Multi-Objective Optimization inspired by Quantum Games

A Thesis

submitted to

Indian Institute of Science Education and Research, Pune

in partial fulfillment of the requirements for the

MS Degree Programme

by

Student: **Rushikesh Sunil Ambekar**

Reg. No. **20232012**

Supervisor: **Dr. Sreejith G. J.**

Expert: **Dr. T. S. Mahesh**

May 2026

Indian Institute of Science Education and Research, Pune

Dr. Homi Bhabha Road, Pashan

PUNE, Maharashtra - 411008, INDIA.

Certificate

This is to certify that this dissertation entitled **"Multi-Objective Optimization inspired by Quantum Games"** towards the partial fulfillment of the **MS degree programme** at the **Indian Institute of Science Education and Research, Pune** represents study/work carried out by **Rushikesh Sunil Ambekar** at IISER Pune under the supervision of **Dr. Sreejith G. J.**, Associate Professor, Department of Physics, during the academic year 2025–2026.

Dr. Sreejith G. J.
Associate Professor
Dept. of Physics
IISER Pune
Supervisor

Dr. T. S. Mahesh
Professor and Chair
Dept. of Physics
IISER Pune
Expert

Declaration

I hereby declare that the matter embodied in the report entitled **"Multi-Objective Optimization inspired by Quantum Games"** are the results of the work carried out by me at the Department of Physics, Indian Institute of Science Education and Research, Pune, under the supervision of **Dr. Sreejith G. J.** and the same has not been submitted elsewhere for any other degree.

Rushikesh Sunil Ambekar
MS student
Reg. No. 20232012
Dept. of Physics
IISER Pune

Contents

Abstract

Game theory provides a mathematical framework for the study of strategic interactions among rational agents whose decisions influence one another. Concepts like *best responses*, *Nash equilibrium*, and *Pareto optimality* have been widely used to study decision-making processes in economics and in computer science as well. At the same time, many problems of practical interest can be formulated into a classical optimization problem, which is generally computationally difficult. Recent developments in *variational quantum algorithms* (VQA) [4], have led to a new technique called *Quantum Approximate Optimization Algorithm* (QAOA) [9], which aims to address certain classical optimization problems using hybrid quantum–classical approaches.

In this work, we introduce a quantum system in which ideas from game theory are applied for the optimization of quantum systems. We formulate a quantum game involving two agents acting on a shared multi-qubit system, where each player seeks to optimize the expectation value of a given payoff operator. We analyze the behavior of this system through numerical simulations and investigate the emergence of equilibrium-like configurations in the parameter space of quantum states.

This perspective connects concepts from game theory, quantum optimization, and variational quantum algorithms. The study may also provide insights into the design of multi-agent reinforcement learning algorithms for quantum-mechanical systems, where multiple learning agents interact through a shared quantum environment.

Chapter 1

Introduction

We explore a mutli-qubit quantum system in particular game settings and study the system behaviour. The study of this problem is presented in the thesis as follows;

1. Chapter 1: We will start with a brief overview of the concepts in *game theory* relevant to our study.
2. Chapter 2: Then we will discuss in brief about Classical Optimization Problems, in particular Boolean Satisfiability Problem (SAT) problem which is the motivation for our study.
3. Chapter 3: Then we will give a brief introduction to *quantum games.*
4. Chapter 4: We formulate the problem for single-qubit system and present the results obtained from numerical simulation.
5. Chapter 5: We formulate the problem now for three-qubit system with access control for players and present the results obtained from numerical simulation.
6. Chapter 6: We propose a *Pareto-Optimal* variant of the quantum game which we studied in this thesis, this seems to be an interesting direction for furthur work which may have potential application in ground state search problems in physics.
7. Chapter 7: The Conclusion of the work.

1.1 Game Theory

Game Theory provides a mathematical framework to study situations involving strategic decision-making among rational players. Concepts such as zero-sum games, payoff matrices, rational behavior, and Nash equilibria have wide applications — from economics to diplomacy between countries, to business strategies, to sports, to politics and also in decision making in personal and professional life. [8]

This understanding of strategic interactions gives a framework for understanding the decision-making processes in various systems. In particular, the ideas of best responses and equilibrium strategies can be extended beyond classical games to quantum systems. In the upcoming chapters, we explore *a quantum system where sequential optimization of two different operators* (called payoff operators) is performed numerically. This can be interpreted as strategic interaction, analogous to players iteratively adjusting their strategies in a game to reach a stable configuration which can be viewed as a Nash equilibrium condition similar to the classical Game Theory. Such ideas could potentially be used to develop multi-agent reinforcement learning algorithms [22] for quantum-mechanical systems.

In this chapter, we first review some fundamental concepts about Game Theory in brief and then their relevance to our problem in the other chapters.

1.2 What is a Game? The Mathematical Definition

A game can be defined as a set having the following properties: [16]

$$G = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}), \tag{1.1}$$

where:

- $N = \{1, 2, \ldots, N\}$ is the set of players,
- S_i is the strategy set available to player i, the elements of this set alter the state of the game,
- $u_i : S_1 \times \cdots \times S_N \to \mathbb{R}$ is the payoff function of player i.

A strategy profile is a joint choice

$$s = (s_1, s_2, \ldots, s_N) \in S_1 \times \cdots \times S_N. \tag{1.2}$$

For example, a two-player game can be defined as follows:

$$G = (\{A, B\}, S_A, S_B, u_A, u_B), \tag{1.3}$$

where:

- A and B denote the two players,

- S_A and S_B are the respective strategy sets of players A and B,
- $u_A : S_A \times S_B \to \mathbb{R}$ and $u_B : S_A \times S_B \to \mathbb{R}$ are the payoff functions for A and B respectively.

A **strategy profile** is a pair

$$s = (s_A, s_B) \in S_A \times S_B,$$

where $s_A \in S_A$ and $s_B \in S_B$. Given a strategy profile $s = (s_A, s_B)$, the payoffs to the players are

$$u_A(s_A, s_B), \qquad u_B(s_A, s_B).$$

For each player, one may define the **best response**. The best response of player A, i.e. one which will maximize $A's$ payoff, to a fixed strategy s_B of player B is

$$\tilde{s}_A = \arg \max_{s_A \in S_A} u_A(s_A, s_B), \tag{1.4}$$

and similarly for player B, i.e. the strategy which will maximize the $B's$ payoff, given the strategy s_A of player A is fixed is,

$$\tilde{s}_B = \arg \max_{s_B \in S_B} u_B(s_A, s_B). \tag{1.5}$$

1.3 Zero-Sum Games

A zero-sum game is one in which one player's gain is exactly equal to the other player's loss.[21] Mathematically, if the payoff matrix (defined in next section) is A, then the second player's payoff matrix will be $-A$. Let us understand this through an example. The game of cricket is a zero-sum game, the batting team has payoffs like (+1,+2,+3,+4,+5,+6) etc as they can score runs, but these will be counted as negative payoffs for the bowling team, similarly we assign certain payoff to a wicket falling for batting team which will be negative, say -20, that will indeed be +20 payoff for the bowling team. In zero-sum games, players are always at odds with each other, i.e. player A's profit is player B's loss; meaning there is no scope for cooperation between the two.

However, there are non zero-sum games also (like Prisoner's Dilemma, explained in next section), wherein cooperative strategies are also possible in certain situations and non cooperative strategies are also possible. Apart from simultaneous game (like Prisoner's Dilemma or Rock-Paper-Scissors) there are also sequential games, where players take a move one after the other (chess or tic-tac-toe). Consider this example from politics, two poltical parties L and R; they will almost always go in opposing each other, but also many a times cooperate silently in decisions where both L and R get benefitted. This is an example of non-cooperative game in general, where there is an occasional chance for cooperation between the players.[8]

1.4 Payoff Matrices

A payoff matrix encodes the payoffs which are the outcomes for all possible strategy combinations in the particular game. The rows represent Player A's strategies' payoff, and the columns represent Player B's strategies' payoff. Therefore, each matrix entry contains the payoff corresponding to the strategies executed by the players. [21] Let us understand this through some examples.

1.4.1 Example 1: Rock-Paper-Scissors

Consider two players, A and B, each choosing between Rock (R), Paper (P), or Scissors (S). The rules are the following:

- Rock beats Scissors
- Scissors beats Paper
- Paper beats Rock

We assign payoffs as: Win $= +1$, Lose $= -1$, Draw $= 0$. The pair (i, j) in the matrix represents the payoffs, i is for player A and j is for player B.

	Rock (B)	Paper (B)	Scissors (B)
Rock (A)	(0,0)	(-1,1)	(1,-1)
Paper (A)	(1,-1)	(0,0)	(-1,1)
Scissors (A)	(-1,1)	(1,-1)	(0,0)

Table 1.1: Payoff matrix for Rock-Paper-Scissors with color coding: green = win for A, red = loss for A, gray = draw.

This is a **zero-sum game** because the pair (i, j) is such that $j = -i$, which means gain of one is exactly equal to the loss of other.

1.4.2 Example 2: Prisoner's Dilemma

The Prisoner's Dilemma is a classic non-zero-sum game. [8] Different variants of Prisoner's Dilemma are encountered in many everyday examples. Imagine two suspects, let us say Alice and Bob, are arrested for a crime and **interrogated separately**. Note that the Police doesn't have a strong evidence against them, so their responses in the Police Interrogation have importance. Each can choose to either **Cooperate** (i.e. stay silent, means deny the charge) or to **Defect** (i.e. betray the other, means confess to the police). Their outcomes depend on their choices i.e. their strategy:

- If both cooperate (i.e. stay silent), they get a light sentence: 1 year of imprisonment each; since there is no strong evidence against them.
- If both defect (i.e. betray), they get a moderate sentence: 3 years of imprisonment each; since now the police have a evidence as their confession.
- If one defects while the other cooperates, the defector goes free, and the cooperator receives a heavy sentence: 5 years of imprisonment.

	Cooperate (Bob)	Defect (Bob)
Cooperate (Alice)	(-1,-1)	(-5,0)
Defect (Alice)	(0,-5)	(-3,-3)

Table 1.2: Payoff matrix for the Prisoner's Dilemma with color coding: green = mutual cooperation (best collective outcome), red = betrayal (one defects, other cooperates), gray = mutual defection.

This illustrates the conflict between *individual rationality* and *collective benefit*: although we see that both would be better off cooperating i.e. $(-1, -1)$, rational self-interest leads the players to defect; since if Alice decides to coperate, then Bob's best response is defect because it gives Bob better payoff $(0 > -1)$. Similarly, if Alice decides to defect, then also Bob's best response is still defect because it gives Bob better payoff $(-3 > -5)$. Same is true vice-versa, since matrix is symmetric. Therefore, both defecting $(-3, -3)$ is the **Nash equlibrium** (defined in section 1.6) i.e. neither one of them has any better payoff by changing their strategy unilaterally. This is an example of a **non-zero-sum game** since, benefit of one is not always the loss of other.

1.5 Rational Behavior and Best Response

Rational agents choose strategies that maximize their payoffs given the opponent's choices. A best response is the strategy that yields the highest payoff under these conditions. **The basic assumption in Game Theory is that all players are rational and individual players will always choose the strategy that gives the best possible payoff** [8]. This rationality assumption allows each player to anticipate or predict the opponent's likely behavior and, in turn, plan their own best possible response accordingly. Hence, the process of strategic decision-making becomes a mutual anticipation of best responses (very much observed in sports or in military tactics), forming the foundation of equilibrium concepts such as the Nash Equilibrium.

1.6 Nash Equilibrium

Within this framework, a Nash equilibrium is defined as a strategy profile

$$s^* = (s_1^*, \dots, s_N^*) \tag{1.6}$$

satisfying the following constraints

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \forall s_i \in S_i, \ \forall i \in N, \tag{1.7}$$

where s_{-i}^* denotes the strategies of all players except player i.

Thus, a Nash equilibrium can be viewed as a point of selfish best-response, wherein each player's strategy is optimal given the strategies of all others. **Importantly, Nash equilibrium conditions impose constraints only on unilateral deviations. No requirement is placed on joint or cooperative deviations.**

As a result, Nash equilibrium need not maximize any global objective and need not be socially efficient. [8] A Nash Equilibrium is a stable outcome in which **no player can improve their payoff by unilaterally changing their strategy, assuming the other players keep their strategies fixed. In other words, each player's strategy is the best response to the strategies chosen by others**. [8]

This concept represents a point of mutual consistency — where expectations and actions align. Even though such an equilibrium does not always guarantee the collectively optimal outcome, it provides a powerful predictive tool to understand how rational individuals or systems behave when their interests are interdependent. In many cases, players' interests are aligned (as in cooperative games, for example, team games where each team within itself tries to perform towards a shared common goal), and sometimes the players' interests are conflicting (i.e. non-cooperative games), resembling a zero-sum situation.

As in our example of Rock–Paper–Scissors (Table 1.3), if Player A chooses Rock, then Player B's best possible response is Paper, since it maximizes B's payoff. Similarly, before A chooses Rock, A will reason that B, being rational, will likely choose Paper. This logical deduction leads to a circular chain of reasoning — from Rock to Paper to Scissors and then again back to Rock. So, which strategy would be executed cannot be deduced logically, but this is for *pure strategies.* [21] A *pure strategy* means anyone of the move Rock, Paper or Scissors is choosen with a probability equal to 1. This also implies that there is no pure strategy Nash Equilibrium for this game.

Table 1.3: Payoff Matrix for Rock–Paper–Scissors (for Player A)

A \ B	**Rock**	**Paper**	**Scissors**
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

However, a Nash Equilibrium does exist in the form of a *mixed strategy.* A *mixed strategy* is where one can choose a move from the available moves with some probabality. So, if Player A decides to mix its strategy between Rock, Paper, and Scissors with probabilities a, b, and c respectively, such that $a + b + c = 1$, then its expected payoff against B's choices would be the following, note that here $E_A(U)$ means the expected payoff for player A when player B executes a strategy U:

$$E_A(Rock) = 0a + 1b - 1c, \quad E_A(Paper) = -1a + 0b + 1c, \quad E_A(Scissors) = 1a - 1b + 0c.$$

Since, we already said that there is no pure strategy Nash equilibrium. We do not know what the opponent is going to play. Therefore, we would like to play such a move that whatever move the opponent chooses, it makes no difference to us i.e. no loss in terms of payoff. If we see the expected payoff equations carefully, we observe that the expected payoffs become indifferent when the probabilities are $a = b = c = \frac{1}{3}$.

We choose them to be indifferent because we cannot logically determine what B will choose — our reasoning loop prevents a definite pure choice. Therefore, there is **no pure strategy Nash Equilibrium**, but there exists a **mixed strategy Nash Equilibrium** at $a = b = c = \frac{1}{3}$. The same argument holds for Player B since the payoff matrix is symmetric. Note that in Table 3, the payoff for B is not given. The payoff for B is just the negative of that of A.

However, in the case of the **Prisoner's Dilemma**, there exists a **pure strategy Nash Equilibrium**. Let us recall the setup of the game. Two prisoners, Alice and Bob, are suspected of committing a crime together. They are interrogated separately and cannot communicate. Each has two choices: *Cooperate* (stay silent) or *Betray* (confess and implicate the other). Their possible outcomes are shown in Table 1.4.

Let us analyze their best responses:

- If B chooses to *Cooperate*, then A's best response is to *Betray* (since $0 > -1$).
- If B chooses to *Betray*, then again A's best response is to *Betray* (since $-3 > -5$).

Table 1.4: Payoff Matrix for Prisoner's Dilemma (Years in Jail; lower is better)

A \ B	**Cooperate**	**Betray**
Cooperate	(-1, -1)	(-5, 0)
Betray	(0, -5)	(-3, -3)

Hence, irrespective of what B does, A's **dominant strategy** is to *Betray*. [8] A *dominant strategy* is a strategy which gives the better payoff than any other strategy irrespective of the opponent's strategy.

By symmetry, the same reasoning holds for B. Therefore, both players choose to *Betray*, leading to the outcome $(-3, -3)$, which is the **Nash Equilibrium**. Interestingly, both players would have been better off if they had chosen to *Cooperate* (resulting in $(-1, -1)$), but rational reasoning forces them into mutual betrayal. This paradox highlights a key insight of Game Theory — **that individual rationality does not always lead to collective optimality**. And many variants of this Prisoner's Dilemma can be seen in our life, like in politics, sports, diplomacy, business, trading, biological systems, etc. Note that, Game Theory doesn't take into account the belief system of a person(s); since any belief which is devoid of any scientific reasoning or rationality is very hard to justify.

1.7 Pareto Optimality

Pareto optimality imposes a fundamentally different set of constraints on the payoff functions. **A strategy profile s^* is Pareto optimal if there does not exist any alternative strategy profile s' such that**

$$u_i(s') \geq u_i(s^*) \quad \forall i \in N, \tag{1.8}$$

$$u_j(s') > u_j(s^*) \quad \text{for at least one } j \in N. \tag{1.9}$$

Equivalently, s^* is Pareto optimal if it is impossible to improve any player's payoff without decreasing the payoff of at least one other player.[5] Unlike Nash equilibrium, Pareto optimality constrains *joint deviations* of strategies rather than unilateral ones.

The set of all Pareto-optimal strategy profiles forms the *Pareto front* in payoff space. Points on this frontier are typically obtained by solving a family of cooperative optimization problems, such as

$$\max_{\mathbf{s}} \sum_{i=1}^{N} \alpha_i u_i(\mathbf{s}), \quad \alpha_i > 0, \tag{1.10}$$

for **different weight vectors α.**

For example, referring to the payoff table of the Prisoner's Dilemma discussed in the previous section, the strategy profile yielding payoff $(-1, -1)$ is Pareto optimal point, since any deviation from that improves one player's payoff necessarily worsens the payoff of the other player. However, note that this outcome is different from $(-3, -3)$, which is Nash equilibrium. This means, in general, a Nash equilibrium is not the Pareto Optimal solution. This illustrates the fundamental distinction between local equilibrium stability (Nash equilibrium) and collective efficiency (Pareto optimality).

Pareto optimality admits a natural geometric interpretation in payoff space. Consider the mapping

$$s \mapsto \big(u_1(s), u_2(s), \ldots, u_N(s)\big) \in \mathbb{R}^N.$$

Each strategy profile s corresponds to a point in $\mathbb{R}^N$, called the *payoff vector*. [7]

The collection of all attainable payoff vectors forms the feasible payoff region. Within this region, a point is Pareto optimal if there exists no other feasible point that lies componentwise above it (with at least one strict improvement). Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\mathbf{x} \neq \mathbf{y}$. If $x_i \leq y_i$ for all $i = 1, \ldots, n$, then we say that $\mathbf{y}$ lies component-wise above $\mathbf{x}$.

The Pareto front therefore forms a boundary surface separating attainable improvements from unattainable ones. Moving along this frontier necessarily involves trade-offs: improving one player's payoff requires sacrificing the payoff of another.

The weighted optimization problem in Eq. (1.10)

$$\max_s \sum_{i=1}^{N} \alpha_i u_i(s)$$

can be interpreted geometrically as maximizing a linear functional over the feasible payoff region. Varying the weights α_i changes the orientation of the supporting hyperplane in payoff space, thereby tracing different points on the Pareto front.

Chapter 2

Classical Optimization Problems

Many problems across physics, computer science, and engineering can be formulated as optimization tasks, where the goal is to minimize or maximize an objective function over a high-dimensional space of configurations. Examples include energy minimization in physical systems, parameter estimation in machine learning, and combinatorial optimization problems in computer science. In practice, these problems often involve complex landscapes with many local minima [14], and many algorithms are typically used to navigate these landscapes and search for optimal solutions.

Game theory provides a useful conceptual framework for understanding such optimization dynamics. In this perspective, the optimization process can be viewed as the outcome of strategic interactions between different agents, each attempting to improve an objective function that depends on the joint configuration of the system. By interpreting optimization algorithms as repeated strategic interactions, one can use tools from game theory to analyze convergence properties, equilibrium structures, and the behavior of learning dynamics. A important class of models in this context is that of *potential games*, [18]. In potential games, the incentives of individual agents are aligned with a global potential function. Any unilateral improvement made by an agent corresponds to an improvement in this global objective. As a result, decentralized best-response dynamics can naturally lead to locally optimal configurations. Many distributed optimization problems can be interpreted within this framework.

Game-theoretic ideas also appear in the study of *multi-agent learning*, where multiple agents adapt their strategies through repeated interactions. [22] Learning rules such as best-response dynamics, fictitious play, and regret minimization provide decentralized mechanisms through which agents can collectively approach equilibrium states. A similar viewpoint arises in machine learning i.e. generative adversarial networks (GANs) [11] explicitly formulate the training process as a game between two competing agents, a generator and a discriminator, whose objectives are coupled. [1] The resulting dynamics can be interpreted as a search for equilibrium in a two-player game, highlighting the deep connection between optimization algorithms and strategic interaction.

In the context of quantum computing, similar ideas appear in *variational quantum algorithms* (VQAs) [4]. In these methods, a parameterized quantum circuit is optimized to minimize the expectation value of a Hamiltonian encoding a computational problem. Classical optimization routines iteratively update the parameters of the circuit in order to drive the quantum state toward low-energy configurations. [9] These algorithms can be viewed as optimization processes on the manifold of quantum states, where different components of the system may effectively play competing or cooperative roles.

The problem discussed in this work adopts a similar perspective. We study a quantum system in which two players apply unitary transformations to a shared quantum state while attempting to optimize expectation values of different operators. The resulting dynamics can be interpreted as a strategic interaction on the quantum state space. By analyzing these dynamics through a game-theoretic lens, we aim to better understand the structure of quantum optimization landscapes in such systems which may have potential application in developing better algorithms in *multi-agent reinforcement learning algorithms* and *variantional quantum algorithms (VQAs).*

2.1 SAT problem i.e. Boolean Satisfiability Problem

Given a Boolean formula $\phi(x, y)$ with variables $x, y \in \{0, 1\}$, the SAT problem asks:

> *Does there exist an assignment of variables that makes ϕ true? Which means for what value of x and y, is $\phi(x, y) = 1$?*

An example of such a formula is shown in Figure 2.1.

Consider the following example Boolean formula, it is a 3-SAT problem because there are three clauses which are dotted with each other (meaning logical and, $\wedge$)

Any k-SAT problem has k such clause, then within each clause there is logical OR among the variables, so if we want the Boolean function ϕ to be true, we have to satisfy all the clauses.

$$\phi(x, y) = (x \vee x \vee y) \ \wedge \ (\neg x \vee y \vee y) \ \wedge \ (\neg x \vee \neg y \vee \neg y). \tag{2.1}$$

- $\vee$ denotes logical OR, $\wedge$ denotes logical AND, $\neg$ denotes logical NOT.

The formula consists of three clauses, each being a disjunction of literals. The full formula is satisfied if and only if all three clauses evaluate to true.

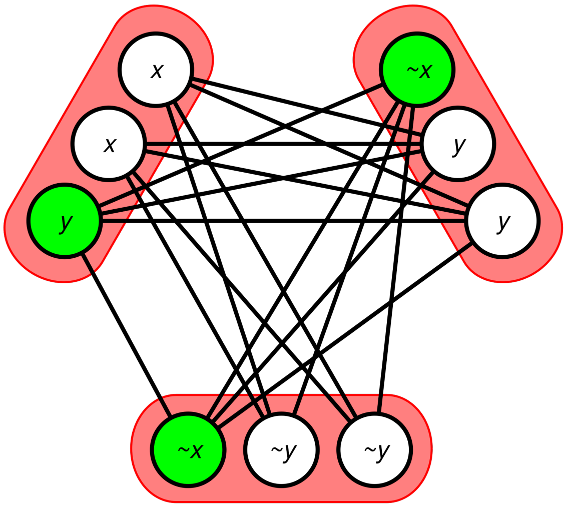

Figure 2.1: Example of a 3-SAT Boolean formula.

For example, consider the assignment:

$$x = 0, \qquad y = 1.$$

We evaluate each clause:

$$(x \lor x \lor y) = (0 \lor 0 \lor 1) = 1,$$

$$(\neg x \lor y \lor y) = (1 \lor 1 \lor 1) = 1,$$

$$(\neg x \lor \neg y \lor \neg y) = (1 \lor 0 \lor 0) = 1.$$

Thus, under this assignment, $\phi(x, y) = 1$, and the formula is satisfiable.

If someone provides such an assignment, we can check correctness in polynomial time by evaluating each clause individually. However, discovering such an assignment may require searching through up to 2^n possible assignments in the worst case. Hence, SAT $\in$ **NP**, SAT is believed to be computationally hard to solve efficiently, yet easy to verify once a candidate solution is given. [10] The Boolean Satisfiability Problem (SAT) is one of the most fundamental problems in theoretical computer science. A Boolean variable takes values in the set

$$\{0, 1\},$$

where 0 represents `False` and 1 represents `True`.

Let $x, y, z \in \{0, 1\}$ be Boolean variables.

We use three basic logical operations:

- $\neg x$: logical NOT

- $x \vee y$: logical OR
- $x \wedge y$: logical AND

Their meanings are:

x	y	$x \vee y$	$x \wedge y$	$\neg x$	$\neg y$
0	0	0	0	1	1
0	1	1	0	1	0
1	0	1	0	0	1
1	1	1	1	0	0

Table 2.1: Truth table for logical OR, AND, and negation operations.

$$x \vee y = 1 \quad \text{if at least one of } x, y \text{ is } 1,$$

$$x \wedge y = 1 \quad \text{iff both } x, y \text{ are } 1,$$

$$\neg x = 1 \quad \text{if } x = 0, \text{ and vice-versa}$$

A **clause** is a logical OR of literals.

A literal is either: a variable (e.g. x), or its negation (e.g. $\neg x$). Some example clauses:

$$(x \vee y), \qquad (\neg x \vee y), \qquad (x \vee \neg y \vee z).$$

A clause is satisfied if at least one of its literals evaluates to 1 i.e. true, because in a clause we take logical OR of literals.

Consider the clause $(x \vee y)$. This clause is satisfied for:

$$(x, y) = (1, 0), (0, 1), (1, 1),$$

and unsatisfied only for:

$$(x, y) = (0, 0).$$

Thus, a clause excludes certain assignments. A Boolean formula is said to be in **Conjunctive Normal Form (CNF)** if it is a logical AND of the clauses:

$$\phi = C_1 \wedge C_2 \wedge \cdots \wedge C_m.$$

Thus, the formula is satisfied if and only if *every clause* is satisfied.

2.1.1 Example 1: Two Clauses

Consider the two clause formula, $\phi(x, y) = (x \vee y) \wedge (\neg x \vee y)$. We check the following assignments:

- $(x, y) = (0, 0)$:
  $$(0 \vee 0) = 0, \quad (1 \vee 0) = 1,$$
  overall: $0 \wedge 1 = 0$.
- $(0, 1)$:
  $$(0 \vee 1) = 1, \quad (1 \vee 1) = 1,$$
  overall: $1 \wedge 1 = 1$.
- $(1, 0)$:
  $$(1 \vee 0) = 1, \quad (0 \vee 0) = 0,$$
  overall: $1 \wedge 0 = 0$.
- $(1, 1)$:
  $$(1 \vee 1) = 1, \quad (0 \vee 1) = 1,$$
  overall: $1 \wedge 1 = 1$.

Thus, the formula is satisfiable, since assignments $(0, 1)$ and $(1, 1)$ satisfy it.

2.1.2 Example 2: Unsatisfiable Formula

Consider $\phi(x) = (x) \wedge (\neg x)$.
If $x = 1$:

$$(1) \wedge (0) = 0.$$

If $x = 0$:

$$(0) \wedge (1) = 0.$$

Thus, no assignment satisfies the formula. This formula is **unsatisfiable**. The SAT problem asks, given a CNF formula in n Boolean variables $x_1, \ldots, x_n \in \{0, 1\}$, does there exist an assignment

$$(x_1, \ldots, x_n) \in \{0, 1\}^n$$

such that the formula evaluates to 1?

Since, it is a CNF form, ϕ is written in terms of logical AND of clauses, there are 2^n possible assignments for n variables, therefore search space is exponentially large as n increases. For small n, we can check all possibilities. But for large n, brute force becomes infeasible.

2.1.3 Why Verification is Easy

Suppose someone gives a candidate assignment x^*. To verify it, we simply evaluate each clause, then check if each clause equals 1. If all are 1, accept otherwise reject. Each clause evaluation takes constant time (if k is fixed). Total verification time is $O(m)$. Thus verification is in polynomial time. If each clause contains at most k literals, the problem is called k-SAT. So, 1-SAT is easy, 2-SAT is solvable in polynomial time and 3-SAT is NP-complete. [2] Each clause removes a region of the hypercube $\{0, 1\}^n$.

In summary, SAT is a constraint satisfaction problem:

- Variables: binary degrees of freedom,
- Constraints: clauses,
- Goal: satisfy all constraints simultaneously.

2.2 MAX k-SAT Problem and Significance

In the previous section, we said that SAT asks whether there exists an assignment that satisfies *all* clauses of a Boolean formula. This is a *decision problem*: the answer is either "yes" or "no". However, in many real-world settings, we are not merely interested in whether perfect satisfaction is possible. Instead, we often ask a more flexible question:

What is the *best* assignment we can find?

This leads naturally from decision problems to optimization problems. Let $\phi(x_1, \ldots, x_n)$ be a Boolean formula in conjunctive normal form (CNF), consisting of m clauses:

$$\phi = C_1 \wedge C_2 \wedge \cdots \wedge C_m,$$

where each clause C_j contains at most k literals.

The MAX k-SAT problem asks:

For $x_i \in \{0,1\}$. What are the maximum number of clauses which are satisfied?

For $k \geq 3$, this problem is NP-complete. Thus, unless P = NP, there is no polynomial-time algorithm that solves all instances. [2]

Now suppose the formula is unsatisfiable. The strict SAT question simply answers "no". But this provides no information about how close we can get to satisfying the formula. Given a k-CNF formula ϕ, find an assignment that maximizes the number of satisfied clauses.

Formally, define the clause indicator function

$$C_j(x) = \begin{cases} 1 & \text{if clause } C_j \text{ is satisfied by } x, \\ 0 & \text{otherwise.} \end{cases}$$

Then define the objective function

$$f(x) = \sum_{j=1}^{m} C_j(x),$$

which counts the total number of satisfied clauses. This is reminiscent of the payoff functions that we assign to players in Game Theory. [8]

The optimization problem becomes

$$\max_{x \in \{0,1\}^n} f(x). \tag{2.2}$$

Thus, MAX k-SAT transforms the binary YES/NO SAT question into a quantitative optimization problem over a discrete space of size 2^n. [10]

The domain $\{0,1\}^n$ can be viewed as the vertices of an n-dimensional hypercube. Each vertex corresponds to a possible assignment. The function $f(x)$ assigns an integer value between 0 and m to each vertex.

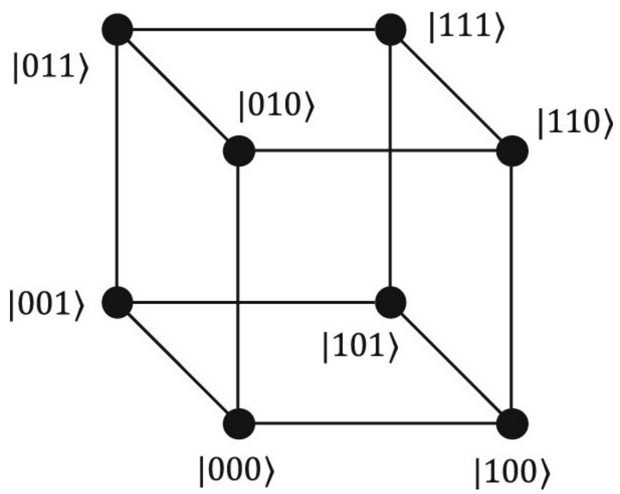

Figure 2.2: A typical 3-SAT state space representation for illustration.

Thus, MAX k-SAT defines a discrete energy landscape:

$$x \mapsto f(x).$$

The difficulty arises because the landscape is highly non-convex and also the number of candidate solutions grows exponentially. Brute-force search requires evaluating all 2^n vertices in the worst case.

For $k \geq 3$:

- **k-SAT is NP-complete and MAX k-SAT is NP-hard**. [10]

That means:

- There is no known polynomial-time algorithm that solves all instances exactly,
- An efficient algorithm for MAX k-SAT would imply $\mathbf{P} = \mathbf{NP}$.

Thus, **MAX k-SAT represents one of the canonical hard optimization problems in theoretical computer science**. [2]

MAX k-SAT is significant for several reasons, it represents a prototypical discrete non-convex optimization problem. Also, many combinatorial problems reduce to it namely graph coloring problem, knapsack problem, scheduling problem, constraint satisfaction. It naturally connects *computational complexity* with *statistical physics models.* As we will see in the next section, MAX k-SAT can be reformulated as the minimization of a classical Ising Hamiltonian. [3] This reformulation transforms a logical constraint problem into an energy minimization problem. Thus, MAX k-SAT serves as a conceptual bridge between optimization problems and statistical mechanics.

2.3 From MAX k-SAT to Ising Hamiltonians

The MAX k-SAT problem can be reformulated as the minimization of a classical Ising Hamiltonian. [3] This establishes a direct bridge between combinatorial optimization and statistical physics.

2.3.1 Boolean Variables as Spin Variables and Energy Clause Representation

Let Boolean variables $x_i \in \{0, 1\}$ be mapped to spin variables

$$s_i \in \{-1, +1\}$$

via the transformation

$$x_i = \frac{1 - s_i}{2}.$$

Under this mapping,

$$x_i = 1 \iff s_i = -1, \qquad x_i = 0 \iff s_i = +1.$$

Thus Boolean assignments correspond to classical spin configurations.

Now, consider a k-clause

$$C_j = (\ell_{j1} \vee \ell_{j2} \vee \cdots \vee \ell_{jk}),$$

where each literal ℓ_{jm} is either x_i or $\neg x_i$.

We define the clause energy

$$E_j(s) = \begin{cases} 0 & \text{if clause } C_j \text{ is satisfied,} \\ 1 & \text{otherwise.} \end{cases}$$

The total energy is

$$E(s) = \sum_{j=1}^{m} E_j(s),$$

which counts the number of violated clauses. We wish to minimize the total energy of the system.

Hence,

$$\text{MAX } k\text{-SAT} \quad \iff \quad \min_{s_i \in \{\pm 1\}} E(s).$$

2.3.2 Example: 2-SAT Clause

Consider the 2-SAT clause

$$(x_i \vee x_j).$$

This clause is violated only when

$$x_i = 0 \quad \text{and} \quad x_j = 0.$$

Using the spin mapping,

$$x_i = \frac{1 - s_i}{2},$$

we see that

$$x_i = 0 \iff s_i = +1.$$

Therefore the clause is violated only when

$$s_i = +1, \quad s_j = +1.$$

An energy term that equals 1 only in this case is

$$E_{ij}(s) = \frac{1}{4}(1 + s_i)(1 + s_j).$$

Expanding,

$$E_{ij}(s) = \frac{1}{4}\left(1 + s_i + s_j + s_i s_j\right).$$

Thus, a 2-SAT clause produces a quadratic polynomial in the spin variables.

2.3.3 General k-SAT Clause and Ising Hamiltonians

More generally, a clause

$$(x_{i_1} \vee x_{i_2} \vee \cdots \vee x_{i_k})$$

is violated only when all k literals are false. [2]

In spin variables, this can be written (up to an overall constant factor) as

$$E_j(s) \propto \prod_{\ell=1}^{k}(1 + s_{i_\ell}).$$

Expanding this product generates:

- a constant term,
- linear terms s_{i_ℓ},
- pairwise interaction terms $s_{i_\ell} s_{i_m}$,
- higher-order interactions,
- and a k-body interaction term

$$J_{i_1 i_2 \dots i_k}\, s_{i_1} s_{i_2} \dots s_{i_k}.$$

Thus each clause contributes a polynomial of degree k in the spin variables.

Collecting all clause contributions, MAX k-SAT becomes the minimization of a classical k-body Ising Hamiltonian:

$$H_{\text{classical}}(s) = \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j + \sum_{i<j<k} J_{ijk} s_i s_j s_k + \cdots . \tag{2.3}$$

Hence,

$$\text{MAX } k\text{-SAT} \quad \Longleftrightarrow \quad \min_{s_i \in \{\pm 1\}} H_{\text{classical}}(s).$$

For $k = 2$, this reduces to the familiar quadratic Ising model:

$$H = \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j,$$

which is already NP-hard in general. [15]

2.3.4 Quantum Extension

Now promoting classical spins to quantum operators,

$$s_i \to \sigma_i^z,$$

we replace classical variables $s_i \in \{-1, +1\}$ with Pauli operators acting on qubit i.

The Pauli-z matrix is

$$\sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.4}$$

with eigenvalues

$$+1, \quad -1,$$

and eigenvectors

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Thus,

$$\sigma^z|0\rangle = +|0\rangle, \qquad \sigma^z|1\rangle = -|1\rangle.$$

For n qubits, the Hilbert space is

$$\mathcal{H} = (\mathbb{C}^2)^{\otimes n}.$$

The operator σ_i^z acts as

$$\sigma_i^z = I \otimes \cdots \otimes I \otimes \sigma^z \otimes I \otimes \cdots \otimes I, \tag{2.5}$$

with σ^z in the i-th position.

Two-body interactions are represented by

$$\sigma_i^z \sigma_j^z,$$

which reproduces the classical interaction $s_i s_j$ on computational basis states.

The quantum Hamiltonian becomes

$$H_{\text{quantum}} = \sum_i h_i \sigma_i^z + \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z + \cdots . \tag{2.6}$$

The ground state energy is the lowest eigenvalue of the Hamiltonian H;

$$\lambda_{\min}(H) = \min_{\|\psi\|=1} \langle \psi | H | \psi \rangle. \tag{2.7}$$

Since global phases are physically irrelevant,

$$|\psi\rangle \sim e^{i\theta} |\psi\rangle ,$$

the optimization takes place on the complex projective space

$$\mathbb{CP}^{2^n-1}.$$

Thus, a discrete optimization over $\{\pm 1\}^n$ is transformed into a continuous but highly non-convex variational problem over a complex manifold. We therefore can map:

$$\text{MAX } k\text{-SAT problem} \iff \text{MIN } k\text{-body Ising Hamiltonian.}$$

2.4 Variational Quantum Algorithms

The mapping of the MAX k-SAT problem to an Ising Hamiltonian provides a natural pathway for using quantum systems to address combinatorial optimization problems. Once the problem is expressed as the minimization of a Hamiltonian, the task reduces to finding the ground state of that Hamiltonian. However, computing the ground state of a general many-body Hamiltonian is itself a difficult problem. [20]

Variational Quantum Algorithms (VQAs) provide a promising framework to approximately solve such optimization problems using near-term quantum devices. [4] These algorithms combine a quantum processor with a classical optimization loop and are particularly suited for the current generation of noisy intermediate-scale quantum (NISQ) devices.

2.4.1 Variational Principle

The foundation of variational quantum algorithms is the *variational principle* of quantum mechanics [19] [12], which states that for any normalized quantum state $|\psi\rangle$,

$$\langle\psi|H|\psi\rangle \geq E_0, \tag{2.8}$$

where E_0 is the ground state energy of the Hamiltonian H.

Thus, minimizing the expectation value of H over all possible quantum states will yield the ground state energy. In practice, instead of searching over the entire Hilbert space, we restrict ourselves to a parameterized family of states,

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|0\rangle\,, \tag{2.9}$$

where $U(\boldsymbol{\theta})$ is a parameterized unitary and $\boldsymbol{\theta}$ is a vector of classical parameters.

The optimization problem therefore becomes

$$\min_{\boldsymbol{\theta}}\langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle. \tag{2.10}$$

This converts the quantum ground-state search problem into a classical optimization problem over the parameter space $\boldsymbol{\theta}$.

2.4.2 Hybrid Quantum-Classical Optimization

Variational quantum algorithms operate in a hybrid quantum-classical loop [4]:

1. A parameterized quantum circuit prepares the state $|\psi(\boldsymbol{\theta})\rangle$.
2. The quantum device measures expectation values of the Hamiltonian.
3. A classical optimizer updates the parameters $\boldsymbol{\theta}$, such that it reduces the expectation value in the next measurement.
4. The process repeats until convergence.

The cost function evaluated during this process is

$$C(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle. \tag{2.11}$$

This hybrid structure is attractive because the quantum device performs the difficult task of representing high-dimensional quantum states, while the classical computer performs the parameter optimization.

2.4.3 Quantum Approximate Optimization Algorithm (QAOA)

One of the most prominent variational algorithms designed for combinatorial optimization problems is the Quantum Approximate Optimization Algorithm (QAOA). [9] In QAOA, the Hamiltonian derived from the optimization problem acts as a *cost Hamiltonian*. A second Hamiltonian, called the *mixer Hamiltonian*, is used to explore the state space in which the statevectors live.

The parameterized state is constructed as follows:

$$|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \prod_{p=1}^{P} e^{-i\beta_p H_M} e^{-i\gamma_p H_C} |+\rangle^{\otimes n}, \tag{2.12}$$

where

- H_C is the cost Hamiltonian defined for the problem instance,
- H_M is a mixing Hamiltonian for exploring the states (often $\sum_i X_i$), where X_i is the Pauli Matrices
- γ_p and β_p are variational parameters,
- P is the circuit depth.

The goal is to maximize

$$\langle H_C \rangle = \langle \psi(\boldsymbol{\gamma}, \boldsymbol{\beta})| H_C |\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle . \tag{2.13}$$

For MAX-SAT, maximizing this expectation corresponds to maximizing the number of satisfied clauses.

The mapping

$$\boldsymbol{\theta} \longrightarrow \langle \psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle$$

defines a continuous optimization landscape over the parameter space.

Unlike the original discrete search over $\{0, 1\}^n$, the variational formulation converts the problem into a continuous optimization problem over a manifold of quantum states.

However, the resulting landscape is typically highly non-convex and may contain many local minima and flat regions for a large number of parameters (often called *barren plateaus*).[**?**] Therefore, the choice of parameterization and optimization algorithm particularly important.

2.5 Motivation to our problem

Now, it seems to be an interesting question to ask, what dynamics would happen if there are two competeing objectives instead of one objective. Instead of optimizing a single Hamiltonian (i.e. $C(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle$), **we consider a scenario in which two different Hamiltonians are optimized alternately by different agents**.

Each player effectively performs a local optimization of their own cost function by applying unitary transformations to the shared quantum state. This leads to a dynamical system in the quantum state space where competing objectives interact.

In this sense, the alternating optimization dynamics studied here can be viewed as a game-theoretic generalization of variational quantum algorithms, where the cost landscape is shaped by multiple interacting objective functions (in our case, two competeing objectives) rather than a single global objective.

Chapter 3

Quantum Games and Non-Local Games

Game theory studies strategic interactions between rational agents. A classical game is defined by the tuple [16]

$$G = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}), \tag{3.1}$$

where:

- $N = \{1, 2, \dots, N\}$ is the set of players,
- S_i is the strategy set available to player i, the elements of this set alter the state of the game,
- $u_i : S_1 \times \cdots \times S_N \to \mathbb{R}$ is the payoff function of player i.

The players select strategies $s_i \in S_i$ and receive payoffs determined by the payoff function

$$u_i(s_1, s_2, \dots, s_n).$$

A central concept in classical game theory is the *Nash equilibrium*, which is a strategy profile in which no player can improve their payoff by unilaterally changing their strategy. [8] Quantum games extend this framework by allowing players to use **quantum mechanical properties** such as **superposition**, **entanglement**, and **quantum measurements**. In a quantum game the players share a quantum state $|\psi\rangle$, or equivalently in the density matrix formulation, [19] a density matrix ρ. Their strategies now correspond to quantum operations, represented by *unitary operators* or more generally, a *completely positive trace-preserving maps* in the density matrix formulation. If player i applies the unitary matrix U_i, the joint state transforms as

$$\rho' = (U_1 \otimes U_2 \otimes \dots)\rho(U_1^\dagger \otimes U_2^\dagger \otimes \dots).$$

A measurement performed on the resulting state produces classical outcomes which determine the payoffs. The expected payoff for player i is

$$\Pi_i = \sum_k p_k u_i(k)$$

where p_k are probabilities obtained from the quantum measurement outcomes.

A *quantum Nash equilibrium* occurs when no player can improve their payoff by changing their quantum strategy unilaterally,

$$\Pi_i(U_i, U_{-i}) \geq \Pi_i(U_i', U_{-i})$$

for all alternative strategies U_i'.

3.1 Quantum Refereed Games

A *quantum refereed game* is a type of game where there are two players Alice and Bob, say; and a referee. The referee sends the question to both of them and checks whether their answers are according to the rules of game or not. A particularly important class of quantum games are *non–local games*. These games are **cooperative games** in which spatially separated players attempt to maximize a joint winning probability **without communicating during the game, although they can strategize their plans before the game**.

Formally, a non–local game is defined by the tuple

$$G = (X, Y, A, B, \pi, V)$$

where

- X, Y are sets of questions sent to Alice and Bob,
- A, B are sets of answers, by Alice and Bob to the questions,
- $\pi(x, y)$ is a probability distribution over questions,
- $V(a, b|x, y)$ is a verification function determining whether the answers are accepted.

The game proceeds as follows:

1. The referee samples questions (x, y), where $x \in X$ and $y \in Y$, according to $\pi(x, y)$.

2. Alice receives x and Bob receives y.

3. Alice outputs answer $a \in A$ and Bob outputs answer $b \in B$.

4. The referee checks the answers are correct or not according to the rules of the game by $V(a, b|x, y)$.

The winning probability is

$$P_{\text{win}} = \sum_{x,y} \pi(x, y) \sum_{a,b} P(a, b|x, y)V(a, b|x, y).$$

In classical strategies the players may share classical randomness but no entanglement. The resulting correlations satisfy [12]

$$P(a, b|x, y) = \sum_{\lambda} p(\lambda)\delta_{a,f(x,\lambda)}\delta_{b,g(y,\lambda)}.$$

The maximum success probability achievable with classical resources is called the *classical value* $\omega_c(G)$.

If the players share an entangled state ρ_{AB} and perform measurements $\{M_a^x\}$ and $\{N_b^y\}$, the correlations become

$$P(a, b|x, y) = \mathrm{Tr}[\rho_{AB}(M_a^x \otimes N_b^y)].$$

The maximum success probability achievable using quantum resources is the *quantum value* $\omega_q(G)$. [6] When $\omega_q(G) > \omega_c(G)$ the game demonstrates a quantum advantage.

3.2 The Mermin–Peres Magic Square Game

+1	+1	+1
+1	1	1
1	+1	?

Figure 3.1: Magic square used in the Mermin–Peres game. Green: +1 and Red: -1

The Mermin–Peres game [17] is a well–known non–local game demonstrating a clear separation between classical and quantum strategies. The game is based on the *Mermin–Peres magic square*, which consists of a 3×3 array of observables acting on two qubits,

$$
\begin{matrix}
Z \otimes I & I \otimes Z & Z \otimes Z \\
I \otimes X & X \otimes I & X \otimes X \\
Z \otimes X & X \otimes Z & Y \otimes Y
\end{matrix}
$$

where X, Y, Z denote the Pauli matrices

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Each operator has eigenvalues ± 1. The square has the remarkable property that the products of operators along rows and columns obey fixed algebraic relations.

Using the tensor product rule

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD),$$

we compute the product of operators in each row.

Row 1

$$(Z \otimes I)(I \otimes Z)(Z \otimes Z) = (Z \otimes Z)(Z \otimes Z) = I \otimes I$$

Thus the product of operators in the first row equals $+I$.

Row 2

$$(I \otimes X)(X \otimes I)(X \otimes X) = (X \otimes X)(X \otimes X) = X^2 \otimes X^2 = I \otimes I$$

Thus the second row also multiplies to $+I$.

Row 3

$$(Z \otimes X)(X \otimes Z)(Y \otimes Y)$$

First multiply the first two operators:

$$(ZX) \otimes (XZ)$$

Using Pauli algebra

$$ZX = iY, \qquad XZ = -iY,$$

we obtain

$$(iY) \otimes (-iY) = Y \otimes Y.$$

Multiplying the third operator gives

$$(Y \otimes Y)(Y \otimes Y) = Y^2 \otimes Y^2 = I \otimes I.$$

Thus every row product equals $+I$.

Next we compute the column products.

Column 1

$$(Z \otimes I)(I \otimes X)(Z \otimes X) = Z^2 \otimes X^2 = I \otimes I$$

Column 2

$$(I \otimes Z)(X \otimes I)(X \otimes Z) = X^2 \otimes Z^2 = I \otimes I$$

Column 3

$$(Z \otimes Z)(X \otimes X)(Y \otimes Y)$$

First multiply the first two terms:

$$(ZX) \otimes (ZX)$$

Since $ZX = iY$,

$$(iY) \otimes (iY) = -Y \otimes Y.$$

Multiplying the final operator gives

$$(-Y \otimes Y)(Y \otimes Y) = -Y^2 \otimes Y^2 = -I \otimes I.$$

Thus the first two column products equal $+I$ while the final column product equals $-I$. These operator relations impose constraints on the outcomes of measurements performed on the corresponding observables. Alice and Bob play the game as follows. The referee selects a row $r \in \{1, 2, 3\}$ for Alice and a column $c \in \{1, 2, 3\}$ for Bob.

Alice outputs three bits corresponding to the entries of the chosen row,

$$(a_{r1}, a_{r2}, a_{r3}),$$

while Bob outputs three bits corresponding to the chosen column,

$$(b_{1c}, b_{2c}, b_{3c}).$$

The players win if

1. Alice's row outputs have even parity i.e.

$$a_{r1} \oplus a_{r2} \oplus a_{r3} = 0$$

2. Bob's column outputs have odd i.e.

$$b_{1c} \oplus b_{2c} \oplus b_{3c} = 1,$$

3. the outputs agree at the overlapping cell

$$a_{rc} = b_{rc}.$$

No classical assignment of predetermined values to the nine cells can satisfy both the row and column parity constraints simultaneously. Multiplying all row constraints yields $+1$, while multiplying all column constraints yields -1, producing a contradiction. Consequently classical players cannot win the game perfectly, and the optimal classical success probability is

$$\omega_c = \frac{8}{9}.$$

However, if Alice and Bob share an entangled pair of qubits and measure the corresponding observables in the magic square, the algebraic relations between the Pauli operators ensure that all parity constraints are automatically satisfied. [17] As a result the players can win the game with certainty,

$$\omega_q = 1.$$

3.3 The GHZ/Mermin Parity Game

Another important non–local game illustrating quantum advantage is the GHZ/Mermin parity game. This game is based on the correlations present in the Greenberger–Horne–Zeilinger (GHZ) state [6] and provides a deterministic contradiction between classical local hidden variable theories and quantum mechanics.

The game involves three spatially separated players: Alice, Bob, and Charlie. Each player receives a binary question

$$x, y, z \in \{0, 1\}$$

with the promise that

$$x \oplus y \oplus z = 0.$$

Thus the allowed input combinations are

$$(0,0,0),(0,1,1),(1,0,1),(1,1,0).$$

Each player outputs a binary answer

$$a,b,c \in \{0,1\}.$$

The players win if their outputs satisfy the parity condition

$$a \oplus b \oplus c = \begin{cases} 0 & \text{if } (x,y,z) = (0,0,0) \\ 1 & \text{otherwise.} \end{cases}$$

Classical Strategies

In a classical deterministic strategy each player assigns fixed outputs for each possible input. Writing these outputs as

$$A_0, A_1, \quad B_0, B_1, \quad C_0, C_1,$$

the four winning conditions become

$$A_0 \oplus B_0 \oplus C_0 = 0,$$

$$A_0 \oplus B_1 \oplus C_1 = 1,$$

$$A_1 \oplus B_0 \oplus C_1 = 1,$$

$$A_1 \oplus B_1 \oplus C_0 = 1.$$

Adding the last three equations modulo two gives

$$(A_0 \oplus B_1 \oplus C_1) \oplus (A_1 \oplus B_0 \oplus C_1) \oplus (A_1 \oplus B_1 \oplus C_0) = 1 \oplus 1 \oplus 1.$$

Simplifying the left side,

$$A_0 \oplus B_0 \oplus C_0 = 1.$$

This contradicts the first equation

$$A_0 \oplus B_0 \oplus C_0 = 0.$$

Thus no classical deterministic assignment can satisfy all four constraints simultaneously. Consequently classical players cannot win the game perfectly and the optimal classical success probability is

$$\omega_c = \frac{3}{4}.$$

Quantum Strategy

Quantum players can achieve perfect success by sharing the GHZ state

$$|\text{GHZ}\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle).$$

Each player performs a measurement depending on their input:

- If the input is 0, measure the Pauli operator X.
- If the input is 1, measure the Pauli operator Y.

The Pauli matrices are

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$

The measurement outcomes correspond to eigenvalues ± 1. We convert them to bits by the rule

$$+1 \to 0, \qquad -1 \to 1.$$

Thus the product of measurement eigenvalues determines the parity of the output bits.

Action of Pauli Operators on Basis States

The Pauli operators act on computational basis states as

$$X|0\rangle = |1\rangle, \qquad X|1\rangle = |0\rangle,$$

$$Y|0\rangle = i|1\rangle, \qquad Y|1\rangle = -i|0\rangle.$$

Using these relations we now compute the relevant operator actions on the GHZ state.

Case 1: Inputs $(0, 0, 0)$

All players measure X. The relevant operator is

$$X \otimes X \otimes X.$$

Acting on the basis states,

$$(X \otimes X \otimes X)|000\rangle = |111\rangle,$$

$$(X \otimes X \otimes X)|111\rangle = |000\rangle.$$

Therefore

$$(X \otimes X \otimes X)|\mathrm{GHZ}\rangle = \frac{1}{\sqrt{2}}(|111\rangle + |000\rangle) = |\mathrm{GHZ}\rangle.$$

Thus

$$X \otimes X \otimes X|\mathrm{GHZ}\rangle = |\mathrm{GHZ}\rangle.$$

The GHZ state is therefore an eigenstate of $X \otimes X \otimes X$ with eigenvalue $+1$. Consequently the product of measurement outcomes is always $+1$, corresponding to even parity of the output bits.

Case 2: Inputs $(0, 1, 1)$

Now Alice measures X while Bob and Charlie measure Y. The operator is

$$X \otimes Y \otimes Y.$$

Acting on the computational basis,

$$(X \otimes Y \otimes Y)|000\rangle = |1\rangle \otimes i|1\rangle \otimes i|1\rangle = -|111\rangle.$$

Similarly,

$$(X \otimes Y \otimes Y)|111\rangle = |0\rangle \otimes (-i|0\rangle) \otimes (-i|0\rangle) = -|000\rangle.$$

Therefore

$$(X \otimes Y \otimes Y)|\text{GHZ}\rangle = -\frac{1}{\sqrt{2}}(|111\rangle + |000\rangle) = -|\text{GHZ}\rangle.$$

Thus

$$X \otimes Y \otimes Y|\text{GHZ}\rangle = -|\text{GHZ}\rangle.$$

Hence the product of measurement outcomes is always -1, corresponding to odd parity of the output bits.

Remaining Cases

By symmetry of the GHZ state the remaining operators satisfy

$$Y \otimes X \otimes Y|\text{GHZ}\rangle = -|\text{GHZ}\rangle,$$

$$Y \otimes Y \otimes X|\text{GHZ}\rangle = -|\text{GHZ}\rangle.$$

Thus for all remaining allowed input combinations the measurement outcomes automatically satisfy the required parity condition.

x	y	z	Measurement Operator	Required Product
0	0	0	$X \otimes X \otimes X$	$+1$
0	1	1	$X \otimes Y \otimes Y$	-1
1	0	1	$Y \otimes X \otimes Y$	-1
1	1	0	$Y \otimes Y \otimes X$	-1

Table 3.1: Measurement operators and required outcome parity in the GHZ/Mermin game.

Table 3.1 summarizes the measurement operators corresponding to each allowed input configuration. The GHZ state is an eigenstate of these operators with eigenvalues $+1$ or -1. Consequently the product of the measurement outcomes automatically satisfies the parity condition required by the game. When the inputs are $(0, 0, 0)$ the product of outcomes is $+1$, corresponding to even parity. For the remaining inputs the product is -1, corresponding to odd parity. Thus the quantum strategy satisfies the winning conditions for every allowed input.

Because the GHZ state is an eigenstate of the above operators with eigenvalues $+1$ or -1, the product of measurement outcomes always matches the required parity relations. Consequently the players always produce outputs consistent with the rules of the game.

Therefore quantum players win the GHZ/Mermin parity game with certainty, $\omega_q = 1$. The GHZ/Mermin game provides a striking demonstration of the difference between classical and quantum correlations. As a result the GHZ/Mermin game plays a central role in the study of quantum nonlocal games.

Chapter 4

A Sequential Quantum Game on a Single-Qubit System

Even before doing the bigger problem of multi-objective quantum optimization. [13] It is important to study the simple problem and gain some intuition about formulating the framework. Therefore, we now discuss single qubit quantum game in the simplest possible setting: a single-qubit system. The goal is again to formulate a strategic interaction between two rational players whose strategies consist of unitary operations acting on a shared quantum state.

4.1 Formulating the problem

We consider a quantum system consisting of a single qubit. The Hilbert space of a single qubit is $\mathcal{H}$. A general pure state of the system is represented by a normalized vector

$$|\psi\rangle \in \mathbb{C}^2, \quad \langle\psi|\psi\rangle = 1.$$

In the computational basis $\{|0\rangle, |1\rangle\}$, the state can be written as

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle,$$

where

$$|\alpha|^2 + |\beta|^2 = 1.$$

Since global phase is physically irrelevant, the true state space of the system is the complex projective manifold $\mathbb{CP}^1$. Geometrically, this space is equivalent to the Bloch sphere. Thus the quantum state of the system corresponds to a point on the surface of the Bloch sphere, and the dynamics of the game can be interpreted as motion on this curved manifold.

4.1.1 Payoff Operators

Each player is assigned a Hermitian operator acting on the single-qubit Hilbert space:

$$H \in \mathbb{C}^{2\times 2}, \quad M \in \mathbb{C}^{2\times 2},$$

with

$$H^\dagger = H, \qquad M^\dagger = M.$$

The payoff (or cost) functions are defined as expectation values of these operators with respect to the quantum state:

$$\Pi_A(\psi) = \langle\psi|H|\psi\rangle,$$

$$\Pi_B(\psi) = \langle\psi|M|\psi\rangle.$$

These functions are real-valued functions defined on $\mathbb{CP}^1$ (the Bloch sphere).

Thus, each player's objective is an energy minimization problem:

$$\min_{|\psi\rangle}\langle\psi|H|\psi\rangle \quad \text{(Player A)},$$

$$\min_{|\psi\rangle}\langle\psi|M|\psi\rangle \quad \text{(Player B)}.$$

Importantly, both objectives depend on the same shared quantum state $|\psi\rangle$, which means the optimization goals of the players are generally coupled.

4.2 Unitary Transformations on the Qubit

The players cannot directly choose the quantum state $|\psi\rangle$. Instead, they act on the system by applying unitary transformations.

Let

$$U \in \mathcal{U}(2)$$

be a 2×2 unitary operator. When applied, the state transforms as

$$|\psi\rangle \mapsto U|\psi\rangle.$$

Under such a transformation, the expectation value of an observable transforms as

$$\langle\psi|H|\psi\rangle \mapsto \langle\psi|U^{\dagger}HU|\psi\rangle.$$

Thus, the players' strategies effectively modify the payoff landscape by conjugating the operators with their chosen unitary transformations.

Strategy Space of Player A

Player A is allowed to apply unitary transformations generated by the Pauli operators. Let σ^{α} denote the Pauli matrices

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The generators of Player A's strategy space are chosen from

$$G_a \in \{I, \sigma^x, \sigma^y, \sigma^z\}.$$

A general strategy of Player A is therefore

$$U_A(\theta) = \exp\left(-i \sum_a \theta_a G_a\right),$$

where the parameters θ_a determine the direction and magnitude of the unitary transformation.

Strategy Space of Player B

Similarly, Player B applies unitary transformations generated by the same set of Pauli operators. The strategy set of Player B therefore takes the form

$$U_B(\phi) = \exp\left(-i \sum_b \phi_b K_b\right),$$

where

$$K_b \in \{I, \sigma^x, \sigma^y, \sigma^z\}.$$

Thus both players act on the same qubit, meaning they share complete access to the quantum system.

4.3 Game Dynamics

A move in the game consists of the sequential application of the players' strategies:

$$|\psi\rangle \mapsto U_B U_A |\psi\rangle.$$

Each player chooses their parameters so as to minimize their own cost function:

$$\theta^* = \arg\min_\theta \langle\psi|U_A^\dagger H U_A|\psi\rangle,$$

$$\phi^* = \arg\min_\phi \langle\psi|U_B^\dagger M U_B|\psi\rangle.$$

Since both players act on the same qubit, their operations generally do not commute:

$$[U_A, U_B] \neq 0.$$

Consequently, the order of operations affects the resulting quantum state and the optimization dynamics. The state evolves on the Bloch sphere under the action of the Lie groups generated by the players' strategy spaces:

$$\mathcal{G}_A = \exp(\mathfrak{g}_A), \qquad \mathcal{G}_B = \exp(\mathfrak{g}_B).$$

Each player updates their parameters using a *stochastic hill-climbing* (or stochastic hill-descent) procedure. In this approach, small random perturbations of the parameters are proposed, and the update is accepted only if it decreases the corresponding cost function. Formally, if a candidate perturbation $\delta\theta$ is proposed, the update rule for Player A takes the form

$$\theta \to \theta + \delta\theta \quad \text{if } \langle H\rangle(\theta + \delta\theta) < \langle H\rangle(\theta),$$

and similarly for Player B,

$$\phi \to \phi + \delta\phi \quad \text{if } \langle M\rangle(\phi + \delta\phi) < \langle M\rangle(\phi).$$

Because the perturbations are sampled randomly from the allowed parameter directions generated by the accessible Lie algebra, the optimization proceeds as a stochastic search over the manifold of quantum states. In the single-qubit case, this corresponds geometrically to a stochastic exploration of the Bloch sphere.

4.4 Numerical Simulation Results

We start with some initial state vector on the Bloch sphere, parameterized by two real angles θ and ϕ. The general single-qubit state is parameterized as follows:

$$|\psi(\theta, \phi)\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle .$$

Here, $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$ specify the point on the Bloch sphere corresponding to $|\psi(\theta, \phi)\rangle$.

We then follow the alternating optimization scheme described earlier — where Player A and Player B take turns applying their respective unitaries from the sets U_A and U_B, each minimizing the expectation value of their Hamiltonian (H and M respectively). After every move, the updated state $|\psi_k\rangle$ evolves on the Bloch sphere.

By iterating this process, we can visualize the trajectory of the state on the Bloch sphere and track how it converges. The evolution of the state vector reflects the competition and cooperation between the two optimization objectives, much like strategic adjustments in a game between rational players.

We record and analyze quantities such as the expectation values $\langle H\rangle_k = \langle\psi_k| H |\psi_k\rangle$ and $\langle M\rangle_k = \langle\psi_k| M |\psi_k\rangle$ at each step, to study whether the system approaches a stable configuration (a local minimum or equilibrium-like state). Such convergence, when it occurs, can be interpreted as the system reaching a quantum analogue of the Nash equilibrium — where neither optimization step can further reduce its corresponding cost function.

In Fig. 4.1, we see that starting from an initial state vector near the higher eigenstate of M, the system evolves towards a cyclic trajectory that lies close to the minimum eigenstates of both H and M. Furthermore, the overlap $|\langle\psi_k|\psi_{k+1}\rangle|$ remains close to unity throughout the evolution, indicating that the change in the state $|\psi\rangle$ is quasi-static. This behavior arises because the strategy sets U_A and U_B were generated by exponentiating random Hermitian matrices as

$$U = e^{iA\epsilon}, \quad \text{where } A^\dagger = A \text{ and } \epsilon \text{ is a small real parameter.}$$

We have also included the 2×2 identity matrix in both unitary sets U_A and U_B to ensure that the option of not changing the state is always available to each player.

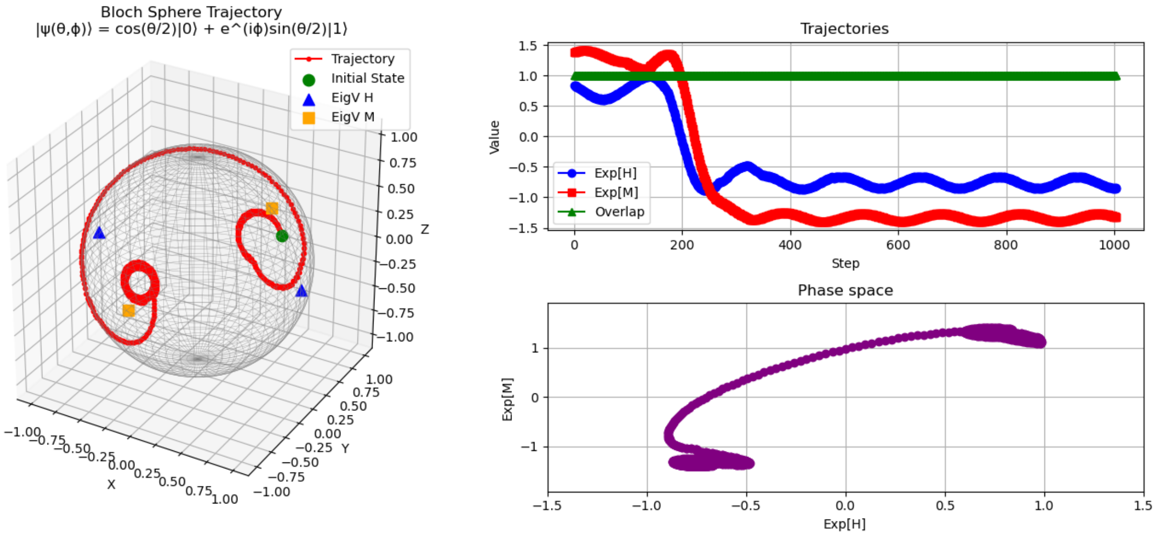

Figure 4.1: Trajectory of the state vector $|\psi(\theta, \phi)\rangle$ on the Bloch sphere. We can see that the state has gone into a cycle and not approached to fixed point. Here $n(U_A) = n(U_B) = 5$.

In Fig. 4.2, we observe a similar behavior; however, in this case, the state does not approach the minimum eigenvectors of H and M, but instead cycles around regions close to their maximum eigenvectors. It is much easier to see (θ, ϕ) plane what path the system follows. Note that, initial states to start with are all different in all of these figures.

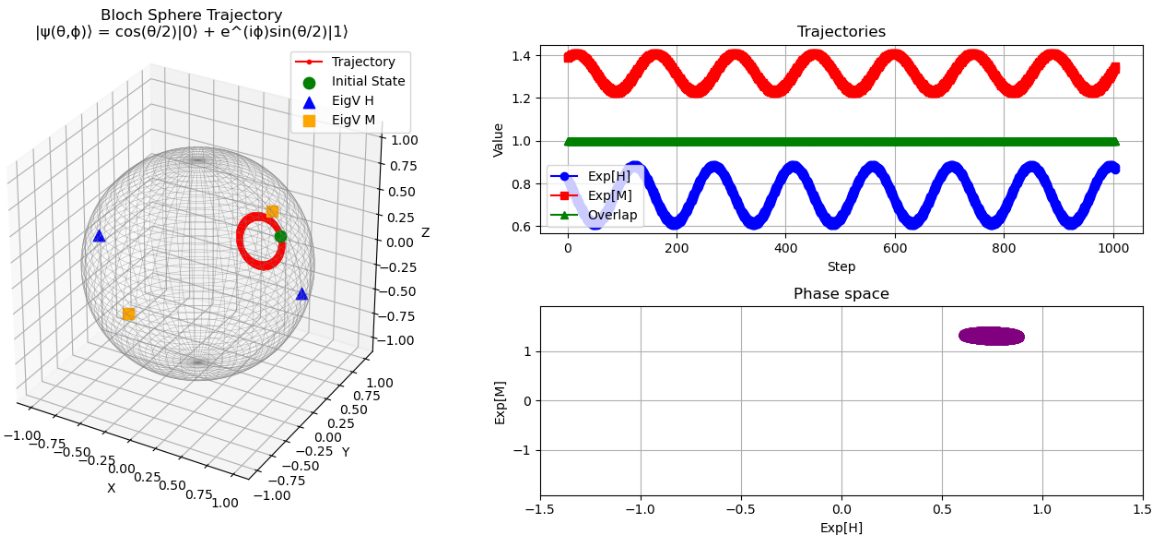

Figure 4.2: Another sample trajectory showing the cyclic behavior. Here also, $n(U_A) = n(U_B) = 5$

To understand why this behavior arises, it is better to analyze the problem in Energy Landscape, meaning we plot the $\langle H\rangle$ and $\langle M\rangle$ in (θ, ϕ) plane (see Fig. 4.3). Refer to the color-bar, blue color means negative values whereas red color means positive values. The

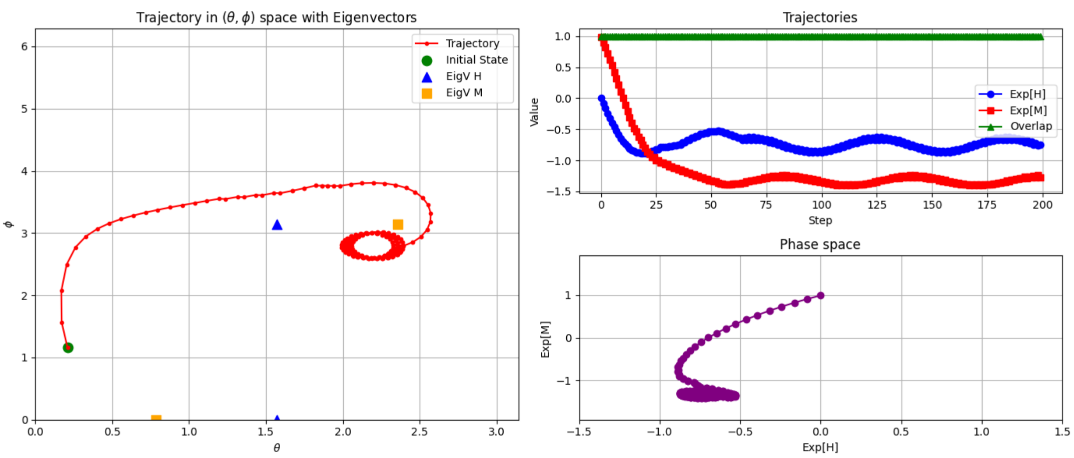

Figure 4.3: Evolution of the expectation values $\langle H \rangle$ and $\langle M \rangle$ in phase space (bottom right) and evolution of $|\psi(\theta, \phi)\rangle$ in (θ, ϕ) plane (left).

dashed line denote equal value contours. Note that we have chosen

$$H = \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{and} \quad M = \sigma_x + \sigma_z = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

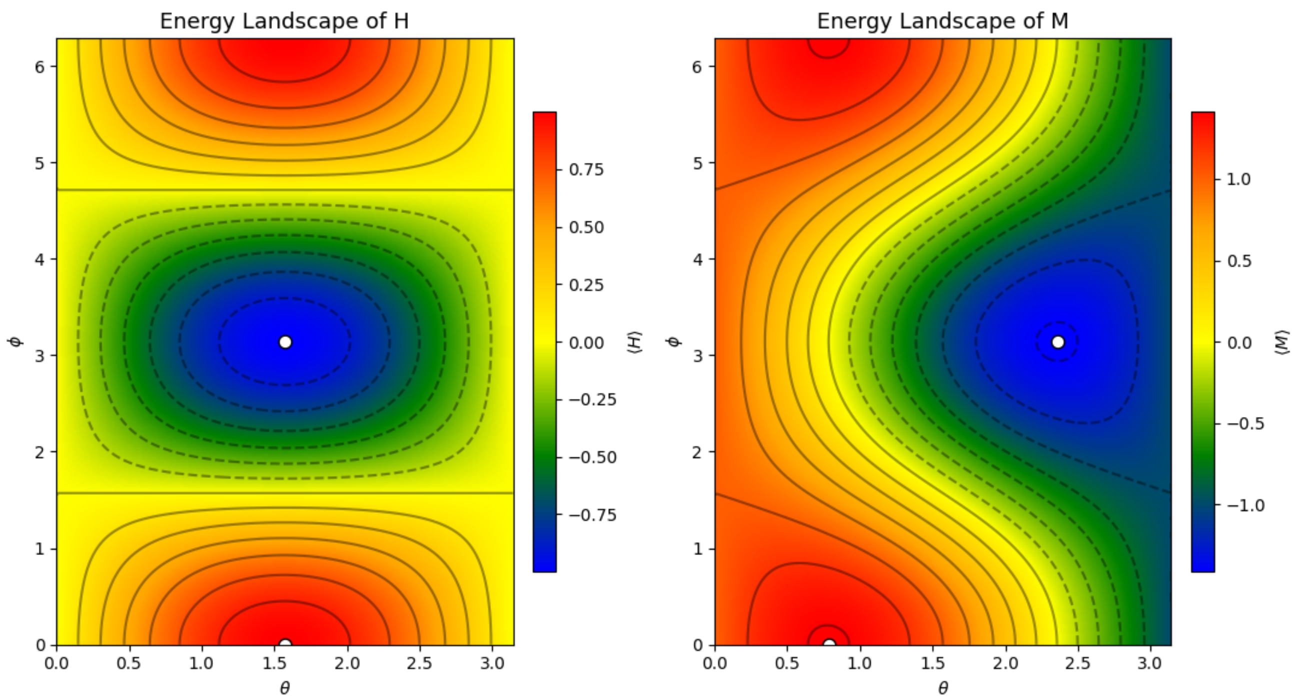

Figure 4.4: Energy Landscape for H and M i.e. $\langle H \rangle$ and $\langle M \rangle$ in (θ, ϕ) plane

The energy landscape makes it clear that we aim to compute the gradient in different possible directions and then move along the direction in which the gradient is most negative. In our setup, the available directions are determined by the allowed unitary operations, i.e., the strategy spaces U_A and U_B. These sets constrain the possible directions in which the state can evolve. If U_A and U_B equal the full unitary group $U(2)$, then all directions on the Bloch sphere would, in principle, be accessible.

It is important to note that in this case, since there are two different Hamiltonians, H and M, the energy landscape effectively changes at each step of the alternating optimization. Player A optimizes with respect to the H-landscape, while Player B optimizes with respect to the M-landscape. This introduces an interesting interplay between cooperative and non-cooperative behavior, reminiscent of classical game theory. Whether the players act in agreement or opposition depends entirely on the relation between the Hamiltonians H and M.

For instance, if $H = -M$, the two energy landscapes are exact opposites, and the players are always in conflict — a purely adversarial scenario. On the other hand, if H and M are nearly identical, differing only slightly, their energy landscapes are largely aligned, leading to mostly cooperative dynamics. It is worth noting that the overall dynamics are not determined by the commutator $[H, M]$ alone: for example, both H and $-H$ commute, yet they produce entirely opposite "tug-of-war" behaviors in this alternating optimization scheme.

Now we can also understand why the trajectories in Fig.4.1, Fig.4.2, and Fig.4.3 sometimes form closed loops instead of converging to a fixed point. This behavior arises because the available strategy spaces U_A and U_B are restricted and do not span all possible directions on the Bloch sphere or (θ, ϕ) plane, equivalently. Consequently, the players can only move within a limited subset of directions determined by these unitaries. Once the state explores all accessible directions within this constrained subspace, the path becomes cyclic, rather than leading to a fixed point.

4.5 Simulation Observations

- We observe chaotic behavior and oscillatory dynamics when the parameter ϵ is increased to large values.

- When a sufficiently large number of unitary operators (i.e., available directions in parameter space) are provided, the state evolves to a final state irrespective of the initial condition.

- The final state is not necessarily the ground state of either Hamiltonian H or M.

- Moreover, the final state is not the ground state of the combined operator, $O = H + M$.

- Interestingly, we observe a kind of **phase transition** as ϵ is varied — the system abruptly switches its behavior from converging to quasi-periodic oscillations, and then back to convergence again.

In summary, the alternating optimization between two competing objectives, H and M, gives rise to a dynamical behavior on the Bloch sphere. There is a balance of opposing drives from the two players. The restriction of the unitary strategy spaces U_A and U_B further constrains the accessible directions, effectively shaping the path in parameter space.

From a game-theoretic perspective, the system's dynamics mimic the iterative adjustment of rational players seeking their best responses under changing conditions. A fixed point in this alternating scheme corresponds to a Nash equilibrium — a state where neither player can improve their objective by unilaterally changing strategy. In contrast, the emergence of cycles signals the absence of a stable equilibrium in pure strategies, reminiscent of the mixed-strategy equilibria in classical games such as Rock–Paper–Scissors.

Extending this analogy to multi-qubit systems or more complex operator sets could provide new insight into equilibrium and convergence behavior, and the structure of optimization landscapes in quantum control and dynamics. We will investigate multi-qubit systems, where Player A has control over one subset of qubits and Player B has control over another subset, with some overlapping qubits forming a shared region.

Chapter 5

A Sequential Quantum Game on a Three-Qubit System

We now describe the three qubit quantum game studied in this work. The goal is to formulate a strategic interaction between two rational players whose strategies consist of local unitary operations acting on a shared quantum system.

5.1 The Underlying Quantum System

Let us now consider a three-qubit quantum system, whose Hilbert space is $\mathcal{H}$. A general pure state of the system is represented by a normalized vector

$$|\psi\rangle \in \mathbb{C}^8, \quad \langle\psi|\psi\rangle = 1.$$

In the computational basis,

$$|\psi\rangle = \sum_{i,j,k\in\{0,1\}} \alpha_{ijk}|ijk\rangle,$$

with

$$\sum_{i,j,k} |\alpha_{ijk}|^2 = 1.$$

5.1.1 Payoff Operators

Each player is assigned a Hermitian operator acting on the full Hilbert space:

$$H \in \mathbb{C}^{8\times 8}, \quad M \in \mathbb{C}^{8\times 8},$$

with

$$H^\dagger = H, \qquad M^\dagger = M.$$

The payoff (or cost) functions are defined as expectation values:

$$\Pi_A(\psi) = \langle\psi|H|\psi\rangle,$$

$$\Pi_B(\psi) = \langle\psi|M|\psi\rangle.$$

These are real-valued functions on $\mathbb{CP}^7$.

Thus, each player's objective is an energy minimization problem:

$$\min_{|\psi\rangle}\langle\psi|H|\psi\rangle \quad \text{(Player A)},$$

$$\min_{|\psi\rangle}\langle\psi|M|\psi\rangle \quad \text{(Player B)}.$$

Importantly, both objectives depend on the same shared state $|\psi\rangle$.

5.2 Strategic Actions as Unitary Transformations

The players cannot directly choose $|\psi\rangle$. Instead, they act on the system via unitary transformations.

Let

$$U \in \mathcal{U}(8)$$

be an 8×8 unitary operator.

When applied,

$$|\psi\rangle \mapsto U|\psi\rangle.$$

The expectation value transforms as:

$$\langle\psi|H|\psi\rangle \mapsto \langle\psi|U^\dagger H U|\psi\rangle.$$

Thus, strategy choices deform the payoff landscape via conjugation.

The important feature of this game is **partial access**.

Player A has access to qubits 1 and 2 but not qubit 3.

Player B has access to qubits 2 and 3 but not qubit 1.

Thus, the allowed strategy sets are restricted.

Strategy Space of Player A

Player A may apply unitaries generated by operators acting nontrivially only on qubits 1 and 2.

Let σ^α denote Pauli matrices:

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Then A's generators are of the form

$$P \otimes P \otimes I,$$

where $P \in \{I, \sigma^x, \sigma^y, \sigma^z\}$.

A general strategy of A is therefore

$$U_A(\theta) = \exp\left(-i \sum_a \theta_a G_a\right),$$

where

$$G_a \in \{\sigma^\alpha \otimes \sigma^\beta \otimes I\}.$$

Strategy Space of Player B

Similarly, B's generators are

$$I \otimes P \otimes P.$$

Thus,

$$U_B(\phi) = \exp\left(-i \sum_b \phi_b K_b\right),$$

with

$$K_b \in \{I \otimes \sigma^\alpha \otimes \sigma^\beta\}.$$

5.3 Game Dynamics

A move in the game consists of sequential application:

$$|\psi\rangle \mapsto U_B U_A |\psi\rangle.$$

Each player chooses parameters to minimize their own cost:

$$\theta^* = \arg\min_{\theta} \langle \psi | U_A^\dagger H U_A | \psi \rangle,$$

$$\phi^* = \arg\min_{\phi} \langle \psi | U_B^\dagger M U_B | \psi \rangle.$$

Because both act on the overlapping qubit 2, their strategy spaces do not commute in general:

$$[U_A, U_B] \neq 0.$$

Thus, the order of operations matters and nontrivial dynamics arise. The state evolves on $\mathbb{CP}^7$ under restricted Lie group actions:

$$\mathcal{G}_A = \exp(\mathfrak{g}_A), \qquad \mathcal{G}_B = \exp(\mathfrak{g}_B).$$

Each player updates their parameters using a *stochastic hill-climbing* (or stochastic hill-descent) procedure. In this approach, small random perturbations of the parameters are proposed, and the update is accepted only if it decreases the corresponding cost function. Formally, if a candidate perturbation $\delta\theta$ is proposed, the update rule for player A takes the form

$$\theta \to \theta + \delta\theta \quad \text{if } \langle H \rangle(\theta + \delta\theta) < \langle H \rangle(\theta),$$

and similarly for player B,

$$\phi \to \phi + \delta\phi \quad \text{if } \langle M \rangle(\phi + \delta\phi) < \langle M \rangle(\phi).$$

Because the perturbations are sampled randomly from the allowed parameter directions generated by the accessible Lie algebra, the optimization proceeds as a stochastic search over the manifold of quantum states.

5.4 Numerical Simulation Results

To investigate the role of controllability in the quantum game described above, we study the dynamics under three different access structures. These settings determine the set of unitary transformations available to the players and therefore strongly influence the optimization dynamics and possible equilibria of the game.

For each of the three access structures described below we perform numerical simulations of the stochastic hill-descent dynamics. The simulations track the evolution of the expectation values $\langle H \rangle$ and $\langle M \rangle$

In the following sections we present representative plots illustrating: the evolution of the players' cost functions, the convergence behavior of the optimization dynamics, and the influence of the access structure on the resulting equilibria. Comparing these three scenarios allows us to understand how the structure of the strategy spaces and the presence or absence of overlapping control affect the strategic behavior of the players in the quantum game. In all cases the players perform their strategies using the stochastic hill-descent procedure described earlier. Starting from a randomly initialized quantum state, the players iteratively update their allowed unitary parameters in order to decrease their respective cost functions. The resulting trajectories of the expectation values

$$\langle H \rangle = \langle \psi | H | \psi \rangle, \qquad \langle M \rangle = \langle \psi | M | \psi \rangle$$

are recorded and analyzed. In the following subsections we describe the three access scenarios studied in this work.

5.4.1 Partial Access i.e. Overlapping Control

This is the primary game defined in the previous section. In this scenario the players have partially overlapping access to the quantum system.

- Player A has access to qubits $(1, 2)$.
- Player B has access to qubits $(2, 3)$.

Thus both players can act on qubit 2, which creates a shared degree of control. The strategy spaces are generated by operators of the form

$$G_A = P \otimes P \otimes I, \qquad G_B = I \otimes P \otimes P,$$

where $P \in \{I, X, Y, Z\}$.

Because both players manipulate the overlapping qubit, their operations generally do not commute,

$$[U_A, U_B] \neq 0.$$

As a result, the order of moves affects the resulting quantum state and the optimization dynamics may exhibit complex behavior such as oscillations, competition between objectives.

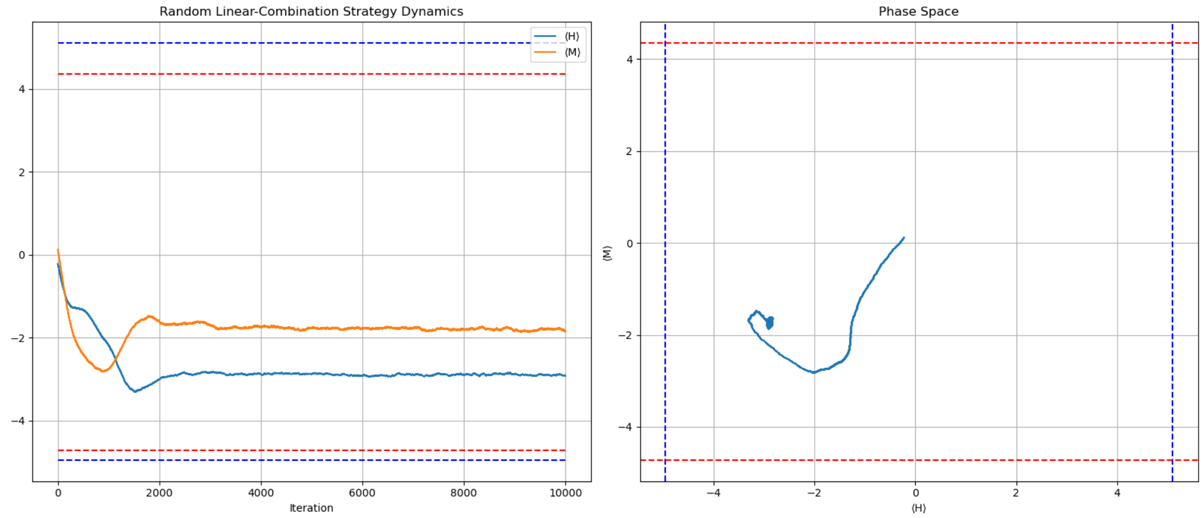

Figure 5.1: Partial Access with Player A on (1, 2) and Player B on (2, 3). Dashed blue lines are min/max eigenvalues for A's cost operator. Dashed red lines are min/max eigenvalues for B's operator. We see that that state has reached into a local minima i.e. Nash equilibrium.

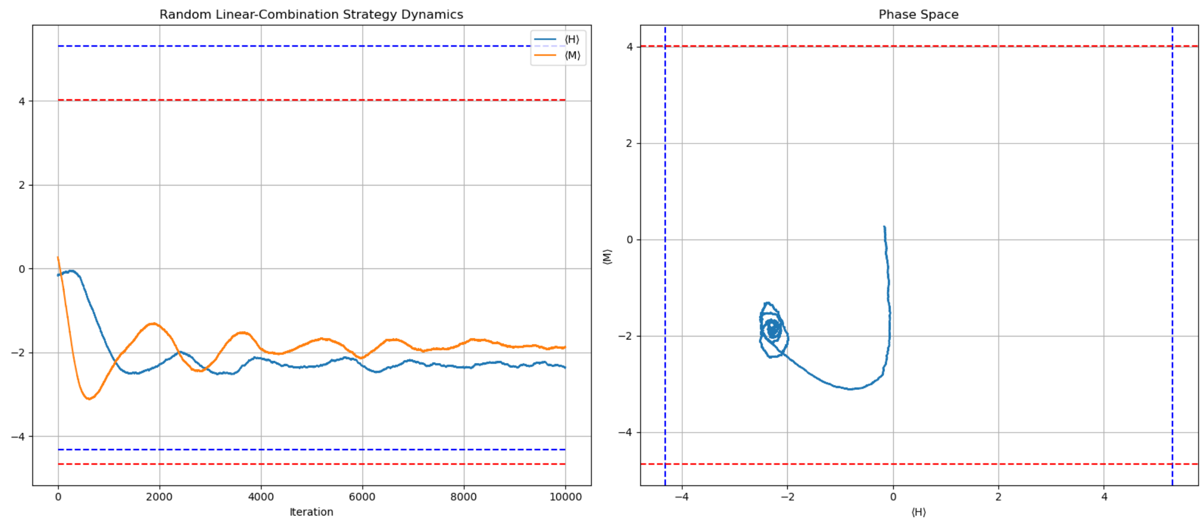

Figure 5.2: Partial Access with Player A on (1, 2) and Player B on (2, 3). Dashed blue lines are min/max eigenvalues for A's cost operator. Dashed red lines are min/max eigenvalues for B's operator. Here state is limited to region, this can also be viewd as a local minima.

This setting captures the essential game-theoretic structure of the problem, where two players attempt to optimize competing objective functions while acting on a partially shared quantum system.

5.4.2 Full Access i.e. U(8)

As a comparison benchmark we also consider a scenario in which both players have access to the entire quantum system.

In this case the allowed strategies belong to the full unitary group

$$U_A, U_B \in \mathcal{U}(8).$$

Equivalently, the generators span the full Lie algebra $\mathfrak{u}(8)$.

Under this setting each player can apply arbitrary unitary transformations to the global state. This represents a situation of complete controllability, where both players have maximal freedom to manipulate the quantum system.

Studying this case allows us to understand the behavior of the optimization dynamics when no structural constraints are imposed on the players' strategies. In particular, it provides a useful baseline against which the restricted-access scenarios can be compared.

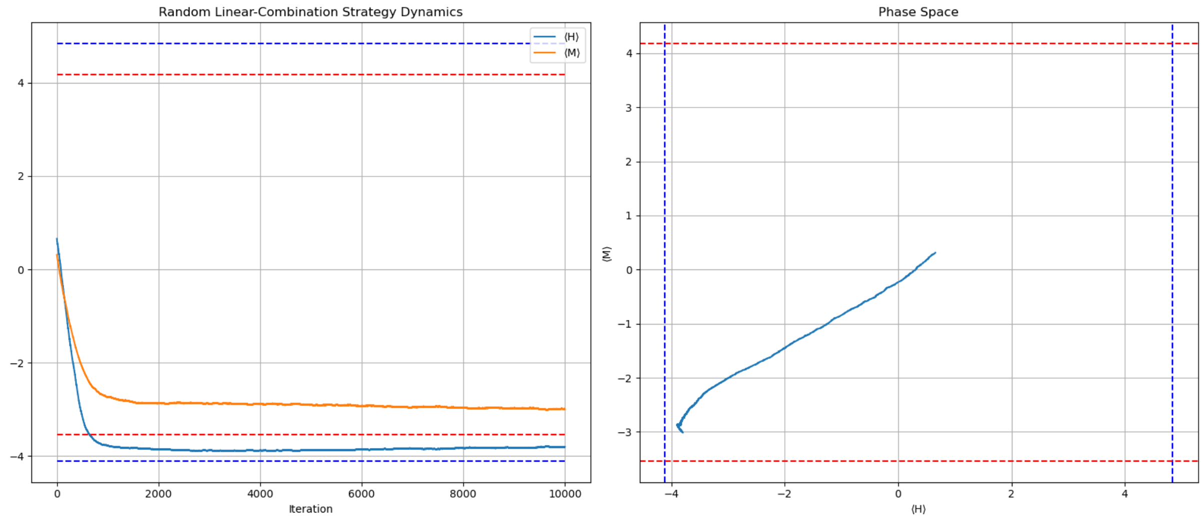

Figure 5.3: Full Access i.e. both players have access to all three qubits. We see that now the state approches much closer to the ground state of the respective cost operators.

Under this setting each player can apply arbitrary unitary transformations to the global state. This represents a situation of complete controllability, where both players have maximal freedom to manipulate the quantum system.

Studying this case allows us to understand the behavior of the optimization dynamics when no structural constraints are imposed on the players' strategies. In particular, it provides a useful baseline against which the restricted-access scenarios can be compared.

5.4.3 No Overlap

Finally we consider a scenario in which the players have completely disjoint access to the system.

- Player A controls qubit 1.

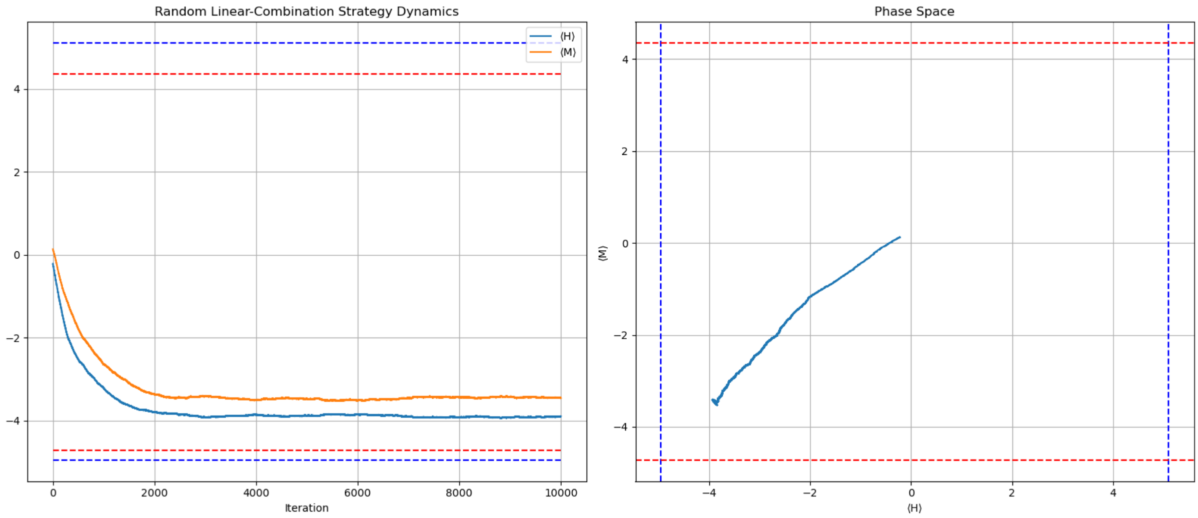

Figure 5.4: Full Access i.e. both players have access to all three qubits. We see that now the state approches much closer to the ground state of the respective cost operators.

- Player B controls qubit 3.

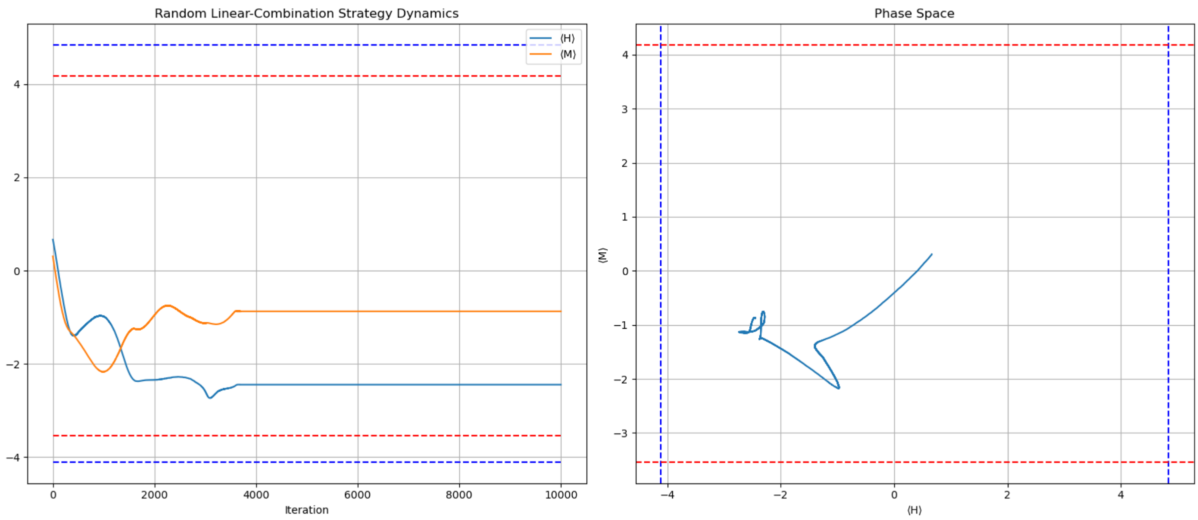

Figure 5.5: No Overlap i.e. both players do not have access to second qubit. We see that system cannot explore many states due to restriced lie algebra, also state is into Nash equilibrium.

In this setting neither player has access to qubit 2. The allowed generators take the form

$$G_A = P \otimes I \otimes I, \qquad G_B = I \otimes I \otimes P,$$

where again $P \in \{I, X, Y, Z\}$.

Because the players act on different qubits, their operations commute,

$$[U_A, U_B] = 0.$$

Consequently the order of operations does not affect the resulting quantum state. This setting therefore removes the direct competition over shared degrees of freedom and allows us to study how the absence of overlap influences the strategic dynamics of the game.

Chapter 6

Pareto-Optimal Variant: Cooperative Quantum Game

So far we have considered a competitive formulation of the quantum game, where Player A and Player B attempt to minimize different cost operators H and M respectively. In many physical and computational settings, however, it is natural to consider a cooperative scenario in which the players aim to jointly minimize multiple objective functions. In this section we formulate such a cooperative version of the quantum game and discuss its implications.

6.1 Motivation for Pareto Optimality

The motivation for studying this cooperative variant comes from a central problem in quantum physics: the search for the ground state of a Hamiltonian. [19] Determining the lowest-energy state of a quantum system is a fundamental task in low-temperature physics, condensed matter theory, and quantum chemistry. Ground states determine many important physical properties of a system, including phase structure, correlation functions, and transport behavior.

In many-body systems the Hamiltonian often consists of multiple interacting components,

$$H_{\mathrm{tot}} = H_1 + H_2 + \cdots + H_k,$$

each describing a different physical interaction or contribution to the total energy. Finding the global ground state therefore corresponds to minimizing the total energy of the system, which can be viewed as a multi-objective optimization problem. [13] From this perspective, the cooperative quantum game introduced here can be interpreted as a distributed optimization process in which different agents attempt to jointly minimize a set of cost operators acting on a shared quantum system. Such formulations are closely related to variational quantum algorithms (VQAs) [4] and quantum control problems.

6.2 Pareto Optimality in Quantum Games

In the cooperative setting, the players no longer compete. Instead, they seek states that simultaneously reduce both cost functions

$$\Pi_A(\psi) = \langle\psi|H|\psi\rangle, \qquad \Pi_B(\psi) = \langle\psi|M|\psi\rangle.$$

In general it may not be possible to minimize both quantities simultaneously. A state that minimizes H may not minimize M, and vice versa. This leads naturally to the concept of *Pareto optimality.*

A quantum state $|\psi^*\rangle$ is said to be **Pareto optimal**, if there exists no other state $|\phi\rangle$ which can be reached by applying the unitary from the strategy space such that

$$\langle\phi|H|\phi\rangle \leq \langle\psi^*|H|\psi^*\rangle$$

and

$$\langle\phi|M|\phi\rangle \leq \langle\psi^*|M|\psi^*\rangle$$

with at least one inequality being strict.

In other words, a Pareto-optimal state is one for which neither objective can be improved without worsening the other. The set of all such states forms the *Pareto frontier* in the space of expectation values

$$(\langle H\rangle, \langle M\rangle).$$

In the cooperative quantum game, the goal of the players is therefore to drive the system toward this Pareto frontier.

6.3 Cooperative Optimization Dynamics

Instead of alternating adversarial updates, the players now coordinate their actions in order to jointly reduce the cost operators. One possible strategy is to minimize a combined cost function such as

$$O = \lambda H + (1 - \lambda) M,$$

where $\lambda \in [0, 1]$ determines the relative importance of the two objectives.

The expectation value of this combined operator is

$$\langle O \rangle = \lambda \langle H \rangle + (1 - \lambda) \langle M \rangle.$$

By varying λ, one can explore different points along the Pareto frontier. [5] Alternatively, the players may perform coordinated stochastic updates that only accept moves if they improve both objectives:

$$\langle H \rangle_{\text{new}} \le \langle H \rangle_{\text{old}} \qquad \textit{and} \qquad \langle M \rangle_{\text{new}} \le \langle M \rangle_{\text{old}}.$$

This corresponds to a cooperative search over the quantum state manifold.

6.4 Access Structures in Strategy Space

As in the competitive game, the strategy space available to each player depends on the subset of qubits to which they have access. The structure of this access strongly influences the cooperative optimization dynamics.

Several natural scenarios arise.

6.4.1 Local Access Only

In the most restrictive case, each player can act only on their own qubits.

- Player A acts on subsystem A.
- Player B acts on subsystem B.

The generators i.e. lie algebra for the corrosponding unitary matrices which form the strategic space, takes the following form

$$G_A = P \otimes I, \qquad G_B = I \otimes P,$$

where P denotes Pauli operators.

In this setting the players can only perform local transformations. Cooperative optimization must therefore occur through local control of the quantum state.

6.4.2 Full Shared Access

In the most permissive case both players can act on the entire system which is a n qubit system,

$$U_A, U_B \in \mathcal{U}(2^n).$$

This corresponds to complete controllability of the quantum system. In this setting the players effectively perform a global cooperative optimization of the state.

Such a scenario is closely related to variational quantum algorithms, where a parameterized unitary circuit is optimized to minimize a cost Hamiltonian.

6.5 Connection to Ground-State Search

The cooperative quantum game can therefore be interpreted as a distributed algorithm for ground-state preparation. If the operators H and M represent different physical interactions within a many-body Hamiltonian,

$$H_{\text{tot}} = H + M,$$

then the global ground state satisfies

$$|\psi_0\rangle = \arg\min_{|\psi\rangle} \langle\psi|H_{\text{tot}}|\psi\rangle.$$

In the cooperative formulation, the players attempt to jointly reduce the expectation values of these operators. When their objectives are aligned, the optimization naturally drives the system toward the ground state of the combined Hamiltonian.

Studying this process within a game-theoretic framework provides insight into how distributed agents, each with partial control over a quantum system, can collectively navigate the energy landscapes. Such ideas may be useful for designing new algorithms for quantum control, variational ground-state search, and cooperative quantum optimization.

Chapter 7

Conclusion

In this thesis, we explored a framework that connects ideas from game theory with the dynamics of quantum systems. Motivated by the role of strategic interactions in classical decision-making problems, we formulated a quantum game in which two agents interact through a shared multi-qubit system. Each player attempts to optimize their own objective function, defined as the expectation value of a Hermitian payoff operator acting on the global Hilbert space.

The strategic actions of the players were modeled as unitary transformations applied to the quantum state. Because the players act on overlapping subsystems in certain configurations, their operations generally do not commute. As a result, the evolution of the system exhibits coupled dynamics, where the actions of one player can influence the optimization landscape experienced by the other. This leads naturally to a game-theoretic interpretation of variational quantum optimization, where competing objectives shape the trajectory of the quantum state in the high-dimensional space.

To study these dynamics, we implemented numerical simulations using a stochastic hill-descent optimization procedure. Three different access structures were considered: partial access with overlapping control of one qubit, full access in which both players can manipulate the entire system, and a no-overlap scenario where the players control disjoint subsets of qubits. These scenarios allowed us to investigate how the structure of the players' strategy spaces affects the resulting optimization behavior.

The results demonstrate that the access structure plays a significant role in determining the dynamics of the game. In the partial-access setting, the shared control of a qubit leads to genuine strategic interaction between the players, often producing complex trajectories and equilibrium-like configurations in the state space. In contrast, when both players have full access to the system, the optimization leads the state much closer to the true ground state. In the no-overlap case, the players act on independent degrees of freedom, and the dynamics become effectively decoupled, reducing the degree of strategic competition. These observations suggest that the interplay between controllability, shared resources, and competing objectives can give rise to rich dynamical behavior in quantum systems. The framework developed here therefore provides a new perspective on the study of multi-agent dynamics in quantum mechanics and highlights a natural connection between game-theoretic reasoning and variational quantum optimization.

Several directions for future work emerge from this study. One possibility is to extend the analysis to larger quantum systems with more qubits and additional players, where the structure of overlapping controls may lead to more intricate strategic interactions. Another direction is to explore alternative optimization schemes and learning algorithms, including reinforcement learning and policy-gradient methods, which may provide more efficient strategies for navigating the quantum state space. It would also be interesting to investigate analytical characterizations of equilibrium states in such quantum games and to relate them to notions of Nash equilibria and Pareto optimality in continuous strategy spaces.

More broadly, people have studied quantum dynamics in various settings such as unitary dynamics of the system, non-unitary dynamics in the context of open quantum systems, measurement-induced dynamics of the quantum system, etc. This work is one such possible approach for thinking about the quantum systems inspired by game theory. The ideas presented in this thesis may help for studying **multi-agent reinforcement learning in quantum environments**. In such settings, multiple learning agents interact through a shared quantum system, and their strategies correspond to quantum operations that modify the state of the environment. Understanding these interactions may provide new insights into the design of quantum-aware learning algorithms and the control of complex quantum systems.

Bibliography

[1] Martin Arjovsky and L'eon Bottou. Towards principled methods for training generative adversarial networks. *International Conference on Learning Representations (ICLR) Workshop*, 2017.

[2] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, 2009.

[3] Stefan Boettcher. Analysis of the relation between quadratic unconstrained binary optimization and the spin-glass ground-state problem. *Physical Review Research*, 1(3):033142, 2019.

[4] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3:625–644, 2021.

[5] Altannar Chinchuluun and Panos M. Pardalos. A survey of recent developments in multiobjective optimization. *Annals of Operations Research*, 154(1):29–50, October 2007.

[6] Richard Cleve, Peter Hoyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. *Proceedings of the 19th IEEE Conference on Computational Complexity*, pages 236–249, 2004.

[7] Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms.* John Wiley & Sons, 2001.

[8] Avinash K. Dixit, Susan Skeath, and David H. Reiley. *Games of Strategy.* W. W. Norton & Company, 3 edition, 2008.

[9] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint*, 2014.

[10] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27, 2014.

[12] David J. Griffiths and Darrell F. Schroeter. *Introduction to Quantum Mechanics*. Cambridge University Press, 3 edition, 2018.

[13] Ayse Kotil, Elijah Pelofske, Stephanie Riedmüller, Daniel J. Egger, Stephan Eidenbenz, Thorsten Koch, and Stefan Woerner. Quantum approximate multi-objective optimization. *npj Quantum Information*, October 2025.

[14] Martín Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 2025.

[15] Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2:5, 2014.

[16] Vladimir Mazalov. *Mathematical Game Theory and Applications*. John Wiley & Sons, 2014.

[17] N. David Mermin. Extreme quantum entanglement in a superposition of macroscopically distinct states. *Physical Review Letters*, 65(15):1838–1840, 1990.

[18] Dov Monderer and Lloyd S Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[19] J. J. Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Addison-Wesley, 2 edition, 2010.

[20] Norbert Schuch, Ignacio Cirac, and Frank Verstraete. Computational difficulty of finding matrix product ground states. *Physical Review Letters*, 100(25), June 2008.

[21] Saul Stahl. *A Gentle Introduction to Game Theory*, volume 13 of *Mathematical World*. American Mathematical Society, 1999.

[22] Thet Htar Su, Shaswot Shresthamali, and Masaaki Kondo. Quantum framework for reinforcement learning: Integrating the markov decision process, quantum arithmetic, and trajectory search. *Physical Review A*, 111(6), June 2025.