Quantum Strategies

by

Gordon B. Dahl

University of California, San Diego

and

Steven E. Landsburg

University of Rochester

ABSTRACT

We investigate the consequences of allowing players to adopt strategies which take advantage of quantum randomization devices. In games of full information, the resulting equilibria are always correlated equilibria, but not all correlated equilibria appear as quantum equilibria. The classical and quantum theories diverge further in games of private information. In the quantum context, we show that Kuhn's equivalence between behavioral and mixed strategies breaks down. As a result, quantum technology allows players to achieve outcomes that would not be achievable with any classical technology short of direct communication; in particular they do not occur as correlated equilibria. In general, in games of private information, quantum technology allows players to achieve outcomes that are Pareto superior to any classical correlated equilibrium, but not necessarily Pareto optimal. A simple economic example illustrates these points.

1. Introduction.

Quantum game theory investigates the behavior of strategic agents with access to quantum technology. Such technology can be employed in randomization devices and/or in communication devices. This paper is concerned strictly with quantum randomization. For other recent papers on quantum randomization in economics and game theory see, for example, [B] or [LaM]. For models of quantum communication see, for example, [EW], [EWL] or [La].

We thank Mark Bils, Alan Stockman, Paulo Barelli, David Miller, David Meyer, Vince Crawford, Val Lambson, Navin Kartik and Joel Sobel for helpful remarks.

Just as a classical mixed strategy can be thought of as a strategy conditioned on the realization of some classical random variable, so a quantum strategy can be thought of as a strategy conditioned on the value of some quantum mechanical observable. This is a non- trivial expansion of the strategy space, because quantum observables are not bound by the laws of classical probability. For example, if X, Y, Z and W are classical binary random variables, it is nearly trivial to prove that

$$\mathrm{Prob}(X \neq W) \leq \mathrm{Prob}(X \neq Y) + \mathrm{Prob}(Y \neq Z) + \mathrm{Prob}(Z \neq W) \tag{1.1}$$

The clear intuition is that the first and last elements of a sequence cannot differ unless two adjacent elements differ along the way. Nevertheless, the existence of observables violating (1.1) is predicted by quantum mechanics and confirmed by laboratory experiments. (The most glaring inconsistencies are avoided by the fact that neither X and Z nor Y and W can be observed simultaneously.)

In Section 2 we will provide a simple informal example illustrating the failure of (1.1) and its consequences for game theory. In the remainder of the paper, we will investigate the consequences of allowing players to adopt quantum strategies. As Levine ([Le]) has observed, the resulting equilibria (at least in games of full information) are always correlated equilibria in the sense of Aumann ([A]). But not all correlated equilibria appear as quantum equilibria. A correlated equilibrium E might not be sustainable in a given quantum environment because no pair of quantum strategies yields the outcome E. Moreover, even in the presence of such quantum strategies, E might fail to be deviation-proof in the quantum context.

The classical and quantum theories diverge further when we turn to games of private information. Here, in the quantum context, Kuhn's equivalence between mixed and behavioral strategies ([K]) breaks down. Of course a game of private information can always be modeled as a game of complete information with a more complex strategy space—but the point is that there is more than one way to construct that more complex strategy space. Classically, the constructions are equivalent; not so in the quantum case. As a result, quantum technology allows players to achieve outcomes that would not be achievable with any classical technology short of direct communication.

Players could in general do even better if they could condition their behavior on each others' signals as well as their own. However, it's important to recognize that our quantum devices allow nothing of the kind. Neither player receives any information about the other player's signal.

As a general rule, in games of private information, quantum technology allows players to achieve outcomes that are Pareto superior to any classical correlated equilibrium, but not necessarily Pareto optimal. Examples in the final section will illustrate both points.

A brief table of contents:

Section 2 offers a brief example to illustrate what we mean by the breakdown of (1.1); we will return to this simple example repeatedly throughout the paper.

Section 3 establishes some notation and vocabulary for talking about classical equilibria (including mixed strategy and correlated equilibria) in a form that is suitable for generalization to the quantum context. A game **G** is imbedded in a larger game by giving players a choice of random variables on which to condition their strategies. When each player is restricted to a single choice, we recover the notion of correlated equilibrium.

Sections 4 and 5 provide the quantum generalization for games of complete information. Here, instead of conditioning their strategies on classical random variables, players are permitted to condition their strategies on quantum observables. We observe that (i) there are quantum strategies that no classical random variables can mimic and (ii) all quantum equilibria are also correlated equilibria (though not vice-versa). Section 4 explains the basic physics and Section 5 applies it to game theory.

Sections 6 and 7 provide the classical and quantum generalizations for games where players receive private signals. We demonstrate that Kuhn's mixed/behavioral strategy equivalence breaks down in quantum environments when there is private information. In fact, while it is easy to define quantum games of behavioral strategies, we show that there is no reasonable way to define a quantum analogue of the Kuhnian games of mixed strategies. We also show that with private information there can be quantum equilibria which are in no sense equivalent to classical correlated equilibria.

In the remaining sections, we present some examples. Section 8 establishes some technical results that will be needed to compute the equilibria in those examples, and

Section 9 establishes the examples themselves. One technical result is deferred to the appendix.

2. Cats and Dogs.

We begin with an informal example that illustrates the failure of (1.1) and its significance for game theory. The analysis here is adapted from [CHTW]. We will return to this example in a more formal context in Section 7.

Example 2.1. Two players, who cannot communicate once the game is underway, are each asked one of two yes/no questions, e.g. "Do you like dogs?" or "Do you like cats?". Each player's question is chosen independently via a fair coin flip. The players both win if and only if their answers agree, unless they both get the "cats" question, in which case they win if and only if their answers disagree.

In a classical environment, it's clear that players can achieve a Pareto optimal outcome by always agreeing, which yields a success rate of 3/4. In particular, the players have nothing to gain by randomizing their responses.

Now equip Player i with a coin, which the player can observe after rotating it through either of two angles, C or D. The outcomes of these observations have the following probabilities:

If both coins are rotated through angle C

		Coin Two	
		H	**T**
Coin One	**H**	.15/2	.85/2
	T	.85/2	.15/2

If either coin is rotated through angle D

		Coin Two	
		H	**T**
Coin One	**H**	.85/2	.15/2
	T	.15/2	.85/2

(2.1.1)

Given these coins, each player can adopt the following strategy:

> **Strategy 2.1.2:** *If I am asked the "cat" question, I will rotate my coin through angle C, and if I am asked the "dog" question I will rotate my coin through angle D. Either way, I will answer "yes" if and only if the coin shows heads.*

A moment's reflection reveals that if both players adopt this strategy, they win 85% of the time, which is a clear improvement over the classical maximum of 75%.

Unfortunately for the players, no such coins exist in a world governed by the classical laws of physics and probability. To see this, define the following random variables:

X is the orientation (i.e. heads or tails) of Coin One, after rotation through Angle C.

Y is the orientation of Coin Two, after rotation through Angle D.

Z is the orientation of Coin One, after rotation through Angle D.

W is the orientation of Coin Two, after rotation through Angle C.

Then chart (2.1.1) reveals that $\mathrm{Prob}(X \neq Y) = \mathrm{Prob}(Y \neq Z) = \mathrm{Prob}(Z \neq W) = .15$, while $\mathrm{Prob}(X \neq W) = .85$, so that (1.1) — an easy theorem of classical probability theory — is violated.

Another way to say essentially the same thing is to note that no joint probability distribution for the random variables X, Y, Z and W can yield the values in Chart (2.1.1).

However, the laws of quantum mechanics do allow for the existence of such "coins" (actually, subatomic particles), which are routinely produced in physics laboratories and could plausibly be incorporated in the machinery of future quantum computers.[1]

Therefore players equipped with quantum coins can do better than players equipped with classical random number generators. On the other hand, they can't do arbitrarily well; in the example at hand, the laws of quantum mechanics set an upper limit of approximately 85% (more precisely, $\cos^2(\pi/8)$) for the expected fraction of wins.

Just as the notion of a mixed strategy captures the options available to players equipped with independent random number generators, we will introduce the notion of a *quantum strategy* to capture the options available to players equipped with "quantum coins".

Remark 2.2. Readers unfamiliar with the relevant physics might be tempted to conclude that some sort of communication must take place between the players in Example 2.1. Note however that no action by either player has any effect on the probability distribution of anything the other player can observe.

[1] The notion of observable quantities with no joint probability distribution is, for most people, highly counterintuitive — so much so that readers without the relevant physics background might suspect they must have misread something here. We want to reassure those readers that we meant what we said, and that this is standard textbook fare for physicists, who often say that one never develops an intuition for quantum mechanics; one just gets used to it.

In particular, beware the following fallacious argument:

Argument 2.2.1. *Suppose both players have agreed to play Strategy 2.1.2. Let p denote the probability that Player One says "yes" conditional on receiving the "dog" question. Then if Player Two receives the cat question, he knows that $p = .85$, and if Player Two receives the dog question, he knows that $p = .15$. Therefore Player Two, on receiving his question, learns something previously known only to Player One, which implies some form of communication.*

One problem with this argument is that Player One never knows the value of p, Therefore nothing that Player Two learns about the value of p can have been "communicated" from Player One.

By analogy, suppose that Players One and Two carry coins that are known to be the same color (because they were prepared that way by a referee). When Player Two first observes his own coin, he immediately learns the color of Player One's coin — but we daresay that nobody would want to call this an instance of communication. Careful reflection will convince the reader that the apparent "communication" in Example 2.1 is of exactly this nature. Quantum technology overcomes classical restrictions on *correlations*, but all of the classical restrictions on *communication* remain intact.[2]

Finally, we remark that the effects of these quantum coins could certainly be mimicked by a mediator who observes both questions, observes Player One's strategy, and hands Player Two a weighted coin that has an 85% chance of yielding the winning response. But the whole point here is to model the capabilities and limits of quantum technology, not of mediators. We can equally well imagine a mediator who coordinates the responses perfectly, but this is beyond the capability of our quantum coins.

3. Classical Game Theory.

[2] We note also that it is a well-established principle of physics that faster-than-light communication is impossible, whereas the apparent "communication" in this example is instantaneous. It follows that the appearance of communication — at least in any sense that a physicist (or for that matter an economist) would recognize — must be deceptive. As in the previous footnote, the fact that these correlations cannot be used to transmit information is, to many people, so counterintuitive that they think they must have misunderstood something, but this is in fact standard physics textbook fare.

In this section, we will review the basic notions of classical game theory (including correlated equilibria) in order to establish some (slightly nontraditional) notation and vocabulary that will be suitable for generalization to the quantum context.

Throughout this section we fix a two-player game $\mathbf{G}$ with strategy sets S_1, S_2 and payoff functions P_1, P_2. For simplicity, we usually assume S_1 and S_2 are finite.

We also fix a probability space Ω, which for concreteness we can take to be the unit interval. A *random variable* always means a random variable with domain Ω.

In the traditional formulation of game theory, a mixed strategy is a probability distribution on the strategy space S_i. It will be more convenient for us to think of a mixed strategy as a random variable with values in S_i. Because a given probability distribution can be induced by many different random variables, this leads to a great proliferation of strategies, many of which are essentially interchangeable, and which we will want to think of as equivalent. The defintions in 3.1 will clarify the notion of equivalence.

Definitions 3.1. Two stategies $s, t \in S_1$ are *equivalent* if:

$$\text{For all } u \in S_2 \text{ we have } P_1(s,u) = P(t,u) \text{ and } P_2(s,u) = P_2(t,u) \tag{3.1.1}$$

Two strategies in S_2 are equivalent if they satisfy the obvious condition symmetric to (3.1.1).

We define the game $\overline{\mathbf{G}}$ by replacing S_i with the set of all equivalence classes of strategies for player i (and retaining the obvious payoff functions).

We say that two games $\mathbf{G}$ and $\mathbf{H}$ are *equivalent* if $\overline{\mathbf{G}}$ is isomorphic to $\overline{\mathbf{H}}$.

Definition 3.2. A *classical environment* is a pair $E = (\mathcal{P}_1, \mathcal{P}_2)$ where each $\mathcal{P}_i$ is a set of measurable partitions of Ω.

In what follows we fix a classical environment $E = (\mathcal{P}_1, \mathcal{P}_2)$. Let $\mathcal{X}_i$ be the set of S_i-valued random variables that are measurable with respect to some partition in $\mathcal{P}_i$.

Definition 3.3. The game $\mathbf{G}(E) = \mathbf{G}(\mathcal{P}_1, \mathcal{P}_2)$ is defined as follows:

- Player i's strategy set is $\mathcal{X}_i$.
- Player i's payoff function is

$$P_i(X,Y) = \int_{S_1 \times S_2} P_i(x,y) d\mu_{X,Y}(x,y)$$

where $\mu_{X,Y}$ is the probability distribution on $S_1 \times S_2$ induced by (X, Y).

(We abuse notation slightly by using the same notation P_i for the payoff functions in $\mathbf{G}$ and in $\mathbf{G}(E)$.)

We view the $\mathbf{G}$-strategy set S_i as contained in the $\mathbf{G}(E)$-strategy set $\mathcal{X}_i$ by identifying $s \in S_i$ with the random variable that takes $\mathbf{s}$ as its only value. We call these the *pure strategies*.

Example 3.4. Suppose that for each i, the random variables in $\mathcal{X}_i$ induce every possible probability distribution on the strategy set S_i. Suppose also that every partition in $\mathcal{P}_1$ is independent from every partition in $\mathcal{P}_2$ (so that every random variable in $\mathcal{X}_1$ is independent of every random variable in $\mathcal{X}_2$). Then $\mathbf{G}(E)$ is equivalent (in the sense of 3.1) to the classical game of mixed strategies associated to $\mathbf{G}$.

We will sometimes abuse language by calling $\mathbf{G}(E)$ *the* game of mixed strategies associated to $\mathbf{G}$, though the various choices for E yield games that are only equivalent, not isomorphic.

Notation 3.5. If $\mathcal{X}$ is any set of random variables, we write $\mathcal{P}(\mathcal{X})$ for the corresponding set of partitions of Ω. If $\mathcal{X}_i$ is a set of S_i-valued random variables, we define the environment $E(\mathcal{X}_1, \mathcal{X}_2) = (\mathcal{P}_1, \mathcal{P}_2)$.

Example 3.6. Let X and Y be S_1-valued and S_2-valued random variables. It is immediate that the probability distribution induced on $S_1 \times S_2$ by (X, Y) is (in the usual sense) a correlated equilbirium in $\mathbf{G}$ if and only if (X, Y) is a Nash equilibrium in the game $\mathbf{G}(\{X\}, \{Y\})$.

Observation 3.7. Let E be any classical environment for $\mathbf{G}$ and suppose that (X, Y) is a Nash equilibrium in the game $\mathbf{G}(E)$. Then (X, Y) is a correlated equilibrium in $\mathbf{G}$.

The converse to 3.7 does not hold:

Example 3.8. Let $S_1 = S_2 = \{\mathbf{H}, \mathbf{T}\}$.

Let $\mathbf{G}$ be the game with the following payoffs:[3]

---

[3] In future examples, players will use coin flips to choose their strategies; therefore we've called the strategies $\mathbf{H}$ and $\mathbf{T}$ for "heads" and "tails".

		Player Two	
		H	**T**
Player One	**H**	$(0,0)$	$(2,1)$
	T	$(1,2)$	$(0,0)$

Let X, Y and W be binary random variables such that:

$$\text{Prob}(X = W = \mathbf{H}) = \text{Prob}(X = W = \mathbf{T}) = 1/8 \qquad \text{Prob}(X \neq W = \mathbf{H}) = \text{Prob}(X \neq W = \mathbf{T}) = 3/8$$

$$\text{Prob}(Y = W = \mathbf{H}) = \text{Prob}(Y = W = \mathbf{T}) = 1/12 \qquad \text{Prob}(Y \neq W = \mathbf{H}) = \text{Prob}(Y \neq W = \mathbf{T}) = 5/12$$

Let $E = E(\{X, Y\}, \{W\})$.

It is easy to check that both (X, W) and (Y, W) yield correlated equilibria in $\mathbf{G}$. But (X, W) is not an equilibrium in the game $\mathbf{G}(E)$, though (Y, W) is.

Remark 3.9. Our modeling assumptions allow Player One to observe the realization of either X or Y but not both. This is because $\mathcal{P}_1$ consists of two partitions. An alternative model would define $\mathcal{P}_1$ to consist of a single partition that refines both of these, thus allowing Player One to play a strategy contingent on the realizations of both X and Y. Depending on the intended real-world implementation, this might or might not be a better model.

Remark 3.10. The moral of this section is this: Every correlated equilibrium (X, Y) occurs as the equilibrium in a game of the form $\mathbf{G}(E)$. (In fact, we just take $E = \mathcal{P}(X, Y)$ as defined in (3.4).) However, the same (X, Y) might *not* be a correlated equilibrium in some other game $\mathbf{G}(E')$.

This will become important when we get to the quantum analogue, because there (unlike here) physical considerations will dictate the choice of the quantum environment. Thus, given a correlated equilibrium (X, Y), the interesting question will not be "Is (X, Y) an equilibrium in *some* environment E?" (to which the answer is trivially yes), but "Is

(X, Y) an equilibrium in a *particular* environment E?" (to which the answer is sometimes yes and sometimes no).

4. Quantum Measurements.

Our next goal is to further expand players' options by allowing them to make quantum measurements. In this section we will explain what that means, and in the next section we will incorporate these quantum measurements into our formal model.

Discussion 4.1: States. A classical coin is in one of two states, "heads" (which we denote $\mathbf{H}$) or "tails" ($\mathbf{T}$). A "quantum coin" can be in any state of the form

$$\alpha\mathbf{H} + \beta\mathbf{T} \tag{4.1.1}$$

where α and β are complex numbers, not both zero.

We view (4.1.1) as an element of the two-dimensional complex vector space spanned by the symbols $\mathbf{H}$ and $\mathbf{T}$. Two such vectors represent identical states if one is a (non-zero) scalar multiple of the other.

Thus, strictly speaking, a *state* is not a vector but an equivalence class of vectors. Nevertheless, we will often abuse language by using the same expression (4.1.1) for both a vector and the state that it represents.

A heads/tails measurement of a coin in state (4.1.1) yields the outcome "heads" or "tails" with probabilities proportional to $|\alpha|^2$ and $|\beta|^2$. When such a measurement is made, the state immediately changes to either $\mathbf{H}$ or $\mathbf{T}$, depending on the measurement's outcome.

It goes without saying that no ordinary-sized coin obeys these laws of quantum physics, but spin-1/2 particles such as electrons do, with "heads" and "tails" replaced by "spin up" and "spin down". For ease of comparison with the classical case, we will continue to speak of "quantum coins".

Discussion 4.2: Transformations. Each physical action (such as rotating the coin through a pariticular angle) corresponds to some *unitary* transformation U of the state space.[4] Performing the action transforms the penny's state from ϕ to $U\phi$.

[4] A unitary transformation is an invertible linear transformation U such that $\overline{U}^T = U^{-1}$. The overbar denotes complex conjugation.

Because states are defined only up to multiplication by nonzero scalars, we can restrict attention to unitary transformations with determinant 1. These are called *special unitary transformations* and are represented by matrices

$$\begin{pmatrix} P & Q \\ -\overline{Q} & \overline{P} \end{pmatrix} \tag{4.2.1}$$

where P and Q are complex numbers satisfying $|P|^2 + |Q|^2 = 1$.

Thus, for example, if a coin in state (4.1.1) is subjected to a physical operation corresponding to the transformation (4.2.1), then it is transformed into the new state

$$\alpha(P\mathbf{H} - \overline{Q}\mathbf{T}) + \beta(Q\mathbf{H} + \overline{P}\mathbf{T}) = (\alpha P + \beta Q)\mathbf{H} + (-\alpha\overline{Q} + \beta\overline{P})\mathbf{T}$$

Discussion 4.3: Randomization. A quantum coin is neither more nor less useful than a classical randomizing device. Given a coin in the state (4.1.1), a player can change the state at will by applying a unitary transformation. Once the coin is in a state $\gamma\mathbf{H} + \delta\mathbf{T}$, it acts just like a weighted coin that comes up heads or tails with probabilities proportional to $|\gamma|^2$ and $|\delta|^2$. All the new phenomena arise not from *single* coins but from interactions between *pairs* of coins.

Discussion 4.4: Entanglement. Once a pair of coins have come into contact, they no longer occupy their own states. Instead, the pair occupies a state jointly represented by a non-zero vector

$$\alpha\mathbf{HH} + \beta\mathbf{HT} + \gamma\mathbf{TH} + \delta\mathbf{TT} \tag{4.4.1}$$

As in the one-coin case, any non-zero multiple of this expression represents the same state. Measurements yield the outcomes (heads,heads), (heads,tails), and so forth with probabilities proportional to $|\alpha|^2$, $|\beta|^2$ and so forth. These probabilities hold even when the coins are examined at physically remote locations.

Coins that occupy a joint state of the form (4.4.1) are called *entangled*.

A physical operation on the first coin is represented by a unitary matrix which we can take to be of the form (4.2.1). This transformation has the following effect on basis elements:

$$\mathbf{HH} \mapsto P\mathbf{HH} - \overline{Q}\mathbf{TH}$$

$$\mathbf{HT} \mapsto P\mathbf{HT} - \overline{Q}\mathbf{TT}$$

$$\mathbf{TH} \mapsto Q\mathbf{HH} + \overline{P}\mathbf{TH}$$

$$\mathbf{TT} \mapsto Q\mathbf{HT} + \overline{P}\mathbf{TT}$$

Its action on a general state of the form (4.4.1) is determined by these rules plus linearity. Likewise, the same operation applied to the second coin has the following effect:

$$\mathbf{HH} \mapsto P\mathbf{HH} - \overline{Q}\mathbf{HT}$$

$$\mathbf{HT} \mapsto Q\mathbf{HH} + \overline{P}\mathbf{HT}$$

$$\mathbf{TH} \mapsto P\mathbf{TH} - \overline{Q}\mathbf{TT}$$

$$\mathbf{TT} \mapsto Q\mathbf{TH} + \overline{P}\mathbf{TT}$$

Notation 4.5. Start with a coin in some state ξ. Let Player One apply the transformation U_1 and let Player Two apply the transformation U_2. The resulting state, computed according to the rules of 4.4, is denoted $(U \otimes 1)\xi(1 \otimes V)$.

(This notation will be familiar to readers familiar with the yoga of multilinear algebra; others can simply take it as a definition.)

Example 4.6. Suppose the two coins start in the *maximally entangled state* **HH**+**TT**. Players One and Two apply transformations U and V to the first and second coins.

The reader may check that the resulting state is given by

$$(U \otimes 1)\xi(1 \otimes V) = \alpha\mathbf{HH} + \beta\mathbf{HT} + \gamma\mathbf{TH} + \delta\mathbf{TT}$$

where

$$UV^T = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$$

In particular, suppose that

$$U = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -sin(\theta) & \cos(\theta) \end{pmatrix} \qquad V = \begin{pmatrix} \cos(\phi) & \sin(\phi) \\ -sin(\phi) & \cos(\phi) \end{pmatrix}$$

for some $\theta, \phi \in [0, 2\pi)$. (These transformations are implemented physically by rotating the coins through the angles 2θ and 2ϕ.) Then the resulting state is

$$(U \otimes 1)(\mathbf{HH} + \mathbf{TT})(1 \otimes V) = \cos(\theta - \phi)\mathbf{HH} + \sin(\theta - \phi)\mathbf{HT} - \sin(\theta - \phi)\mathbf{TH} + \cos(\theta - \phi)\mathbf{TT}$$

so that the resulting probability distribution over strategy pairs is

$$\text{Prob}(\mathbf{H}, \mathbf{H}) = \text{Prob}(\mathbf{T}, \mathbf{T}) = \cos^2(\theta - \phi)/2 \qquad \text{Prob}(\mathbf{H}, \mathbf{T}) = \text{Prob}(\mathbf{T}, \mathbf{H}) = \sin^2(\theta - \phi)/2$$

Note that the unitary operators U and V affect the join probability distribution of the two coins, but not the individual probability distributions, so that each coin always turns up $\mathbf{H}$ with probability $1/2$. More generally, the reader can check that regardless of the initial state, no choice of U can affect the probability distribution of Player Two's outcomes (nor, of course, can the choice of V affect the probability distribution for Player One). This follows from the mathematics, or, if you prefer, from the physical principle that no influence can travel faster than light.

Even if Player One chooses to ignore his coin (planning, perhaps, to play a pure strategy or a classical mixed strategy) then Player Two's coin reverts to a randomization device ala Discussion 4.3.

5. Quantum Environments.

In this section we will formalize the options available to players with quantum coins. We view the possession of a quantum coin as analogous to the "possession" of a set of measurable partitions as in Section 3.2. That is, it leads naturally to an expansion of the player's strategy set.

For ease of exposition, we will restrict ourselves to games $\mathbf{G}$ in which the strategy spaces S_1 and S_2 each have cardinality two. Everything generalizes easily to the case of arbitrarily large finite strategy spaces, and somewhat less easily to infinite strategy spaces.

Definition 5.1. Let $\mathbf{G}$ be a two-player game with strategy spaces $S_1 = S_2 = \{\mathbf{H}, \mathbf{T}\}$. Then a *quantum environment* for $\mathbf{G}$ is a non-zero vector in the complex vector space spanned by $\mathbf{HH}, \mathbf{HT}, \mathbf{TH}, \mathbf{TT}$. ($\xi$ represents the initial state of a pair of quantum coins, and is thus analogous to the pair of measurable partitions introduced in Section 3.)

In what follows, we fix a quantum environment ξ.

Notation and Conventions 5.2. We write $\mathcal{U}$ for the set of unitary transformations of the vector space spanned by $\mathbf{H}$ and $\mathbf{T}$.

Discussion 5.3. We want to define a game $\mathbf{G}(\xi)$ that expands each player's strategy space by allowing him to randomize over pure strategies via the selection of a unitary operator. This will be the analogue of the game $\mathbf{G}(E)$ in Definition 3.3. One small difference is that in the classical environment of 3.3, pure strategies are automatically included in the strategy sets $\mathcal{X}_i$ whereas they are not automatically included in the set $\mathcal{U}$; therefore we have to append them.[5] Thus $\mathbf{G}(\xi)$ will be a game in which Player i's strategy set is

$$\mathcal{U}_i = \mathcal{U} \cup S_i \tag{5.3.1}$$

.

The payoff functions are defined in the obvious way; a strategy pair imposes a probability distribution on $S_1 \times S_2$ and we compute payoffs as expected values with respect to this probability distribution. The next two definitions will make this precise.

Definition 5.4. With notation as above, let $(U, V) \in \mathcal{U}_1 \times \mathcal{U}_2$.

We define a probability distribution μ_{UV} on $S_1 \times S_2$ as follows:

a) If U and V are both unitary operators, write

$$(U \otimes 1)\xi(1 \otimes V) = \alpha\mathbf{HH} + \beta\mathbf{HT} + \gamma\mathbf{TH} + \delta\mathbf{TT}$$

(where the left hand side is defined in Section 4.5). Then the probabilities associated to $\mathbf{HH}, \mathbf{HT}$, and so forth are proportional to $|\alpha|^2, |\beta|^2$ and so forth.

b) If U is a unitary operator and $V = s$ is a pure strategy, write

$$(U \otimes 1)\xi(1 \otimes I) = \alpha\mathbf{HH} + \beta\mathbf{HT} + \gamma\mathbf{TH} + \delta\mathbf{TT}$$

(where the I on the left hand side is the identity transformation). Then $\mu_{U,s}(\mathbf{H}, s)$ and $\mu(U, s)(\mathbf{T}, s)$ add to one and are proportional to the squared norms of the coefficients on $\mathbf{H}s$ and $\mathbf{T}s$.

[5] An alternative formalism would require players to observe their quantum coins and play accordingly. We think it is far more natural to allow players the option of discarding their quantum coins and adopting pure strategies.

c) If $U = s$ is a pure strategy and V is a unitary operator, we define μ_{UV} analogously to part b).

d) If $U = s$ and $V = t$ are pure strategies, the probability distribution μ_{UV} is concentrated on (s,t).

Definition 5.5. Given a game $\mathbf{G}$ and a quantum environment ξ, define a new game $\mathbf{G}(\xi)$ as follows:

a) Player i's strategy set is $\mathcal{U}_i = \mathcal{U} \cup S_i$ (as in 5.3.1).

b) The payoff functions are given by

$$P_i(X,Y) = \int_{S_1 \times S_2} P_i(X_1, X_2) d\mu_{XY}(s,t)$$

where $\mu_(XY)$ is the probability distribution associated to X and Y according to Definition 5.4.

Example 5.6. Let $\mathbf{G}$ be the game of Example 3.8 and let $\xi = \mathbf{HH} + \mathbf{TT}$.

In the game $\mathbf{G}(\xi)$, suppose that Player One chooses a unitary operator U. Then, in light of Example 4.6, Player 2 clearly optimizes by choosing a unitary operator V so that

$$UV^T = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{5.6.1}$$

(thereby putting all the probability weight on the outcomes **HT** and **TH**). The same is true with the players reversed, so the equilibria are precisely the pairs (U_1, U_2) satisfying (5.6.1). In any equilibrium, the outcomes $(2,1)$ and $(1,2)$ are each realized with probability 1/2.

It is easy to check that this is a correlated equilibrium, but it is more than that. For example, the correlated equilibria (X,W) and (Y,W) of Example 3.8 are not sustainable as quantum equilibria in this environment.

Remarks 5.7. We want to single out those quantum environments that can be mimicked by classical technology.

To that end, continue to let S_i be Player i's initial strategy set and let Ω be the unit interval, thought of as a probability sample space. Let $\mathcal{F}_i$ be the set of all S_i-valued random variables defined on Ω. Given $(F,G) \in \mathcal{F}_1 \times \mathcal{F}_2$, let μ_{FG} be the probability distribution on $S_1 \times S_2$ induced by the random variable $F \times G$.

We'll say that the quantum enviroment ξ is *classical* if each strategy in $\mathcal{U}_i$ can be mimicked by some random variable in $\mathcal{F}_i$. More precisely:

Definition 5.8. With notation as in 5.7, we say that ξ is *classical* if there are maps

$$\phi_i : \mathcal{U}_i \to \mathcal{F}_i$$

such that for any $(U, V) \in \mathcal{U}_1 \times \mathcal{U}_2$, the probability distributions μ_{UV} and $\mu_{\phi_1(U)\phi_2(V)}$ are identical.

(Reminder: μ_{UV} is defined in 5.4.)

Example 5.9. The quantum environment $\xi = \mathbf{HH} + \mathbf{TT}$ is not classical.

To see this, define, for each $\theta \in [0, 2\pi)$, the following unitary operator:

$$M(\theta) = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$$

(Physically, the operator $M(\theta)$ corresponds to rotating the coin through the angle 2θ.)

Construct the $\{\mathbf{H}, \mathbf{T}\}$-valued "random variable" $\hat{M}_i(\theta)$ by applying the operation $M(\theta)$ to coin i and then observing the coin's orientation.

Now let

$$X = \hat{M}_1(-\pi/8) \quad Y = \hat{M}_2(0) \quad Z = \hat{M}_1(\pi/8) \quad W = \hat{M}_2(\pi/4)$$

Then by the computations in 4.6, we have

$$\text{Prob}(X \neq Y) = \text{Prob}(Y \neq Z) = \text{Prob}(Z \neq W) = \sin^2(\pi/8) \approx .15$$

$$\text{Prob}(X \neq W) = \sin^2(3\pi/8) \approx .85$$

so that condition (1.1) is violated. Thus no classical random variables X, Y, Z, W can mimic the effects of the strategies $M(0), M(\pi/8), M(\pi/4), M(3\pi/8)$.

Remark 5.10. It is not hard to prove (and will be immediately obvious to readers with appropriate physics backgrounds) that ξ is classical if and only if it is of the form

$$\xi = \alpha\mathbf{HH} + \beta\mathbf{HT} + \gamma\mathbf{TH} + \delta\mathbf{TT}$$

where $\alpha\delta - \beta\gamma = 0$.

Remark and Notation 5.11. A quantum environment ξ expands players' strategy sets by allowing them to make quantum observations. We can expand the strategy sets still further by allowing players to randomize over those observations: If E is a classical environment (as defined in 3.2) we can consider the game $\mathbf{G}(\xi)(E)$. We will abbreviate this game by $\mathbf{G}(\xi, E)$.

Example 5.12. Let $\mathbf{G}$ be a two by two game and let ξ be the quantum environment $\xi = \mathbf{HH} + \mathbf{TT}$. Then if both players adopt quantum strategies, either player can force a probability distribution that is concentrated either on the main diagonal or the off-diagonal (with the the two boxes on that diagonal represented equiprobably). (If Player One plays U, player Two can play $\overline{U}$ to force the main diagonal, or $\overline{U} \cdot \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ to force the off-diagonal.) Unless the players agree on which diagonal is more desirable, this means there can be no equilibrium in pure quantum strategies, but there can certainly be equilibria in mixed quantum strategies, and/or correlated equilbria in quantum strategies i.e. equilibria in games of the form $\mathbf{G}(\xi, E)$.

6. Games of Private Information.

We begin this section with a slight reformulation of the classical theory of games with private information. Our setup is clearly equivalent to the standard formulation, but better suited to quantum generalization.

Definition 6.1. A (two-player) *game of private information* consists of two *strategy sets* S_i, two *signal sets* $\mathcal{A}_i$, a probability distribution ν on $\mathcal{A}_1 \times \mathcal{A}_2$, and two payoff functions

$$P_i : \mathcal{A}_1 \times \mathcal{A}_2 \times S_1 \times S_2 \longrightarrow R$$

Unless stated otherwise, we will assume the strategy and signal sets are finite, though all this could be generalized.

We call these "games of private information" because all of our equilibrium concepts will assume (in effect) that Nature selects a signal pair $(a, b) \in \mathcal{A}_1 \times \mathcal{A}_2$ and that each player knows only his own signal.

Notation 6.2. For any sets X, Y, we write $\mathrm{Hom}(X, Y)$ for the set of all functions from X to Y.

In the discussion to follow, we fix a game of private information $\mathbf{G} = (S_1, S_2, \mathcal{A}_1, \mathcal{A}_2, \nu, P_1, P_2)$.

Definition 6.3. The *associated game* $\mathbf{G}^\#$ has strategy sets $S_i^\# = \mathrm{Hom}(\mathcal{A}_i, S_i)$ and payoff functions

$$P_i^\#(F_1, F_2) = \int_{\mathcal{A}_1 \times \mathcal{A}_2} P_i(A_1, A_2, F_1(A_1), F_2(A_2))d\nu$$

We think of $\mathbf{G}^\#$ as the game that results from $\mathbf{G}$ when we introduce contingent strategies.

A *Nash equilibrium* in $\mathbf{G}$ is (by definition) a Nash equilibrium in $\mathbf{G}^\#$.

Reminders 6.4. Now we will allow players to randomize. As in Section 3, we work with a fixed sample space Ω. Recall from Definition 3.2 that a *classical environment* $E = (\mathcal{P}_1, \mathcal{P}_2)$ is a pair of sets of measurable partitions of Ω, and that $\mathcal{X}_i$ is the set of S_i-valued random variables that are measurable with respect to some partition in $\mathcal{P}_i$.

In the discussion to follow we fix an environment $E = (\mathcal{P}_1, \mathcal{P}_2)$

Definition 6.5. Given $\mathbf{G}$ and E as above, we define a new game of private information $\mathbf{G}(E)$ as follows:

- The information sets are the $\mathcal{A}_i$
- The strategy sets are the $\mathcal{X}_i$
- The payoff functions are given by

$$P_i(A_1, A_2, X, Y) = \int_{S_1 \times S_2} P_i(A_1, A_2, x, y)d\mu_{X,Y}(x, y)$$

  where $\mu_{X,Y}$ is the probability distribution on $S_1 \times S_2$ induced by (X, Y).

Remarks 6.6. Starting with the game of private information $\mathbf{G}$ and the environment E, we can first introduce contingent strategies, yielding the ordinary game $\mathbf{G}^\#$ (see 6.3), and then allow players to randomize, forming the game $\mathbf{G}^\#(E)$ (see 3.3).

Alternatively, we can first allow players to randomize, forming the game of private information $\mathbf{G}(E)$ (see 6.5) and then introduce contingent strategies, forming the game $\mathbf{G}(E)^\#$ (see 6.3 again).

These games are equivalent (in the sense of 3.1) to Kuhn's games of mixed and behavioral strategies ([K]) and are therefore equivalent to each other.

In fact, more is true; they are isomorphic. Explicitly: in $\mathbf{G}^{\#}(E)$, a strategy is an element of $\mathrm{Hom}(\Omega, \mathrm{Hom}(\mathcal{A}_i, S_i))$, whereas in $\mathbf{G}(E)^{\#}$, a strategy is an element of $\mathrm{Hom}(\mathcal{A}_i, \mathrm{Hom}(\Omega, S_i))$. The strategy sets are carried back and forth to each other by the inverse bijections

$$\Phi : \mathrm{Hom}(\Omega, \mathrm{Hom}(\mathcal{A}_i, S_i)) \longleftrightarrow \mathrm{Hom}(\mathcal{A}_i, \mathrm{Hom}(\Omega, S_i)) : \Psi$$

defined by $((\Phi(f))(a))(\omega) = (f(\omega))(a)$ and $((\Psi(g))(\omega))(a) = (g(a))(\omega)$, and these bijections induce an isomorphism of games.

(It's straightforward to check that the maps Ψ and Φ preserve the necessary measurability conditions.)

Thus we have:

Theorem 6.7. The games $\mathbf{G}^{\#}(E)$ and $\mathbf{G}(E)^{\#}$ are isomorphic. [6]

Remarks 6.8. We will see in the next section that in the quantum context, this simple and natural isomorphism does not exist. The proof of 6.7 no longer works, essentially because it's not possible to think of all the quantum observables as random variables originating in the same sample space. Indeed, we will see in Theorems 7.4 and 7.6 that no analogue of Theorem 6.7 can possibly hold in the quantum case.

Remark 6.9. Figure 1 below might be useful for keeping track of the notation in Remark 6.6 and Theorem 6.7:

[6] Readers of a certain mathematical bent will note that we've constructed not just an isomorphism but a *natural* isomorphism, which entitles us to think of these games as essentially "the same". Other readers won't miss much if they ignore this footnote.

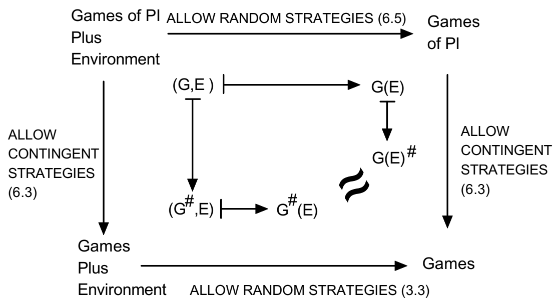

FIGURE ONE

(The section numbers indicate where these concepts are defined.)

7. Quantum Environments for Games of Private Information.

Here we generalize as much of the previous section as possible to the quantum context. We will see that in games of private information, there can be quantum equilibria that are in no sense equivalent to any classical correlated equilibrium.

We fix a game of private information $\mathbf{G}$. For concreteness we take the strategy sets to be $S_1 = S_2 = \{\mathbf{H}, \mathbf{T}\}$. We also fix a quantum environment ξ as in 5.1.

Definition 7.1. (This is the quantum analogue of Definition 6.5). Given $\mathbf{G}$ and ξ as above, we define a new game of private information $\mathbf{G}(\xi)$ as follows:

- The signal sets are the $\mathcal{A}_i$
- The strategy sets are the sets $\mathcal{U}_i$ defined in 5.2 and 5.3.
- The payoff functions are given by

$$P_i(A_1, A_2, U, V) = \int_{S_1 \times S_2} P_i(A_1, A_2, s, t) d\mu_{UV}(s, t)$$

where μ_{UV} is the probability distribution defined in 5.4.

Remarks 7.2. Having formed the game of private information $\mathbf{G}(\xi)$, we can apply 6.3 to form the game $\mathbf{G}(\xi)^{\#}$, which is analogous to a game of behavioral strategies. (In this game, a strategy consists of a map from the signal set A_i to the set of quantum strategies.) It then becomes natural to inquire about an analogue of Theorem 6.7.

Therefore, let (G, ξ) be a pair consisting of a game of private information, and a quantum environment. Figure Two is the quantum analogue of Figure One:

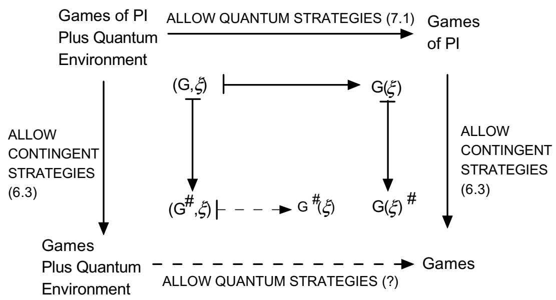

FIGURE TWO

The first minor difficulty is that we haven't defined the bottom map. That's because ξ, as defined in (5.1), is a quantum observable that takes only the two values $\mathbf{H}$ and $\mathbf{T}$, and so serves as an environment only for games where the strategy sets S_i have cardinality two. In the game $\mathbf{G}^{\#}$, the strategy sets are typically much larger.

So if we want to define the bottom horizontal arrow, we'll need a new construction that converts the pair $(\mathbf{G}^{\#}, \xi)$ to a game $\mathbf{G}^{\#}(\xi)$. It's not difficult to imagine reasonable constructions; we could, for example, allow players to observe multiple quantum coins so that the space of possible outcomes has the appropriate cardinality.

However, we don't have to worry about the details, because we're about to prove that *no* reasonable construction can work.

Definition 7.3. Let $\mathbf{H}$ be any (two-player but not necessarily two-by-two) game with with strategy sets S_i and payoff functions P_i. Let $\mathbf{H}'$ be any other such game, with strategy sets S_i' and payoff functions P_i'.

We say that $\mathbf{H}'$ is a *stochastic extension* of $\mathbf{H}$ if for any $(s', t') \in S_1' \times S_2'$, there is a probability distribution μ on $S_1 \times S_2$ such that

$$P_i'(s', t') = \int_{S_1 \times S_2} P_i(s, t) d\mu(s, t)$$

For example, if $\mathbf{G}$ is any game, E any classical environment, and ξ any quantum environment, then $\mathbf{G}(E)$ and $\mathbf{G}(\xi)$ are both stochastic extensions of $\mathbf{G}$.

Theorem 7.4. There exists a game of private information $\mathbf{G}$ and a quantum environment ξ such that $\mathbf{G}(\xi)^{\#}$ is not a stochastic extension of $\mathbf{G}^{\#}$.

Remark 7.4.1. Because the (still undefined) bottom arrow in Figure Two is supposed to be the analogue of "allowing random (or quantum) strategies", it is reasonable to require it to map any pair $(\mathbf{H}, \xi)$ to some stochastic extension of $\mathbf{H}$. Theorem 7.4 says that if we impose this requirement, then no matter how we define the bottom map, the analogue of Theorem 6.7 must fail.

Proof of 7.4. Let $\mathbf{G}$ be the following game of private information:

a) The strategy sets are $S_1 = S_2 = \{\mathbf{H}, \mathbf{T}\}$
b) The signal sets are $\mathcal{A}_1 = \mathcal{A}_2 = \{\text{cat,dog}\}$
c) The probability distribution on $\mathcal{A}_1 \times \mathcal{A}_2$ is uniform
d) The payoff functions are given by (2.1.1).

Let ξ be the quantum environment $\mathbf{HH} + \mathbf{TT}$. Then in the game $\mathbf{G}^{\#}$, the maximum possible payoff is .75. But it follows from example 5.9 that in the game $\mathbf{G}(\xi)^{\#}$, it is possible to achieve a payoff of approximately .85. Therefore $\mathbf{G}(\xi)^{\#}$ cannot be a stochastic extension of $\mathbf{G}^{\#}(\xi)$.

Remark 7.5. Far more generally, if ξ is any non-classical quantum environment, then there are (by definition) finite families $\mathcal{U}_1$ and $\mathcal{U}_2$ of unitary operators such that the various strategy pairs (U_1, U_2) $(U_i \in \mathcal{U}_i)$ yield a family of probability distributions on $S_1 \times S_2$ that cannot be mimicked by any classical random variables. Thus if the signal sets in some game $\mathbf{G}$ have cardinalities at least as large as those of the $\mathcal{U}_i$, then $\mathbf{G}(\xi)^{\#}$ cannot be a stochastic extension of $\mathbf{G}^{\#}$. Thus we can state:

Theorem 7.6. Let ξ be a quantum environment, and suppose that for every game of private information $\mathbf{G}$, that $\mathbf{G}(\xi)^{\#}$ is a stochastic extension of $\mathbf{G}^{\#}(\xi)$. Then ξ is classical (as defined in 5.8).

Less formally: Theorem 7.4 shows that there exists a quantum environment for which there is no reasonable analogue of the Kuhnian construction for games of mixed strategies; Theorem 7.6 says that the same is true in any non-classical quantum environment.

Remark 7.7. If $\mathbf{G}$ is a game of private information, then any quantum equilibrium in the game $\mathbf{G}^{\#}$ is a correlated equilibrium in $\mathbf{G}^{\#}$ and hence *a fortiori* (see 3.6) a Nash equilibrium in some stochastic extension of $\mathbf{G}^{\#}$.

But if ξ is a quantum environment, then $\mathbf{G}(\xi)^{\#}$ is *not* in general a stochastic extension of $\mathbf{G}^{\#}$, so it's plausible that it will have new equilibria which do not come from correlated equilibria in the game $\mathbf{G}^{\#}$. In Section 9, we will see explicit examples of exactly this phenomenon.

8. Prelude to the Examples.

In Section 9, we will offer two examples to illustrate the value of quantum strategies in games of private information. To compute equilibria, we will exploit some special features of those examples, which we highlight here.

General Setup 8.1. In our examples, the strategy sets and signal sets will all have cardinality two. We can take both signal sets to be $\{\mathbf{C}, \mathbf{D}\}$ and both strategy sets to be $\{\mathbf{H}, \mathbf{T}\}$. Then a game of private information is specified by a probability distribution over the four signal pairs **CC**, **CD**, **DC**, **DD**, together with four two-by-two game matrices, one for each of these signal pairs. In our examples, we will always take the last three of these four game matrices to be the same. Thus the payoff structure is specified as follows:

If both players receive signal **C**				If either player receives signal **D**			
		Player Two				**Player Two**	
		H	**T**			**H**	**T**
Player One	**H**	(A,B)	(C,D)	**Player One**	**H**	(I,J)	(K,L)
	T	(E,F)	(G,H)		**T**	(M,N)	(P,Q)

(8.1.1)

Let $\mathbf{G}$ be a game of private information with payoff structure (8.1.1) and let ξ be the quantum environment $\mathbf{HH}+\mathbf{TT}$. We will study the game of behavioral strategies $\mathbf{G}(\xi)^{\#}$. Here a strategy for Player One is a pair $(U_{\mathbf{C}}, U_{\mathbf{D}})$ where each U_i is either a special unitary matrix or one of the pure strategies **H** and **T**. (A special unitary matrix corresponds to physically manipulating one's coin before observing it; a pure strategy involves throwing one's coin away.)

Likewise, a strategy for Player Two is a pair $(V_{\mathbf{C}}, V_{\mathbf{D}})$.

We will (temporarily) restrict attention to the subgame $\mathbf{G}(\xi)^{\#}_0$, in which $U_{\mathbf{C}}, U_{\mathbf{D}}, V_{\mathbf{C}}$ and $V_{\mathbf{D}}$ are restricted to be special unitary matrices — that is, pure strategies are not allowed. (This is unlikely to be a very interesting game in practice, but computing equilibria in $\mathbf{G}(\xi)_0^{\#}$ will be a stepping stone to computing equilibria in $\mathbf{G}(\xi)^{\#}$.)

In the game $\mathbf{G}(\xi)_0^{\#}$, Example (4.6) implies that Player One's payoff is computed as follows:

$$\begin{aligned} &\mathrm{Prob}(\mathbf{CC})\left(s\Big(U_{\mathbf{C}}V_{\mathbf{C}}^T\Big)\frac{A+G}{2}+t\Big(U_{\mathbf{C}}V_{\mathbf{C}}^T\Big)\frac{C+E}{2}\right)\\ +&\mathrm{Prob}(\mathbf{CD})\left(s\Big(U_{\mathbf{C}}V_{\mathbf{D}}^T\Big)\frac{I+P}{2}+t\Big(U_{\mathbf{C}}V_{\mathbf{D}}^T\Big)\frac{K+M}{2}\right)\\ +&\mathrm{Prob}(\mathbf{DC})\left(s\Big(U_{\mathbf{D}}V_{\mathbf{C}}^T\Big)\frac{I+P}{2}+t\Big(U_{\mathbf{D}}V_{\mathbf{C}}^T\Big)\frac{K+M}{2}\right)\\ +&\mathrm{Prob}(\mathbf{DD})\left(s\Big(U_{\mathbf{D}}V_{\mathbf{D}}^T\Big)\frac{I+P}{2}+t\Big(U_{\mathbf{D}}V_{\mathbf{D}}^T\Big)\frac{K+M}{2}\right) \end{aligned} \tag{8.1.2}$$

where, for any special unitary matrix $U=\begin{pmatrix} x & y \\ \overline{y} & \overline{x}\end{pmatrix}$, we set $s(U)=|x|^2$ and $t(U)=|y|^2$.

We define the payoff structure (8.1.1) to be *balanced* if the following four equations hold:

$$\begin{array}{ll} A+G=B+H & C+E=D+F \\ I+P=J+Q & K+M=L+N \end{array} \tag{8.1.3}$$

If the original game $\mathbf{G}$ is balanced, then in the quantum game $\mathbf{G}(\xi)_0^{\#}$ the two players' payoff functions are identical. This implies:

Proposition 8.2. If $\mathbf{G}$ is balanced, then there is a Pareto optimal equilbrium in the game $\mathbf{G}(\xi)_0^{\#}$.

Proof. Because the space of special unitary matrices is compact, there exist matrices $U_{\mathbf{C}}, U_{\mathbf{D}}, V_{\mathbf{C}}, V_{\mathbf{D}}$ that maximize the value of (8.1.2).

Discussion 8.3. If we assume that players achieve a Pareto optimal equilibrium in the game $\mathbf{G}(\xi)_0^{\#}$, then the discussion above reduces the computation of that equilibrium to a maximization problem over 4-tuples of special unitary matrices. This can be computationally quite cumbersome. Our next result will demonstrate that we can restrict our attention to a small subset of those 4-tuples, which will make it possible to do these computations by hand.

Definition 8.4. For $\theta \in [0, 2\pi)$, set (as in Example 5.9)

$$M(\theta) = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$$

Theorem 8.5. In the setup of (8.1), suppose that $\mathbf{G}$ is balanced. Then there exists a real number θ such that the maximum of (8.1.2) is achieved at

$$U_{\mathbf{C}} = M(-\theta) \quad U_{\mathbf{D}} = M(\theta) \quad V_{\mathbf{C}} = M(2\theta) \quad V_{\mathbf{D}} = M(0) \tag{8.5.1}$$

Proof. See the appendix to this paper.

Corollary 8.6. In the setup of (8.1), if $\mathbf{G}$ is balanced, then in $\mathbf{G}(\xi)_0^{\#}$, there is a Pareto optimal equilibrium of the form (8.5.1).

9. Examples.

Example 9.1: Cats and Dogs Revisited. In section 2, we introduced the cat/dog game (adapted from [CHTW]) and claimed that players could improve their payoffs via quantum strategies. In this section, we will return to that game and compute the equilibrium outcome when quantum strategies are available.

As in Section 2, each player is asked "Do you like cats" or "Do you like dogs", with the questions chosen independently via fair coin flips. Players then answer yes or no, and receive the following payoffs:

If both players are asked about cats

		Player Two	
		YES	**NO**
Player One	**YES**	$(0,0)$	$(1,1)$
	NO	$(1,1)$	$(0,0)$

If either player is asked about dogs

		Player Two	
		YES	**NO**
Player One	**YES**	$(1,1)$	$(0,0)$
	NO	$(0,0)$	$(1,1)$

Note that this game has the form of (8.1.1) and is balanced in the sense of (8.2). Thus we can apply Corollary 8.6 to conclude that there is a Pareto-optimal equilibrium in $\mathbf{G}(\xi)_0^{\#}$ of the form (8.5.1):

$$U_{\mathbf{C}} = M(-\theta) \quad U_{\mathbf{D}} = M(\theta) \quad V_{\mathbf{C}} = M(2\theta) \quad V_{\mathbf{D}} = M(0)$$

(Here $\mathbf{C}$ and $\mathbf{D}$, of course, stand for "cats" and "dogs".)

At this equilibrium, expression (8.1.2), for the common payoff to the two players, reduces to

$$\frac{1}{4}\sin^2(3\theta) + \frac{3}{4}\cos^2(\theta) \tag{9.1.1}$$

This expression is maximized at $\theta = \pi/8$, where it takes the approximate value .85. This, then, is a Pareto optimal equilibrium in $\mathbf{G}(\xi)_0^{\#}$.

It is easy to verify that this remains a Pareto optimal equilibrium in the full quantum game $\mathbf{G}(\xi)^{\#}$ (that is, it remains both deviation-proof and Pareto-optimal when players are allowed to choose pure strategies). In example 9.2, we will carry out such a verification in detail.

Remark 9.1.1. The physical interpretation of this equilibrium is that Player One rotates his coin through an angle $-\pi/4$ or $\pi/4$ depending on whether he gets the cat question or the dog question, while Player Two rotates his coin through angle $\pi/2$ or 0. (Note that these are the same operations we considered in Example 5.9, where we used them to illustrate that no classical apparatus could mimic the resulting probability distribution.)

Remark 9.1.2. Not only do players win with probability .85 overall, they also win with probability .85 conditional on being asked any particular pair of questions.

Remark 9.1.3. We have assumed here that players choose behavioral strategies — maps from the signal set $\{\mathbf{C}, \mathbf{D}\}$ to the quantum strategy set $[0, 2\pi)$. Suppose instead that players were required to choose (the quantum analogoue of) mixed strategies, so that each player throws a four-sided quantum coin yielding one of the four possible maps $\{\mathbf{C}, \mathbf{D}\} \rightarrow \{\text{Yes,No}\}$. The result would be a correlated equilibrium in the four-by-four game with these strategies. No such correlated equilibrium can yield an expected payoff greater than 3/4. This, reiterates the point of Theorems 7.4 and 7.5, i.e. the non-equivalence of mixed and behavioral strategies in the quantum context.

Remark 9.1.4. One might imagine the players pooling their information to achieve a better outcome. It's important to recognize, however, that our quantum devices *in no way* allow such information pooling. (Indeed, with information pooling, players could easily earn a certain payoff of 1.) As we stressed in Remark 2.2, players do not communicate any aspect of their private information either to each other or to anyone else.

Example 9.2: Airline Pricing. In Example 9.1, players have a shared goal. Our next example illustrates similar phenomena in a game of greater economic interest, arising from a model of price competition with uncertain demand.

Two identical airlines serve two types of customers. Low-demand customers have a reservation price L; high- demand customers have a reservation price H. There is a fixed population of $2x$ low-demand customers. There is an uncertain population of high-demand customers.

First the airlines receive imperfect signals about the population of high-demand customers. Then they set prices. Then the high-demand customers, if any, arrive and buy seats. Finally, the low-demand customers arrive and buy any remaining available seats at a low price.

The airlines' signals — either **N** (negative) or **P** (postive) are drawn independently, with **N** and **P** equally probable. If either signal is **N**, there are $2y$ high-demand customers, for some $y < x$; otherwise there are none.

Each airline has a capacity constraint equal to the number of low-demand customers. Their only costs are the fixed costs F of running a flight. This defines a game **G** with the following payoff structure (where **L** and **H** stand for "announce the low price L" and "announce the high price H"):

If both firms receive signal **N**

		Firm Two: **L**	Firm Two: **H**
Firm One	**L**	(A, A)	$(B, 0)$
	H	$(0, B)$	$(0, 0)$

If either firm receives signal **P**

		Firm Two: **L**	Firm Two: **H**
Firm One	**L**	(C, C)	$(B, 0)$
	H	$(0, B)$	(D, D)

(9.2.1)

Here $A = xL - F$, $B = 2xL - F$, $C = (x + y)L - F$, and $D = yH - F$. In particular, $A < C < B$. Note that if $D < B$ then **L** is always a dominant strategy for both players, so to keep things interesting we will assume $B < D$.

The analysis of this game depends heavily on the values of A, B, C and D. To avoid a proliferation of cases, and to focus on a particularly interesting example, we take $x = 49$, $y = 19$, $L = 1$, $H = 108/19 \approx 5.68$, and $F = 48$. This gives payoffs of $A = 1$, $B = 50$, $C = 20$, $D = 60$.

9.2.2. Classical Equilibrium. A (mixed) strategy for Firm One is a pair of probabilities (p_N, p_P), with p_i the probability of playing L after receiveing signal i.

Similarly, Firm Two's strategy is a pair (q_N, q_P).

When Firm One receives signal **N**, it is equiprobable that Firm Two has received either signal **N** or signal **P**. Firm One's expected payoff is then

$$\frac{1}{2}\Big(p_N q_N \cdot 1 + p_N(1-q_N)\cdot 50\Big) + \frac{1}{2}\Big(p_N q_P \cdot 20 + p_N(1-q_P)\cdot 50 + (1-p_N)(1-q_P)\cdot 60\Big)$$

When Firm One receives a positive signal, the expected payoff is

$$\frac{1}{2}\Big(p_P q_N{\cdot}20{+}p_P(1{-}q_N){\cdot}50{+}(1{-}p_P)(1{-}q_N)60\Big){+}\frac{1}{2}\Big(p_P q_P{\cdot}20{+}p_P(1{-}q_P){\cdot}50{+}(1{-}p_P)(1{-}q_P){\cdot}60\Big)$$

Firm One chooses p_N and p_P to maximize these expressions and Firm Two behaves symmetrically. One checks that the unique equilibrium is at $p_N = p_P = q_N = q_P = 1$; that is, both firms always play **L**. This guarantees them payoffs of 1 when both receive signal **N** (1/4 of the time) and 20 when either receives a signal **P** (3/4 of the time), for an expected payoff of 15.25.

9.2.3. Quantum Equilibrium. Next we equip our players with a pair of coins in the state $\xi = \mathbf{HH} + \mathbf{TT}$. Because the payoff structure (9.2.1) satisfies the assumptions of Theorem (8.5), the game $\mathbf{G}(\xi)_0^{\#}$ (that is, the game in which players are required to adopt quantum strategies) has a Pareto optimal equilibrium of the form (8.5.1).

At this equilibrium, when both players receive signal **N**, their payoff is

$$\cos^2(3\theta)\frac{A}{2} + \sin^2(3\theta)\frac{B}{2} \tag{9.2.3.1}$$

and otherwise their payoff is

$$\cos^2(\theta)\frac{C+D}{2} + \sin^2(\theta)\frac{B}{2} \tag{9.2.3.2}$$

Adding (9.2.3.1) to 3 times (9.2.3.2) and substituting the assumed values for A, B, C, and D, the payoff becomes

$$\cos^2(3\theta)\frac{1}{2} + \sin^2(3\theta)\frac{50}{2} + 3\cos^2(\theta)\frac{20+60}{2} + 3\sin^2(\theta)\frac{50}{2}$$

It is easy to verify that this expression is maximized at

$$\theta = \mathrm{ArcCos}\left(\frac{1}{2}\sqrt{\frac{14+\sqrt{79}}{7}}\right) \tag{9.2.3.3}$$

where it takes the value

$$\frac{3087+79\sqrt{79}}{112} \approx 33.83$$

This beats the classical payoff of 15.25.

Thus we have found a quantum equilibrium in $\mathbf{G}(\xi)_0^{\#}$. We want to show that it remains an equilibrium in the full quantum game $\mathbf{G}(\xi)^{\#}$, where players are allowed to adopt pure strategies. That is, we want to check that neither player wants to deviate from the quantum equilibrium by playing classically.

Write (s_1, s_2) for the strategy pair described by (8.5.1) and (9.2.3.3)

Claim 1: If Firm Two plays the quantum strategy $\mathbf{s}_2$, and if Firm One receives the negative signal N, then Firm One plays the quantum strategy $\mathbf{s}_1$.

Proof. Note that Firm Two plays **L** and **H** with equal probability. Therefore, when Firm One receives a negative signal, it can earn any of the following returns, with equal probability:

Strategy	Return
L	$\frac{1}{4}(1+50+20+50)=30.25$
H	$\frac{1}{4}(0+0+0+60)=15$
s_1	$\frac{181+7\sqrt{79}}{8}\approx 30.40$

Because $30.40 > 30.25$ (and because the quantum strategy s maximizes both players' payoffs over all alternative quantum strategies), Firm One chooses strategy $\mathbf{s}_1$.

Claim 2: If Firm Two plays the quantum strategy $\mathbf{s}_2$, and if Firm One receives the positive signal P, then Firm One plays the quantum strategy $\mathbf{s}_1$.

Proof. In this case, Firm One's payoffs are known to be

Strategy	Return
L	$\frac{1}{2}(20+50)=35$
H	$\frac{1}{2}(0+60)=30$
s_1	$\frac{65}{2}+\frac{15\sqrt{79}}{28}\approx 37.26$

The result follows because $37.26 > 35$.

Combining Claims 1 and 2, we see that when Firm Two plays $\mathbf{s}_2$, Firm One does not want to deviate; in view of the symmetries of the game, the same is of course true with the firms reversed. Thus $(\mathbf{s}_1, \mathbf{s}_2)$ is genuinely an equilibrium, and, as we have already seen, it is Pareto superior to the unique classical equilibrium where both firms always play $\mathbf{L}$.

9.2.3. Correlated Equilibrium. Consider the game in which a strategy is a behavioral strategy, i.e. a map from the set of signals $\{\mathbf{N}, \mathbf{P}\}$ to the set of actions $\{\mathbf{L}, \mathbf{H}\}$. The conclusion of section (9.2) is equivalent to the statement that the only mixed strategy equilibrium in this game is the pair $(1_{\mathbf{L}}, 1_{\mathbf{L}})$ where $1_{\mathbf{L}}$ is the constant map "always play $\mathbf{L}$". It's not hard to check that this is also the only correlated equilibrium in the sense of Aumann [A].

In other words, in a world governed by the usual laws of probability theory, players who condition their strategies on both the signals they receive and their observations of (possibly correlated) random variables can still do no better than in the classical equilibrium of section (9.2.1). In still other words, a referee who dictates conditional strategies (subject to a deviation-proofness criterion) cannot improve the outcome. This makes it all the more striking that they *can* improve the outcome in the quantum world of section (9.2.2).

9.2.4. Welfare. Consumer surplus occurs only when high demand customers pay low prices. In any of these cases, 38 high-demand customers earn surpluses of 89/19, for a total surplus of 178.

In classical equilibrium, both firms play $\mathbf{L}$, ending up in the upper left corner of the left-hand payoff matrix 1/4 of the time and the upper left corner of the right-hand payoff matrix 3/4 of the time. Thus producer surplus is $(1/4)\cdot 2 + (3/4)\cdot 40 = 30.5$ and consumer surplus is $(3/4)\cdot 178 = 133.5$.

If firms could collude, they would play either $(\mathbf{L}, \mathbf{H})$ or $(\mathbf{H}, \mathbf{L})$ in the low-demand state of the world and $(\mathbf{H}, \mathbf{H})$ in the high demand state of the world for a producer surplus of $(1/4)\cdot 50 + (3/4)\cdot 120 = 102.5$, while consumer surplus would fall to zero. Thus it makes sense that a regulator would want to prohibit collusion.

In quantum equilibrium, one computes that producer surplus is approximately 67.66 and consumer surplus is approximately 78.94.

In summary:

	Classical	Quantum	Classical with Collusion
Consumer Surplus	133.5	78.94	0
Producer Surplus	30.5	67.66	102.5
Total	164	146.6	102.5

Thus our regulator would want to prohibit the use of quantum technology, though not as much as he wants to prohibit collusion (and firms would want to use quantum technology, though not as much as they want to collude). The quantum technology, however, is quite undetectable (after a firm announces it's strategy, how do you know whether it randomized by looking at a classical or a quantum coin?), and hence presumably impossible to regulate.

Appendix: Proof of Theorem 8.5.

Given a special unitary matrix

$$U = \begin{pmatrix} P & Q \\ -\overline{Q} & \overline{P} \end{pmatrix}$$

set $t(U) = |Q|^2$.

Given four special unitary matrices A, B, S, T, set

$$X = t(AS) \quad Y = t(BS) \quad Z = t(BT) \quad W = t(AT) \tag{A.1}$$

Let p, q and r be any real numbers.

Consider the following function

$$f(A, B, S, T) = pW + q(X + Y + Z) + r \tag{A.2}$$

Theorem A.3. There exists a real number θ such that the maximum of (A.2) is achieved at

$$A = M(\theta) \quad B = M(-\theta) \quad S = M(0) \quad T = M(2\theta) \tag{A.4}$$

At this maximum we have

$$X = Y = Z = \sin^2(\theta) \qquad W = \sin^2(3\theta) \tag{A.5}$$

Proof. From [Land] (unnumbered proposition on page 455) it follows that $(X, Y, Z, W) \in [0,1]^4$ is in the image of f if and only if

$$|XY - X - Y - ZW + Z + W| \leq 2(\sqrt{X - X^2}\sqrt{Y - Y^2} + \sqrt{Z - Z^2}\sqrt{W - W^2}) \quad (A.6)$$

We seek to maximize $p(X + Y + Z) + qW$ over $[0,1]^4$ subject to the constraint (A.6). Note that if (X, Y, Z, W) satisfies the constraint, then so does (M, M, M, W) where M is the average of X, Y and Z. Thus to find a maximum, we can assume that $X = Y = Z$. This reduces the problem to a constrained maximization over two variables X and W which can be solved (somewhat laboriously) by hand. This gives X and W as explicit functions of p and q; by inspection, we find that $\mathrm{ArcSin}(\sqrt{W}) = 3\mathrm{ArcSin}(\sqrt{X})$; thus we can set $X = \sin^2(\theta)$ and $W = \sin^2(3\theta)$. Therefore (A.5) is a sufficient condition for a maximum, and this condition holds at (A.4).

Corollary A.7: Theorem 8.5 is true.

Proof. Note that expression (8.1.2) is of the form A.2 (use the fact that s(x) = 1-t(x)). Therefore Theorem A.3 applies.

References

[A] Aumann, "Subjectivity and Correlation in Randomized Strategies", *J. Math Econ* 1 (1974).

[B] Adam Brandenburger, "The relationship between quantum and classical correlation in games", *Games and Economic Behavior* 69 (2010), 175-183.

[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

[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. Rev. Letters* 83, 3077 (1999).

[K] H. Kuhn, "Extensive Games and the Problem of Information", in *Contributions to the Theory of Games II*, Annals of Math Studies 28 (1953).

[L] S. Landsburg, “Nash Equilibria in Quantum Games”, to appear in *Proceedings of the American Mathematical Society*, December 2011.

[LaM] La Mura, P., “Correlated Equilibria of Classical Strategic Games with Quantum Signals,” *International Journal of Quantum Information*, 3, 2005, 183-188.

[Land] L.J. Landau, “Empirical Two-Point Correlation Functions”, *Foundations of Physics* 18 (1988).

[Le] D. Levine, “Quantum Games Have No News for Economists”, working paper.