Quantum Strategies in Noncooperative Games

by

Gordon B. Dahl

and

Steven E. Landsburg

University of Rochester

Abstract. In a game of price-setting under uncertain demand, we allow players to adopt mixed strategies that violate the laws of classical probability theory, in ways that are consistent with the laws of quantum mechanics. By adopting such mixed strategies, firms can improve their profits without collusion, and partially thwart the intentions of regulators. Despite confusion elsewhere in the literature, the resulting quantum equilibria are not the same as correlated equilibria, for at least two reasons which we explain in detail. Finally, to highlight these differences, we present a general model in which quantum and correlated equilibria appear as (distinct) special cases.

1. Introduction.

Traditional game theory assumes the laws of classical probability. However, some of those laws are violated in the real world. Take, for example, the near-triviality:

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

for binary random variables X, Y, Z, W. The clear intuition is that the first and last elements of a sequence cannot differ unless two adjacent elements differ along the way. But despite the clear intuition and easy proof, (1.1) simply isn't true. The existence of random variables violating (1.1) is predicted by quantum mechanics and has been repeatedly verified in physics laboratories.

Note that the classical proof of (1.1) requires writing down expressions like $\mathrm{Prob}(X = Z)$ and $\mathrm{Prob}(Y = W)$; the most glaring inconsistencies are avoided by the fact that,

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

according to quantum mechanics, neither X and Z nor Y and W can be observed simultaneously. The failure of (1.1) is essentially equivalent to the statement that there can be no joint probability distribution for the random variables X, Y, Z and W (though the physicist Richard Feynman observed that one can still write down a joint probability distribution if negative probabilities are allowed—a neat suggestion, but one that still deviates from classical probability theory.)[1]

In some games, players can improve their performance by conditioning their strategies on random variables that violate (1.1). This phenomenon requires no communication between the players. One player might have a choice of observing either X or Z while another player in a distant location has the choice of observing either Y or W. No information is shared.

In the next few decades, technological advances could plausibly make such "quantum strategies" available to any player with a desktop computer. This paper presents a simple example (a two-by-two game of price setting in the face of uncertain demand) to illustrate that quantum strategies matter. In particular, in our example, producers earn more in the quantum equilibrium than in the (unique) classical mixed strategy equilibrium. These gains are more than offset by consumers' losses.

Like collusion, then, quantum technology can increase producers' surplus while lowering social welfare. Unlike collusion, the quantum technology requires no communication and is therefore more difficult to detect and prohibit.

In the final few sections of the paper, we will detail the ways in which quantum equilibria differ from the correlated equilibria of Aumann [A], by presenting a general model which includes quantum and correlated equilibria as distinct special cases.

2. Historical Background.

Quantum mechanics was first applied to game theory by Meyer [M], who describes a

[1] Levine [L] observes— correctly—that even in quantum mechanics, 1) all probabilities are non-negative and 2) the probabilities of the two events $X = 0$ and $X \neq 0$ must add to one. It does not follow, as Levine attempts to conclude, that the classical laws of probability can not be violated; for example, 1) and 2) do not suffice to refute (1.1).

coin-flipping game where one player gains an advantage via access to quantum technology.

Eisert and Wilkens [EW] consider an arbitrary two by two game in which players communicate their strategies to the referee by either flipping or not flipping pennies. Both players have access to arbitrary quantum moves a la Meyer's game. In many games (e.g. the Prisoner's Dilemma), this can lead to Nash equilibria that are Pareto-superior to the classical equilibrium. The Eisert/Wilkens setup requires the referee to observe the players' messages in a particular way. If the referee makes a different kind of observation, the results change.

Cleve, Hoyer, Toner and Watrous [CHTW] consider a pure cooperation game, where players have private signals about the state of the world. Here players can improve the outcome by using quantum particles to randomize their strategies. No cooperation by the referee is required.

Our own example is a non-cooperative variant of the [CHTW] game. We believe this is economically more interesting than any of the previous examples for these reasons:

Meyer's game improves one player's outcome by giving that player (and only that player) access to a more sophisticated technology. Our example (like the Eisert/Wilkens and [CHTW] examples) treats both players symmetrically

Eisert and Wilkens's game requires quantum communication with the referee, who must agree to make a particular sort of observation. Our example (like the Meyer and [CHTW] examples) requires only classical communication and does not require the referee to behave any differently than in a classical game.

[CHTW]'s example is purely cooperative; that is, either both players win or both players lose. Our example (like the Meyer example and the Eisert/Wilkens family of examples) allows for competition between the players.

Also, like the Eisert/Wilkens version of the Prisoner's Dilemma and the [CHTW] example (but unlike Meyer's), our example demonstrates that quantum technology can lead to Pareto-superior outcomes.

3. A Cooperative Game.

To illustrate the value of quantum strategies, we summarize the purely cooperative example from [CHTW].

Two players, who can coordinate strategies in advance but 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 according to a fair coin flip. The players both win if their answers agree, unless they both got the "cats" question, in which case they win if their answers disagree.

It is clear that the players' optimal strategy is to always agree (e.g. by conspiring in advance to always answer "yes" regardless of the question) and that this strategy yields a success rate of 3/4. In a world of classical random variables, there is no advantage to be gained from randomizing.

Now suppose that Player One can observe either of two yes/no–valued random variables X, Z, that Player Two can observe either of two yes/no–valued random variables Y, W, and that the inequality (1.1) is violated. Player One conditions his answer on the value of X if he's asked about cats, or Z if he's asked about dogs. Player Two does the same with Y for dogs and W for cats.

Then the probability of a win is

$$\begin{aligned} & \frac{1}{4}\Big(\text{Prob}(X=Y)+\text{Prob}(Y=Z)+\text{Prob}(Z=W)+\text{Prob}(X\neq W)\Big) \\ = \;& \frac{1}{4}\left(3-\Big(\text{Prob}(X\neq Y)+\text{Prob}(Y\neq Z)+\text{Prob}(Z\neq W)\Big)+\text{Prob}(X\neq W)\right) \\ > \;& 3/4 \end{aligned} \tag{3.1}$$

In other words, quantum technology really does improve the outcome.

4. A Non-Cooperative Game.

Our goal is to demonstrate that the advantages of [CHTW]-like quantum strategies are not restricted to cooperative games. We will illustrate the same phenomenon in 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 seats if they're available for a low price.

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

We assume that each airline has a capacity constraint equal to the number of low-demand customers. Their only costs are the fixed cost F of running a flight. This generates payoff matrices of the following form:

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 (4.0.1)

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

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 $A = 1$, $B = 50$, $C = 20$, $D = 60$.

4.1. Classical Equilibrium. A (mixed) strategy for Firm One is a pair of probabilities (p_N, p_P), with p_N the probability of playing **L** when the signal is negative and p_P the probability of playing **L** when the signal is positive. Similarly, Firm Two's strategy is a pair (q_N, q_P).

When Firm One receives a negative signal, there is a 50/50 chance that Firm Two has received a negative signal and a 50/50 chance that Firm Two has received a positive signal. 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

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

Firm One chooses p_N and p_P to maximize these expressions and Firm B 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.

4.2. Quantum Equilibrium. Next we ask what happens when players can choose strategies contingent on the observation of random variables violating (1.1).

In this case, unlike the cooperative example of Section 3, it will not do simply to demonstrate, a la (3.1), that the quantum payoff exceeds the classical payoff. To illustrate that we've really got an equilibrium (and in particular that players won't want to revert to classical strategies), we'll need to compute payoffs exactly.

Thus we must specify the random variables X, Y, Z and W that are available to be observed. We assume that each of firms A and B is in possession of a spherical "Ouija ball" that can be held at any angle and then asked to recommend a strategy (either **L** or **H**). (Our Ouija balls are metaphors for real-world quantum particles whose spin (which is a binary variable) can be measured at any angle. Quantum theory predicts, and many experiments confirm, that spin measurements satisfy the probabilities given in (4.2.1) below.) The balls are assumed to be *maximally entangled*, a condition that is easy to achieve in a physics laboratory today and possibly within the confines of a desktop quantum computer in the not too distant future.

We will describe how these balls behave. For a detailed description of the underlying mathematical formalism (aimed at an audience of economists) see the appendix to [NE].

If the first ball is rotated through an angle $\phi/2$ and the second through an angle $\theta/2$, then the probabilities of various answer-pairs are:

$$\left.\begin{array}{ll} (\mathbf{L},\mathbf{L}) & \cos^2(\phi-\theta)/2 \\ (\mathbf{L},\mathbf{H}) & \sin^2(\phi-\theta)/2 \\ (\mathbf{H},\mathbf{L}) & \sin^2(\phi-\theta)/2 \\ (\mathbf{H},\mathbf{H}) & \cos^2(\phi-\theta)/2 \end{array}\right\} \qquad (4.2.1)$$

A *quantum strategy* for Firm One consists of a pair of angles $(\phi_N/2, \phi_P/2)$ at which the firm will hold its Ouija ball, depending on whether it receives a negative or positive signal about demand. Similarly, a quantum strategy for Firm Two is a pair of angles $(\theta_N/2, \theta_P/2)$.

In principle, the firms could rotate their balls in three dimensions rather than two; in the language of the appendix [NE], a firm could in principle apply any special unitary matrix whereas we are restricting them to special unitary matrices with real entries. However, firms have nothing to gain from using the extra dimension, so economic modelers have nothing to lose from prohibiting it; we've included a proof in the appendix.

Let X, Y, Z and W denote the observations of Firm One's ball at angle $\phi_N/2$, of Firm Two's ball at angle $\theta_N/2$, of Firm One's ball at angle $\phi_P/2$ and of Firm Two's ball at angle $\theta_P/2$. Then depending on the angles chosen, inequality (1.1) can be violated; thus there is at least the possibility of a non-classical outcome. To see whether such a possibility is realized, we must let the players choose their angles optimally.

When both firms receive negative signals, they each receive a payoff of

$$\cos^2(\phi_N - \theta_N) \cdot \frac{1}{2} + \sin^2(\phi_N - \theta_N) \cdot \frac{50}{2} \tag{4.2.2a}$$

When Firm One receives a negative signal and Firm Two receives a positive, they each receive a payoff of

$$\cos^2(\phi_N - \theta_P) \cdot \frac{20 + 60}{2} + \sin^2(\phi_N - \theta_P) \cdot \frac{50}{2} \tag{4.2.2b}$$

When Firm One receives a positive signal and Firm Two receives a negative signal, they each receive a payoff of

$$\cos^2(\phi_P - \theta_N) \cdot \frac{20 + 60}{2} + \sin^2(\phi_P - \theta_N) \cdot \frac{50}{2} \tag{4.2.2c}$$

When both firms receive positive signals, they each receive a payoff of

$$\cos^2(\phi_P - \theta_P) \cdot \frac{20 + 60}{2} + \sin^2(\phi_P - \theta_P) \cdot \frac{50}{2} \tag{4.2.2d}$$

The firms choose $\phi_N, \phi_P, \theta_N, \theta_P$ to maximize the sum of these four expressions. A maximum is achieved at

$$\phi_N = 0 \quad \theta_P = \mathrm{ArcCos}\left(\frac{1}{2}\sqrt{\frac{14 + \sqrt{79}}{7}}\right) \quad \phi_P = 2\theta_P \quad \theta_N = 3\theta_P \tag{4.2.3}$$

This strategy (call it **Q**) yields a payoff (computed as the average of expressions (4.2.2a) through (4.2.2d)) of

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

which beats the classical payoff of 15.25.

We still have to check that neither player wants to deviate from the quantum equilibrium by playing classically.

Claim 1: If Firm Two plays the quantum strategy **Q**, and if Firm One receives the negative signal N, then Firm One plays the quantum strategy **Q**.

Proof. It follows from (4.2.1) that Firm Two plays **L** and **H** with equal probability. Therefore, when Firm One receives a negative signal (making either of the two payoff matrices in (4.0.1) equally likely), it can earn any of the following returns:

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

Here the payoff to the quantum strategy **Q** is calculated as the average of (5.a) and (5.b) (because Firm Two might have received either a positive or a negative signal), evaluated at (6). Because $30.40 > 30.25$ (and because the quantum strategy **Q** maximizes both players' payoffs over all alternative quantum strategies), Firm One chooses strategy **Q**.

Claim 2: If Firm Two plays the quantum strategy **Q**, and if Firm One receives the positive signal P, then Firm One plays the quantum strategy **Q**.

Proof. In this case, the payoff matrix is surely the right-hand matrix in (4.0.1), so Firm One's payoffs are

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

where the payoff to **Q** is calculated as the average of expressions (4.2.2c) and (4.2.2d). The result follows because $37.26 > 35$.

Combining Claims 1 and 2, we see that when both firms play **Q**, Firm One does not want to deviate. Neither, of course, does Firm Two. Thus **Q** is genuinely an equilibrium, and, as we have already seen, it is Pareto superior to the unique classical equilibrium where both firms always play **L**.

Remark. Although quantum technology can create new equilibria, it can never destroy classical equilibria (such as the equilibrium of (4.1)). To see this, we need only consider deviations by a single player. But for a single player, consulting a quantum ball is equivalent to flipping a fair coin, a strategy which is already available to the classical player.

4.3. Correlated Equilibrium. Consider the game in which a strategy is a map from the set of signals $\{\mathbf{N}, \mathbf{P}\}$ to the set of actions $\{\mathbf{L}, \mathbf{H}\}$. (Call such a map a "conditional strategy".) The conclusion of section (4.1) 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 **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 (4.1). 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 (4.2).

By contrast, one might imagine two-way communication with a referee who observes (or is told) both players' signals and *then* dictates their actions. (In the game theory literature, this is sometimes called a *generalized correlated equilibrium*.) Such a referee could indeed improve the outcome—he always dictates the strategy pair $(\mathbf{H}, \mathbf{H})$ unless both players have received signal **N**, in which case he dicates the strategy pair $(\mathbf{L}, \mathbf{L})$.

It's important to recognize, however, that our quantum devices *in no way* resemble such a referee. Our players do not communicate their signals to each other or to anyone else.

4.4. Welfare. Consumer surplus occurs only when high demand customers pay low prices, i.e. in the upper-left, upper-right and lower-left corners of the right-hand payoff matrix in (4.0.1). 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 **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 Ouija ball?), and hence presumably impossible to regulate.

4.5. Further Remarks. [CHTW] provides an example of a cooperative game with a quantum equilibrium that Pareto dominates any classical equilibrium. Our example (4.0.1) shows that the same phenomenon can occur in a non-cooperative game.

Our example has something of the flavor of a prisoner's dilemma in the sense that in classical Nash equilibrium, both players always play **L**, even though it would be Pareto-superior for them to always play **H**. Using quantum technology, however, they achieve an equilibrium **Q** in which the various strategy pairs occur with the following probabilities:

If both firms receive signal N

		Firm Two L	Firm Two H
Firm One	**L**	.03	.47
	H	.47	.03

If either firm receives signal P

		Firm Two L	Firm Two H
Firm One	**L**	.41	.09
	H	.09	.41

These probabilities are calculated by plugging (4.2.3) into (4.2.1). Note that according to classical probability theory, no mixed strategies can achieve these probabilities, which is why quantum technology can change the outcome. **Q** is really an equilibrium (that is, it is deviation-proof) and it is Pareto-superior to the classical equilibrium $(\mathbf{L}, \mathbf{L})$.

Note also that our result is highly dependent on parameter values (and we had to search for a while to find parameter values that work). For alternative values of A, B, C, D in (4.0.1), we can get no quantum equilibrium, a quantum equilibrium that is Pareto dominated by a classical equilibrium, or, as above, a Pareto-superior quantum equilibrium.

5. Comparison With Correlated Equilibrium.

Levine [L] has suggested that quantum equilibria are merely special cases of the correlated equilibria introduced by Aumann in [A]. We disagree for (at least) two reasons. This section describes those reasons informally; the next section makes them precise in the context of a formal model.

5.1. Deviation-Proofness. In a correlated equilibrium, each player observes an (exogenously determined) random variable and conditions his strategy on its realization. The outcome is deviation-proof in the sense that each player must be willing to obey the dictates of his random variable.

For a player with strategy set $\mathcal{S}$, the set of potential deviations can be identified with the set of maps $\sigma : \mathcal{S} \to \mathcal{S}$. (The player observes the given random variable X and plays the realization of $\sigma \circ X$.) The equilibrium condition is that each player is willing to choose the identity map (taking the other player's behavior as given).

In a more general model, each player might be offered a choice of random variables to observe. Equilibria would then have to be deviation-proof in two senses: In equilibrium, nobody wants to switch to a different random variable, and nobody wants to compose with a non-identity map as in the preceding paragraph. Thus this notion of equilibrium is strictly stronger than that of correlated equilibrium.

Quantum equilibria are like that. Players choose angles at which to measure their Ouija balls (each angle is a different "random variable"). *Given* those angles, they must want to obey the balls' advice (therefore every quantum equilibrium is a correlated equilibrium) but they also must be satisfied with their choice of angles. (Therefore not every correlated equilibrium is a quantum equilibrium.)

Quantum equilibrium is not just a special case of a correlated equilibrium where players can choose among random variables. That's because the "random variables" available to the players are not random variables in the classical sense; in particular they have no joint probability distribution. (This follows from the violation of (1.1).)

In other words, quantum equilibria must be deviation- proof under a particular set of allowable deviations that would never arise in a classical context.

5.2. Private Information. When players have private information, another set of issues arises. Suppose players have information sets A_i and strategy sets S_i. Player i observes a signal $a_i \in A_i$ and chooses a strategy $s_i \in S_i$; payoffs depend on the quadruple (a_1, a_2, s_1, s_2).

There are two ways to model Player i's problem: Either he chooses a *mixed strategy*, i.e. a probability distribution over maps $A_i \to S_i$, or he chooses a *behavioral strategy*, i.e., a map from A_i to probability distributions on S_i.

Explicitly, with information set $A = \{\mathbf{N}, \mathbf{P}\}$ and strategy set $S = \{\mathbf{L}, \mathbf{H}\}$, a mixed strategy assigns probabilities to each of the four strategies

$$\begin{matrix} \mathbf{N} \to \mathbf{L} & \mathbf{N} \to \mathbf{L} & \mathbf{N} \to \mathbf{H} & \mathbf{N} \to \mathbf{H} \\ \mathbf{P} \to \mathbf{L} & \mathbf{P} \to \mathbf{H} & \mathbf{P} \to \mathbf{L} & \mathbf{P} \to \mathbf{H} \end{matrix}$$

while a behavioral strategy assigns probabilities $p_\mathbf{N}$ and $p_\mathbf{P}$ of playing $\mathbf{L}$.

It was shown by Kuhn [K] that this modeling choice is of no consequence, in that any equilibrium under mixed strategies is equivalent (in an appropriate sense) to an equilibrium under behavioral strategies and vice-versa.

In what follows, we will prefer to think of players as choosing random variables, rather than probability distributions. In this case (at least as long as the information and strategy sets are finite) the equivalent of Kuhn's result is immediate: Let T be some fixed sample space; then a mixed strategy is a map from T to S^A, whereas a behavioral strategy is a map from A to S^T, and there's a clear natural equivalence between the two sets of maps. (The finiteness assumption is so that we don't have to worry about measurability.)

However, in the quantum context, this simple and natural equivalence disappears. The proof just given no longer works, essentially because it's not possible to think of all the quantum random variables as maps originating in the same sample space. Players can still follow behavioral strategies, but they are equivalent to mixed strategies *if and only if* all the random variables are classical. This, then, is another fundamental difference between the classical and quantum setups.

6. Games with Oracles

This section presents a formal model of "equilibria with oracles", having correlated equilibria and quantum equilibria as special cases; this will clarify and make precise the arguments of Section 5.1.

Intuitively, an "oracle" is something that helps players choose strategies—such as, for example, a weighted coin. We envision the following situation: Two players set out to play an ordinary game. Their strategy spaces are enhanced by giving each a set of oracles. An equilibrium occurs when each player chooses an oracle, plays a strategy contingent on the oracle's advice, and nobody wants to deviate. The extended game is *classical* if the oracles can all be physically realized without resorting to quantum mechanics and *quantum* if they can be realized via quantum technology. In this section, we will make these notions precise.

Notation 6.1. If X is any set, write $\Omega(X)$ for the set of all measurable functions $[0,1] \to X$, thought of as random variables with values in X.

Definitions 6.2. A *game* consists of the following:

a) Two sets S_1, S_2, called the *strategy spaces*. *We will always assume the S_i are finite.*

b) Two functions $P_1, P_2 : S_1 \times S_2 \to \mathbf{R}$ called the *payoff functions*.

(That is, we use the word "game" to mean a two-player game. This is strictly for expositional convenience, and everything to follow can easily be generalized to n players.)

An *equilibrium* is a pair $(s_1, s_2) \in S_1 \times S_2$ such that s_1 maximizes $P_1(-, s_2)$ and s_2 maximizes $P_2(s_1, -)$.

Definition 6.3. A *game with oracles* $\mathbf{G}$ consists of

a) A game $\mathbf{G}_0$, with strategy sets S_i and payoff functions P_i. $\mathbf{G}_0$ is called the *underlying game*.

b) Two sets $\mathcal{X}_i$, called the sets of *oracles*

c) A function $F : \mathcal{X}_1 \times \mathcal{X}_2 \to \Omega(S_1 \times S_2)$, such that the functions $F(X, -)$ and $F(-, Y)$ uniquely determine the oracles X and Y (the *irredundancy condition*).

We also require the following condition:

Condition 6.3.1. For each oracle $X \in \mathcal{X}_i$ and for each map $\sigma : S_i \to S_i$, there must exist an oracle $X^\sigma \in X_i$ with the following property:

Given $X \in \mathcal{X}_1$, $Y \in \mathcal{X}_2$, $\sigma : \mathcal{X}_1 \to \mathcal{X}_1$ and $\tau : \mathcal{X}_2 \to \mathcal{X}_2$, we have

$$F(X^\sigma, Y^\tau) = (\sigma \times \tau) \circ F(X, Y)$$

Notations 6.3.2. Given a game with oracles as in (6.3), we write $\mathbf{G} = (\mathbf{G}_0, \mathcal{X}_1, \mathcal{X}_2, F)$.

Given $X \in \mathcal{X}_1, Y \in \mathcal{X}_2$, and $s_i \in S_i$ we write $F_{XY}(s_1, s_2)$ for the Lebesgue measure of $F(X, Y)^{-1}(s_1, s_2)$.

Remarks 6.4. The idea is that each player consults an oracle X for advice on how to play; having received that advice players should be free to modify it arbitrarily; the assumed existence of X^σ models that freedom.

Definition 6.5. We will say that the set $\mathcal{X}$ of oracles is *generated by* the subset $\mathcal{X}'$ if every element of $\mathcal{X}$ is of the form X^σ for some $X \in \mathcal{X}'$.

Intuitively, this means that consulting any oracle at all is equivalent to consulting some oracle in $\mathcal{X}'$ and then modifying its advice.

Definition 6.6. Let $\mathbf{G}$ be a game with oracles. We define the *associated game* $\mathbf{G}'$ as follows:

a) The strategy sets are $\mathcal{X}_1$ and $\mathcal{X}_2$

b) The payoff functions are

$$\begin{array}{rccc} P_i: & \mathcal{X}_1 \times \mathcal{X}_2 & \longrightarrow & \mathbf{R} \\ & (X,Y) & \mapsto & \sum_{(s_1,s_2)\in S_1\times S_2} F_{XY}(s_1,s_2)P_i(s_1,s_2) \end{array}$$

Note that we abuse notation by using P_i both for the payoff function in the underlying game $\mathbf{G}_0$ and the payoff function in the associated game $\mathbf{G}'$.

Definition 6.7. Let $\mathbf{G}$ be a game with oracles. Then an *equilibrium* in $\mathbf{G}$ is an (ordinary) equilibrium in the associated game $\mathbf{G}'$.

Example 6.8. Let $\mathcal{X}_i = S_i$, and let $F(s_1, s_2)$ be the random variable that always takes the value (s_1, s_2). Then an equilibrium in $\mathbf{G}$ is the same thing as an equilibirium in the underlying game $\mathbf{G}_0$.

Example 6.9. Let $\mathcal{X}_i$ be the set of all random variables $[0,1] \to S_i$ with the property that the inverse image of any point is connected. for $x \in [0,1]$, let

$$F(X_1, X_2)(x) = \Big(X_1(x), X_2((1-a)x/(b-a))\Big) \tag{6.9.1}$$

where a and b are the least and greatest upper bounds of the set $\{t \in [0,1] | X_1(t) = X_1(x)\}$.

Then an equilibrium in $\mathbf{G}'$ is the same thing as a mixed strategy equilibrium in the underlying game $\mathbf{G}_0$.

Remark 6.9.2. The point of this example is to show that ordinary mixed strategy equilibria occur as a special case of our construction. The random variables X_1 and X_2 induce probability distributions on S_1 and S_2, and we want a random variable that induces the product distribution on $S_1 \times S_2$.

The most natural way to do this would be to define $F(X_1, X_2)$ to be the random variable

$$X_1 \times X_2 : I \times I \to S_1 \times S_2$$

However, we have required all random variables to have I, not $I \times I$ as their domains. Thus we've employed the somewhat *ad hoc* formula (6.9.1); to make this work, we've required the players to choose random variables with connected inverse images. This, of course, does not restrict their freedom to choose arbitrary probability distributions over their strategy sets, which is all that matters.

Example 6.10. Let $\mathbf{G}_0$ be a game. Let $X : [0,1] \to S_1 \times S_2$ be a fixed random variable and let μ be the induced probability distribution on $S_1 \times S_2$. Let $\mathcal{X}_i = Hom(S_i, S_i)$. Let $F(\sigma, \tau) = (\sigma \times \tau) \circ X$. Let $\mathbf{G} = (\mathbf{G}_0, \mathcal{X}_1, \mathcal{X}_2, F)$. Then $(1_{S_1}, 1_{S_2})$ is an equilibrium in $\mathbf{G}$ if and only if μ is a correlated equilibrium in $\mathbf{G}_0$.

More generally:

Example 6.11. In any game with oracles $\mathbf{G}$, we have:

If (X, Y) is an equilibrium in $\mathbf{G}$, then F_{XY} is a correlated equilibrium in $\mathbf{G}_0$.

The converse is in general false:

Example 6.12. Consider the game $\mathbf{G}_0$ with the following payoffs:

		Player Two C	Player Two D
Player One	C	$(0,0)$	$(7,2)$
Player One	D	$(2,7)$	$(6,6)$

Let $\mathbf{G} = (\mathbf{G}_0, \mathcal{X}_1, \mathcal{X}_2, F)$ be a game with oracles. Suppose $\mathcal{X}_1$ is generated by two oracles X, Y and $\mathcal{X}_2$ is generated by two oracles Z, W, which induce the following probability distributions:

$$\overset{(X,Z)}{\begin{pmatrix} 0 & 1/3 \\ 1/3 & 1/3 \end{pmatrix}} \quad \overset{(X,W)}{\begin{pmatrix} 1/6 & 1/6 \\ 1/3 & 1/3 \end{pmatrix}} \quad \overset{(Y,Z)}{\begin{pmatrix} 0 & 1/2 \\ 1/3 & 1/6 \end{pmatrix}} \quad \overset{(Y,W)}{\begin{pmatrix} 0 & 1/2 \\ 1/2 & 0 \end{pmatrix}}$$

It is easy to check that (X, Z) and (Y, W) are both correlated equilibria in $\mathbf{G}_0$. However, only the second of these is an equilibrium in $\mathbf{G}$. For if the players obey oracles X and Z, then Player One can improve his payoff (from 5 to 31/6) by switching to oracle Y.

Remark 6.13. In Example 6.12, (Y, W) is an equilibrium in $\mathbf{G}$. It is also a correlated equilibrium in $\mathbf{G}_0$. But it would be a mistake to say that this is "just an example of a correlated equilibrium", because that observation fails to distinguish between correlated equilibria like (Y, B) (which is a correlated equilibrium in $\mathbf{G}$) and correlated equilibria like (X, Z) (which is *not* a correlated equilibrium in $\mathbf{G}$).

Definition 6.14. A game with oracles is *classical* if each $\mathcal{X}_i$ is a subset of $\Omega(S_i)$ and the inclusion maps ι_i make the following diagram commute:

$$\begin{array}{ccccc} \Omega(S_1) & \longleftarrow & \Omega(S_1 \times S_2) & \longrightarrow & \Omega(S_2) \\ \uparrow \iota_1 & & \uparrow F & & \uparrow \iota_2 \\ \mathcal{X}_1 & \longleftarrow & \mathcal{X}_1 \times \mathcal{X}_2 & \longrightarrow & \mathcal{X}_2 \end{array}$$

In other words, the extended game is classical if an extended strategy consists of observing a random variable (where the two observed random variables might be correlated).

Example 6.15. In Example 6.12, $\mathbf{G}$ is classical. We can take X, Y, A, B to be random variables that take the values (C, C, D, D), (C, D, D, C), (D, D, C, C), (D, C, D, D) with probabilities $1/6, 1/6, 1/3, 1/3$.

"Definition" 6.16. We *loosely* define a set of oracles to be *quantum* if they are physically realizable according to the laws of quantum mechanics. We loosely define an extended game to be *quantum* if $\mathcal{X}_1 \cup \mathcal{X}_2$ is a quantum set of oracles.

Example 6.17. Suppose $\mathcal{X}_1 = \mathcal{X}_2 = [0, 2\pi]$. For $X, Y \in [0, 2\pi]$ we take $F(X, Y)$ to be a random variable with the probability distribution

$$\begin{pmatrix} \cos^2(X-Y)/2 & \sin^2(X-Y)/2 \\ \sin^2(X-Y)/2 & \cos^2(X-Y)/2 \end{pmatrix}$$

These oracles are non-classical. However, they are realizable quantum mechanically by measuring the spins of maximally entangled particles in the chosen directions. (That is, Player One chooses X and measures the spin of Particle One in the X direction, while Player Two chooses Y and measures the spin of Particle Two in the Y direction.)

This is, of course, precisely the set of oracles introduced in Section 4.

Reminder 6.18. Even when two games-with-oracles $\mathbf{G}$ and $\mathbf{H}$ have the same underlying game $\mathbf{G}_0$, a given pair of oracles might be an equilibrium in one but not the other.

In Example 6.12, let $\mathcal{Y}_1$ be the subset of $\mathcal{X}_1$ generated by $\{X\}$ and $\mathcal{Y}_2$ the subset of $\mathcal{X}_2$ generated by $\{A\}$). Then (X, A) is an equilibrium in $\mathbf{H}$ but not in $\mathbf{G}$.

This is of course unsurprising; after all, equilibria in $\mathbf{G}$ and $\mathbf{H}$ are just equilibria in the ordinary games $\mathbf{G}'$ and $\mathbf{H}'$. Because $\mathbf{G}'$ has larger strategy sets than $\mathbf{H}'$, the equilibria in $\mathbf{H}'$ might not survive passage to $\mathbf{G}'$.

This pedestrian observation implies that the problem of finding equilibria in quantum games with oracles cannot be reduced to the problem of finding equilibria in classical games with oracles. It's all very well to point out that a given equilibrium in a given quantum game is also an equilibrium in some classical game, but doing so does not obviate the need to study the quantum game in its own right.

7. Games Of Private Information.

In this section we will consider (generalized) games of private information. Throughout, players are allowed to observe (possibly correlated) random variables, but they are *not* allowed to communicate with each other or with a referee.

Definition 7.1. A *game of private information* consists of two strategy spaces S_i, two information sets $\mathcal{A}_i$, a probability distribution on $\mathcal{A}_1 \times \mathcal{A}_2$, and two payoff functions

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

Definition 7.2. Given a game $\mathbf{G}$ of private information, the *associated game* $\mathbf{G}^\#$ has strategy sets $\mathcal{T}_i = Hom(\mathcal{A}_i, S_i)$ and payoff functions

$$P_i^\#(T_1, T_2) = \sum_{(A_1, A_2) \in \mathcal{A}_1 \times \mathcal{A}_2} \mathrm{Prob}(A_1, A_2) P_i(A_1, A_2, T_1(A_1), T_2(A_2))$$

An *equilibrium* in the game of private information $\mathbf{G}$ is an equilibrium in the associated game $\mathbf{G}^\#$.

Definition 7.3. A *game of private information with oracles* consists of:

a) A game of private information $\mathbf{G}_0$

b) Two sets of oracles $\mathcal{X}_i$

c) A function $F : \mathcal{X}_1 \times \mathcal{X}_2 \to \Omega(S_1 \times S_2)$ satisfying the irredundancy condition of 6.3(c)

We also require condition (6.3.1) to hold.

A game of private information with oracles is *classical* if the condition of Definiton 6.14 holds.

Definition 7.4. Let $\mathbf{G}$ be a game of private information with oracles. The *associated game of private information* $\mathbf{G}'$ is defined as follows:

a) The strategy spaces are $\mathcal{X}_1$ and $\mathcal{X}_2$

b) The payoff functions are

$$P_i: \begin{array}{ccc} \mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{X}_1 \times \mathcal{X}_2 & \to & \mathbf{R} \\ (A_1, A_2, X_1, X_2) & \mapsto & \sum_{(s_1,s_2)\in S_1\times S_2} F_{XY}(s_1,s_2) P_i(A_1,A_2,s_1,s_2) \end{array}$$

Remarks 7.5. Our various "associated game" definitions can be represented schematically as follows:

$$\begin{array}{ccc} \textbf{Games of private information with oracles} & \xrightarrow{\text{Def. 7.4}} & \textbf{Games of private information} \\ \Big\downarrow & & \Big\downarrow \text{Def. 7.2} \\ \textbf{Games with oracles} & \xrightarrow{\text{Def. 6.6}} & \textbf{Games} \end{array}$$

Starting with $\mathbf{G}$ in the upper left corner, we can construct $\mathbf{G}'$ in the upper right and then $\mathbf{G}^\#$ in the lower right.

It is natural to ask whether there is a construction on the left, taking $\mathbf{G}$ to some game with oracles $\mathbf{G}^\#$ in the lower left, with the property that

$$(\mathbf{G}^\#)' = (\mathbf{G}')^\#$$

As we shall see the answer is yes when $\mathbf{G}$ is classical, but not otherwise.

Observation 7.6. Given a game of private information with oracles, there are obvious maps

$$\begin{array}{ccc} \mathcal{X}_1^{\mathcal{A}_1} \times \mathcal{X}_2^{\mathcal{A}_2} \longrightarrow (\mathcal{X}_1 \times \mathcal{X}_2)^{(\mathcal{A}_1 \times \mathcal{A}_2)} \longrightarrow \Omega(S_1 \times S_2)^{(\mathcal{A}_1 \times \mathcal{A}_2)} \approx & \Omega\left((S_1 \times S_2)^{(\mathcal{A}_1 \times \mathcal{A}_2)}\right) \\ & \cup \\ & \Omega\left(S_1^{\mathcal{A}_1}\right) \times \Omega\left(S_2^{\mathcal{A}_2}\right) \end{array}$$

We use Θ to denote the composition along the top row.

Theorem 7.7. The map

$$\Theta : \mathcal{X}_1^{\mathcal{A}_1} \times \mathcal{X}_2^{\mathcal{A}_2} \longrightarrow \Omega\left((S_1 \times S_2)^{(\mathcal{A}_1 \times \mathcal{A}_2)}\right)$$

factors through

$$\Omega\left(S_2^{\mathcal{A}_2}\right) \times \Omega\left(S_2^{\mathcal{A}_2}\right)$$

if and only if **G** is classical.

Proof. If **G** is classical, the map $\mathcal{X}_1 \times \mathcal{X}_2 \to \Omega(S_1 \times S_2)$ factors through $\Omega(S_1) \times \Omega(S_2)$; applying this to the constrution in 7.6 yields the result immediately.

Conversely, suppose Θ factors as indicated. Then, identifying an oracle with the constant function that has that oracle as its value, it's easy to check that the diagram

$$\begin{array}{ccc} \mathcal{X}_1 \times \mathcal{X}_2 & \hookrightarrow & \left(\mathcal{X}_1^{\mathcal{A}_1}\right) \times \left(\mathcal{X}_2^{\mathcal{A}_2}\right) \\ \downarrow & & \downarrow \\ \Omega(S_1 \times S_2) & \rightarrow & \Omega(S_1 \times S_2)^{(\mathcal{A}_1 \times \mathcal{A}_2)} \end{array}$$

commutes, so that

$$F(\mathcal{X}_1 \times \mathcal{X}_2) \subset \left[\Omega\left(S_1^{\mathcal{A}_1}\right) \times \Omega\left(S_2^{\mathcal{A}_2}\right)\right] \cap \Omega(S_1 \times S_2) = \Omega(S_1) \times \Omega(S_2)$$

as needed.

Definition 7.8. Let **G** be a *classical* game of private information with oracles. We define the *associated game with oracles* $\mathbf{G}^{\#}$ as follows:

a) The strategy sets are the $S_i^{\mathcal{A}_i}$

b) The payoff functions are

$$P_i(f, g) = \sum_{(A_1, A_2) \in \mathcal{A}_1 \times \mathcal{A}_2} \text{Prob}(A_1, A_2) P_i(A_1, A_2, f(A_1), g(A_2))$$

c) The oracles are the sets $\mathcal{X}_i^{\mathcal{A}_i}$

d) The required map

$$F : \mathcal{X}_1^{\mathcal{A}_2} \times \mathcal{X}_2^{\mathcal{A}_2} \to \Omega(S_1^{\mathcal{A}_1} \times S_2^{\mathcal{A}_2})$$

is given by taking $F = \Theta$, which makes sense by Theorem 7.7.

Theorem 7.9. If $\mathbf{G}$ is a classical game of private information with oracles, then

$$(\mathbf{G}^{\#})' \quad = \quad (\mathbf{G}')^{\#}$$

Remark 7.10. In spirit, and in the language of [K], $(\mathbf{G}^{\#})'$ is like the associated game with *mixed strategies*, where players choose a probability distribution over maps $A_i \to S_i$ and $(\mathbf{G}')^{\#}$ is the associated game with *behavioral strategies*, where players choose, for each possible piece of information, a probability distribution over strategies. In [K], these games are equivalent in an appropriate sense; in our setup they're actually the same game. The difference is that in [K], players choose probability distributions, whereas here, players choose random variables. Obviously, the two approaches are fundamentally equivalent.

However, by Theorem 7.7, we have no hope of proving Theorem 7.9 (or even defining $(\mathbf{G}^{\#})$ in the non-classical case. In other words, the quantum setup requires us to focus on behavioral strategies.

Appendix: Real Strategies Suffice

The model of Section 4.2 envisions agents equipped with entangled particles that they can rotate before observing. We assume that the players rotate their particles in a single direction. This note is to prove that they have no reason to rotate in any other direction.

What follows assumes familiarity with the yoga of quantum mechanics at the level of the appendix to [NE].

Arbitrary rotations are represented by unitary matrices; rotations about the Y-axis are represented by unitary matrices with real entries.

Thus if arbitrary rotations are allowed, Player One's strategy consists of two unitary matrices M_N and M_P, to be played in the event of signal N or signal P. Player Two's strategy consists of two unitary matrices M'_N and M'_P. Our goal is to show that all of these matrices can be taken to have real entries. More precisely: For any choice of M_N and M_P, there exist optimal responses M'_N and M'_P all of whose entries are real (and similarly, of course, with the players reversed).

We can multiply all four matrices on the right by $(M'_N)^{-1}$ and assume $M'_N = I$.

We can multiply all four matrices by constants of norm 1 to assume that the upper left entries are real.

Write $(X, Y + iZ)$, $(U, V + iW)$, and $(R, S + iT)$ for the top rows of M_N, M_P, M'_P.

Then expressions (4.2.2.a) through (4.2.2.b) become

$$X^2 \cdot \frac{1}{2} + (Y^2 + Z^2) \cdot \frac{50}{2} \tag{4.2.2.a'}$$

$$\begin{aligned}&\big(X^2 D^2 (Y^2 + Z^2)(S^2 + T^2) + 2RX(SY - TZ)\big) \cdot \frac{20 + 60}{2}\\ +&\big(X^2(S^2 + T^2) + (Y^2 + Z^2)R^2 + 2RX(TZ - SY)\big) \cdot \frac{50}{2}\end{aligned} \tag{4.2.2.b'}$$

$$U^2 \cdot \frac{20 + 60}{2} + (V^2 + W^2) \cdot \frac{50}{2} \tag{4.2.2.c'}$$

$$\begin{aligned}&\big(U^2 D^2 (V^2 + W^2)(S^2 + T^2) + 2RU(SV - TW)\big) \cdot \frac{20 + 60}{2}\\ +&\big(U^2(S^2 + T^2) + (V^2 + W^2)R^2 + 2RU(TW - SV)\big) \cdot \frac{50}{2}\end{aligned} \tag{4.2.2.d'}$$

The players act to maximize the sum of these expressions.

Hold fixed X, U and R (and therefore $Y^2 + Z^2$, $V^2 + W^2$ and $S^2 + T^2$). Then the sum of (4.2.2.a′) through (4.2.2.d′) is equal to a constant plus

$$30R\big(S(UV + XY) - T(UW + XZ)\big) \tag{$*$}$$

so that Y, Z, V, W, S, T must be chosen to maximize $(*)$ subject to the constancy of Y^2+Z^2, $V^2 + W^2$ and $S^2 + T^2$. The maxima clearly occur at $Y = V = S = 0$ and $Z = W = T = 0$, and both maxima are equal. Thus players might as well choose the latter, in which case all of their strategy matrices have all real entries.

References

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

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

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

[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] D. Levine, "Quantum Games Have No News for Economists"

[M] D. Meyer, "Quantum Strategies", *Phys. Rev. Lett.* 82 (1999), 1052-1055

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