Lecture 4: Min-max theorem for zero-sum quantum games, maximum output probabilities, an efficient min-max algorithm

We begin this lecture with another application of the link product and characterization of strategies: the min-max theorem for zero-sum quantum games. This theorem asserts the existence of *optimal strategies* for the two players in a zero-sum quantum game.

We then analyze the SDP for maximum output probabilities introduced at the end of Lecture 3 and derive a useful formula for this quantity. This formula allows us to frame the problem of finding optimal strategies for players in a zero-sum quantum game as an SDP, from which we obtain an efficient min-max *algorithm* for finding optimal strategies.

1 Min-max theorem for zero-sum quantum games

Game theory is the study competition and cooperation among independent entities. A very simple type of game is the *normal form* game, in which two players P_0, P_1 each select a strategy s_0, s_1 among a discrete set of strategies (usually a finite set) and are each awarded a payout $v_0(s_0, s_1), v_1(s_0, s_1)$. Such a game can be specified by two matrices V_0, V_1 whose rows and columns are indexed by all possible strategies s_0, s_1 of the players and whose entries are the payouts associated with those strategies.

A normal form game is *zero-sum* if $v_1 = -v_0$, in which case the game is a competition between the two players—Player 1's gain is player 0's loss. Write $v = v_1$, so that player 1 wishes to maximize v while player 0 wishes to minimize v.

Von Neumann's famous min-max theorem for zero-sum games (established in 1928) asserts that there exist *optimal* strategies for each player, independent of the strategy chosen by the other. That is, each zero-sum game v has a *value* λ_v such that there exist a strategies $s_0^\star, s_1^\star$ such that

$$\begin{aligned} v(s_0^\star, s_1) &\leq \lambda_v \qquad \forall s_1 \\ v(s_0, s_1^\star) &\geq \lambda_v \qquad \forall s_0 \end{aligned}$$

This result could also be written

$$\min_{s_0} \max_{s_1} v(s_0, s_1) = \max_{s_1} \min_{s_0} v(s_0, s_1) = \lambda_v.$$

This min-max theorem has been extended to more general classes of games. For example, one could consider a game in which two players exchange *multiple rounds* of messages with a *referee*, who declares payouts at the end of the interaction. Games of this form are called *refereed games*.

Any refereed game can be simulated (albeit with exponential blow-up) by a normal form game and so von Neumann's min-max theorem applies also to refereed games.

Simliar min-max theorems have been proven in analysis. For example, if A, B are compact convex sets and $f : A \times B \to \mathbb{R}$ is bilinear then there exists a value λ_f with

$$\min_{a \in A} \max_{b \in B} f(a, b) = \max_{b \in B} \min_{a \in A} f(a, b) = \lambda_f.$$

Using this result, one could extend von Neumann's min-max theorem to games in which the referee and players may act probabilistically, as opposed to deterministically.

Refereed games generalize in a straightforward way to *quantum refereed games*, in which a the players and referee are free to manipulate and exchange quantum information instead of classical information.

One expects an analogous min-max theorem to hold. But such a theorem was not proven until the advent of quantum strategies.

Any quantum refereed game can be assumed to take the following form.

[*** draw a picture. ***]

Player 0 implements an r-round strategy for input spaces $\mathcal{A}_1, \ldots, \mathcal{A}_r$ and output spaces $\mathcal{C}_1, \ldots, \mathcal{C}_r$. Player 1 implements an r-round strategy for input spaces $\mathcal{B}_1, \ldots, \mathcal{B}_r$ and output spaces $\mathcal{D}_1, \ldots, \mathcal{D}_r$. The referee implements an r-round measuring co-strategy $\{R_a\}$ for input spaces $\mathcal{A}_1 \otimes \mathcal{B}_1, \ldots, \mathcal{A}_r \otimes \mathcal{B}_r$ and output spaces $\mathcal{C}_1 \otimes \mathcal{D}_1, \ldots, \mathcal{C}_r \otimes \mathcal{D}_r$.

According to the link product, the probability with which the referee produces outcome a given that player 0 acts according to A and player 1 acts according to B is $\langle R_a, A \otimes B \rangle$.

Associated with each outcome a is a payout $v(a)$ awarded to player 1. Letting $R = \sum_a v(a) R_a$, the expected payout to player 1 is thus

$$\sum_a v(a) \langle R_a, A \otimes B \rangle = \langle R, A \otimes B \rangle.$$

(The Hermitian operator R could be called a *payoff observable*.) Because $\langle R, A \otimes B \rangle$ is a real-valued bilinear function of A, B and because the sets of all strategies for the players are compact and convex (a simple consequence of the characterization of strategies) it follows that

$$\min_A \max_B \langle R, A \otimes B \rangle = \max_B \min_A \langle R, A \otimes B \rangle,$$

which is precisely the statement of a quantum min-max theorem.

Theorem 1 (Min-max theorem for zero-sum quantum games)**.** *Let R be any payoff observable as defined above. There exists a value $\lambda_R \in \mathbb{R}$ and optimal strategies $A^\star, B^\star$ for the players such that*

$$\begin{aligned} \langle R, A^\star \otimes B \rangle &\leq \lambda_R \qquad \forall B \\ \langle R, A \otimes B^\star \rangle &\geq \lambda_R \qquad \forall A \end{aligned}$$

Equivalently,

$$\min_A \max_B \langle R, A \otimes B \rangle = \max_B \min_A \langle R, A \otimes B \rangle = \lambda_R.$$

2 Formula for maximum output probabilities

At the end of Lecture 3 we noted that the link product and characterization of strategies allow us to apply algorithms for solving SDPs to the task of computing maximum output probabilities.

Specificlly, given an r-round measuring strategy $\{Q_a\}$ we were interested in the problem

$$\begin{array}{ll} \text{maximize} & \Pr[\{Q_a\} \text{ produces outcome } a \text{ when interacting with } R] \\ \text{subject to} & R \text{ is an } r\text{-round non-measuring co-strategy} \end{array}$$

We used the link product and characterization of strategies to express this problem as an SDP

$$\begin{array}{ll} \text{maximize} & \mathrm{Tr}(Q_a R) \\ \text{subject to} & R = I_{\mathcal{Y}_r} \otimes R_r \\ & \mathrm{Tr}_{\mathcal{X}_r}(R_r) = I_{\mathcal{Y}_{r-1}} \otimes R_{r-1} \\ & \quad \vdots \\ & \mathrm{Tr}_{\mathcal{X}_2}(R_2) = I_{\mathcal{Y}_1} \otimes R_1 \\ & \mathrm{Tr}_{\mathcal{X}_1}(R_1) = 1 \\ & R_i \in \mathrm{Pos}(\mathcal{Y}_{1\ldots i-1} \otimes \mathcal{X}_{1\ldots i}) \quad (1 \le i \le r) \\ & R \in \mathrm{Pos}(\mathcal{Y}_{1\ldots r} \otimes \mathcal{X}_{1\ldots r}) \end{array}$$

A fruitful concept in the study of SDPs is *duality*. We will not discuss SDP duality in detail here. Interested readers are referred to the lecture notes of John Watrous for a good introduction to SDPs and duality.

Suffice it to say that associated with every SDP is another SDP called the *dual* of the original SDP. (The "original" SDP is often called *primal* in this context. Use of primal-dual terminology is well-motivated. *e.g.* the dual of the dual is the primal, *etc.*)

Given the primal it is "easy" to compute the dual, meaning that there are fast algorithms that compute the dual. Computing dual SDPs by hand can be laborious, as in the present case. I will spare you the details.

The dual of the above SDP is

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & S_r \succeq Q_a \\ & \mathrm{Tr}_{\mathcal{Y}_r}(S_r) = I_{\mathcal{X}_r} \otimes S_{r-1} \\ & \quad \vdots \\ & \mathrm{Tr}_{\mathcal{Y}_2}(S_2) = I_{\mathcal{X}_2} \otimes S_1 \\ & \mathrm{Tr}_{\mathcal{Y}_1}(S_1) = p I_{\mathcal{X}_1} \\ & p \ge 0, S_i \in \mathrm{Pos}(\mathcal{Y}_{1\ldots i} \otimes \mathcal{X}_{1\ldots i}) \quad (1 \le i \le r) \end{array}$$

which can be written more succinctly as

$$\begin{array}{ll} \text{minimize} & p \\ \text{subject to} & pS \succeq Q_a \\ & p \ge 0,\ S \text{ is an } r\text{-round non-measuring strategy} \end{array}$$

Every SDP and its dual obeys a fascinating condition known as *weak duality*: if the primal is a maximization problem with optimal value p then the dual is a minimization problem with optimal value $d \ge p$. (And *vice versa* if the primal is a minimization problem.)

Thus, any feasible solution T to a dual SDP with objective value $d(T)$ serves as a succinct "certificate" that the optimal value p of the primal is at most $d(T)$.

Under certain conditions and SDP and its dual will obey an even stronger condition known as *strong duality*: the optimal primal and dual values p, d are *equal*.

It just so happens that the above primal-dual SDPs meet the conditions for strong duality. We have the following theorem.

Theorem 2 (Formula for maximum output probabilities)**.** *Let $\{Q_a\}$ be an r-round measuring strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$. For each outcome a the maximum probability with which $\{Q_a\}$ can be forced to produce a after an interaction with an r-round co-strategy for the same input and output spaces is equal to the* minimum *p for which there exists an r-round non-measuring strategy S for the same input and output spaces with $Q_a \preceq pS$. In symbols,*

$$\max_{\textit{co-strategies } R} \langle Q_a, R\rangle = \min \{p \geq 0 : Q_a \preceq pS \textit{ for some strategy } S\}$$

3 Efficient min-max algorithm

Theorem 2 can be used to phrase the value of a zero-sum quantum game as and SDP. Then, by the existence of efficient algorithms for SDPs, we claim that there is an efficient, classical, deterministic algorithm for computing the value of a zero-sum quantum game (and for the associated search problem of *finding* optimal strategies for the players).

Let us start with a somewhat restricted form of zero-sum game called a *win-lose* game. In this game the referee is a two-outcome r-round measuring co-strategy $\{R_{\text{Alice}}, R_{\text{Bob}}\}$ with the referee's outcome dictating the winner of the game.

We are interested in computing the probability with which the referee declares, say, Bob the winner. In order to use Theorem 1 we define a payout function $v(\text{Alice}) = 0, v(\text{Bob}) = 1$ so that the payout observable R obeys

$$R = v(\text{Alice})R_{\text{Alice}} + v(\text{Bob})R_{\text{Bob}} = R_{\text{Bob}}.$$

By Theorem 1 we know that there exists a fixed value λ_{Bob} with

$$\lambda_{\text{Bob}} = \min_A \max_B \langle R_{\text{Bob}}, A \otimes B\rangle.$$

In the case of a win-lose game, λ_{Bob} is simply the maximum probability with which Bob can force the referee to declare him the winner, taken over all possible strategies of Alice. We wish to compute λ_{Bob}.

Using the link product we can bundle any fixed strategy A for Alice together with the referee to form an r-round measuring co-strategy $\{R_{\text{Alice}} * A, R_{\text{Bob}} * A\}$ for input spaces $\mathcal{B}_1, \dots, \mathcal{B}_r$ and output spaces $\mathcal{D}_1, \dots, \mathcal{D}_r$. That is, we may view the three-party interaction between the referee, Alice, and Bob as a two-party interaction between the referee-Alice combination and Bob.

Under this notation, we have

$$\lambda_{\text{Bob}} = \min_A \max_B \langle R_{\text{Bob}} * A, B\rangle$$

For each fixed strategy A for Alice we know from Theorem 2 that

$$\max_B \langle R_{\text{Bob}} * A, B\rangle = \min \{p : R_{\text{Bob}} * A \preceq pT \text{ for some } T\}$$

and so

$$\lambda_{\text{Bob}} = \min\{p : R_{\text{Bob}} * A \preceq pT \text{ for some } T, A\}$$

where the minimum is taken over $p \geq 0$, r-round co-strategies T for input spaces $\mathcal{B}_1, \ldots, \mathcal{B}_r$ and output spaces $\mathcal{D}_1, \ldots, \mathcal{D}_r$ (*i.e.* the referee-Alice party), and r-round strategies A for input spaces $\mathcal{A}_1, \ldots, \mathcal{A}_r$ and output spaces $\mathcal{C}_1, \ldots, \mathcal{C}_r$ (*i.e.* Alice).

This formula for λ_{Bob} is an SDP. More explicitly:

$$\begin{array}{ll} \text{minimize} & \operatorname{Tr}(T_1) \\ \text{subject to} & R_{\text{Bob}} * A_r \preceq T_r \otimes I_{\mathcal{D}_r} \\ & \operatorname{Tr}_{\mathcal{B}_r}(T_r) = I_{\mathcal{D}_{r-1}} \otimes T_{r-1} \\ & \quad \vdots \\ & \operatorname{Tr}_{\mathcal{B}_2}(T_2) = I_{\mathcal{D}_1} \otimes T_1 \\ & T_i \in \operatorname{Pos}(\mathcal{D}_{1\ldots i-1} \otimes \mathcal{B}_{1\ldots i}) \quad (1 \leq i \leq r) \\ & \operatorname{Tr}_{\mathcal{C}_r}(A_r) = I_{\mathcal{A}_r} \otimes A_{r-1} \\ & \quad \vdots \\ & \operatorname{Tr}_{\mathcal{C}_2}(A_2) = I_{\mathcal{A}_2} \otimes A_1 \\ & \operatorname{Tr}_{\mathcal{C}_1}(A_1) = I_{\mathcal{A}_1} \\ & A_i \in \operatorname{Pos}(\mathcal{C}_{1\ldots i} \otimes \mathcal{A}_{1\ldots i}) \quad (1 \leq i \leq r) \end{array}$$

Thus, the value λ_{Bob} of an r-round win-lose quantum game can be computed efficiently by algorithms for SDPs. Indeed, those algorithms for SDPs can also solve the related search problem of *finding* feasible solutions to the SDP that achieve the optimal value. Those solutions indicate optimal strategies for the players.

Perhaps the easiest way to extend the above analysis from the special case of win-lose games to fully general zero-sum games is as follows.

Given an arbitrary Hermitian payout observable R, translate and scale R to obtain a different payout observable R' with $0 \preceq R' \preceq Q$ for some referee co-strategy Q. (A "referee" co-strategy is an r-round non-measuring co-strategy for input spaces $\mathcal{A}_1 \otimes \mathcal{B}_1, \ldots, \mathcal{A}_r \otimes \mathcal{B}_r$ and output spaces $\mathcal{C}_1 \otimes \mathcal{D}_1, \ldots, \mathcal{C}_r \otimes \mathcal{D}_r$.)

By the characterization of measuring strategies from Lecture 3, we know that $\{R', Q - R'\}$ is a referee measuring co-strategy. As such, we know that the above analysis applies with R' in place of R_{Bob}. In particular, efficient SDP solvers can be used to compute the value λ' and find optimal strategies $A^\star, B^\star$ for the players.

These optimal strategies $A^\star, B^\star$ are also optimal strategies for the original game with payout observable R. To recover the optimal value λ of the original game, simply invert the translation and scaling operations used to construct R' from R and apply that inverted operation to λ'.