Nash Equilibria in Quantum Games

by

Steven E. Landsburg

Quantum game theory investigates the behavior of strategic agents with access to quantum technology. In some models (see [CHTW] or [DL] for example), players can randomize their strategies by observing quantum mechanical variables (such as the spins of electrons). In others, players use quantum devices to communicate with each other or with the referee. The present paper is exclusively concerned with models of quantum communication.

The most widely studied such model is the *EWL model* introduced by Eisert, Wilkens and Lewenstein ([EW], [EWL]). The EWL model supplements classical game theory by specifying a protocol for the communication of strategies and computation of payoffs. For example, starting with an ordinary two-by-two game, EWL envisions a referee who issues each player a penny in the state **H** ("heads up"). The player indicates his choice of strategy by returning the penny either unflipped or flipped. The referee observes the pair of returned pennies, which are in one of four states $\mathbf{H} \otimes \mathbf{H}$, $\mathbf{H} \otimes \mathbf{T}$, etc., and calculates payoffs accordingly.

Now replace the pennies with subatomic particles (though I will continue to call them pennies) in the maximally entangled state $\mathbf{H} \otimes \mathbf{H} + \mathbf{T} \otimes \mathbf{T}$. Each player returns his penny after acting on it by the special unitary operator of his choice. The pennies are returned to the referee, who makes an observation with one of four possible outcomes, and calculates payoffs accordingly.

Players are free to discard their entangled pennies and substitute classical pennies, but in equilibrium they won't want to.

This model is of interest because it captures, in a fairly general setting, the behavior of players who use quantum technology to manipulate their communications, taking as given the protocol by which those communications will be deciphered.

In the EWL model, each player's strategy space naturally expands from the two-point space to the group SU_2 , which we identify with the three-sphere S^3 . The space of *mixed quantum strategies*, then, is the space of probability distributions on S^3 . In principle, the vastness of this space makes it difficult to

find equilibria or to establish that they've all been found. Quite a bit of attention has been lavished on identifying equilibria in particular two-by-two EWL games.

The chief contribution of this paper is to classify all Nash equilibria in mixed quantum strategies up to a natural notion of equivalence, and to show that they all take quite simple forms. In equilibrium, each player's mixed strategy is supported on at most four points. Moreover, these points must lie in certain quite restrictive geometric configurations. This transforms the search for equilibria from a potentially intractable problem into an almost mechanical one.

I call a game "generic" if, for each player, the four possible payoffs and the six pairwise sums of those payoffs are all distinct. In this paper, I will state and prove the main theorem for generic games, referring the reader to my unpublished working paper [NE] for the (considerably uglier) generalization to the non-generic case.

Section 1 lays out 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.

A classical penny can be either in the state \mathbf{H} ("heads") or \mathbf{T} ("tails"). The states of a quantum penny are represented by expressions $\alpha \mathbf{H} + \beta \mathbf{T}$ where α and β are complex numbers, not both zero. Two such expressions represent the same state if (and only if) one is a (complex) scalar multiple of the other.

An entangled pair of pennies is in a state represented by a non-zero expression

$$\alpha \mathbf{H} \otimes \mathbf{H} + \beta \mathbf{H} \otimes \mathbf{T} + \gamma \mathbf{T} \otimes \mathbf{H} + \delta \mathbf{T} \otimes \mathbf{T}$$

where the coefficients are complex numbers, and where a scalar multiple represents the same state.

Start with a two-by-two game \mathbf{G} where each player's strategy space is $\{\mathbf{C}, \mathbf{D}\}$. Each player acquires (or is issued) one of two entangled pennies, which start out in the maximally entangled state

$$\mathbf{H}\otimes \mathbf{H} + \mathbf{T}\otimes \mathbf{T}$$

They return their pennies unflipped to indicate a play of \mathbf{C} or flipped to indicate a play of \mathbf{D} . After the pennies are returned, the referee observes whether they've been flipped and makes payoffs accordingly.

As long as players obediently play either \mathbf{C} or \mathbf{D} , the game is unchanged from its old, non-quantum version. But quantum mechanics allows players to act on their pennies by arbitrary special unitary matrices. Acting by the identity matrix means leaving the penny unflipped. Acting by the matrix

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{1.0.1}$$

means flipping the penny. Other unitary matrices correspond to physical operations with no classical analogues.

If Players One and Two act by the unitary matrices

$$\begin{pmatrix} A & B \\ -\overline{B} & \overline{A} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} P & Q \\ -\overline{Q} & \overline{P} \end{pmatrix}$$

then the pennies come back to the referee in the state

$$(A\mathbf{H} - \overline{B}\mathbf{T}) \otimes (P\mathbf{H} - \overline{Q}\mathbf{T}) + (B\mathbf{H} + \overline{A}\mathbf{T}) \otimes (Q\mathbf{H} + \overline{P}\mathbf{T})$$

which expands to

$$(AP + BQ)\mathbf{H} \otimes \mathbf{H} + (-A\overline{Q} + B\overline{P})\mathbf{H} \otimes \mathbf{T} + (-\overline{B}P + \overline{A}Q)\mathbf{T} \otimes \mathbf{H} + (\overline{AP} + \overline{BQ})\mathbf{T} \otimes \mathbf{T}$$
(1.0.2)

If the players choose classical strategies (either the identity matrix or the matrix (1.0.1)), the pennies come back to the referee in one of four states

$$\mathbf{C}\mathbf{C} = \mathbf{H} \otimes \mathbf{H} + \mathbf{T} \otimes \mathbf{T} \tag{1.0.3a}$$

$$\mathbf{C}\mathbf{D} = \mathbf{H} \otimes \mathbf{T} + \mathbf{T} \otimes \mathbf{H} \tag{1.0.3b}$$

$$\mathbf{DC} = \mathbf{H} \otimes \mathbf{T} - \mathbf{T} \otimes \mathbf{H} \tag{1.0.3c}$$

$$\mathbf{D}\mathbf{D} = \mathbf{H} \otimes \mathbf{H} - \mathbf{T} \otimes \mathbf{T} \tag{1.0.3d}$$

and the referee makes payoffs accordingly. If players adopt more general strategies, returning the pennies in state (1.0.2), the referee's observation causes them to collapse into one of the states (1.0.3a- d) with probabilities calculated (according to the laws of quantum mechanics) as follows: First write (1.0.2) as a linear combination

$$\alpha_1 \mathbf{C} \mathbf{C} + \alpha_2 \mathbf{D} \mathbf{D} + \alpha_3 \mathbf{C} \mathbf{D} + \alpha_4 \mathbf{D} \mathbf{C}$$
(1.0.4)

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

Notice that the referee cannot detect (and therefore cannot prohibit) the play of nonclassical strategies.

All the referee ever observes is a pair of pennies in one of the states (1.0.3a-d).

Now identify Player One's strategy space with the unit quaternions by mapping the matrix with top row (A, B) to the quaternion A + Bj; identify Player Two's strategy space with the unit quaternions by mapping the matrix with top row (P, Q) to the quaternion P - jQ. From (1.0.2) one readily calculates the coefficients in (1.0.4) 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

$$|\alpha_t| = 2 \left| \pi_t(\mathbf{pq}) \right|$$

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 **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

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

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 **G** is a mixed strategy in the game \mathbf{G}^{Q} , i.e. a probability distribution on the space of unit quaternions.

If **p** is a unit quaternion, **I** will sometimes identify **p** with the mixed strategy supported entirely on **p**. If ν and μ are mixed strategies, **I** will write $P_i^Q(\nu, \mu)$ for the corresponding expected playoff to player *i*; that is:

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

Example 1.3. The Prisoners' Dilemma. In [EW], Eisert and Wilkens [EW] analyze the quantum version of the Prisoners' Dilemma:

Player Two

$$C D$$

$$C D$$

$$C (3,3) (0,5)$$
Player One
$$D (5,0) (1,1)$$

$$(1.3.1)$$

which has the following equilibrium in mixed quantum strategies:

Player 1 plays the quaternions 1 and
$$i$$
, each with probability 1/2.
Player 2 plays the quaternions j and k , each with probability 1/2.
$$\left.\right\}$$
 (1.3.2)

To check that this is indeed a Nash equilibrium, first take Player 1's strategy as given and suppose Player 2 plays the quaternion $\mathbf{q} = \mathbf{q}_1 + \mathbf{q}_2 i + \mathbf{q}_3 j + \mathbf{q}_4 k$, so that $i\mathbf{q} = (-\mathbf{q}_2) + \mathbf{q}_1 i + (-\mathbf{q}_4)j + \mathbf{q}_3 k$. Then Player 2's expected payoff is

$$\begin{aligned} \frac{1}{2}P_2(1,\mathbf{q}) + \frac{1}{2}P_2(i,\mathbf{q}) &= \frac{1}{2} \Big(\mathbf{q}_1^2 \cdot 3 + \mathbf{q}_2^2 \cdot 1 + \mathbf{q}_3^2 \cdot 5 + \mathbf{q}_4^2 \cdot 0 \Big) + \frac{1}{2} \Big((-\mathbf{q}_2)^2 \cdot 3 + \mathbf{q}_1^2 \cdot 1 + (-\mathbf{q}_4)^2 \cdot 5 + \mathbf{q}_3^2 \cdot 0 \Big) \\ &= 2\mathbf{q}_1^2 + 2\mathbf{q}_2^2 + (5/2)\mathbf{q}_3^2 + (5/2)\mathbf{q}_4^2 \end{aligned}$$

which is to be maximized subject to the condition that \mathbf{q} is a unit quaternion, i.e.

$$\mathbf{q}_1^2 + \mathbf{q}_2^2 + \mathbf{q}_3^2 + \mathbf{q}_4^2 = 1$$

Clearly this is maximized when $\mathbf{q}_0 = \mathbf{q}_1 = 0$, e.g. at j and k. Thus Player 2 is optimizing, and similarly for Player 1.

Note that this equilibrium, in which each player receives an expected payoff of 5/2, is Pareto-superior to the classical equilibrium (though still Pareto-supoptimal).

Our goal is to classify the 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 1.4. Two mixed strategies μ and μ' are equivalent if

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

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 1.5. 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 1.6. Let ν be a mixed strategy and \mathbf{u} a unit quaternion. The *right translate* of ν by \mathbf{u} is the measure $\nu \mathbf{u}$ definied 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 1.7. Let (ν, μ) be a pair of mixed strategies and **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 1.8. Two pairs of mixed strategies (ν, μ) and (ν', μ') are *equivalent* if there exists a unit quaternion **u** such that ν' is equivalent to ν **u** and μ' is equivalent to $\mathbf{u}^{-1}\mu$. Note that this definition is independent of any particular game.

Proposition 1.9. 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 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}_{α} 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

$$X(\mathbf{p}, \mathbf{q}) = \sum_{\alpha=1}^{4} \pi_{\alpha}(\mathbf{p}) \pi_{\alpha}(\mathbf{q}) X_{i}$$
(2.1.1)

Then for any **p** we have

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

To conclude that μ is equivalent to ν it is sufficient (and necessary) to choose the \mathbf{q}_{α} 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 $B: {\bf R}^4 \times {\bf R}^4 \rightarrow {\bf R}$ 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}_{α} to be an orthonormal basis with respect to which *B* is diagonal. Then we have (for $\alpha \neq \beta$)

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

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 $\mathbf{p} \in S^3$ to maximize

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

Expression (2.2.1) is a (real) quadratic form in the coefficients $\pi_i(\mathbf{p})$ and hence is maximized (over S^3) on the intersection of 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:

$$(\alpha - \beta)(\alpha - \gamma)(\alpha - \delta)K(\mathbf{pq}_1) + (\beta - \alpha)(\beta - \delta)(\beta - \gamma)K(\mathbf{pq}_2)$$

+(\gamma - \alpha)(\gamma - \beta)(\gamma - \beta)(\beta - \beta)(\delta - \beta)(\delta

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

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

If **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 **G** be a two-by-two game with payoff pairs $(X_1, Y_1), \ldots, (X_4, Y_4)$ (listed in arbitrary order). **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 non-generic 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 **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, v; ν is supported on two orthogonal points p, pu, and the quadruple (p, pv, pu, 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 (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 (1.8) 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 (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 non-diagonal coefficients are all zero. Computing these coefficients explicitly and dividing by (non-zero) 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 $\alpha = \beta = \gamma = \delta = 1/4$, proving the theorem.

Case II. 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

$$\begin{pmatrix} \alpha & \beta & E^2\gamma + F^2\delta & E^2\delta + F^2\gamma \\ \beta & \alpha & E^2\delta + F^2\gamma & E^2\gamma + F^2\delta \\ E^2\gamma + F^2\delta & E^2\delta + F^2\gamma & \alpha & \beta \\ E^2\delta + F^2\gamma & E^2\gamma + F^2\delta & \beta & \alpha \end{pmatrix} \cdot \begin{pmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \end{pmatrix} = \begin{pmatrix} \lambda \\ \lambda \\ \lambda \\ \lambda \end{pmatrix}$$
(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, one 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, **u**, **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+Ejplayed 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 **G**) gives

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

so that by (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 $EF(\psi - \xi) = 0$. 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 (3.5), μ is supported on at most three points; we can rename so those points are 1, **u**, **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 **p** is any optimal response by Player One, apply (2.4) with $\delta = 0$ (and possibly $\gamma = 0$) to get

$$\sigma_1 K(\mathbf{p}) + \sigma_2 K(\mathbf{pu}) + \sigma_3 K(\mathbf{pv}) + \sigma_4 K(\mathbf{puv}) = 0$$
(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 is equal to zero. Then $\gamma \neq 0$ so a) holds.

By (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

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

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

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

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

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

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

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

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 **u** and **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

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

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

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 (3.5) we can assume that μ is supported on $\{1, \mathbf{u}, \mathbf{v}\}$. By (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 (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) we have

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

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 $P_1(j,\mu) = P_1(k,\mu)$, 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, **u** and ν is supported on two quaternions **p**, **pv** where either

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

Proof. Suppose 1 and **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) with $\mathbf{q}_1 = 1, \mathbf{q}_2 = \mathbf{u}, \gamma = \delta = 0$ gives

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

. This, plus the identical observation with the players reversed, estabilishes 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.

Example 3.11. We apply Theorem 3.3 to find all mixed strategy quantum equilibria in the Prisoner's Dilemma (1.3).

First, there are certainly equilibria of the form (3.3a): Each player chooses any four orthogonal quaternions and plays each with probability 1/4. Any such pair of strategies is an equilibrium; in any such equilibrium, each player earns an expected payoff of 9/4 = 2.25.

There are also equilibria of the form (3.3d); (1.3.2) is such an equilibrium. Up to equivalence, this is the only one. Sketch of proof: Player 2 plays 1 and \mathbf{v} equiprobably where $\mathbf{v} = Si + Tj + Uk$. Write down the first order conditions for Player One's problem to find a circle of optimal responses. Player One chooses \mathbf{p} and \mathbf{pv} on this circle; from the explicit first order conditions we can take \mathbf{p} proportional to

$$\left(U(21 - 24S^2 - 20T^2 + 3\sqrt{49 - 48S^2 - 40T^2}, T(23 - 24S^2 - 20T^2 + \sqrt{49 - 48S^2 - 40T^2}, 4S(6S^2 + 5t^2 - 6), 0\right)$$

(We can assume this choice of **p** because any two points on the optimal circle yield the same maximization problem for Player Two—this is essentially because two strategies are played equiprobably.) Now we can write down the first order conditions for Player Two's problem and require that they be satisfied at 1. For $S \neq 0, 1$, the resulting equations imply $T^2 < 0$ or $U^2 < 0$, contradiction. For S = 0, the first order conditions are automatically satisfied but the quaternion 1 minimizes rather than maximizes Player Two's payoff (an optimal response is i); this rules out S = 0. That leaves S = 1, which is (1.3.2).

An analogous search for equilibria of the form (3.3c) would require substantially more computational resources than are available to the author. (For one thing, there are now six independent first order conditions instead of three; for another, it now matters which points on his optimal circle Player One chooses.) However, such equilibria can be ruled out on the basis of [I], which characterizes all fully intertwined quadruples. Checking each of the possibilities is straightforward, though (very) tedious.

There are no equilibria of the form (3.3e) thanks to the remarks in (1.2).

Finally, I claim there are no equilibria of the form (3.3b). Suppose for example that Player Two's strategy is supported on $\{1, i, j\}$ with probabilities p, q, r. Then three of the strategies 1, i, j, k must return the same payoff to Player One; that is, three of the expressions

$$3p+q$$
 $p+3q+5r$ $5q+3r$ $5p+r$

must be equal. Together with p, q, r > 0 and p + q + r = 1, this implies (p, q, r) = (14/25, 9/25, 2/25) or (p, q, r) = (14/31, 12/31, 5/31). In the former case, Player 1's unique optimal play is k (for a payoff of 72/25, as opposed to 51/25 for plays of 1, i, j); thus we can rule this out. In the latter case, Player One's optimal responses are i, j, k so he plays these with some probabilities p', q', r'. But then Player One's payoffs for 1, i, j are p' + 5q', 3p' + 5r', 3q' + r', and for these to be equal we must have p' < 0 (specifically p' = -5/11); contradiction. Thus Player Two's strategy is not supported on $\{1, i, j\}$. The same calculation rules out the other three-element subsets of $\{1, i, j, k\}$ as well.

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. 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.

I have listed these fifteen families in the appendix; again, see [I] for the (tedious) proof that this list is exhaustive.

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 its corollary 4.1A apply to all two by two games (whether generic or not) and are of independent interest:

Theorem 4.1. Let **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 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 **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 (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.

Appendix: Intertwining

The following result is proved in [I]:

Theorem 4.1. Suppose $(\mathbf{p}, \mathbf{pv}, \mathbf{pu}, \mathbf{pvu})$ is fully interwined, where \mathbf{p} is an arbitrary unit quaternion and \mathbf{u}, \mathbf{v} are square roots of -1. Then, up to permuting i, j and k, and writing $\langle \mathbf{p}, \mathbf{q} \rangle$ for the circle generated by \mathbf{p} and \mathbf{q} , at least one of the following holds:

- i) $\mathbf{u} = i$, $\mathbf{v} \in \{\overline{\mathbf{p}}j\mathbf{p}, \overline{\mathbf{p}}k\mathbf{p}\}\$
- ii) $\mathbf{u} = i$, $\mathbf{v} \in \langle i, \overline{\mathbf{p}}i\mathbf{p} \rangle$
- ii') $\mathbf{u} \in \langle i, \overline{\mathbf{p}}i\mathbf{p} \rangle, \quad \mathbf{v} = \overline{\mathbf{p}}i\mathbf{p}$
- iii) $\mathbf{p} \in <1, j><1, i>$, $\mathbf{u} = i$, $\mathbf{v} \in \overline{\mathbf{p}} < i, k>\mathbf{p}$
- iii') $\mathbf{p} \in <1, i ><1, j >$, $\mathbf{v} = \overline{\mathbf{p}}i\mathbf{p}$, $\mathbf{u} \in <i, k >$
- iv) $\mathbf{u} = i$ and \mathbf{v} is (uniquely) determined by one of the following three conditions:

$$\mathbf{pv} \sim \left(-A(C^2 + D^2), -B(C^2 + D^2), C(A^2 + B^2), D(A^2 + B^2) \right)$$
$$\mathbf{pv} \sim \left(-C(BC + AD), D(BC + AD), B(AC - BD), A(AC - BD) \right)$$
$$\mathbf{pv} \sim \left(D(AC - BD), C(AC - BD), -A(BC + AD), B(BC + AD) \right)$$

where $\mathbf{p} = (A, B, C, D)$

iv') $\mathbf{v} = \overline{\mathbf{p}}i\mathbf{p}$ and \mathbf{u} is uniquely determined by one of the following three conditions:

$$\begin{aligned} \mathbf{pu} &\sim \left(-A(C^2 + D^2), -B(C^2 + D^2), C(A^2 + B^2), D(A^2 + B^2) \right) \\ \mathbf{pu} &\sim \left(-D(BD + AC), C(BD + AC), A(AD - BC), B(AD - BC) \right) \\ \mathbf{pu} &\sim \left(C(AD - BC), D(AD - BC), -B(BD + AC), A(BD + AC) \right) \end{aligned}$$

where $\mathbf{p} = (A, B, C, D)$

- v) $\mathbf{p} \in <1, i > \cup < j, k > \mathbf{u}, \mathbf{v} \in \{i\} \cup < j, k >$
- vi) $\mathbf{p} \in \langle i, j \rangle \langle 1, \mathbf{v} \rangle$, $\mathbf{u} = \mathbf{v} \in \langle i, j \rangle$
- vii) For some (A, B, C) with $A^2 + B^2 + C^2 = 1$ we have:

$$\mathbf{p} \sim (A, A, 0, 2C)$$
 or $\mathbf{p} \sim (A, -A, -2B, 0)$

and \mathbf{u}, \mathbf{v} given by one of the following pairs of expressions:

$$\mathbf{u} \sim (0, C - B, A, -A)$$
 and $\mathbf{v} \sim (0, A^2 + 2BC, A(B - C), A(C - B))$
or $\mathbf{u} \sim (0, -B - C, A, A)$ and $\mathbf{v} \sim (0, A^2 - 2BC, A(B + C), A(B + C))$

viii) $\mathbf{p} = (0, A, B, \pm B), \mathbf{u} = (0, X, Y, \pm Y)$ and $\mathbf{v} \sim (0, 2ABX - Y, 2B^2X, \pm 2B^2X)$

- ix) $\mathbf{p} = (AY 2BX, 0, BY, \mp BY)$, with **u** and **v** as in (viii).
- x) $\mathbf{u} = \mathbf{v} \perp \mathbf{p} \perp 1$
- xi) $\mathbf{u} = \mathbf{v} \perp i\mathbf{p} \perp 1$
- xii) $\mathbf{v} = (j \pm k)/\sqrt{2}$ $\mathbf{p} \in <1, i\mathbf{v} > \cup < i, i\mathbf{v} > \mathbf{u} = \overline{\mathbf{p}}i\mathbf{p}$
- xii') $\mathbf{u} = (j \pm k)/\sqrt{2}$ $\mathbf{p} \in <1, i\mathbf{u} > \cup < i, i\mathbf{u} > \mathbf{v} = i$
- xiii) $(\mathbf{p}, \mathbf{v}, \mathbf{u})$ is a real point on a certain one-dimensional algebraic variety \mathcal{X}_0 .

References

- [A] R. J. Aumann, "Subjectivity and Correlation in Randomized Strategies", J. of Mathematical Economics 1 (1974).
- [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 in Noncooperative Games", preprint, available at http://www.landsburg.com/dahl6.pdf
 - [EW] J. Eisert and M. Wilkens, "Quantum Games", J. of Modern Optics 47 (2000), 2543-2556
 - [EWL] J. Eisert, M. Wilkens and M. Lewenstein, "Quantum Games and Quantum Strategies", Phys. Reve. Lett. 83 (1999), 3077-3080
 - [I] S. Landsburg, "Intertwining", working paper available at

http://www.landsburg.com/twining4.pdf

- [M] P. La Mura, "Correlated Equilibria of Classical Strategic Games with Quantum Signals": September 2003
- [NE] S. Landsburg, "Nash Equilibrium in Quantum Games", RCER Working Paper #524, Rochester Center for Economic Research, February 2006