ -Nash Equilibrium in $\mathcal{O}_d(1/\epsilon)$ steps. Note, however, that the $\mathcal{O}_d(\cdot)$ notation hides the dimension-dependence of the result, incurred by the $\frac{1}{\eta}\mathcal{BD}_h(\Phi_0\|Z)$ term. In order to achieve an explicit bound, we will demonstrate how different OMMP regularizers affect the Bregman divergence $\mathcal{BD}_h(\Phi_0\|Z)$ and step-size η . In particular, we consider two regularizers—the Frobenius norm and the von Neumann entropy—which are quantum analogs of two of the most widely used classical regularizers (entropy and Euclidean norm). Note that OMMP instantiated with the Frobenius norm implements the *Optimistic Matrix Extra-Gradient (OMEG) algorithm* (updates given in Figure 2), while the von Neumann entropy regularizer implements the *Optimistic Matrix Multiplicative Weights Update (OMMWU) algorithm*.

We will first consider the Bregman divergence term, noting that in OMMP we initialize $\Phi_0 = \left(\frac{1}{2^n}\mathbb{1}_A, \frac{1}{2^m}\mathbb{1}_B\right)$. We offer an explicit proof for the Frobenius norm regularizer:

Proposition 5 *The Frobenius norm regularizer induces a maximum Bregman divergence that is constant, i.e. $\max_{X \in \mathcal{C}} \mathcal{BD}_h(\Phi_0\|X) = 2 = \mathcal{O}(1)$.*

Proof. Using the derivation of the Bregman divergence for the Frobenius norm in Equation (52) and the fact that $\forall X \in \mathcal{C}, \text{Tr}(X) = \text{Tr}(\alpha) + \text{Tr}(\beta) \leq 1 + 1 = 2$,

$$\begin{aligned} \max_{X \in \mathcal{C}} \mathcal{BD}_h(\Phi_0\|X) &= \max_{X \in \mathcal{C}} \frac{1}{2} \|\Phi_0 - X\|_F^2 = \frac{1}{2} \max_{X \in \mathcal{C}} \left(\|\Phi_0\|_F^2 - 2\text{Tr}(\Phi_0^\dagger X) + \|X\|_F^2 \right) \\ &\leq \frac{1}{2} \max_{X \in \mathcal{C}} \left(2 + \|X\|_F^2 \right) \leq \frac{1}{2} (2 + 2) = 2. \end{aligned}$$

□

In the case of the von Neumann entropy regularizer, as discussed in [71], for a $(4^d - 1)$ -dimensional spectraplex, the corresponding bound on the Bregman divergence is equal to the maximum quantum entropy for a $2d$ -qubit system, and hence linear in d —more precisely, equal to $\log(4^d - 1) = \mathcal{O}(d)$.

We now consider the step-size. Following from Equation (112), the optimal step-size is dependent on h 's strong convexity parameter (μ_h) and \mathcal{F} 's Lipschitz constant ($\gamma_{\mathcal{F}}$), as

$$\eta = \Theta(\mu_h/\gamma_{\mathcal{F}}). \quad (117)$$

In scenarios utilizing the Frobenius regularizer, the modulus of strong convexity is 1, while γ_{Lipsch} , the Lipschitz constant with respect to $\|\cdot\|_F$ of the gradient for games with bounded utilities in $[-1, 1]$, is of the order of $\mathcal{O}(\sqrt{4^d}) = \mathcal{O}(2^d)$. In contrast, for the von Neumann entropy, the modulus of strong convexity is 1 with respect to the Shatten 1-norm (the matrix analog of the ℓ_1 -norm) [72]. Moreover, for games with bounded utilities, the relative Lipschitz constant with respect to Shatten 1-norm $\|\cdot\|_1$ is of the order of $\mathcal{O}(1)$.

In light of the preceding discussions, we conclude with the dimension-dependent iteration complexities of OMMWU and OMEG for achieving ϵ -Nash equilibria of quantum zero-sum games:

Corollary 1 (OMMWU Iteration Complexity) *In a 4^d -dimensional spectraplex, OMMWU computes average-iterate ϵ -Nash equilibria in $\mathcal{O}(d/\epsilon)$ iterations.*

Corollary 2 (OMEG Iteration Complexity) *In the 4^d -dimensional spectraplex, OMEG computes average-iterate ϵ -Nash equilibria in $\mathcal{O}(2^d/\epsilon)$ iterations.*

Not only does OMMWU match the dimension-dependence of MMWU, but it also logarithmically outperforms the dimension dependence of OMEG. Thus, between the Frobenius norm (OMEG) and von Neumann entropy (OMMWU), the latter is the clear choice for efficient implementation of the OMMP method for quantum zero-sum games played in the spectraplex.

5 Conclusions

In this work, we proposed a “design” of classical algorithms for finding ϵ -Nash equilibria of finite-valued quantum zero-sum games, served to unify and clarify the landscape of quantum zero-sum game algorithms. We demonstrated that the best prior method—Jain and Watrous’ Matrix Multiplicative Weight Updates (MMWU) algorithm [29]—fits within the framework as an instantiation of our proposed Matrix Dual Averaging (MDA) method. To improve upon MMWU, we leveraged a gradient-based view, which enabled us to draw from the large classical game theory and optimization literature for solving variational inequalities of monotone, Lipschitz operators. Via this gradient-based perspective, proximal steps, and “optimism” we proposed and proved convergence guarantees of the Optimistic Matrix Mirror Prox (OMMP) method. When instantiated with the von Neumann entropy regularizer, OMMP implements the Optimistic Matrix Multiplicative Weights Update (OMMWU) algorithm, which we proved to have dimensionality dependence equivalent to that of MMWU. However, OMMWU establishes a new benchmark in convergence rate, achieving a quadratic speed-up relative MMWU, while still requiring only one gradient call per iteration.

With these contributions several challenges remain and interesting open problems are introduced:

1. **Tightness of Convergence Rates:** Our first conjecture posits that the $O(1/\epsilon)$ convergence rate of 1-MMP is tight, but empirical or theoretical confirmation remains elusive. Establishing the tightness would clarify the limitations of our method.

Conjecture 1 *The rates of [Theorem 1](#) are tight: specifically, there exists a quantum two-player, zero-sum game for which OMMP methods require $\Omega_d(1/\epsilon)$ iterations to compute ϵ -Nash equilibria.*

2. **Dimensionality of Quantum Games:** Our work focuses on finite quantum zero-sum games. Extending the applicability of our algorithms to games in infinite dimensional Hilbert spaces remains an open challenge that could have theoretical applications to operator analysis and practical implications to derive new optimal transport techniques.
3. **Robustness to Noise:** Quantum systems are inherently prone to noise. Important for practical implementations is how our method performs with different kind of feedback under noisy measurement conditions.
4. **Computational Scaling Challenges:** Our proposal requires knowledge of the entire system density matrix in each iteration, which is computationally prohibitive for large system sizes. Therefore, it would be valuable to explore efficient classical representations of the system state or propose entirely quantum protocols/algorithms which can eliminate the need to classically learn/store the system state.

5. **Going Beyond 2-Player Zero-Sum:** Early forays into quantum potential games [73] and quantum zero-sum polymatrix games [45] are suggestive of exciting possibilities for further research with novel classes of games and solution concepts.

6 Acknowledgements

Part of this work was done while the authors were visiting the Simons Institute for the Theory of Computing and the Institute for Mathematical Sciences at the National University of Singapore. We thank Rahul Jain for useful discussions. F.V. is supported by NSF grant DGE-2146752 and the Vannevar Bush Faculty Fellowship Program grant N00014-21-1-2941. E.V. is grateful for financial support by the Post-Doctoral FODSI-Simons Fellowship. P.M. is a member of the Archimedes Unit—Athena RC—Department of Mathematics, National & Kapodistrian University of Athens, and gratefully acknowledges financial support by the French National Research Agency (ANR) in the framework of the “Investissements d’avenir” program (ANR-15- IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), MIAI@Grenoble Alpes (ANR-19- P3IA-0003), and project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program. G.P. would like to acknowledge National Research Foundation, Singapore and DSO National Laboratories under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016), NRF 2018 Fellowship NRF-NRFF2018- 07, NRF2019-NRF-ANR095 ALIAS grant, grant PIESGP-AI-2020-01, AME Programmatic Fund (Grant No.A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and Provost’s Chair Professorship grant RGEPPV2101. M.J. is partially supported by NSF Grant IIS-1901252 and Vannevar Bush Faculty Fellowship program under grant number N00014-21-1-2941.

References

- [1] John Von Neumann. “Zur Theorie der Gesellschaftsspiele”. *Mathematische Annalen* **100**, 295–320 (1928).
- [2] John Nash. “Non-Cooperative Games”. *Annals of Mathematics* **54**, 286–295 (1951). url: <http://www.jstor.org/stable/1969529>.
- [3] Parvathy S. Pillai and Shrisha Rao. “Resource Allocation in Cloud Computing Using the Uncertainty Principle of Game Theory”. *IEEE Systems Journal* **10**, 637–648 (2016).
- [4] Steven J. Brams. “Game theory and politics”. Free Press. (1975). url: <https://cir.nii.ac.jp/crid/1130000793730346368>.
- [5] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. “Mastering the game of Go with deep neural networks and tree search”. *Nature* **529**, 484–489 (2016).
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. “Generative Adversarial Nets”. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*. Volume 27. Curran Associates, Inc. (2014). url: https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
- [7] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. “Robust Adversarial Reinforcement Learning”. In *International Conference on Machine Learning*. Pages 2817–2826. PMLR (2017).
- [8] J. S. Bell. “On the Einstein Podolsky Rosen paradox”. *Physics Physique Fizika* **1**, 195–200 (1964).
- [9] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. “Proposed Experiment to Test Local Hidden-Variable Theories”. *Phys. Rev. Lett.* **23**, 880–884 (1969).
- [10] N. David Mermin. “Simple unified form for the major no-hidden-variables theorems”. *Phys. Rev. Lett.* **65**, 3373–3376 (1990).
- [11] Lucien Hardy. “Nonlocality for two particles without inequalities for almost all entangled states”. *Phys. Rev. Lett.* **71**, 1665–1668 (1993).
- [12] P.K. Aravind. “Bell’s theorem without inequalities and only two distant observers”. In *Optical Fiber Communication Conference and International Conference on Quantum Information*. Page EAPA7. Optica Publishing Group (2001).
- [13] Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. “Going beyond bell’s theorem” (2007). [arXiv:0712.0921](https://arxiv.org/abs/0712.0921).
- [14] Ben W. Reichardt, Falk Unger, and Umesh Vazirani. “Classical command of quantum systems”. *Nature* **496**, 456–460 (2013).
- [15] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. “MIP* = RE”. *Commun. ACM* **64**, 131–138 (2021).

- [16] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. “Quantum Advantage from Any Non-local Game”. In Barna Saha and Rocco A. Servedio, editors, Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023. Pages 1617–1628. ACM (2023).
- [17] G. Gutoski. “Upper bounds for quantum interactive proofs with competing provers”. In 20th Annual IEEE Conference on Computational Complexity (CCC’05). Pages 334–343. (2005).
- [18] Gus Gutoski and John Watrous. “Quantum interactive proofs with competing provers”. In Volker Diekert and Bruno Durand, editors, STACS 2005. Pages 605–616. Berlin, Heidelberg (2005). Springer Berlin Heidelberg.
- [19] Gus Gutoski and John Watrous. “Toward a general theory of quantum games”. In Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing. Page 565–574. STOC ’07New York, NY, USA (2007). Association for Computing Machinery.
- [20] Hirotada Kobayashi and Keiji Matsumoto. “Quantum multi-prover interactive proof systems with limited prior entanglement”. *Journal of Computer and System Sciences* **66**, 429–450 (2003).
- [21] R. Cleve, P. Hoyer, B. Toner, and J. Watrous. “Consequences and limits of non-local strategies”. In Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004. Pages 236–249. (2004).
- [22] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. “Using Entanglement in Quantum Multi-Prover Interactive Proofs”. *computational complexity* **18**, 273–307 (2009).
- [23] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. “Entangled Games Are Hard to Approximate”. *SIAM Journal on Computing* **40**, 848–877 (2011).
- [24] Andris Ambainis. “A New Protocol and Lower Bounds for Quantum Coin Flipping”. In Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing. Page 134–142. STOC ’01New York, NY, USA (2001). Association for Computing Machinery.
- [25] R. W. Spekkens and Terry Rudolph. “Quantum Protocol for Cheat-Sensitive Weak Coin Flipping”. *Phys. Rev. Lett.* **89**, 227901 (2002).
- [26] Carlos Mochon. “Quantum weak coin flipping with arbitrarily small bias” (2007). [arXiv:0711.4114](https://arxiv.org/abs/0711.4114).
- [27] Carl A. Miller. “The Impossibility of Efficient Quantum Weak Coin Flipping”. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. Page 916–929. STOC 2020New York, NY, USA (2020). Association for Computing Machinery.
- [28] David A. Meyer. “Quantum Strategies”. *Phys. Rev. Lett.* **82**, 1052–1055 (1999).
- [29] Rahul Jain and John Watrous. “Parallel Approximation of Non-interactive Zero-sum Quantum Games”. In 2009 24th Annual IEEE Conference on Computational Complexity. Pages 243–253. (2009).
- [30] Pierre-Luc Dallaire-Demers and Nathan Killoran. “Quantum generative adversarial networks”. *Phys. Rev. A* **98**, 012324 (2018).

- [31] Shouvanik Chakrabarti, Yiming Huang, Tongyang Li, Soheil Feizi, and Xiaodi Wu. “Quantum Wasserstein Generative Adversarial Networks”. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. Pages 6778–6789. (2019). url: <https://proceedings.neurips.cc/paper/2019/hash/f35fd567065af297ae65b621e0a21ae9-Abstract.html>.
- [32] John Bostanci and John Watrous. “Quantum game theory and the complexity of approximating quantum Nash equilibria”. *Quantum* **6**, 882 (2022).
- [33] Kyriakos Lotidis, Panayotis Mertikopoulos, Nicholas Bambos, and Jose Blanchet. “Payoff-based learning with matrix multiplicative weights in quantum games”. In *NeurIPS ’23: Proceedings of the 37th International Conference on Neural Information Processing Systems*. (2023).
- [34] Koji Tsuda, Gunnar Rätsch, and Manfred K. Warmuth. “Matrix Exponentiated Gradient Updates for On-line Learning and Bregman Projection”. *Journal of Machine Learning Research* **6**, 995–1018 (2005). url: <http://jmlr.org/papers/v6/tsuda05a.html>.
- [35] Arkadi Semen Nemirovski and David Berkovich Yudin. “Problem complexity and method efficiency in optimization”. Wiley. New York, NY (1983).
- [36] Amir Beck and Marc Teboulle. “Mirror descent and nonlinear projected subgradient methods for convex optimization”. *Operations Research Letters* **31**, 167–175 (2003).
- [37] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. “QIP = PSPACE”. *J. ACM* **58**, 30:1–30:27 (2011).
- [38] Sanjeev Arora and Satyen Kale. “A Combinatorial, Primal-Dual Approach to Semidefinite Programs”. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*. Page 227–236. STOC ’07New York, NY, USA (2007). Association for Computing Machinery.
- [39] Li Chen, Rasmus Kyng, Maximilian Probst Gutenberg, and Sushant Sachdeva. “A Simple Framework for Finding Balanced Sparse Cuts via APSP”. In *Symposium on Simplicity in Algorithms (SOSA)*. Pages 42–55. SIAM (2023).
- [40] Zeyuan Allen-Zhu, Zhenyu Liao, and Lorenzo Orecchia. “Spectral Sparsification and Regret Minimization Beyond Matrix Multiplicative Updates”. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. Page 237–245. STOC ’15New York, NY, USA (2015). Association for Computing Machinery.
- [41] Panayotis Mertikopoulos, E. Veronica Belmega, and Aris L. Moustakas. “Matrix exponential learning: Distributed optimization in MIMO systems”. In *2012 IEEE International Symposium on Information Theory Proceedings*. Pages 3028–3032. (2012).
- [42] Panayotis Mertikopoulos and Aris L. Moustakas. “Learning in an Uncertain World: MIMO Covariance Matrix Optimization With Imperfect Feedback”. *IEEE Transactions on Signal Processing* **64**, 5–18 (2016).
- [43] Olivier Bilenne, Panayotis Mertikopoulos, and Elena Veronica Belmega. “Fast Optimization With Zeroth-Order Feedback in Distributed, Multi-User MIMO Systems”. *IEEE Transactions on Signal Processing* **68**, 6085–6100 (2020).
- [44] Sham M Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. “Regularization techniques for learning with matrices”. *The Journal of Machine Learning Research* **13**, 1865–1890 (2012).

- [45] Wayne Lin, Georgios Piliouras, Ryann Sim, and Antonios Varvitsiotis. “No-regret learning and equilibrium computation in quantum games” (2023). [arXiv:2310.08473](#).
- [46] Sanjeev Arora, Elad Hazan, and Satyen Kale. “The Multiplicative Weights Update Method: a Meta-Algorithm and Applications”. *Theory of Computing* **8**, 121–164 (2012).
- [47] Arkadi Nemirovski. “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**, 229–251 (2004).
- [48] Alfred Auslender and Marc Teboulle. “Interior projection-like methods for monotone variational inequalities”. *Mathematical Programming* **104**, 39–68 (2005).
- [49] Yuyuan Ouyang and Yangyang Xu. “Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems”. *Mathematical Programming* **185**, 1–35 (2021).
- [50] Galina M Korpelevich. “The extragradient method for finding saddle points and other problems”. *Matecon* **12**, 747–756 (1976). url: <https://cir.nii.ac.jp/crid/1571698600143951616>.
- [51] Boris Teodorovich Polyak and Anatoli Juditsky. “Acceleration of stochastic approximation by averaging”. *SIAM Journal on Control and Optimization* **30**, 838–855 (1992).
- [52] Leonid Denisovich Popov. “A modification of the Arrow-Hurwicz method for search of saddle points”. *Mathematical notes of the Academy of Sciences of the USSR* **28**, 845–848 (1980).
- [53] Sasha Rakhlin and Karthik Sridharan. “Optimization, learning, and games with predictable sequences”. *Advances in Neural Information Processing Systems* **26** (2013).
- [54] Kyriakos Lotidis, Panayotis Mertikopoulos, and Nicholas Bambos. “Learning in quantum games” (2023). [arXiv:2302.02333](#).
- [55] Ilayda Canyakmaz, Wayne Lin, Georgios Piliouras, and Antonios Varvitsiotis. “Multiplicative updates for online convex optimization over symmetric cones” (2023).
- [56] Alina Ene and Huy Lê Nguyễn. “Adaptive and Universal Algorithms for Variational Inequalities with Optimal Convergence”. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**, 6559–6567 (2022).
- [57] Gerard Debreu. “A Social Equilibrium Existence Theorem”. *Proceedings of the National Academy of Sciences* **38**, 886–893 (1952).
- [58] Gesualdo Scutari, Daniel P. Palomar, Francisco Facchinei, and Jong-shi Pang. “Convex Optimization, Game Theory, and Variational Inequality Theory”. *IEEE Signal Processing Magazine* **27**, 35–49 (2010).
- [59] Stephen Boyd and Lieven Vandenbergh. “Convex Optimization”. *Cambridge University Press*. (2004).
- [60] Shai Shalev-Shwartz. “Online Learning and Online Convex Optimization”. *Foundations and Trends in Machine Learning* **4**, 107–194 (2012).
- [61] John Watrous. “Theory of quantum information”. 2.3 Norms of operators, lecture notes, University of Waterloo (2011).
- [62] Carl D. Meyer and Ian Stewart. “Matrix Analysis and Applied Linear Algebra, Second Edition”. *Society for Industrial and Applied Mathematics*. Philadelphia, PA (2023).

- [63] Xingyu Zhou. “On the Fenchel Duality between Strong Convexity and Lipschitz Continuous Gradient” (2018). [arXiv:1803.06573](https://arxiv.org/abs/1803.06573).
- [64] D. Bertsekas, A. Nedic, and A. Ozdaglar. “Convex Analysis and Optimization”. Athena Scientific optimization and computation series. Athena Scientific. (2003).
- [65] Constantin Zalinescu. “Convex analysis in general vector spaces”. [World scientific](https://www.worldscientific.com/). (2002).
- [66] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. “Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile”. In International Conference on Learning Representations. (2019). url: <https://openreview.net/forum?id=Bkg8jjc9KQ>.
- [67] Sébastien Bubeck. “Convex optimization: Algorithms and complexity”. [Found. Trends Mach. Learn.](https://www.foundtrends.com/) **8**, 231–357 (2015).
- [68] Yurii Nesterov. “Primal-dual subgradient methods for convex problems”. [Mathematical Programming](https://www.optimization-online.org/DB_HTML/2009/01/2211.html) **120**, 221–259 (2009).
- [69] Andersen Ang. “Projected gradient algorithm”. https://angms.science/doc/CVX/CVX_PGD.pdf (2023). Version: July 13, 2023. First draft: August 2, 2017.
- [70] Satyen Kale. “Efficient algorithms using the Multiplicative Weights Update method”. PhD thesis. Princeton University. (2007). url: <https://www.proquest.com/dissertations-theses/efficient-algorithms-using-multiplicative-weights/docview/304824121/se-2>.
- [71] Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. “Solving Variational Inequalities with Stochastic Mirror-Prox Algorithm”. [Stochastic Systems](https://www.stochasticsystems.com/) **1**, 17–58 (2011).
- [72] Yao-Liang Yu. “The strong convexity of von neumann’s entropy”. Lecture Notes for MSE 213: Optimization Theory (2020). Available online: https://web.stanford.edu/~sidford/courses/20fa_opt_theory/sidford_mse213_2020fa_chap_5_extensions.pdf.
- [73] Wayne Lin, Georgios Piliouras, Ryann Sim, and Antonios Varvitsiotis. “Quantum Potential Games, Replicator Dynamics, and the Separability Problem” (2023). [arXiv:2302.04789](https://arxiv.org/abs/2302.04789).
- [74] D. Petz. “Bregman divergence as relative operator entropy”. [Acta Mathematica Hungarica](https://www.ams.org/journals/acta-math-hungarica) **116**, 127–131 (2007).
- [75] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. “Fundamentals of convex analysis”. [Pages X](https://www.springer.com/), 259. Grundlehren Text Editions. Springer Berlin, Heidelberg. (2001). 1 edition.

A Complex Matrix Differentiation and Gradients

In this appendix section, we review complex matrix differentiation and gradients - an invaluable tool for framing our quantum game feedback and analyzing our algorithm. Specifically, we will show why it is natural to define the gradient of a real-valued function $u : \mathbb{C}^d \rightarrow \mathbb{R}$ with respect to complex matrix $Z \in \mathbb{C}^d$ as

$$\nabla_{Z^\top} u(Z) \quad (118)$$

and how this relates to a directional derivative. Concretely, our goal is to compute the directional derivative of $u(Z(t))$ at Z in the direction of matrix dZ/dt .

In order to derive this derivative, we will start by introducing a few important properties of Z and u . To begin, note that Z is a Hermitian matrix of the form

$$Z = X + iY, \quad (119)$$

where X is a real symmetric matrix and Y is a real skew-symmetric matrix. Thus, can use the chain rule to derive du/dt :

$$\frac{du}{dt} = \sum_{i,j} \frac{\partial u}{\partial X_{ij}} \frac{dX_{ij}}{dt} + \frac{\partial u}{\partial Y_{ij}} \frac{dY_{ij}}{dt}. \quad (120)$$

From the definition of the Wirtinger derivative, we also have that

$$\frac{\partial u}{\partial Z_{ij}} = \frac{\partial u}{\partial X_{ij}} - i \frac{\partial u}{\partial Y_{ij}} \quad (121)$$

Finally, we note the following property of symmetric and skew-symmetric matrices:

Proposition 6 *For symmetric matrix A and skew-symmetric matrix B ,*

$$\text{Tr}(AB) = \text{Tr}(A^\top B) = \sum_{i,j} A_{ij} B_{ij} = 0. \quad (122)$$

Proof. $\text{Tr}(AB) = \text{Tr}(A^\top B) = \text{Tr}(AB^\top) = -\text{Tr}(AB) \implies \text{Tr}(AB) = 0 \quad \square$ Since

X is symmetric, $\partial u/\partial X$ and dX/dt are symmetric. Similarly, since Y is skew-symmetric, $\partial u/\partial Y$ and dY/dt are symmetric. Therefore, by [Proposition 6](#),

$$\sum_{i,j} \frac{\partial u}{\partial X_{ij}} \frac{dY_{ij}}{dt} = \text{Tr} \left(\frac{\partial u}{\partial X} \frac{dY}{dt} \right) = 0 \quad (123)$$

$$\sum_{i,j} \frac{\partial u}{\partial Y_{ij}} \frac{dX_{ij}}{dt} = \text{Tr} \left(\frac{\partial u}{\partial Y} \frac{dX}{dt} \right) = 0 \quad (124)$$

Now, we can combine the previous properties to obtain du/dt expressed in terms of the complex matrix Z . Starting from [Equation \(120\)](#), then leveraging [Equation \(123\)](#) and [Equation \(124\)](#), we get that

$$\frac{du}{dt} = \sum_{i,j} \frac{\partial u}{\partial X_{ij}} \frac{dX_{ij}}{dt} + \frac{\partial u}{\partial Y_{ij}} \frac{dY_{ij}}{dt} \quad (125)$$

$$= \sum_{i,j} \left(\frac{\partial u}{\partial X_{ij}} \frac{dX_{ij}}{dt} + \frac{\partial u}{\partial Y_{ij}} \frac{dY_{ij}}{dt} \right) + i \text{Tr} \left(\frac{\partial u}{\partial X} \frac{dY}{dt} \right) - i \text{Tr} \left(\frac{\partial u}{\partial Y} \frac{dX}{dt} \right) \quad (126)$$

$$= \sum_{i,j} \frac{\partial u}{\partial X_{ij}} \frac{dX_{ij}}{dt} + i \frac{\partial u}{\partial X_{ij}} \frac{dY_{ij}}{dt} - i \frac{\partial u}{\partial Y_{ij}} \frac{dX_{ij}}{dt} + \frac{\partial u}{\partial Y_{ij}} \frac{dY_{ij}}{dt} \quad (127)$$

$$= \sum_{i,j} \left(\frac{\partial u}{\partial X_{ij}} - i \frac{\partial u}{\partial Y_{ij}} \right) \left(\frac{dX_{ij}}{dt} + i \frac{dY_{ij}}{dt} \right). \quad (128)$$

By Equation (119) and Equation (121), this achieves the desired directional derivative, which is expressed in terms of the gradient $\nabla_{Z^\top} u$:

$$\frac{du}{dt} = \sum_{i,j} \left(\frac{\partial u}{\partial X_{ij}} - i \frac{\partial u}{\partial Y_{ij}} \right) \left(\frac{dX_{ij}}{dt} + i \frac{dY_{ij}}{dt} \right) = \sum_{i,j} \frac{\partial u}{\partial Z_{ij}} \frac{dZ_{ij}}{dt} \quad (129)$$

$$= \text{Tr} \left(\left(\frac{\partial u}{\partial Z} \right)^\top \frac{dZ}{dt} \right) = \text{Tr} \left(\frac{\partial u}{\partial Z^\top} \frac{dZ}{dt} \right) = \text{Tr} \left(\nabla_{Z^\top} u \cdot \frac{dZ}{dt} \right). \quad (130)$$

B Mirror & Proximal Maps

In this appendix we offer further intuition for definitions and provide intuitive proofs of key lemmas regarding the mirror and proximal maps/steps introduced in Section 2.5.

Definition 6 (Mirror Map) For a state D in the dual space \mathcal{D} , the mirror map, or regularized best response,

$$\text{Mir}\Pi_{\mathcal{P}}^h(D) := \arg \max_{P \in \mathcal{P}} \{ \text{Tr}[D^\dagger P] - h(P) \} \quad (43)$$

implements the inverse mirror map, $\partial h^{-1} : \mathcal{D} \rightarrow \mathcal{P}$, mapping $D \in \mathcal{D}$ to the corresponding state P in the feasible region \mathcal{P} of the primal space.

Proof. Since \mathcal{H}^d is the set of all complex matrices, this is an unconstrained minimization problem. We can thus solve for the maximum by setting the gradient of the minimization term to zero, as

$$0 = \nabla_{P^\top} \left(\text{Tr}[D^\dagger P] - h(P) \right) = D - \partial h(P), \quad (131)$$

which implies the solution

$$\partial h(P) = D \implies P = \partial^{-1} h(D). \quad (132)$$

□

Lemma 4 (Mirror Step) For state X_t in the feasible set \mathcal{H}_+^d of primal space \mathbb{C}^d ,

$$\text{MirProx}\Pi_{\mathcal{X}_+^d}^{h,\eta,f}(X_t) := \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(X_t)^\dagger X_{t+1}] + \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1} || X_t) \right\}, \quad (45)$$

implements the mirror step mapping to an updated feasible set primal state $X_{t+1} \in \mathcal{H}_+^d$.

Proof. The constrained minimization problem has Lagrangian

$$\mathcal{L} = \text{Tr}[\nabla f(X_t)^\dagger X_{t+1}] - \frac{1}{\eta} \left(h(X_{t+1}) - h(X_t) - \text{Tr}[\nabla h(X_t)^\dagger (X_{t+1} - X_t)] \right) - \lambda (\text{Tr}[X_{t+1}] - 1), \quad (133)$$

with gradients

$$\begin{cases} \nabla_{X_{t+1}^\top} \mathcal{L} &= \nabla f(X_t) - \frac{1}{\eta} (\nabla h(X_{t+1}) - \nabla h(X_t)) - \lambda \mathbb{1} \\ \nabla_\lambda \mathcal{L} &= \text{Tr}[X_{t+1}] - 1. \end{cases} \quad (134)$$

Setting both of these gradients to zero gives us the equations

$$\begin{cases} \partial h(X_{t+1}) &= \partial h(X_t) + \eta \nabla f(X_t) - \eta \lambda \mathbb{1} \\ \text{Tr}[X_{t+1}] &= 1, \end{cases} \quad (135)$$

which implies the solution

$$X_{t+1} = \partial^{-1} h(\partial h(X_t) + \eta \nabla f(X_t) - \eta \lambda \mathbb{1}), \quad (136)$$

with λ chosen to ensure that X_t lies in the feasible set \mathcal{X}^d ; i.e., it satisfies the constraint

$$\text{Tr}[X_{t+1}] = \text{Tr}[\partial^{-1} h(\partial h(X_t) + \eta \nabla f(X_t) - \eta \lambda \mathbb{1})] = 1. \quad (137)$$

□

Lemma 5 (Proximal Step) For state $X_t \in \mathcal{H}_+^d$ in the feasible set of the primal space,

$$\begin{aligned} \text{Prox}\Pi_{\mathcal{H}_+^d}^{h,\eta}(X_t, \nabla f(Y_{t+1})) &:= \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(Y_{t+1})^\dagger (X_{t+1} - X_t)] + \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1} \| X_t) \right\} \\ &= \arg \min_{X_{t+1} \in \mathcal{H}_+^d} \left\{ \text{Tr}[\nabla f(Y_{t+1})^\dagger X_{t+1}] + \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1} \| X_t) \right\} \end{aligned} \quad (47)$$

implements the proximal step mapping to an updated feasible set primal state $X_{t+1} \in \mathcal{H}_+^d$.

Proof. The constrained minimization problem has Lagrangian

$$\mathcal{L} = \text{Tr}[\nabla f(Y_{t+1})^\dagger (X_{t+1} - X_t)] - \frac{1}{\eta} \mathcal{B}\mathcal{D}_h(X_{t+1} \| X) - \lambda (\text{Tr}[X_{t+1}] - 1), \quad (138)$$

where

$$\mathcal{B}\mathcal{D}_h(X_{t+1} \| X) = h(X_{t+1}) - h(X_t) - \text{Tr}[\partial h(X_t)^\dagger (X_{t+1} - X_t)]. \quad (139)$$

The Lagrangian has gradients

$$\begin{cases} \nabla_{X_{t+1}^\top} \mathcal{L} &= \nabla f(Y_{t+1}) - \frac{1}{\eta} (\partial h(X_{t+1}) - \partial h(X_t)) - \lambda \mathbb{1} \\ \nabla_\lambda \mathcal{L} &= \text{Tr}[X_{t+1}] - 1. \end{cases} \quad (140)$$

Setting both of these gradients to zero gives us the equations

$$\begin{cases} \partial h(X_{t+1}) &= \partial h(X_t) + \eta \nabla f(Y_{t+1}) - \eta \lambda \mathbb{1} \\ \text{Tr}[X_{t+1}] &= 1, \end{cases} \quad (141)$$

which implies the solution

$$X_{t+1} = \partial^{-1} h(\partial h(X_t) + \eta \nabla f(Y_{t+1}) - \eta \lambda \mathbb{1}), \quad (142)$$

with λ chosen to ensure that X_t lies in the feasible set \mathcal{X}^d , i.e., it satisfies the constraint

$$\text{Tr}[X_t] = \text{Tr}[\partial^{-1} h(\partial h(X_t) + \eta \nabla f(Y_{t+1}) - \eta \lambda \mathbb{1})] = 1. \quad (143)$$

□

Lemma 6 (Bregman Divergence of the Frobenius Norm) *If h is the squared Frobenius norm,*

$$h(X) = \frac{1}{2}\|X\|_F^2 = \frac{1}{2}\text{Tr}[X^\dagger X], \quad (51)$$

then the Bregman divergence is the squared Frobenius distance,

$$\mathcal{BD}_h(X\|Y) = \frac{1}{2}\|X - Y\|_F^2, \quad (52)$$

the corresponding mirror map (regularized best-response) is the orthogonal projection,

$$\text{Mir}\Pi_C^h(Y) = \arg \min_{C \in \mathcal{C}} \|Y - C\|_F^2 = \text{Orth}\Pi_C(Y), \quad (53)$$

and the corresponding proximal map is

$$\text{Prox}\Pi_C^h(X, Y) = \arg \min_{C \in \mathcal{C}} \|X + Y - C\|_F^2 = \text{Orth}\Pi_C(X + Y). \quad (54)$$

Proof. We begin by noting that

$$\nabla h(Y) = \frac{1}{2}\nabla\|Y\|_F^2 = Y, \quad (144)$$

which implies that the mirror map ∇h is the identity in this case. Therefore,

$$\nabla^{-1}h(Y) = \nabla h(Y) = Y. \quad (145)$$

Furthermore, since

$$\|X - Y\|_F^2 = \|X\|_F^2 + \|Y\|_F^2 - 2\text{Tr}[X^\dagger Y] \quad (146)$$

$$= \|X\|_F^2 + (2\|Y\|_F^2 - \|Y\|_F^2) - 2\text{Tr}[X^\dagger Y] \quad (147)$$

$$= \|X\|_F^2 - \|Y\|_F^2 + 2\text{Tr}[Y^\dagger Y] - 2\text{Tr}[X^\dagger Y] \quad (148)$$

$$= \|X\|_F^2 - \|Y\|_F^2 - 2\text{Tr}[(X - Y)^\dagger Y], \quad (149)$$

the Bregman divergence, as desired, is

$$\mathcal{BD}_h(X\|Y) = h(X) - h(Y) - \text{Tr}[(X - Y)^\dagger \nabla h(Y)] \quad (150)$$

$$= \frac{1}{2}\|X\|_F^2 - \frac{1}{2}\|Y\|_F^2 - \text{Tr}[(X - Y)^\dagger \nabla h(Y)] \quad (151)$$

$$= \frac{1}{2}\|X\|_F^2 - \frac{1}{2}\|Y\|_F^2 - \text{Tr}[(X - Y)^\dagger Y] \quad (152)$$

$$= \frac{1}{2}(\|X\|_F^2 - \|Y\|_F^2 - 2\text{Tr}[(X - Y)^\dagger Y]) \quad (153)$$

$$= \frac{1}{2}\|X - Y\|_F^2. \quad (154)$$

The corresponding regularized best-response is

$$\text{Mir}\Pi_C^h(Y) = \arg \min_{C \in \mathcal{C}} \left\{ \text{Tr}[Y^\dagger C] - \frac{1}{2}\|C\|_F^2 \right\},$$

which has Lagrangian

$$\mathcal{L} = \text{Tr}[Y^\dagger C] - \frac{1}{2}\|C\|_F^2 - \lambda(\text{Tr}[C] - 1), \quad (155)$$

with gradients

$$\nabla \mathcal{L} = \begin{cases} \nabla_{C^\top} \mathcal{L} & = Y - C - \lambda \mathbf{1} \\ \nabla_\lambda \mathcal{L} & = -\text{Tr}[C] + 1. \end{cases} \quad (156)$$

Note, however, that these are also gradients of the Lagrangian

$$\mathcal{L}' = \frac{1}{2} \|Y - C\|_F^2 - \lambda(\text{Tr}[C] - 1), \quad (157)$$

corresponding to the minimization of the orthogonal projection

$$\text{Orth}\Pi_C(Y) = \arg \min_{C \in \mathcal{C}} \|Y - C\|_F^2.$$

Therefore, as desired,

$$\text{Mir}\Pi_C^h(Y) = \arg \min_{C \in \mathcal{C}} \|Y - C\|_F^2 = \text{Orth}\Pi_C(Y).$$

Finally, the corresponding proximal projection is given by

$$\text{Prox}\Pi_C^h(X, Y) = \arg \min_{C \in \mathcal{C}} \left\{ \text{Tr}[Y^\dagger(X - C)] + \frac{1}{2} \|C - X\|_F^2 \right\}, \quad (158)$$

which has Lagrangian

$$\mathcal{L} = \text{Tr}[Y^\dagger(X - C)] + \frac{1}{2} \|C - X\|_F^2 - \lambda(\text{Tr}[C] - 1), \quad (159)$$

with gradients

$$\nabla \mathcal{L} = \begin{cases} \nabla_{C^\top} \mathcal{L} & = X + Y - C - \lambda \mathbf{1} \\ \nabla_\lambda \mathcal{L} & = -\text{Tr}[C] + 1. \end{cases} \quad (160)$$

Note, however, that these are also gradients of the Lagrangian

$$\mathcal{L}' = \frac{1}{2} \|X + Y - C\|_F^2 - \lambda(\text{Tr}[C] - 1), \quad (161)$$

corresponding to the minimization of the orthogonal projection

$$\text{Orth}\Pi_C(X + Y) = \arg \min_{C \in \mathcal{C}} \|X + Y - C\|_F^2.$$

Therefore, as desired,

$$\text{Prox}\Pi_C^h(X, Y) = \arg \min_{C \in \mathcal{C}} \|X + Y - C\|_F^2 = \text{Orth}\Pi_C(X + Y).$$

□

Lemma 7 (Bregman Divergence of the von Neumann Entropy) *If h is the von Neumann entropy,*

$$h(X) = \text{Tr}[X \log X], \quad (55)$$

then the Bregman divergence is the quantum relative entropy

$$\mathcal{BD}_h(X||Y) = \text{Tr}[X(\log X - \log Y)], \quad (56)$$

the corresponding mirror map (regularized best-response) is the logit map,

$$\text{Mir}\Pi_C^h(Y) = \frac{\exp(Y)}{\text{Tr}[\exp(Y)]} = \Lambda(Y), \quad (57)$$

and the corresponding proximal map is

$$\text{Prox}\Pi_C^h(X, Y) = \frac{\exp(\log X + Y)}{\text{Tr}[\exp(\log X + Y)]} = \Lambda(\log X + Y). \quad (58)$$

Proof. We will begin with an informal proof of the Bregman divergence of the von Neumann entropy of density matrices (for the original, rigorous proof refer to [74]). In doing so, we will leverage the following facts about density matrices X and Y :

1. $\nabla_{X^\top} \text{Tr}[X \log X] = \mathbb{1} + \log X$
2. $\text{Tr}[X] = \text{Tr}[Y] = 1$
3. Density matrices are Hermitian, which implies $X^\dagger = X$ and $Y^\dagger = Y$.

Thus, the Bregman divergence is the relative quantum entropy as follows,

$$\begin{aligned} \mathcal{BD}_h(X||Y) &= h(X) - h(Y) - \text{Tr}[(X - Y)^\dagger \nabla h(Y)] \\ &= \text{Tr}[X \log X] - \text{Tr}[Y \log Y] - \text{Tr}[(X - Y)^\dagger \nabla \text{Tr}[Y \log Y]] \\ &= \text{Tr}[X \log X] - \text{Tr}[Y \log Y] - \text{Tr}[(X - Y)^\dagger (\mathbb{1} + \log Y)] \\ &= \text{Tr}[X \log X] - \text{Tr}[Y \log Y] - \text{Tr}[X] + \text{Tr}[Y] - \text{Tr}[X \log Y] + \text{Tr}[Y \log Y] \\ &= \text{Tr}[X(\log X - \log Y)]. \end{aligned}$$

The regularized best-response,

$$\text{Mir}\Pi_{\mathcal{C}}^h(Y) = \arg \min_{C \in \mathcal{C}} \{ \text{Tr}[Y^\dagger C] - \text{Tr}[C \log C] \}$$

has Lagrangian

$$\mathcal{L}(Y) = \text{Tr}[Y^\dagger C] - \text{Tr}[C \log C] - \lambda \cdot (\text{Tr}[C] - 1),$$

with gradients

$$\nabla \mathcal{L} = \begin{cases} \nabla_{C^\top} \mathcal{L}(Y) &= Y - (\mathbb{1} + \log C) - \lambda \mathbb{1} = Y - \log C - \alpha \mathbb{1} \\ \nabla_\lambda \mathcal{L}(Y) &= -\text{Tr}[C] + 1, \end{cases}$$

where $\alpha = \lambda + 1$. Setting the gradients to zero results in the system of equations

$$\begin{cases} \log C &= Y - \alpha \mathbb{1} \\ \text{Tr}[C] &= 1 \end{cases} \implies \begin{cases} C &= \exp(Y - \alpha \mathbb{1}) = \exp(Y) \cdot e^{-\alpha} \mathbb{1} \\ \text{Tr}[C] &= 1. \end{cases}$$

We can solve for α by plugging the first equation into the second and using the properties of matrix exponentials as follows.

$$\text{Tr}[\exp(Y) \cdot e^{-\alpha} \mathbb{1}] = e^{-\alpha} \text{Tr}[\exp(Y)] = 1 \implies e^{-\alpha} = \frac{1}{\text{Tr}[\exp(Y)]}.$$

Substituting this back into our equation for C gives

$$C = \frac{\exp(Y)}{\text{Tr}[\exp(Y)]}.$$

Therefore, as desired, the regularized best response is the logit map

$$\text{Mir}\Pi_{\mathcal{C}}^h(Y) = \frac{\exp(Y)}{\text{Tr}[\exp(Y)]} = \Lambda(Y).$$

Finally, the corresponding proximal projection is given by

$$\text{Prox}\Pi_{\mathcal{C}}^h(X, Y) = \arg \min_{C \in \mathcal{C}} \left\{ \text{Tr}[Y^\dagger (X - C)] + \text{Tr}[C(\log C - \log X)] \right\}, \quad (162)$$

which has Lagrangian

$$\mathcal{L}(Y) = \text{Tr}[Y^\dagger(X - C)] + \text{Tr}[C(\log C - \log X)] - \lambda \cdot (\text{Tr}[C] - 1),$$

with gradients

$$\nabla \mathcal{L} = \begin{cases} \nabla_{C^\top} \mathcal{L}(Y) &= -Y + \mathbb{1} + \log C - \log X - \lambda \mathbb{1} = \log C - Y - \log X + \alpha \mathbb{1} \\ \nabla_\lambda \mathcal{L}(Y) &= -\text{Tr}[C] + 1, \end{cases}$$

where $\alpha = \lambda + 1$. Setting the gradients to zero results in the system of equations

$$\begin{cases} \log C &= Y + \log X - \alpha \mathbb{1} \\ \text{Tr}[C] &= 1 \end{cases} \implies \begin{cases} C &= \exp(Y + \log X) \cdot e^{-\alpha \mathbb{1}} \\ \text{Tr}[C] &= 1 \end{cases}$$

We can solve for α by plugging the first equation into the second and using the properties of matrix exponentials as follows.

$$\text{Tr}[\exp(Y + \log X) \cdot e^{-\alpha \mathbb{1}}] = e^{-\alpha} \text{Tr}[\exp(Y + \log X)] = 1 \implies e^{-\alpha} = \frac{1}{\text{Tr}[\exp(Y + \log X)]}.$$

Substituting this back into our equation for C gives

$$C = \frac{\exp(Y + \log X)}{\text{Tr}[\exp(Y + \log X)]}.$$

Therefore, as desired, the regularized best response is the logit map

$$\text{Prox}\Pi_C^h(X, Y) = \frac{\exp(Y + \log X)}{\text{Tr}[\exp(Y + \log X)]} = \Lambda(Y + \log X).$$

□

C Properties of Finite-Valued Zero-Sum Quantum Games

In this appendix we will prove several important properties of finite-valued zero-sum quantum games.

C.1 Properties of Spectraplexes

Lemma 10 *The spectraplex $\mathcal{X} = \{X \in \mathcal{H}_+^k : \text{Tr}(X) = 1\}$ is a convex set.*

Proof. By the definition of convex set, \mathcal{X} is a convex set iff, for any $X_1, X_2 \in \mathcal{X}$ and $t \in (0, 1]$,

$$\tilde{X} = (1 - t) \cdot X_1 + t \cdot X_2 \tag{163}$$

such that $\tilde{X} \in \mathcal{H}_+^k$ and $\text{Tr}(\tilde{X}) = 1$.

We begin by proving that \tilde{X} is positive semi-definite, $\tilde{X} \in \mathcal{H}_+^k$, which is true iff $\forall |v\rangle \in \mathbb{C}^k$,

$$\langle v | \tilde{X} | v \rangle \geq 0. \tag{164}$$

Given that $X_1, X_2 \in \mathcal{H}_+^k$,

$$\forall |v\rangle \in \mathbb{C}^k : \langle v|X_1|v\rangle \geq 0 \quad \text{and} \quad \langle v|X_2|v\rangle \geq 0. \quad (165)$$

Therefore, since $t \geq 0$ and $(1-t) \geq 0$,

$$\langle v|\tilde{X}|v\rangle = \langle v|(1-t) \cdot X_1 + t \cdot X_2|v\rangle = (1-t) \langle v|X_1|v\rangle + t \langle v|X_2|v\rangle \geq 0. \quad (166)$$

We now prove that $\text{Tr}(\tilde{X}) = 1$. Since $\text{Tr}[X_1] = \text{Tr}[X_2] = 1$,

$$\text{Tr}(\tilde{X}) = \text{Tr}[(1-t) \cdot X_1 + t \cdot X_2] = (1-t) \cdot \text{Tr}[X_1] + t \cdot \text{Tr}[X_2] = 1 - t + t = 1 \quad (167)$$

□

Lemma 11 *The spectraplex $\mathcal{X} = \{X \in \mathcal{H}_+^k : \text{Tr}(X) = 1\}$ is a compact set.*

Proof. See page 162 of [75].

□

C.2 Properties of \mathcal{F}

Lemma 8 (Mononicity of \mathcal{F}) *\mathcal{F} is monotone or, equivalently,*

$$\text{Tr}[(\mathcal{F}(X) - \mathcal{F}(Y))(X - Y)] \geq 0, \quad \forall X, Y \in \mathcal{C}. \quad (73)$$

Proof. For this proof, we will leverage the Pauli decompositions of the involved operators. Note that every matrix M can be decomposed in terms of the Pauli matrices $\mathcal{P} = \{I, X, Y, Z\}$, as

$$M = \sum_{P \in \mathcal{P}^{\otimes n}} \widehat{M}(P)P,$$

where the tensored Pauli matrices P constitute an orthonormal basis, with respect to trace inner product, and possess corresponding Pauli coefficients $\widehat{M}(P)$. Note that for all Paulis, we have $P = P^\dagger$. The Pauli coefficients of an n -qubit system can also be explicitly defined as

$$\widehat{M}(P) = \frac{1}{2^n} \text{Tr}(P^\dagger M).$$

The payoff observable U can be decomposed as

$$U = \sum_{R \in \mathcal{P}^{\otimes(nm)}} \widehat{U}(R)R = \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)P \otimes Q,$$

where $\widehat{U}(P, Q) = \widehat{U}(P \otimes Q)$. Plugging this into the payoff gradients defined in [Equation \(16\)](#) and [Equation \(17\)](#), we get that

$$\begin{aligned}
\mathcal{F}_\alpha(\alpha, \beta) &= \nabla_{\alpha^\top} \text{Tr} \left(U^\dagger(\alpha \otimes \beta) \right) \\
&= \nabla_{\alpha^\top} \text{Tr} \left(\left(\sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q) P \otimes Q \right)^\dagger (\alpha \otimes \beta) \right) \\
&= \nabla_{\alpha^\top} \text{Tr} \left(\sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* P^\dagger \alpha \otimes Q^\dagger \beta \right) \\
&= \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot \nabla_{\alpha^\top} \text{Tr} \left(P^\dagger \alpha \otimes Q^\dagger \beta \right) \\
&= \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot \nabla_{\alpha^\top} \text{Tr} \left(P^\dagger \alpha \right) \text{Tr} \left(Q^\dagger \beta \right) \\
&= \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot 2^m \widehat{\beta}(Q) P^\dagger
\end{aligned}$$

and, similarly,

$$\begin{aligned}
\mathcal{F}_\beta(\alpha, \beta) &= -\nabla_{\beta^\top} \text{Tr} \left(U^\dagger(\alpha \otimes \beta) \right) \\
&= -\sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot \nabla_{\beta^\top} \text{Tr} \left(P^\dagger \alpha \right) \text{Tr} \left(Q^\dagger \beta \right) \\
&= -\sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot 2^n \widehat{\alpha}(P) Q^\dagger.
\end{aligned}$$

Thus, plugging these results into [Equation \(19\)](#), the operator \mathcal{F} can be decomposed as

$$\mathcal{F}(\alpha, \beta) = \begin{pmatrix} 2^m & \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot \widehat{\beta}(Q) P^\dagger \\ -2^n & \sum_{P \in \mathcal{P}^{\otimes n}} \sum_{Q \in \mathcal{P}^{\otimes m}} \widehat{U}(P, Q)^* \cdot \widehat{\alpha}(P) Q^\dagger \end{pmatrix}. \quad (168)$$

Now, let us consider the two distinct games states, $X = (x_1, x_2)$ and $Y = (y_1, y_2)$, where $x_1, y_1 \in \mathcal{A}$ and $x_2, y_2 \in \mathcal{B}$, with difference

$$X - Y = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} x_1 - y_1 \\ x_2 - y_2 \end{pmatrix}.$$

Following from [Equation \(168\)](#), the difference of the payoff gradients is

$$F(X) - F(Y) = F(x_1, x_2) - F(y_1, y_2) = \begin{pmatrix} 2^m \sum_{P, Q} \widehat{U}(P, Q)^* \cdot (\widehat{x}_2(Q) - \widehat{y}_2(Q)) P^\dagger \\ 2^n \sum_{P, Q} \widehat{U}(P, Q)^* \cdot (\widehat{y}_1(P) - \widehat{x}_1(P)) Q^\dagger \end{pmatrix}$$

Plugging these into the left-hand side of Equation (73), we get that

$$\begin{aligned}
& \text{Tr} [(\mathcal{F}(X) - \mathcal{F}(Y))(X - Y)] = \\
&= \text{Tr} \left[\begin{pmatrix} 2^m \sum_{P,Q} \widehat{U}(P, Q)^* \cdot (\widehat{x}_2(Q) - \widehat{y}_2(Q)) & P^\dagger \\ 2^n \sum_{P,Q} \widehat{U}(P, Q)^* \cdot (\widehat{y}_1(P) - \widehat{x}_1(P)) & Q^\dagger \end{pmatrix} \begin{pmatrix} x_1 - y_1 \\ x_2 - y_2 \end{pmatrix} \right] \\
&= \text{Tr} \left[\sum_{P,Q} \widehat{U}(P, Q)^* \left(2^m (\widehat{x}_2(Q) - \widehat{y}_2(Q)) P^\dagger (x_1 - y_1) + 2^n (\widehat{y}_1(P) - \widehat{x}_1(P)) Q^\dagger (x_2 - y_2) \right) \right] \\
&= \sum_{P,Q} \widehat{U}(P, Q)^* \left(2^m (\widehat{x}_2(Q) - \widehat{y}_2(Q)) \text{Tr} [P^\dagger (x_1 - y_1)] + 2^n (\widehat{y}_1(P) - \widehat{x}_1(P)) \text{Tr} [Q^\dagger (x_2 - y_2)] \right) \\
&= \sum_{P,Q} \widehat{U}(P, Q)^* \left(2^m (\widehat{x}_2(Q) - \widehat{y}_2(Q)) 2^n (\widehat{x}_1(P) - \widehat{y}_1(P)) + 2^n (\widehat{y}_1(P) - \widehat{x}_1(P)) 2^m (\widehat{x}_2(Q) - \widehat{y}_2(Q)) \right) \\
&= 2^{n+m} \sum_{P,Q} \widehat{U}(P, Q)^* \left((\widehat{x}_2(Q) - \widehat{y}_2(Q)) (\widehat{x}_1(P) - \widehat{y}_1(P)) - (\widehat{x}_1(P) - \widehat{y}_1(P)) (\widehat{x}_2(Q) - \widehat{y}_2(Q)) \right) \\
&= 0.
\end{aligned}$$

□

Lemma 9 (Lipschitz Continuity of \mathcal{F}) For finite-valued quantum zero-sum games, \mathcal{F} is a Lipschitz continuous operator, meaning there exists a Lipschitz constant $\gamma \in \mathbb{R}$ such that

$$\|\mathcal{F}(Z) - \mathcal{F}(Z')\|_* \leq \gamma \|Z - Z'\|, \quad \forall Z, Z' \in \mathcal{C}. \quad (108)$$

For the case of $(\|\cdot\|_F, \|\cdot\|_F)$ and $(\|\cdot\|_\infty, \|\cdot\|_1)$, $\gamma_{F,F} = \mathcal{O}(2^d)$ while $\gamma_{\infty,1} = \mathcal{O}(1)$.

Proof. In Lemma 12, we proved that \mathcal{F} is linear, which implies that for any $Z, Z' \in \mathcal{C}$,

$$\mathcal{F}(Z) - \mathcal{F}(Z') = \mathcal{F}(Z - Z').$$

In order to prove that \mathcal{F} is Lipschitz continuous, we simply need to prove that there exists a $\gamma \in \mathbb{R}$ such that, for all $Z \neq Z'$,¹⁰

$$\frac{\|\mathcal{F}(Z - Z')\|_F}{\|Z - Z'\|_F} \leq \gamma.$$

Letting $Y = Z - Z'$, this is equivalent to proving

$$\sup_{Y \neq 0} \frac{\|\mathcal{F}(Y)\|_F}{\|Y\|_F} \leq \gamma. \quad (169)$$

Vectorizing Y , as a vector of size $|\mathcal{C}|$, the Frobenius norm can be expressed as

$$\|Y\|_F = \sqrt{\sum_i |Y_i|^2}.$$

Furthermore, since \mathcal{F} is a linear operator, it maps Y to some new joint state $Y' = \mathcal{F}(Y) \in \mathcal{C}$, with entries

$$Y'_j = \sum_i c_{ij} Y_i.$$

¹⁰In the case that $Z = Z'$, trivially $Z - Z' = 0$ and $\mathcal{F}(Z - Z') = 0$, meaning any value for γ suffices.

for some linear coefficients $c_{ij} \in \mathbb{C}$. Note that since the setting is restricted to finite-valued quantum zero-sum games, with a finite-valued utility function \mathcal{U} , the coefficients c_{ij} are guaranteed to be finite-valued. Therefore, the right-hand side of Equation (169) can be re-expressed as

$$\sup_{Y \neq 0} \frac{\|\mathcal{F}(Y)\|_F}{\|Y\|_F} = \sup_{Y \neq 0} \frac{\|Y'\|_F}{\|Y\|_F} = \sup_{Y \neq 0} \sqrt{\frac{\sum_j |\sum_i c_{ij} Y_i|^2}{\sum_i |Y_i|^2}}.$$

Letting $c^* = \max(\{c_{i,j}\})$,

$$\sup_{Y \neq 0} \sqrt{\frac{\sum_j |\sum_i c_{ij} Y_i|^2}{\sum_i |Y_i|^2}} \leq \sup_{Y \neq 0} \sqrt{\frac{\sum_j |\sum_i c^* Y_i|^2}{\sum_i |Y_i|^2}} = \sup_{Y \neq 0} \sqrt{\frac{c^* \sum_j |\sum_i Y_i|^2}{\sum_i |Y_i|^2}}$$

and leveraging the triangle inequality,

$$\sup_{Y \neq 0} \sqrt{\frac{c^* \sum_j |\sum_i Y_i|^2}{\sum_i |Y_i|^2}} \leq \sup_{Y \neq 0} \sqrt{\frac{c^* \sum_j \sum_i |Y_i|^2}{\sum_i |Y_i|^2}} = \sup_{Y \neq 0} \sqrt{\frac{c^* |\mathcal{C}| \cdot \sum_i |Y_i|^2}{\sum_i |Y_i|^2}} = \sqrt{c^* |\mathcal{C}|}.$$

Thus, for any $\gamma \geq \sqrt{c^* |\mathcal{C}|}$, Equation (169) is satisfied. Recall that $c^* = \Theta(1)$ for games in $[-1, 1]$. Therefore, the linear operator \mathcal{F} is Lipschitz continuous where $\gamma_{F,F} = \mathcal{O}(\sqrt{4^d})$. Applying similar argumentation for the case of real vector and norm pairs (ℓ_∞, ℓ_1) , we get $\gamma_{\infty,1} = \mathcal{O}(1)$. Indeed, it is easy to see

$$\sup_{Y \neq 0} \frac{\|\mathcal{F}(Y)\|_\infty}{\|Y\|_1} = \sup_{Y \neq 0} \frac{\|Y'\|_\infty}{\|Y\|_1} = \sup_{Y \neq 0} \sqrt{\frac{\max_j |\sum_i c_{ij} Y_i|}{\sum_i |Y_i|}} \leq c^* = \mathcal{O}(1).$$

□

Lemma 1 *If \mathcal{F} is monotone, then a strong solution is a weak solution.*

Proof. Denote a strong solution, satisfying Equation (20), as Ψ^S . Since \mathcal{F} is monotone, from the definition of monotonicity, $\forall \Psi \in \mathcal{C}$,

$$\text{Tr} \left[(\mathcal{F}(\Psi) - \mathcal{F}(\Psi^S)) (\Psi - \Psi^S) \right] \geq 0, \quad (170)$$

which implies that,

$$\text{Tr} \left[\mathcal{F}(\Psi) (\Psi - \Psi^S) \right] - \text{Tr} \left[\mathcal{F}(\Psi^S) (\Psi - \Psi^S) \right] \geq 0. \quad (171)$$

Since Ψ^S is a strong solution, by Equation (20),

$$\text{Tr} \left[\mathcal{F}(\Psi) (\Psi - \Psi^S) \right] \geq \text{Tr} \left[\mathcal{F}(\Psi^S) (\Psi - \Psi^S) \right] \geq 0. \quad (172)$$

Therefore, $\forall \Psi \in \mathcal{C}$,

$$\text{Tr} \left[\mathcal{F}(\Psi) (\Psi^S - \Psi) \right] \leq 0, \quad (173)$$

meaning Ψ^S satisfies Equation (21) and is, thus, also a weak solution. □

Lemma 2 *If \mathcal{F} is Lipschitz continuous, then a weak solution is a strong solution.*

Proof. Denote a weak solution, satisfying Equation (21), as Ψ^W . In Lemma 10 we proved that spectraplexes are convex, which implies that \mathcal{A} , \mathcal{B} , and, thus, \mathcal{C} are convex sets. Therefore, by the definition of convex set, for any $\Psi \in \mathcal{C}$ and $t \in (0, 1]$,

$$\Psi^W + t \cdot (\Psi - \Psi^W) \in \mathcal{C}. \quad (174)$$

Plugging this state into Equation (21) gives

$$\text{Tr} \left[\mathcal{F}(\Psi^W + t \cdot (\Psi - \Psi^W)) \left(\Psi^W + t \cdot (\Psi - \Psi^W) - \Psi^W \right) \right] \geq 0 \quad (175)$$

$$t \cdot \text{Tr} \left[\mathcal{F}(\Psi^W + t \cdot (\Psi - \Psi^W)) \left(\Psi - \Psi^W \right) \right] \geq 0 \quad (176)$$

$$\text{Tr} \left[\mathcal{F}(\Psi^W + t \cdot (\Psi - \Psi^W)) \left(\Psi - \Psi^W \right) \right] \geq 0. \quad (177)$$

Since \mathcal{F} is Lipschitz continuous, we can take the following limit,

$$\lim_{t \rightarrow 0} \text{Tr} \left[\mathcal{F}(\Psi^W + t \cdot (\Psi - \Psi^W)) \left(\Psi - \Psi^W \right) \right] \geq 0. \quad (178)$$

The resulting expression,

$$\text{Tr} \left[\mathcal{F}(\Psi^W) \left(\Psi - \Psi^W \right) \right] \geq 0, \quad \forall \Psi \in \mathcal{C}, \quad (179)$$

implies that Ψ^W satisfies Equation (20) and is also a strong solution. \square

Lemma 12 (Linearity of \mathcal{F}) $\mathcal{F} : \mathcal{A} \times \mathcal{B} \mapsto \mathcal{C}$ is a linear operator. Therefore, \mathcal{F} must satisfy the following properties:

1. For any $\lambda \in \mathbb{C}$, $\mathcal{F}(\lambda\alpha, \lambda\beta) = \lambda\mathcal{F}(\alpha, \beta)$.
2. For $\alpha, \alpha' \in \mathcal{A}$ and $\beta, \beta' \in \mathcal{B}$, $\mathcal{F}(\alpha + \alpha', \beta + \beta') = \mathcal{F}(\alpha, \beta) + \mathcal{F}(\alpha', \beta')$.

Proof. From Equation (16) and Equation (17), it follows that

$$\begin{aligned} \mathcal{F}_\alpha(\alpha, \beta) &= \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \beta) \right] = \mathcal{F}_\alpha(\beta) \in \mathcal{A} \\ \mathcal{F}_\beta(\alpha, \beta) &= -\text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha \otimes \mathbb{1}_m) \right] = \mathcal{F}_\beta(\alpha) \in \mathcal{B} \\ \mathcal{F}(\alpha, \beta) &= (\mathcal{F}_\alpha(\alpha, \beta), \mathcal{F}_\beta(\alpha, \beta)) \in \mathcal{C} \end{aligned}$$

such that $\mathcal{F} : \mathcal{A} \times \mathcal{B} \mapsto \mathcal{C}$.

We will now prove that \mathcal{F} satisfies the *multiplicative property*. If $\lambda \in \mathbb{C}$, then

$$\begin{aligned} \mathcal{F}_{\lambda\alpha}(\lambda\alpha, \lambda\beta) &= \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \lambda\beta) \right] = \lambda \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \beta) \right] = \lambda \mathcal{F}_\alpha(\alpha, \beta) \\ \mathcal{F}_{\lambda\beta}(\lambda\alpha, \lambda\beta) &= -\text{Tr}_{\mathcal{A}} \left[U^\dagger(\lambda\alpha \otimes \mathbb{1}_m) \right] = -\lambda \text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha \otimes \mathbb{1}_m) \right] = \lambda \mathcal{F}_\beta(\alpha, \beta), \end{aligned}$$

which implies, as desired,

$$\mathcal{F}(\lambda\alpha, \lambda\beta) = \left(\mathcal{F}_{\lambda\alpha}(\lambda\alpha, \lambda\beta), \mathcal{F}_{\lambda\beta}(\lambda\alpha, \lambda\beta) \right) = \left(\lambda \mathcal{F}_\alpha(\alpha, \beta), \lambda \mathcal{F}_\beta(\alpha, \beta) \right) = \lambda \mathcal{F}(\alpha, \beta).$$

We will now prove that \mathcal{F} satisfies the *additive property*. If $\alpha, \alpha' \in \mathcal{A}$ and $\beta, \beta' \in \mathcal{B}$, then

$$\begin{aligned}
\mathcal{F}_{\alpha+\alpha'}(\alpha + \alpha', \beta + \beta') &= \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes (\beta + \beta')) \right] \\
&= \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \beta) + U^\dagger(\mathbb{1}_n \otimes \beta') \right] \\
&= \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \beta) \right] + \text{Tr}_{\mathcal{B}} \left[U^\dagger(\mathbb{1}_n \otimes \beta') \right] \\
&= \mathcal{F}_\alpha(\alpha, \beta) + \mathcal{F}_{\alpha'}(\alpha', \beta') \\
\mathcal{F}_{\beta+\beta'}(\alpha + \alpha', \beta + \beta') &= -\text{Tr}_{\mathcal{A}} \left[U^\dagger((\alpha + \alpha') \otimes \mathbb{1}_m) \right] \\
&= -\text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha \otimes \mathbb{1}_m) + U^\dagger(\alpha' \otimes \mathbb{1}_m) \right] \\
&= -\text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha \otimes \mathbb{1}_m) \right] - \text{Tr}_{\mathcal{A}} \left[U^\dagger(\alpha' \otimes \mathbb{1}_m) \right] \\
&= \mathcal{F}_\beta(\alpha, \beta) + \mathcal{F}_{\beta'}(\alpha', \beta'),
\end{aligned}$$

which implies, as desired,

$$\begin{aligned}
\mathcal{F}(\alpha + \alpha', \beta + \beta') &= \left(\mathcal{F}_\alpha(\alpha, \beta) + \mathcal{F}_{\alpha'}(\alpha', \beta'), \mathcal{F}_\beta(\alpha, \beta) + \mathcal{F}_{\beta'}(\alpha', \beta') \right) \\
&= \left(\mathcal{F}_\alpha(\alpha, \beta), \mathcal{F}_\beta(\alpha, \beta) \right) + \left(\mathcal{F}_{\alpha'}(\alpha', \beta'), \mathcal{F}_{\beta'}(\alpha', \beta') \right) \\
&= \mathcal{F}(\alpha, \beta) + \mathcal{F}(\alpha', \beta').
\end{aligned}$$

□

D Convergence Analysis Proofs

Proposition 7 *If $g(X)$ is a μ -strongly convex function and $h(X)$ is an α -strongly convex function, then $f(X) = g(X) + h(X)$ is a $\mu + \alpha$ -strongly convex function.*

Proof. By the definition of strong convexity (Definition 2), $\forall X, Y \in \mathcal{C}$

$$g(X) \geq g(Y) + \langle \nabla g(Y), X - Y \rangle + \frac{\mu}{2} \|X - Y\|_F^2, \quad (180)$$

$$h(X) \geq h(Y) + \langle \nabla h(Y), X - Y \rangle + \frac{\alpha}{2} \|X - Y\|_F^2. \quad (181)$$

Summing together these two expressions gives

$$g(X) + h(X) \geq g(Y) + h(Y) + \langle \nabla g(Y) + \nabla h(Y), X - Y \rangle + \frac{\mu + \alpha}{2} \|X - Y\|_F^2 \quad (182)$$

$$\implies f(X) \geq f(Y) + \langle \nabla f(Y), X - Y \rangle + \frac{\mu + \alpha}{2} \|X - Y\|_F^2, \quad (183)$$

which implies f is $(\mu + \alpha)$ -strongly convex. □

Proposition 3 *For μ -strongly convex regularization function h with respect to norm $\|\cdot\|$, ϕ_t is $\frac{\mu}{\eta}$ -strongly convex with respect to norm $\|\cdot\|$.*

Proof. To prove that $\phi_t(Z)$ is $\frac{\mu}{\eta}$ -strongly convex, we will leverage [Proposition 7](#). Namely, by proving that $\text{Tr} [\mathcal{F}(\Psi_t) (Z - \Phi_{t-1})]$ is convex (i.e., 0-strongly convex) and that $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ is $\frac{\mu}{\eta}$ -strongly convex, then

$$\phi_t(Z) = \text{Tr} [\mathcal{F}(\Psi_t) (Z - \Phi_{t-1})] + \frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1}), \quad (184)$$

must be $0 + \frac{\mu}{\eta} = \frac{\mu}{\eta}$ -strongly convex.

We begin by proving that $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ is $\frac{\mu}{\eta}$ -strongly convex. From the definition of strong convexity ([Definition 2](#)), if h is μ -strongly convex, then the Bregman divergence \mathcal{BD}_h must also be μ -strongly convex. This implies that

$$\frac{1}{\eta}\mathcal{BD}_h(X\|Y) \geq \frac{\mu}{2\eta}\|X - Y\|_F^2, \quad \forall X, Y \in \mathcal{C}, \quad (185)$$

meaning the function $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ must be $\frac{\mu}{\eta}$ -strongly convex.

We now prove that $f(Z) = \text{Tr} [\mathcal{F}(\Psi_t) (Z - \Phi_{t-1})]$ is convex. For all $X, Y \in \mathcal{C}$,

$$f(X) - f(Y) = \text{Tr} [\mathcal{F}(\Psi_t) (X - \Phi_{t-1})] - \text{Tr} [\mathcal{F}(\Psi_t) (Y - \Phi_{t-1})] = \text{Tr} [\mathcal{F}(\Psi_t) (X - Y)]. \quad (186)$$

Since $\forall Z \in \mathcal{C}$, $\nabla f(Z) = \mathcal{F}(\Psi_t)$, this implies that

$$f(X) - f(Y) = \text{Tr} [\mathcal{F}(\Psi_t) (X - Y)] = \text{Tr} [\nabla f(Y) (X - Y)], \quad (187)$$

meaning, by [Definition 1](#), f is convex. \square

Proposition 4 For μ -strongly convex regularization function h with respect to norm $\|\cdot\|$, ψ_t is $\frac{\mu}{\eta}$ -strongly convex with respect to norm $\|\cdot\|$.

Proof. To prove that $\psi_t(Z)$ is $\frac{\mu}{\eta}$ -strongly convex, we will leverage [Proposition 7](#). Namely, by proving that $\text{Tr} [\mathcal{F}(\Psi_{t-1}) (Z - \Phi_{t-1})]$ is convex (i.e. 0-strongly convex) and that $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ is $\frac{\mu}{\eta}$ -strongly convex, then

$$\psi_t(Z) = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (Z - \Phi_{t-1})] + \frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1}), \quad (188)$$

must be $0 + \frac{\mu}{\eta} = \frac{\mu}{\eta}$ -strongly convex.

We begin by proving that $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ is $\frac{\mu}{\eta}$ -strongly convex. From the definition of strong convexity ([Definition 2](#)), if h is μ -strongly convex, then the Bregman divergence \mathcal{BD}_h must also be μ -strongly convex. This implies that

$$\frac{1}{\eta}\mathcal{BD}_h(X\|Y) \geq \frac{\mu}{2\eta}\|X - Y\|_F^2, \quad \forall X, Y \in \mathcal{C}, \quad (189)$$

meaning the function $\frac{1}{\eta}\mathcal{BD}_h(Z\|\Phi_{t-1})$ must be $\frac{\mu}{\eta}$ -strongly convex.

We now prove that $f(Z) = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (Z - \Phi_{t-1})]$ is convex. For all $X, Y \in \mathcal{C}$,

$$f(X) - f(Y) = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (X - \Phi_{t-1})] - \text{Tr} [\mathcal{F}(\Psi_{t-1}) (Y - \Phi_{t-1})] = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (X - Y)]. \quad (190)$$

Since $\forall Z \in \mathcal{C}$, $\nabla f(Z) = \mathcal{F}(\Psi_{t-1})$, this implies that

$$f(X) - f(Y) = \text{Tr} [\mathcal{F}(\Psi_{t-1}) (X - Y)] = \text{Tr} [\nabla f(Y) (X - Y)], \quad (191)$$

meaning, by [Definition 1](#), f is convex. \square