Quantum Game Theory

Steven E. Landsburg

An n-person game consists of

1. n sets S_i $(i = 1, \ldots, n)$.
2. n real-valued functions

$$P_i : S_1 \times \cdots \times S_n \to \mathbf{R} \qquad (i = 1, \ldots, n).$$

The set S_i is called *Player i's strategy space*. The function P_i is called *Player i's payoff function*.

This formulation is general enough to model pretty much any real-world strategic interaction: we take S_i to be the set of actions available to the character called Player i, we imagine that each player must choose some action, we imagine that those actions have some joint consequence and that P_i measures Player i's assessment of the value of that consequence. Given such a model, one can ask questions that tend to fall into two broad categories: First, what do we think the players *will* do (assuming, ordinarily, that they are sublimely selfish and sublimely rational)? Second, what do we think the players *ought* to do (according to some standard of fairness or justice or morality)? Questions of the first sort call for *solution concepts*; questions of the second sort call for *normative criteria*.

The most thoroughly studied solution concept is *Nash equilibrium*, an outcome that results when players maximize their own payoffs, taking other players' behavior as given. More precisely, an n-tuple of strategies $(s_1, \ldots, s_n)$ is a Nash equilibrium if for every i and for every $s \in S_i$,

$$P_i(s_1, \ldots, s_{i-1}, s_i, s_{i+1} \ldots s_n) \geq P_i(s_1, \ldots, s_{i-1}, s, s_{i+1}, \ldots, s_n).$$

Steven E. Landsburg is adjunct associate professor of economics at the University of Rochester. His email address is `landsbur@troi.cc.rochester.edu`.

Nash equilibrium either is or is not a good predictor of actual behavior depending on exactly what sort of behavior is being modeled.

Turning now to normative criteria (criteria intended to judge the *desirability* of outcomes), the least controversial is the criterion of *Pareto-Optimality*. Given two n-tuples of strategies $s = (s_1, \ldots, s_n)$ and $s' = (s_1, \ldots, s_n)$, we say that s (weakly) *Pareto dominates* s' if $P_i(s) \geq P_i(s')$ for all i; it is easy to verify that Pareto dominance is a partial order, and we say that s is Pareto optimal if it is maximal for this order.

In general, we have no guarantees of existence or uniqueness for Nash equilibria or for Pareto optima. When they do exist, Nash equilibria and Pareto optima might or might not coincide. The single most famous example in the history of game theory is the *Prisoner's Dilemma*, one form of which can be represented by the following matrix:

(1)

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

Here the rows are indexed by Player 1's strategy set $\{\mathbf{C}, \mathbf{D}\}$, the columns are indexed by Player 2's strategy set $\{\mathbf{C}, \mathbf{D}\}$, and the (i,j) entry is $(P_1(i,j), P_2(i,j))$. The associated story, which the reader can see enacted every Tuesday night on *N.Y.P.D. Blue*, goes as follows: Two prisoners have jointly committed a crime. They are separated and invited to testify against each other. Each receives the same menu of options: You're facing a ten-year sentence for this crime. You can either cooperate with your buddy (strategy **C**) by refusing to testify against him or defect from your criminal partnership (strategy **D**) by testifying. If you both cooperate (i.e., if you both stay silent), we'll have to convict you of a lesser crime, which will take three years off each of your sentences. But if you defect while your buddy cooperates, we'll take five years off your sentence (while he serves a full term). And if you both defect, we'll take one year off each of your sentences.

As the reader can (and should) verify, the unique Nash equilibrium $(\mathbf{D},\mathbf{D})$ is also the unique outcome that is *not* Pareto optimal. Rational selfish prisoners always choose the one strategy pair that both can agree is undesirable—in the sense that they would both prefer $(\mathbf{C},\mathbf{C})$.[1]

The simplest game without a Nash equilibrium is "odds and evens", represented by the game matrix

(2)

		Player Two	
		O	**E**
Player One	**O**	$(1,-1)$	$(-1,1)$
	E	$(-1,1)$	$(1,-1)$

Suppose there is a Nash equilibrium in which Player 1 plays Even. Then Player 2 plays Odd, so Player 1 plays Odd—contradiction; and similarly with Odd and Even reversed. Thus there is no Nash equilibrium, and without an alternative solution concept we are unable to predict anything other than paralysis on the part of both players. But of course anyone who has ever played this game knows what *actually* happens: players randomize their strategies, and each wins half the time.

One might be tempted to conclude that Nash equilibrium is the wrong solution concept for this game. A better conclusion is that the mathematical structure (2) is a poor model for the real-world game of odds and evens. A better model would allow for *mixed* (i.e. randomized) *strategies*. So we replace the strategy space $S_i = \{\text{Odd}, \text{Even}\}$ with the unit interval $S_i^+ = [0,1]$, using $p \in [0,1]$ to model the strategy "play Odd with probability p and Even with probability $1-p$". It is then natural to define new payoff functions

$$\begin{aligned} P_i^+(p,q) &= pqP_i(\text{Odd},\text{Odd}) \\ &\quad + p(1-q)P_i(\text{Odd},\text{Even}) \\ &\quad + (1-p)qP_i(\text{Even},\text{Odd}) \\ &\quad + (1-p)(1-q)P_i(\text{Even},\text{Even}). \end{aligned}$$

More generally, given any game **G** with finite strategy sets S_i, we define a new game $\mathbf{G}^+$ as follows: Let S_i^+ be the set of all probability distributions on S_i, and define

$$P_i^+ : S_1^+ \times \cdots \times S_n^+ \to \mathbf{R}$$

by

$$(3)\quad \begin{aligned} &P_i^+(s_1^+,\ldots,s_n^+) \\ &= \int_{S_1\times\cdots\times S_n} P_i(s_1,\ldots,s_n)ds_1^+(s_1)\ldots ds_n^+(s_n). \end{aligned}$$

(The restriction to games with finite strategy sets is so that we do not have to worry about convergence issues in (3).) One proves via standard fixed point theorems that the game $\mathbf{G}^+$ has at least one Nash equilibrium, the key point being that each S_i^+, unlike the original S_i, can be identified with a convex subset of a Euclidean space on which the P_i^+ are continuous.

Thus in the case of Odds and Evens, $\mathbf{G}^+$ is a better model of reality than **G** is. I want to argue that the same thing is true more generally: If **G** is any game with finite strategy spaces intended to model some real-world interaction, then $\mathbf{G}^+$ is always a better model of that same interaction. Here's why: In the real world, players must *communicate* their strategies either to each other or to a referee or to an interrogating detective, who then computes the payoffs. And as a practical matter, it is quite impossible for a referee or anyone else to prohibit the use of mixed strategies. Player 1 announces "I defect!" How can the referee know whether Player 1 arrived at this strategy through a legal deterministic process or an illegal random one?

Because there is no way to prohibit mixed strategies in practice, we might as well allow them in the model. More generally, whenever the real world imposes limits on referees' ability to observe and/or calculate, we should improve the model by adjusting the strategy spaces and payoff functions accordingly.

Quantum game theory begins with the observation that the technology of the (near?) future is likely to dictate that much communication will occur through quantum channels. For example, players might communicate their strategies to a referee via email composed on quantum computers. Such communication creates the prospect of new strategies whose use the referee cannot detect and

[1] *In fact, things are even worse than that. Each player has* **D** *as a dominant strategy, which means that* **D** *is the optimal play regardless of whether the other player chooses* **C** *or* **D**. *This is a far more powerful reason to anticipate an outcome of* $(\mathbf{D},\mathbf{D})$ *than the mere fact that* $(\mathbf{D},\mathbf{D})$ *is a Nash equilibrium.*

therefore cannot prohibit: Instead of cooperating or defecting (or randomizing between the two), a player might send a message that is some quantum superposition of the messages "I cooperate" and "I defect". To read the message, the referee must destroy the superposition, along with any evidence that the superposition ever existed, which makes superpositions effectively impossible to prohibit. What cannot be prohibited must be allowed; therefore, if we want to model accurately the behavior of games in which players have access to "quantum moves", we should expand our strategy spaces accordingly.

One might guess that a quantum move is just one more way to implement a mixed strategy, so that there is nothing new here for game theory. The physicist David Meyer [M] was the first to publish a counterexample to that guess. In Meyer's example a single coin is passed back and forth between two blindfolded players. The coin starts out heads up (call this state **H**). Player One has the option either to flip the coin or to return it unflipped. Then Player Two (still blindfolded so he doesn't know Player One's first move) has the same option: Flip or don't flip. And finally, Player One gets another turn. If the coin ends up in its initial state **H**, Player One wins. If it ends up in the opposite state **T** (tails up), Player Two wins.

Here Player One has four strategies ("flip, flip", "flip, don't flip", etc.). Player Two has two strategies ("flip" and "don't flip"). In any mixed strategy Nash equilibrium, Player Two flips with probability .5, Player One flips an even number of times with probability .5, and each wins half the games.

Now suppose we treat the coin as a quantum object. Its state is an equivalence class of nonzero vectors in the complex vector space spanned by **H** ("heads") and **T** ("tails"); two vectors are equivalent if one is a scalar multiple of the other. A physical operation on the coin corresponds to a unitary operation on the state space; in particular, we can set things up so the operation "not flip" is represented by the identity transformation and the operation "flip" is represented (with respect to the basis $\{\mathbf{H},\mathbf{T}\}$) by the unitary matrix

$$F = \begin{pmatrix} 0 & 1 \\ -i & 0 \end{pmatrix}.$$

Now suppose that Player One (but not Player Two) has access to a full array of quantum moves; that is, instead of simply flipping or not flipping, he can apply any unitary matrix he chooses. In particular, if he is clever, Player One will choose the matrix

$$U = \begin{pmatrix} \frac{-1-i}{2} & -\frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2} & \frac{-1+i}{2} \end{pmatrix}$$

on his first turn and the matrix U^{-1} on his second. Here's why that's clever: If Player Two fails to flip, then the net result of the three operations is

$$U^{-1} \circ I \circ U = I; \tag{4}$$

whereas if Player Two flips, the net result is

$$U^{-1} \circ F \circ U = \begin{pmatrix} \frac{-\sqrt{2}+i\sqrt{2}}{2} & 0 \\ 0 & \frac{\sqrt{2}-i\sqrt{2}}{2} \end{pmatrix}. \tag{5}$$

This is great for Player One, because both (4) and (5) map the state represented by **H** into itself (remember that any scalar multiple of **H** is equivalent to **H**). Thus *whether or not Player Two flips the coin*, Player One is *guaranteed* a win.

Meyer's example shows that quantum moves can be more powerful than mere mixed strategies, at least in a context where quantum moves are available to only one player. But of course it is more natural to ask what happens if both players are given access to a full set of quantum moves.

The first example of a full-fledged quantum game is due to Jens Eisert, Martin Wilkens, and Maciej Lewenstein [EWL]. Let **G** be a two-by-two game, that is, a game with two players, each of whom has a two-point strategy space, say $\{\mathbf{C},\mathbf{D}\}$. (The reader will have no difficulty extending this construction to *n*-by-*m* games.) Each player is given a coin that he returns to the referee either in its original state (to indicate a play of **C**) or flipped (to indicate a play of **D**). A player with access to quantum moves can act on his coin with any unitary matrix and therefore return it in a state

$$\alpha\mathbf{H} + \beta\mathbf{T}$$

where α and β are arbitrary complex numbers, not both zero. (Here and in everything that follows I will freely abuse notation by writing $\alpha\mathbf{H} + \beta\mathbf{T}$ both for an element of the vector space $\mathbf{C}^2$ and for the state represented by that vector.) When the referee observes the coin, it will appear to be unflipped or flipped with probabilities proportional to $|\alpha|^2$ and $|\beta|^2$. As long as the coins can be treated as independent quantum entities, then indeed all we have is a fancy way to implement a mixed strategy—in other words, nothing new for game theory.

Meyer's example was more interesting because both players acted on a single coin. Eisert, Wilkens, and Lewenstein (referred to henceforth as EWL) make their example more interesting by assuming the players' coins are *entangled* so that there is a single state space for the *pair* of coins. Explicitly, let $\mathbf{C}^2$ be the two-dimensonal complex vector space spanned by symbols **H** and **T**; then the state space for an entangled pair is

$$\mathbf{S} = (\mathbf{C}^2 \otimes \mathbf{C}^2 - \{0\})/\sim \tag{6}$$

where $\sim$ is the equivalence relation that identifies a vector with all its nonzero scalar multiples. As before, I will write, for example, $\mathbf{H} \otimes \mathbf{H}$ both for a vector in the space $\mathbf{C}^2 \otimes \mathbf{C}^2$ and for the state it represents. A physical operation on the first coin is represented by a two-by-two unitary matrix U acting on the state space $\mathbf{S}$ as $U \otimes 1$. A physical operation on the second coin is represented by a two-by-two unitary matrix V acting on the state space as $1 \otimes V^T$.

Now EWL conjure the following scenario: A pair of coins starts in the state[2]

$$(\mathbf{H} \otimes \mathbf{H}) + (\mathbf{T} \otimes \mathbf{T}).$$

As before, each player is handed one of the coins and invited to indicate a play of **C** by applying the identity matrix ("leaving the coin untouched") or to indicate a play of **D** by applying the matrix

$$F = \begin{pmatrix} 0 & 1 \\ -i & 0 \end{pmatrix}$$

("flipping the coin"). As long as the players restrict themselves to the two-point strategy space $\{\mathbf{C}, \mathbf{D}\}$, the pair of coins lands in one of the four states

(7a) $$\mathbf{CC} = (\mathbf{H} \otimes \mathbf{H}) + (\mathbf{T} \otimes \mathbf{T}),$$

(7b) $$\mathbf{CD} = (\mathbf{H} \otimes \mathbf{T}) - i(\mathbf{T} \otimes \mathbf{H}),$$

(7c) $$\mathbf{DC} = (\mathbf{H} \otimes \mathbf{T}) + i(\mathbf{T} \otimes \mathbf{H}),$$

(7d) $$\mathbf{DD} = (\mathbf{H} \otimes \mathbf{H}) - (\mathbf{T} \otimes \mathbf{T}).$$

The referee now performs an observation to determine which of the states (7a–d) the coins occupy and makes appropriate payoffs.

If players cannot be trusted to restrict themselves to the two strategies **C** and **D**, then the mathematical modeler should replace the game **G** with a new game $\mathbf{G}^Q$ that expands the strategy spaces accordingly. Player One's strategy space should consist of the operations that can be effected on the state space **S** (defined in (6)) via the action of unitary matrices on the first variable. Let $\mathbf{U}_2$ be the group of two-by-two unitary matrices. The matrices that fix **S** are the scalar matrices, which form a subgroup $\mathbf{S}^1 \subset \mathbf{U}_2$. Therefore, we define Player One's strategy space to be the group $\mathbf{U}_2/\mathbf{S}^1$.

Let $\mathbf{SU}_2 \subset \mathbf{U}_2$ be the subgroup of matrices with determinant one; then inclusion induces an isomorphism

$$\mathbf{SU}_2/\{\pm 1\} \to \mathbf{U}_2/\mathbf{S}^1,$$

so we can just as well define Player One's strategy space to be $\mathbf{SU}_2/\{\pm 1\}$. Moreover, the group $\mathbf{SU}_2$ can be identified with the group $\mathbf{S}^3$ of unit quaternions via the map

$$\mathbf{SU}_2 \to \mathbf{S}^3$$
$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \mapsto A + Bj.$$

The same analysis applies to Player Two. Thus we define the strategy spaces

(8) $$S_1^Q = S_2^Q = \mathbf{S}^3/\{\pm 1\} = \mathbf{RP}^3.$$

Using language loosely, I will often identify a strategy with either of the two quaternions that represent it.

Next we motivate definitions for the payoff functions. Start with the game **G**:

(9)

		Player Two	
		C	**D**
Player One	**C**	(X_1, Y_1)	(X_2, Y_2)
	D	(X_3, Y_3)	(X_4, Y_4)

Suppose Player One plays the quaternion **p** and Player Two plays the quaternion **q**. Write the product as

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

where the π_α are real numbers unique up to a sign, because **p** and **q** are defined up to a sign. Using the notation of (7a–d), a chase through the isomorphisms reveals that the coin is transformed from the initial state **CC** to a final state

$$\pi_1(\mathbf{pq})\mathbf{CC} + \pi_2(\mathbf{pq})\mathbf{CD} + \pi_3(\mathbf{pq})\mathbf{DC} + \pi_4(\mathbf{pq})\mathbf{DD}.$$

When the referee observes the coins' joint state, he observes each of the four outcomes with probabilities

$$\text{Prob}(\mathbf{CC}) = \pi_1(\mathbf{pq})^2 \quad \text{Prob}(\mathbf{CD}) = \pi_2(\mathbf{pq})^2$$
$$\text{Prob}(\mathbf{DC}) = \pi_3(\mathbf{pq})^2 \quad \text{Prob}(\mathbf{DD}) = \pi_4(\mathbf{pq})^2.$$

Thus we should define the payoff functions by

(11) $$P_1^Q(\mathbf{p}, \mathbf{q}) = \sum_{\alpha=1}^{4} \pi_i(\mathbf{pq})^2 X_i,$$

$$P_2^Q(\mathbf{p}, \mathbf{q}) = \sum_{\alpha=1}^{4} \pi_i(\mathbf{pq})^2 Y_i.$$

[2] *Everything to follow depends heavily on the assumption that the coins start in the "maximally entangled" state* $(\mathbf{H} \otimes \mathbf{H}) + (\mathbf{T} \otimes \mathbf{T})$. *If they were to start in a state of the form* $\mathbf{s} \otimes \mathbf{t}$, *then the construction that follows would only reconstruct the classical mixed strategy game* $\mathbf{G}^+$. *[EWL] studies families of quantum games parameterized by the choice of initial state.*

Equations (8) and (11) define the quantum game $\mathbf{G}^Q$ associated to the game $\mathbf{G}$ of equation (9). The quantum game $\mathbf{G}^Q$ is not at all the same as the mixed strategy game $\mathbf{G}^+$. Using mixed strategies, the players can jointly effect some but not all probability distributions over the four possible outcomes; there is, for example, no pair of mixed strategies that can effect the probability distribution

$$\text{Prob}\,(\mathbf{CC}) = 1/2 \quad \text{Prob}\,(\mathbf{CD}) = 0$$
$$\text{Prob}\,(\mathbf{DC}) = 0 \quad \text{Prob}\,(\mathbf{DD}) = 1/2.$$

By contrast, in the game $\mathbf{G}^Q$ any probability distribution at all is realizable; in fact, more is true: taking Player One's strategy as given, Player Two can choose a strategy that effects any desired probability distribution. (Proof: Let Player One choose strategy $\mathbf{p}$, and let $\mathbf{r}$ be an arbitrary unit quaternion; then Player Two can play $\mathbf{p}^{-1}\mathbf{r}$.)

Thus Nash equilibria must be a great rarity in quantum games; in fact, Nash equilibria exist only when there exists $\alpha \in \{1,2,3,4\}$ that maximizes both X_α and Y_α; in other words, Nash equilibria exist only when there is no conflict between the players.

But nobody would ever actually play the game $\mathbf{G}^Q$ anyway. Just as there is nothing to stop our players from adopting quantum strategies, there is also nothing to stop them from adopting *mixed* quantum strategies. So we really want to study the game

$$\mathbf{G}^! = (\mathbf{G}^Q)^+. \tag{12}$$

So far, we have defined the game $\mathbf{G}^+$ only when the game $\mathbf{G}$ has finite strategy spaces, which is certainly not the case for $\mathbf{G}^Q$. So to turn (12) into a definition, we must first give a more general definition of $\mathbf{G}^+$.

Definition. *A measurable game consists of*

1. n measure spaces $(S_i, \mathcal{O}_i)$ $(i = 1, \dots, n)$. (That is, S_i is a set and $\mathcal{O}_i$ is a σ-algebra on S_i.)
2. n bounded measurable functions

$$P_i : S_1 \times \cdots \times S_n \to \mathbf{R} \qquad (i = 1, \dots, n).$$

Definition. *Let G be a measurable game. Then a mixed strategy for Player i consists of a probability measure on the space $(S_i, \mathcal{O}_i)$.*

Definition. *Let G be a measurable game. Then we define a game $\mathbf{G}^+$ as follows: Let S_i^+ be the set of all probability measures on S_i and define*

$$P_i^+ : S_1^+ \times \cdots \times S_n^+ \to \mathbf{R}$$

by

$$P_i^+(s_1^+, \dots, s_n^+) = \int_{S_1 \times \cdots \times S_n} P_i(s_1, \dots, s_n)\, ds_1^+(s_1) \dots ds_n^+(s_n).$$

(We identify the original strategy space S_i with a subset of S_i^+ by identifying the point s with the probability measure concentrated on s.) Now if we equip $\mathbf{RP}^3$ with its Borel σ-algebra, then $\mathbf{G}^Q$ acquires, quite naturally, the structure of a measurable game. Thus equation (12) becomes a meaningful definition.

The payoff functions in $\mathbf{G}^!$ should be called $P_i^!$ or P_i^{Q+}, but I will just call them P_i.

In the game $\mathbf{G}^!$ there is always at least one Nash equilibrium, namely (μ, μ) where μ is the uniform probability distribution on $\mathbf{RP}^3$. There are usually more interesting equilibria as well. For example, let us return to the Prisoner's Dilemma (1). It is easy to verify that the following pair of mixed strategies constitutes a Nash equilibrium:

(13)
μ: Player 1 plays the quaternions 1 and k, each with probability 1/2.
ν: Player 2 plays the quaternions i and j, each with probability 1/2.

Indeed, for any unit quaternion $q = a + bi + cj + dk$, we have

$$P_2(\mu, \mathbf{q}) = 2a^2 + \frac{5}{2}b^2 + \frac{5}{2}c^2 + 2d^2,$$

so that $\mathbf{q}$ is an optimal response to μ if and only if $a = d = 0$. Thus $\mathbf{q} = i$, $\mathbf{q} = j$ are optimal responses; whence so is the strategy ν. Similarly with the players reversed.

In the Nash equilibrium (13), each player's payoff is 5/2, so (13) Pareto dominates the unique classical equilibrium $(\mathbf{D}, \mathbf{D})$ (where the payoffs are both 1). Nevertheless, (13) is still Pareto suboptimal, being dominated by $(\mathbf{C}, \mathbf{C})$.

More generally, we would like to classify the Nash equilibria in $\mathbf{G}^!$ where $\mathbf{G}$ is an arbitrary two-by-two game. The results that follow are from the forthcoming article [L].

Given a strategy μ, we define the *optimal response sets*

$$O_1(\mu) = \{\mathbf{p} \in \mathbf{RP}^3 \mid P_1(\mathbf{p}, \mu) \text{ is maximized}\},$$

$$O_2(\mu) = \{\mathbf{q} \in \mathbf{RP}^3 \mid P_2(\mu, \mathbf{q}) \text{ is maximized}\}.$$

Thus (μ, ν) is a Nash equilibrium if and only if ν is supported on $O_2(\mu)$ and μ is supported on $O_1(\nu)$. This leads us to ask: Which subsets of $\mathbf{RP}^3$ can occur as $O_i(\mu)$? The answer is:

Theorem. *For any μ, each $O_i(\mu)$ is a projective hyperplane in $\mathbf{RP}^3$.*

Proof. *We have*

$$P_1(\mathbf{p}, \mu) = \int P_1(\mathbf{p}, \mathbf{q}) d\mu(\mathbf{q}), \tag{14}$$

which, for fixed μ*, is a quadratic form in the coefficients* $\pi_\alpha(\mathbf{p})$ *and hence maximized (over* $\mathbf{S}^3$*) on the intersection of* $\mathbf{S}^3$ *with the linear subspace of* $\mathbf{R}^4$ *corresponding to the maximum eigenvalue of that form.* $O_1(\mu)$ *is the image in* $\mathbf{RP}^3$ *of that linear subspace. Similarly, of course, for* O_2*.*

The theorem (or more precisely, the proof of the theorem) has some immediate corollaries of considerable interest.

Corollary. *Let G be the arbitrary two-by-two game of expression (9). Then in any mixed strategy quantum Nash equilibrium, Player One earns a payoff of at least* $(X_1 + X_2 + X_3 + X_4)/4$*, and Player Two earns a payoff of at least* $(Y_1 + Y_2 + Y_3 + Y_4)/4$*.*

Proof. *The quadratic form (14) has trace* $X_1 + X_2 + X_3 + X_4$ *and hence a maximum eigenvalue of at least one fourth that amount. Similarly, of course, for Player Two.*

Corollary. *Suppose the game (9) is zero-sum, meaning that* $X_\alpha + Y_\alpha = 0$ *for all* α*. Then in any mixed strategy quantum equilibrium, Player One earns exactly* $(X_1 + X_2 + X_3 + X_4)/4$ *and Player Two earns exactly* $(Y_1 + Y_2 + Y_3 + Y_4)/4$*.*

To describe Nash equilibria in general, we need to describe probability measures on $\mathbf{RP}^3$ in general. Of course, there are a huge number of such measures, but fortunately they fall quite naturally into large equivalence classes. In particular, we say that two mixed quantum strategies μ and μ' are *equivalent* if for all mixed quantum strategies ν and for all $\alpha \in \{1, 2, 3, 4\}$, we have

$$\int_{\mathbf{RP}^3 \times \mathbf{RP}^3} \pi_\alpha(\mathbf{pq})^2 d\mu(\mathbf{p}) d\nu(\mathbf{q}) = \int_{\mathbf{RP}^3 \times \mathbf{RP}^3} \pi_\alpha(\mathbf{pq})^2 d\mu'(\mathbf{p}) d\nu(\mathbf{q})$$

where the π_α are the coordinate functions on the quaternions as in (10). That is, two strategies are equivalent if one can be substituted for the other without changing either player's payoffs in any game.

It turns out that equivalence classes of strategies are quite large. I proved in [L] that every mixed quantum strategy is equivalent to a strategy supported on (at most) four mutually orthogonal points (that is, four points that are the images of mutually orthogonal points in $\mathbf{S}^3$). Thus an equivalence class of measures can be identified with a "weighted frame" in $\mathbf{R}^4$, that is, a set of mutually orthogonal vectors, each assigned a positive weight in such a way that the weights add up to one.

That cuts the set of (equivalence classes of) potential Nash equilibria way down to size. But we can cut it down much further. Call a pair of weighted frames (μ, ν) *realizable* if it is a Nash equilibrium for some quantum game. A reasonable first guess is that every pair (μ, ν) is realizable. But the truth is quite the opposite: the main theorem of [L] establishes highly restrictive (and easily checkable) necessary conditions for realizability. Modulo some minor technical provisos, the theorem implies that for each μ there are only a small finite number of ν such that (μ, ν) is realizable. In many cases, that small finite number is zero; in the remaining cases, the relevant strategies ν are easy to describe. Given a particular game, this makes the search for Nash equilibria a tractable problem.

One can view quantum game theory as an exercise in pure mathematics: Given a game $\mathbf{G}$, we create a new game $\mathbf{G}^!$ and we study its properties. But game theory has historically been interesting primarily for its applications. As with quantum computing, the applications of quantum game theory lie in the future.[3] The immediate task is to prove theorems that we expect will be useful a generation from now.

References

[EWL] J. Eisert, M. Wilkens, and M. Lewenstein, Quantum games and quantum strategies, *Phys. Rev. Lett.* **83** (1999), 3077.

[L] S. Landsburg, Nash equilibria in quantum games, preprint available at `http://www.landsburg.com/qgt.pdf`.

[M] D. Meyer, Quantum strategies, *Phys. Rev. Lett.* **8** (1999), 1052.

[3] *It has been suggested that quantum games might have immediate applications in the theory of evolution, assuming that genetic mutations are driven by quantum events. To my knowledge there is no evidence to support this admittedly intriguing speculation.*