PROCEEDINGS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 139, Number 12, December 2011, Pages 4423–4434
S 0002-9939(2011)10838-4

NASH EQUILIBRIA IN QUANTUM GAMES

STEVEN E. LANDSBURG

(Communicated by Richard C. Bradley)

Abstract. For any two-by-two game $\mathbf{G}$, we define a new two-player game $\mathbf{G}^Q$. The definition is motivated by a vision of players in game $\mathbf{G}$ communicating via quantum technology according to the protocol introduced by J. Eisert and M. Wilkins.

In the game $\mathbf{G}^Q$, each player's (mixed) strategy set consists of the set of all probability distributions on the 3-sphere $\mathbf{S}^3$. Nash equilibria in the game can be difficult to compute.

Our main theorems classify all possible mixed-strategy equilibria. First, we show that up to a suitable definition of equivalence, any strategy that arises in equilibrium is supported on at most four points; then we show that those four points must lie in one of a small number of allowable geometric configurations.

Introduction

Quantum game theory models the behavior of strategic agents with access to quantum technology. For example, agents might use observations of entangled particles to randomize their strategies as in [CHTW] or [DL]; alternatively they might use quantum devices to communicate with each other or with the referee.

Among models of quantum communication, the most widely studied is the *EWL model* ([EW], [EWL]) of Eisert, Wilkens and Lewenstein. Here we start with an ordinary game $\mathbf{G}$, envision players who communicate with a referee via a specific quantum protocol, and, motivated by this vision, construct a new game $\mathbf{G}^Q$ with greatly enlarged strategy spaces.

If $\mathbf{G}$ is a two-by-two game (that is, a game with two players, each of whom has a two point strategy space), then the game $\mathbf{G}^Q$ has strategy spaces that are naturally identified with the 3-sphere $\mathbf{S}^3$. A mixed strategy in the game $\mathbf{G}^Q$ is therefore an arbitrary probability distribution on $\mathbf{S}^3$.

Partly because the space of mixed strategies is so large, mixed-strategy Nash equilibria in the game $\mathbf{G}^Q$ can be difficult to find. The present paper ameliorates this difficulty by proving that up to a natural notion of equivalence, all Nash equilibria have a particularly simple form. More specifically (and still up to equivalence, as defined in Section 1):

1) Every mixed strategy that occurs in a Nash equilibrium is supported on at most four points of $\mathbf{S}^3$.

2) Those four points must lie in one of a small number of allowable geometric configurations.

Received by the editors October 24, 2009 and, in revised form, October 17, 2010.
2010 *Mathematics Subject Classification.* Primary 91A05, 81P45.

A two-by-two game is called "generic" if, for each player, the four possible payoffs are all distinct and the six pairwise sums of those payoffs are all distinct. In this paper, we will state and prove the main theorem for generic games, referring the reader to the author's unpublished working paper [NE] for the (considerably uglier but no more enlightening) generalization to the nongeneric case.

Section 1 lays out the motivation and the details of the EWL model. Section 2 presents the main technical lemmas. Section 3 contains the main classification theorem (3.3). Section 4 addresses some natural questions raised by the statement of the main theorem. Section 5 collects a few additional remarks and applications; the most striking is that in any mixed strategy quantum equilibrium of any two-by-two zero sum game, each player earns exactly the average of the four possible payoffs.

1. The Eisert-Wilkens-Lewenstein model

Let $\mathbf{G}$ be a two-player game with strategy sets $\mathcal{S}_1$, $\mathcal{S}_2$. The EWL construction is motivated by specifying a purely classical communication protocol: A referee issues each player a penny in one of two states, $\mathbf{H}$ ("heads up") or $\mathbf{T}$ ("tails up"). A player indicates his choice of strategy by returning his penny either flipped or unflipped. The referee observes the returned pennies and computes payoffs accordingly.

Now replace the pennies with subatomic particles whose state spaces are complex vector spaces with basis $\{\mathbf{H},\mathbf{T}\}$. A state is a line through the origin, which we will often denote by specifying some nonzero point on that line. The state space for the pair of particles is the tensor product of the two individual state spaces.

The referee prepares two pennies in the *maximally entangled state*

$$\mathbf{H}\otimes\mathbf{H}+i\,\mathbf{T}\otimes\mathbf{T}.$$

Each player returns the penny after acting on its state by the special unitary operator of his choice. The two classical strategies (i.e. the strategies in the game $\mathbf{G}$) are identified with the operators

$$\mathbf{C}=\begin{pmatrix}1&0\\0&1\end{pmatrix},\qquad \mathbf{D}=\begin{pmatrix}0&1\\-1&0\end{pmatrix}. \tag{1.0.1}$$

If Players One and Two select the unitary operators $\mathbf{U}$ and $\mathbf{V}$, we denote the resulting state by

$$\mathbf{UV}=(\mathbf{U}\otimes\mathbf{V})\ \ (\mathbf{H}\otimes\mathbf{H}+\mathbf{T}\otimes\mathbf{T}). \tag{1.0.2}$$

The referee then performs an observation with the four possible outcomes $\mathbf{CC}$, $\mathbf{CD}$, $\mathbf{DC}$ and $\mathbf{DD}$, and makes payoffs accordingly; if we write

$$\mathbf{UV}=\alpha_1\mathbf{CC}+\alpha_2\mathbf{DD}+\alpha_3\mathbf{DC}+\alpha_4\mathbf{CD} \tag{1.0.3}$$

(with complex scalar coefficients), then the probabilities of the four states are proportional to $|\alpha_1|^2,|\alpha_2|^2,|\alpha_3|^2,|\alpha_4|^2$.

Now identify Player One's strategy space with the unit quaternions by mapping the unitary matrix with top row (A,B) to the quaternion $A+Bj$; identify Player Two's strategy space with the unit quaternions by mapping the unitary matrix with top row (P,Q) to the quaternion $P+kQ$. From (1.0.2) one readily calculates the coefficients in (1.0.3) and discovers the following remarkably simple formula:

Proposition 1.1. *Suppose Player One plays the quaternion* $\mathbf{p}$ *and Player Two plays the quaternion* $\mathbf{q}$*. Then for* $t = 1, \dots, 4$*, we have*

$$|\alpha_t| = 2\Big|\pi_t(\mathbf{pq})\Big|,$$

where the π_t *are the coordinate functions defined by*

$$\mathbf{p} = \pi_1(\mathbf{p}) + \pi_2(\mathbf{p})i + \pi_3(\mathbf{p})j + \pi_4(\mathbf{p})k.$$

Motivated by Proposition 1.1 and the preceding discussion, we make the following definitions:

Definitions and Remarks 1.2. Let $\mathbf{G}$ be a two-by-two game with strategy spaces $S_i = \{\mathbf{C}, \mathbf{D}\}$ and payoff functions $P_i : S_1 \times S_2 \to \mathbf{R}$. Then the *associated quantum game* $\mathbf{G}^Q$ is the two-player game in which each strategy space is the unit quaternions, and payoffs are calculated as

$$\begin{aligned} P_i^Q(\mathbf{p}, \mathbf{q}) =& \pi_1(\mathbf{pq})^2 P_i(\mathbf{C}, \mathbf{C}) + \pi_2(\mathbf{pq})^2 P_i(\mathbf{D}, \mathbf{D}) \\ &+ \pi_3(\mathbf{pq})^2 P_i(\mathbf{D}, \mathbf{C}) + \pi_4(\mathbf{pq})^2 P_i(\mathbf{C}, \mathbf{D}). \end{aligned}$$

Note that for any strategy $\mathbf{p}$ chosen by Player 1 and for any probability distribution whatsoever over the four outcomes $(\mathbf{C}, \mathbf{C})$, etc., Player 2 can always adopt a strategy $\mathbf{q}$ that effects this probability distribution: Let a^2, b^2, c^2, d^2 be the desired probabilities, let $\mathbf{r} = a + bi + cj + dk$ and set $\mathbf{q} = \mathbf{p}^{-1}\mathbf{r}$. Therefore, in the game $\mathbf{G}^Q$, there can never be an equilibrium in pure strategies unless one of the four strategy pairs leads to an optimal outcome for both players. Thus in $\mathbf{G}^Q$, pure-strategy equilibria are both rare and uninteresting.

Next we consider mixed strategies. A *mixed quantum strategy* for $\mathbf{G}$ is a mixed strategy in the game $\mathbf{G}^Q$, i.e. a probability distribution on the space of unit quaternions. If $\mathbf{p}$ is a unit quaternion, we will sometimes identify $\mathbf{p}$ with the mixed strategy supported entirely on $\mathbf{p}$. If ν and μ are mixed strategies, we will write $P_i^Q(\nu, \mu)$ for the corresponding expected playoff to player i; that is,

$$P_i^Q(\nu, \mu) = \int P_i(\mathbf{p}, \mathbf{q}) d\nu(\mathbf{p}) d\mu(\mathbf{q}).$$

This gives rise to a new game $\mathbf{G}^Q_{\mathbf{mixed}}$, in which the strategy sets are the sets of mixed quantum strategies in $\mathbf{G}$. A *mixed-strategy Nash equilibrium* in $\mathbf{G}^Q$ is a Nash equilibrium in $\mathbf{G}^Q_{\mathbf{mixed}}$.

Our goal is to classify the mixed-strategy Nash equilibria in $\mathbf{G}^Q$. The definitions that occupy the remainder of this section kick off that process by partitioning the set of Nash equilibria into natural equivalence classes.

Definition 1.3. Two mixed strategies μ and μ' are *equivalent* if

$$\int \pi_t(\mathbf{pq}) d\mu(\mathbf{q}) = \int \pi_t(\mathbf{pq}) d\mu'(\mathbf{q})$$

for all unit quaternions $\mathbf{p}$ and all $t = 1, 2, 3, 4$.

In other words, μ and μ' are equivalent if in every quantum game and for every quantum strategy $\mathbf{p}$, we have $P_1(\mathbf{p}, \mu) = P_1(\mathbf{p}, \mu')$ and $P_2(\mathbf{p}, \mu) = P_2(\mathbf{p}, \mu')$.

Example 1.4. The strategy supported on the singleton $\{\mathbf{p}\}$ is equivalent to the strategy supported on the singleton $\{-\mathbf{p}\}$ and to no other singleton.

Definition 1.5. Let ν be a mixed strategy and $\mathbf{u}$ a unit quaternion. The *right translate* of ν by $\mathbf{u}$ is the measure $\nu\mathbf{u}$ defined by $(\nu\mathbf{u})(A) = \nu(A\mathbf{u})$, where A is any subset of the unit quaternions and $A\mathbf{u} = \{\mathbf{xu} | x \in A\}$. Similarly, the left translate of ν by $\mathbf{u}$ is defined by $(\mathbf{u}\nu)(A) = \nu(\mathbf{u}A)$. The following proposition is immediate:

Proposition 1.6. *Let (ν, μ) be a pair of mixed strategies and $\mathbf{u}$ a unit quaternion. Then in any game $\mathbf{G}^Q$, (ν, μ) is a mixed strategy Nash equilibrium if and only if $(\nu\mathbf{u}, \mathbf{u}^{-1}\mu)$ is.*

Definition 1.7. Two pairs of mixed strategies (ν, μ) and (ν', μ') are *equivalent* if there exists a unit quaternion $\mathbf{u}$ such that ν' is equivalent to $\nu\mathbf{u}$ and μ' is equivalent to $\mathbf{u}^{-1}\mu$. Note that this definition is independent of any particular game.

Proposition 1.8. *In a given game, a pair of mixed strategies is a Nash equilibrium if and only if every equivalent pair of mixed strategies is also a Nash equilibrium.*

2. Preliminary results

Theorems 2.1, 2.2 and 2.4 are the main results which will be used in Section 3 to classify Nash equilibria.

Theorem 2.1. *Every mixed strategy μ is equivalent to a mixed strategy supported on (at most) four points. Those four points can be taken to form an orthonormal basis for $\mathbf{R}^4$.*

Proof. First, choose any orthonormal basis $\{\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3, \mathbf{q}_4\}$ for $\mathbf{R}^4$. For any quaternion $\mathbf{p}$, write (uniquely)

$$\mathbf{p} = \sum_{\alpha=1}^{4} A_\alpha(p)\mathbf{q}_\alpha,$$

where the $A_\alpha(p)$ are real numbers.

Define a probability measure ν supported on the four points $\mathbf{q}_\alpha$ by

$$\nu(\mathbf{q}_\alpha) = \int_{\mathbf{S}^3} A_\alpha(\mathbf{q})^2 d\mu(\mathbf{q}).$$

For any two quaternions $\mathbf{p}$ and $\mathbf{q}$, define

$$X(\mathbf{p}, \mathbf{q}) = \sum_{\alpha=1}^{4} \pi_\alpha(\mathbf{p})\pi_\alpha(\mathbf{q})X_i. \tag{2.1.1}$$

Then for any $\mathbf{p}$ we have

$$\begin{aligned}
P(\mathbf{p}, \mu) &= \int_{\mathbf{S}^3} P(\mathbf{pq}) d\mu(\mathbf{q}) \\
&= \int_{S^3} P\left(\sum_{\alpha=1}^{4} A_\alpha(\mathbf{q})\mathbf{pq}_\alpha\right) d\mu(\mathbf{q}) \\
&= \sum_{\alpha=1}^{4} P(\mathbf{pq}_\alpha) \int_{S^3} A_\alpha(\mathbf{q})^2 d\mu(\mathbf{q}) + 2\sum_{\alpha\neq\beta} X(\mathbf{pq}_\alpha, \mathbf{pq}_\beta) \int_{\mathbf{S}^3} A_\alpha(\mathbf{q})A_\beta(\mathbf{q}) d\mu(\mathbf{q}) \\
&= P(\mathbf{p}, \nu) + 2\sum_{\alpha\neq\beta} X(\mathbf{pq}_\alpha, \mathbf{pq}_\beta) \int_{\mathbf{S}^3} A_\alpha(\mathbf{q})A_\beta(\mathbf{q}) d\mu(\mathbf{q}).
\end{aligned}$$

To conclude that μ is equivalent to ν it is sufficient (and necessary) to choose the $\mathbf{q}_\alpha$ so that for each $\alpha \neq \beta$ we have

$$\int_{\mathbf{S}^3} A_\alpha(\mathbf{q}) A_\beta(\mathbf{q}) d\mu(\mathbf{q}) = 0.$$

For this, consider the function $B : \mathbf{R}^4 \times \mathbf{R}^4 \to \mathbf{R}$ defined by

$$B(\mathbf{a}, \mathbf{b}) = \int_{\mathbf{S}^3} \pi_1(\overline{\mathbf{a}}\mathbf{q}) \pi_1(\overline{\mathbf{b}}\mathbf{q}) d\mu(\mathbf{q}).$$

B is a bilinear symmetric form and so can be diagonalized; take the $\mathbf{q}_\alpha$ to be an orthonormal basis with respect to which B is diagonal. Then we have (for $\alpha \neq \beta$)

$$\begin{aligned} \int_{\mathbf{S}^3} A_\alpha(\mathbf{q}) A_\beta(\mathbf{q}) d\mu(\mathbf{q}) &= \int_{\mathbf{S}^3} \pi_1(\overline{\mathbf{q}_\alpha}\mathbf{q}) \pi_1(\overline{\mathbf{q}_\beta}\mathbf{q}) d\mu(\mathbf{q}) \\ &= B(\mathbf{q}_\alpha, \mathbf{q}_\beta) = 0. \end{aligned}$$ □

Theorem 2.2. *Taking Player 2's (mixed) strategy μ as given, Player 1's optimal response set is equal to the intersection of $\mathbf{S}^3$ with a linear subspace of $\mathbf{R}^4$.*

(Recall that we identify the unit quaternions with the three-sphere $\mathbf{S}^3$.)

Proof. Player One's problem is to choose $\mathbf{p} \in \mathbf{S}^3$ to maximize

$$P_1(\mathbf{p}, \mu) = \int P_1(\mathbf{pq}) d\mu(\mathbf{q}). \tag{2.2.1}$$

Expression (2.2.1) is a (real) quadratic form in the coefficients $\pi_i(\mathbf{p})$ and hence is maximized (over $\mathbf{S}^3$) on the intersection of $\mathbf{S}^3$ with the real linear subspace of $\mathbf{R}^4$ corresponding to the maximum eigenvalue of that form. □

Definition 2.3. We define the function $K : \mathbf{S}^3 \to \mathbf{R}$ by $K(A + Bi + Cj + Dk) = ABCD$. Thus in particular $K(\mathbf{p}) = 0$ if and only if $\mathbf{p}$ is a linear combination of at most three of the fundamental units $\{1, i, j, k\}$.

Theorem 2.4. *Let μ be a mixed strategy supported on four orthogonal points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3, \mathbf{q}_4$ played with probabilities $\alpha, \beta, \gamma, \delta$. Suppose $\mathbf{p}$ is an optimal response to μ in some game where it is not the case that $X_1 = X_2 = X_3 = X_4$. Then $\mathbf{p}$ must satisfy:*

$$\begin{aligned} &(\alpha - \beta)(\alpha - \gamma)(\alpha - \delta) K(\mathbf{pq}_1) + (\beta - \alpha)(\beta - \delta)(\beta - \gamma) K(\mathbf{pq}_2) \\ &+ (\gamma - \alpha)(\gamma - \beta)(\gamma - \delta) K(\mathbf{pq}_3) + (\delta - \alpha)(\delta - \beta)(\delta - \gamma) K(\mathbf{pq}_4) = 0. \end{aligned} \tag{2.4.1}$$

Proof. Set $\mathbf{p}_n = \pi_n(\mathbf{p})$ and consider the function

$$\begin{array}{rccc} \mathcal{P}: & \mathbf{S}^3 \times \mathbf{R}^4 & \to & \mathbf{R}, \\ & (\mathbf{p}, \mathbf{x}) & \mapsto & \sum_{n=1}^4 \mathbf{p}_n^2 \mathbf{x}_n d\mu(\mathbf{q}). \end{array}$$

In particular, if we let $X = (X_1, X_2, X_3, X_4)$, then $\mathcal{P}(\mathbf{p}, X) = P_1(\mathbf{p}, \mu)$.

The function $\mathcal{P}$ is quadratic in $\mathbf{p}$ and linear in $\mathbf{x}$; explicitly we can write

$$\mathcal{P}(\mathbf{p}, \mathbf{x}) = \sum_{i,j,k} t_{ijk} \mathbf{p}_i \mathbf{p}_j \mathbf{x}_k$$

for some real numbers t_{ijk}.

Set

$$M_{ij}(\mathbf{x}) = \sum_{k=1}^{4} t_{ijk}\mathbf{x}_k,$$
$$N_{ij}(\mathbf{p}) = \sum_{k=1}^{4} t_{ikj}\mathbf{p}_j,$$

so that

$$M(\mathbf{x}) \cdot \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \mathbf{p}_3 \\ \mathbf{p}_4 \end{pmatrix} = N(\mathbf{p}) \cdot \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \mathbf{x}_4 \end{pmatrix}. \tag{2.4.2}$$

If $\mathbf{p}$ is an optimal response to the strategy μ, then $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4)^T$ must be an eigenvector of $M(X)$, say with associated eigenvalue λ. From this and (2.4.2) we conclude that

$$N(\mathbf{p}) \cdot \begin{pmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \end{pmatrix} = \lambda \cdot \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \mathbf{p}_3 \\ \mathbf{p}_4 \end{pmatrix} = N(\mathbf{p}) \cdot \begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \lambda \end{pmatrix},$$

where the second equality holds by an easy calculation.

Thus $N(\mathbf{p})$ must be singular. But it follows from a somewhat less easy calculation that the determinant of $N(\mathbf{p})/2$ is given by the left side of (2.4.1). □

3. CLASSIFICATION

Definition 3.1. Let $\mathbf{G}$ be a two-by-two game with payoff pairs $(X_1, Y_1), \ldots, (X_4, Y_4)$ (listed in arbitrary order). $\mathbf{G}$ is a *generic game* if the X_i are all distinct, the Y_i are all distinct, the twofold sums $X_i + X_j$ are all distinct and the twofold sums $Y_i + Y_j$ are all distinct.

Theorem 3.3 will classify Nash Equilibria in $\mathbf{G}^Q$, where $\mathbf{G}$ is any generic two-by-two game. Subtler versions of the same arguments work for nongeneric games (yielding somewhat messier results); see [NE].

To state Theorem 3.3 we need a definition:

Definition 3.2. Let $\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}$ be quaternions; write $\mathbf{p} = p_1 + p_2 i + p_3 j + p_4 k$, etc. Then the quadruple $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s})$ is *intertwined* if there is a nonzero constant α such that

$$\alpha(X\mathbf{p} + Y\mathbf{q}) = X\mathbf{r} + Y\mathbf{s}$$

identically in the polynomial variables X and Y.

Thus if the components of $\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}$ are all nonzero, then $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s})$ is intertwined if and only if the four quotients $\frac{p_1}{q_1}, \frac{p_2}{q_2}, \frac{p_3}{q_3}, \frac{p_4}{q_4}$ are equal (in some order) to the four quotients $\frac{r_1}{s_1}, \frac{r_2}{s_2}, \frac{r_3}{s_3}, \frac{r_4}{s_4}$.

The intertwined quadruple $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s})$ is *fully intertwined* if $(\mathbf{p}, \mathbf{r}, \mathbf{q}, \mathbf{s})$ is also intertwined.

We can now state the main theorem.

Theorem 3.3. *Let* $\mathbf{G}$ *be a generic game. Then up to equivalence, every equilibrium in* $\mathbf{G}^Q$ *is of one of the following types:*

a) *Each player plays each of four orthogonal quaternions with probability* $1/4$.
b) *Each player's strategy is supported on three of the four quaternions* $1, i, j, k$.
c) μ *is supported on two orthogonal points* $1, \mathbf{v}$*;* ν *is supported on two orthogonal points* $\mathbf{p}, \mathbf{pu}$*; and the quadruple* $(\mathbf{p}, \mathbf{pv}, \mathbf{pu}, \mathbf{pvu})$ *is fully intertwined.*
d) *Each of* μ *and* ν *is supported on two orthogonal points, each played with probability* $1/2$. *Moreover, the supports of* μ *and* ν *lie in parallel planes.*
e) *Each player plays a pure strategy from the four point set* $\{1, i, j, k\}$.

Proof. Let (ν, μ) be an equilibrium. By Theorem 2.1 we can assume that each of ν and μ is supported on a set of at most four orthogonal points. Applying a translation as in Definition 1.7 we can assume that the support of μ contains the quaternion 1. Then from standard facts about orthogonality in the space of quaternions, the support of μ is contained in a set of the form $\{1, \mathbf{u}, \mathbf{v}, \mathbf{uv}\}$, where $\mathbf{u}^2 = \mathbf{v}^2 = -1$ and $\mathbf{uv} + \mathbf{vu} = 0$, played with probabilities of $\alpha, \beta, \gamma, \delta \geq 0$. We will maintain these assumptions and this notation while proving Theorems 3.4, 3.5, 3.9, and 3.10, which together imply Theorem 3.3. ☐

Theorem 3.4. ν *is a pure strategy if and only if* μ *is a pure strategy.*

Proof. If ν is a pure strategy, Player Two can guarantee any desired probability distribution over four outcomes; by genericity his optimal probability distribution is unique. ☐

Theorem 3.5. *If the support of* ν *contains four points, then* μ *assigns probability* $1/4$ *to each of four strategies.*

Proof. Explicitly write $\mathbf{u} = Ai + Bj + Ck, \mathbf{v} = Di + Ej + Fk, \mathbf{uv} = Gi + Hj + Ik$. Write

$$\mathcal{M} = \begin{pmatrix} AB & DE & GH \\ AD & DF & GI \\ BC & EF & HI \end{pmatrix}.$$

By Theorem 2.2 the quadratic form

$$\mathbf{p} \mapsto P_1(\mathbf{p}, \mu) \tag{3.5.1}$$

is constant on the unit sphere $\mathbf{S}^3$. Therefore its nondiagonal coefficients are all zero. Computing these coefficients explicitly and dividing by (nonzero) expressions of the form $(X_i - X_j)$, we get

$$\mathcal{M} \cdot (\beta, \gamma, \delta)^T = (0, 0, 0)^T. \tag{3.5.2}$$

But $\mathcal{M}$ also kills the column vector $(1, 1, 1)^T$. Thus we have two cases:

Case I. $\beta = \gamma = \delta$. Then the four diagonal terms of (3.5.1) (which must all be equal) are given by $(X_1 + X_2 + X_3 + X_4)\beta + X_i(\alpha - \beta)$, with $i = 1, 2, 3, 4$. Since the X_i are not all equal, it follows that $\alpha = \beta = \gamma = \delta = 1/4$, proving the theorem.

Case II. $\mathcal{M}$ has rank at most one. From this and the orthogonality of $\mathbf{u}, \mathbf{v}, \mathbf{uv}$, we have $\{\mathbf{u}, \mathbf{v}, \mathbf{uv}\} \cap \{i, j, k\} \neq \emptyset$. Assume $\mathbf{u} = i$ (the other cases are similar). Then

$A=1$, $B=C=D=G=0$, $H=-F$ and $I=E$. The four diagonal entries of (3.3.1) are now equal; call their common value λ so that we have

$$\begin{pmatrix} \alpha & \beta & E^2\gamma+F^2\delta & E^2\delta+F^2\gamma \\ \beta & \alpha & E^2\delta+F^2\gamma & E^2\gamma+F^2\delta \\ E^2\gamma+F^2\delta & E^2\delta+F^2\gamma & \alpha & \beta \\ E^2\delta+F^2\gamma & E^2\gamma+F^2\delta & \beta & \alpha \end{pmatrix} \cdot \begin{pmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \end{pmatrix} = \begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \lambda \end{pmatrix}. \tag{3.5.3}$$

Combining (3.5.2), (3.5.3), the conditions $\alpha+\beta+\gamma+\delta=E^2+F^2=1$ and the genericity conditions, we get $\alpha=\beta=\gamma=\delta$ as required. □

Corollary 3.5A. *If either player's strategy has a four-point support, then each player plays each of four orthogonal quaternions with probability* $1/4$.

Proof. Apply Theorem 3.5 twice, once as stated and once with the players reversed. □

Theorem 3.9, dealing with the case where ν is supported on exactly three points, requires some preliminary lemmas:

Lemma 3.6. *It is not the case that Player Two plays* $1, \mathbf{u}, \mathbf{v}$ *each with probability* $1/3$.

Proof. If $1, \mathbf{u}, \mathbf{v}$ are played with probability $1/3$, then one computes that the eigenvalues of the form (2.2.1) are $X_1+X_2+X_3$, $X_1+X_2+X_4$, $X_1+X_3+X_4$, $X_2+X_3+X_4$, which are all distinct by genericity. Thus Player One responds with a pure strategy, and Theorem 3.4 provides a contradiction. □

Lemma 3.7. *Suppose the support of* ν *is contained in the linear span of* $1, i, j$, *and suppose that* 1 *and* i *are both optimal responses for Player Two. Then one of the following is true:*

a) *The support of* ν *is contained in the three point set* $\{1, i, j\}$.
b) *The support of* ν *is contained in a set of the form* $\{1, Ei+Fj, -Fi+Ej\}$ *with* $Ei+Fj$ *and* $-Fi+Ej$ *played equiprobably.*

Moreover, if b) holds and either j *or* k *is also an optimal response for Player Two, then* 1 *is played with probability zero.*

Proof. Suppose ν is supported on three orthogonal quaternions $\mathbf{q}_1=A+Bi+Cj$, $\mathbf{q}_2=D+Ei+Fj$, $\mathbf{q}_3=G+Hi+Ij$, played with probabilities ϕ, ψ, ξ. The first-order conditions for Player Two's maximization problem must be satisfied at both 1 and i; this (together with genericity for the game $\mathbf{G}$) gives

$$\begin{pmatrix} AC & DF & GI \\ BC & EF & HI \end{pmatrix} \begin{pmatrix} \phi \\ \psi \\ \xi \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} AC & DF & GI \\ BC & EF & HI \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \tag{3.7.1}$$

so that by Lemma 3.6 with the players reversed, the matrix on the left has rank at most one. This (together with the orthogonality of $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$) gives $\{\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3\} \cap \{1, i, j\} \neq \emptyset$. We can assume $\mathbf{q}_1=1$ (all other cases are similar); thus $A=1$, $B=C=D=G=0$, $H=-F$, $I=E$. Now (3.7.1) says that $EF(\psi-\xi)=0$. If $EF=0$, then a) holds, and if $\psi-\xi=0$, then b) holds.

Now suppose j is also an optimal response for Player Two. Then $0=P_2(\nu, i)-P_2(\nu, j)=\phi(Y_2-Y_3)$, so that by genericity $\phi=0$. A similar argument works if k is optimal. □

Lemma 3.8. *Suppose ν is supported on exactly three points and continue to assume that μ is supported on a subset of $\{1, \mathbf{u}, \mathbf{v}, \mathbf{uv}\}$. Then at least two of the four quaternions $1, \mathbf{u}, \mathbf{v}, \mathbf{uv}$ are optimal responses for Player One.*

Proof. By Theorem 3.5, μ is supported on at most three points; we can rename so that those points are $1, \mathbf{u}, \mathbf{v}$. These are played with probabilities α, β, γ, and we can rename again so that α lies (perhaps not strictly) between β and γ.

If $\mathbf{p}$ is any optimal response by Player One, apply (2.4.1) with $\delta = 0$ (and possibly $\gamma = 0$) to get

$$\sigma_1 K(\mathbf{p}) + \sigma_2 K(\mathbf{pu}) + \sigma_3 K(\mathbf{pv}) + \sigma_4 K(\mathbf{puv}) = 0, \tag{3.8.1}$$

where $\sigma_1 = (\alpha - \beta)(\alpha - \gamma)\alpha$, etc., so that

$$\sigma_1, \sigma_4 \le 0 \qquad \text{and} \qquad \sigma_2, \sigma_3 \ge 0. \tag{3.8.2}$$

Case I. Suppose none of the σ_i are equal to zero. Then $\gamma \neq 0$ so a) holds.

By Theorem 2.2, the support of ν spans a three-dimensional hyperplane in $\mathbf{R}^4$ and thus must include some quaternion of the form $A + B\mathbf{u}$ ($A, B \in \mathbf{R}$). Inserting $\mathbf{p} = A + B\mathbf{u}$ into (3.8.1) gives

$$AB(\sigma_1 B^2 - \sigma_2 A^2) K(1 + \mathbf{u}) = 0. \tag{3.8.3}$$

Thus either $AB = 0$ (in which case either $\mathbf{p} = 1$ or $\mathbf{p} = \mathbf{u}$) or $K(1 + \mathbf{u}) = 0$. This and similar arguments establish the following:

(3.8.3a) If 1 and $\mathbf{u}$ are both suboptimal responses, then $K(1 + \mathbf{u}) = 0$.

(3.8.3b) If 1 and $\mathbf{v}$ are both suboptimal responses, then $K(1 + \mathbf{v}) = 0$.

(3.8.3c) If $\mathbf{u}$ and $\mathbf{uv}$ are both suboptimal responses, then $K(1 + \mathbf{v}) = 0$.

(3.8.3d) If $\mathbf{v}$ and $\mathbf{uv}$ are both suboptimal responses, then $K(1 + \mathbf{u}) = 0$.

Taken together, these imply that if the lemma fails, then $K(1+\mathbf{u}) = K(1+\mathbf{v}) = 0$. From this it follows that $\{\mathbf{u}, \mathbf{v}, \mathbf{uv}\} \cap \{\pm i, \pm j, \pm k\} \neq \emptyset$; assume without loss of generality that $\mathbf{u} = i$ and therefore $\mathbf{v}$ is in the linear span of $\{j, k\}$. (Generality is not lost because the argument to follow works just as well, with obvious modifications, in all the remaining cases.)

Now we have

$$\begin{aligned} P_1(A + B\mathbf{v}, \mu) &= \alpha P_1(A + B\mathbf{v}) + \beta P_1(Ai + B\mathbf{v}i) + \gamma P_1(A\mathbf{v} - B) \\ &= A^2\Big(\alpha P_1(1) + \beta P_1(i) + \gamma P_1(\mathbf{v})\Big) + B^2\Big(\alpha P_1(\mathbf{v}) + \beta P_1(\mathbf{v}i) + \gamma P_1(1)\Big), \end{aligned}$$

which is maximized at an endpoint, so either 1 or $\mathbf{v}$ is an optimal response for Player One. Similarly, at least one from each pair $\{1, \mathbf{uv}\}$, $\{\mathbf{u}, \mathbf{v}\}$, and $\{\mathbf{u}, \mathbf{uv}\}$ is an optimal response, from which b) (and therefore the lemma) follows.

Case II. Suppose at least one of the σ_i is equal to zero. Up to renaming $\mathbf{u}$ and $\mathbf{v}$, there are three ways this can happen:

Subcase IIA. $\alpha = \beta$, $\gamma = 0$. As above, Player One's optimal response set contains a quaternion of the form $(A + B\mathbf{u})$. But $P_1(A + B\mathbf{u}, \mu)$ is independent of A and B, so both 1 and $\mathbf{u}$ are optimal, proving the theorem. (Note that $\mathbf{v}$ and $\mathbf{uv}$ are also both optimal, so that in fact by (3.5A) this case never occurs.)

Subcase IIB. $\alpha = \beta$, $\gamma \neq 0$. By Lemma 3.6, $\gamma \neq \alpha, \beta$. Thus σ_3 and σ_4 are nonzero, so (3.8.3b), (3.8.3c) and (3.8.3d) (but not (3.8.3a)) still hold. But $\sigma_1 = \sigma_2 = 0$, so the same techniques now yield:

(3.8.3e) If 1 and $\mathbf{u}$ are both suboptimal responses, then $K(1+\mathbf{u}) = 0$.

(3.8.3f) If 1 and $\mathbf{v}$ are both suboptimal responses, then $K(1+\mathbf{v}) = 0$.

We can now repeat the argument from Case I.

Subcase IIC. $\alpha \neq \beta$, $\gamma = 0$. Now we have $\sigma_1, \sigma_2 \neq 0$, $\sigma_3 = \sigma_4 = 0$, so that (3.8.3a) through (3.8.3c) still hold, along with (3.8.3e) and (3.8.3f). We can now repeat the argument from Case I. □

Theorem 3.9. *If ν is supported on exactly three points, then up to equivalence, both μ and ν are supported on three-point subsets of $\{1, i, j, k\}$.*

Proof. By Theorem 3.5 we can assume that μ is supported on $\{1, \mathbf{u}, \mathbf{v}\}$. By Lemma 3.8 we can assume without much loss of generality that 1 and $\mathbf{u}$ are optimal responses for Player One. (The argument below works equally well, with obvious modifications, for other pairs.) Let $\mathbf{w}$ be a quaternion orthogonal to 1 and $\mathbf{u}$ such that the support of ν is contained in the linear span of $1, \mathbf{u}$ and $\mathbf{w}$.

By Theorem 2.2, any quaternion of the form $X + Y\mathbf{u} + Z\mathbf{w}$ is an optimal response for Player One, so by (2.4.1) we have

$$\sigma_1 K(X + Y\mathbf{u} + Z\mathbf{w}) + \sigma_2 K(X\mathbf{u} - Y + Z\mathbf{wu}) + \sigma_3 K(X\mathbf{v} + Y\mathbf{uv} + Z\mathbf{wv}) \\ + \sigma_4 K(X\mathbf{uv} - Y\mathbf{v} + Z\mathbf{wuv}) = 0$$

identically in X, Y, Z. Writing out the left side as a polynomial in these three variables, the coefficients, all of which must vanish, can be expressed in terms of the components of $\mathbf{u}, \mathbf{v}, \mathbf{w}$. Setting all these expressions equal to zero and solving, we find that $\{\mathbf{u}, \mathbf{v}, \mathbf{w}\} \in \{\pm i, \pm j, \pm k\}$. (The details of this tedious but straightforward calculation can be found on pages 32-33 of [NE].) We assume $\mathbf{u} = i$, $\mathbf{w} = j$.

Claim. Player Two's strategy is not supported just on 1 and i. □

Proof. If so, the fact that $P_1(1, \mu) = P_1(i, \mu)$ implies that μ assigns equal weights to 1 and i, which implies that $P_1(j, \mu) = P_1(k, \mu)$, contradicting the fact that j but not k is optimal for Player One.

Thus the support of μ is a three-point subset of $\{1, i, j, k\}$. It now follows from Lemma 3.8 (together with the assumption that the support of ν contains three points) that the support of ν is $\{1, i, j\}$, completing the proof. □

Theorem 3.10. *Suppose ν is supported on two points. Then μ is supported on $1, \mathbf{u}$ and ν is supported on two quaternions $\mathbf{p}, \mathbf{pv}$ where either*

a) *the quadruple $(\mathbf{p}, \mathbf{pu}, \mathbf{pv}, \mathbf{pvu})$ is fully intertwined or*
b) *$\mathbf{u} = \mathbf{v}$ and each player plays each strategy with probability $1/2$.*

Proof. Suppose 1 and $\mathbf{u}$ are played with probabilities α and β.

Any unit quaternion of the form $X\mathbf{p} + Y\mathbf{pv}$ is an optimal response for Player One; thus (2.4.1) with $\mathbf{q}_1 = 1, \mathbf{q}_2 = \mathbf{u}, \gamma = \delta = 0$ gives

$$(\alpha - \beta)\Big(\alpha^2 K(X\mathbf{p} + Y\mathbf{pv}) - \beta^2 K(X\mathbf{pv} + Y\mathbf{pvu})\Big) = 0.$$

This, plus the identical observation with the players reversed, establishes full intertwining except when $\alpha = \beta = 1/2$. In that case, $P_1(\mathbf{p}, \mu) = P_1(\mathbf{pu}, \mu)$, so $\mathbf{pu}$ must be optimal; i.e. we can take $\mathbf{v} = \mathbf{u}$. □

This completes the proof of Theorem 3.3.

Remark. The statement of Theorem 3.3 makes it natural to ask for a classification of fully intertwined quadruples of the form $(\mathbf{p}, \mathbf{pv}, \mathbf{pu}, \mathbf{pvu})$ with $\mathbf{u}$, $\mathbf{v}$ square roots of -1. That classification is provided in [I]. The thrust of the result is this: All such quadruples fall into one of approximately 15 families. Each of these families is at most four-dimensional (inside the twelve-dimensional manifold of all four-tuples). For all but one of the families, it is easy to tell by inspection whether a given quadruple satisifies the membership condition. The exceptional family is one-dimensional.

In short: Condition b) of Theorem 3.3 allows only four dimensions' worth of possible equilibria, all of which are easily identifiable except for a one-dimensional subset.

4. Minimal payoffs and opting out

Theorem 3.3 classifies all mixed strategy Nash equilibria in generic games. Here we briefly address the issue of whether these equilibria survive in a larger game where the players can opt out of the assigned communication protocol.

A key tool is the very simple Theorem 4.1; this and Corollary 4.1A apply to all two-by-two games (whether generic or not) and are of independent interest:

Theorem 4.1. *Let* $\mathbf{G}$ *be a game with payoff pairs* $(X_1, Y_1), \ldots, (X_4, Y_4)$*. Then in any mixed strategy quantum equilibirum, Player One earns at least* $(X_1 + X_2 + X_3 + X_4)/4$.

Proof. Player One maximizes the quadratic form (2.2.1) over the sphere $\mathbf{S}^3$. The trace of this form is $X_1 + X_2 + X_3 + X_4$, so the maximium eigenvalue must be at least $(X_1 + X_2 + X_3 + X_4)/4$. □

Corollary 4.1A. *If, in Theorem* 4.1*, the game* $\mathbf{G}$ *is zero-sum, then in any mixed strategy quantum equilibrium, Player One earns exactly* $(X_1 + X_2 + X_3 + X_4)/4$.

Proof. Apply Theorem 4.1 to both players. □

4.2. *Remarks on opting out.* A player can throw away his entangled penny and substitute an unentangled penny (or for that matter a purely classical penny, but this offers no additional advantage, because the unentangled quantum penny can always be returned in one of the two classical states $\mathbf{H}$ or $\mathbf{T}$). However, a simple quantum mechanical calculation shows that if Player One unilaterally substitutes an unentangled penny, then no matter what strategies the players follow from there, the result is a uniform distribution over the four possible outcomes. By Theorem 4.1, Player One considers this weakly inferior to any $\mathbf{G}^Q$ equilibrium. Thus, even if we allow players to choose their pennies, all of the $\mathbf{G}^Q$ equilibria survive.

References

[CHTW] R. Cleve, P. Hoyer, B. Toner and J. Watrous, "Consequences and Limits of Nonlocal Strategies", *Proc. of the 19th Annual Conference on Computational Complexity* (2004), 236-249.

[DL] G. Dahl and S. Landsburg, "Quantum Strategies", preprint, available at http://www.landsburg.com/dahlcurrent.pdf.

[EW] J. Eisert and M. Wilkens, "Quantum Games", *J. of Modern Optics* 47 (2000), 2543-2556. MR1801549 (2001m:81047)

[EWL] J. Eisert, M. Wilkens and M. Lewenstein, "Quantum Games and Quantum Strategies", *Phys. Rev. Lett.* 83 (1999), 3077-3080. MR1720182 (2001c:81027)

[I] S. Landsburg, "Intertwining", working paper available at http://www.landsburg.com/twining4.pdf.

[NE] S. Landsburg, "Nash Equilibrium in Quantum Games", RCER Working Paper #524, Rochester Center for Economic Research, February 2006.

Department of Economics, University of Rochester, Rochester, New York 14627