

How Well Do People Play a Quantum Prisoner's Dilemma?

Kay-Yut Chen*and Tad Hogg
Hewlett Packard Labs
1501 Page Mill Road
Palo Alto, CA 94304

January 11, 2006

Abstract

Game theory suggests quantum information processing technologies could provide useful new economic mechanisms. For example, using shared entangled quantum states can alter incentives so as to reduce the free-rider problem inherent in economic contexts such as public goods provisioning. However, game theory assumes players understand fully the consequences of manipulating quantum states and are rational. Its predictions do not always describe human behavior accurately. To evaluate the potential practicality of quantum economic mechanisms, we experimentally tested how people play the quantum version of the prisoner's dilemma game in a laboratory setting using a simulated version of the underlying quantum physics. Even without formal training in quantum mechanics, people nearly achieve the payoffs theory predicts, but do not use mixed-strategy Nash equilibria predicted by game theory. Moreover, this correspondence with game theory for the quantum game is closer than that of the classical game.

keywords: quantum economic mechanisms, prisoner's dilemma, free-riding, experimental economics

*corresponding author: kay-yut.chen@hp.com

1 Introduction

Recent developments in quantum computing and demonstrated exchanges of entangled quantum states over distances of tens of kilometers have led to increased interests in quantum games. Extending classical games into the quantum realm broadens the range of strategies [20]. Examples of quantum games analyzed with game theory methods include the prisoner’s dilemma [12, 13, 10, 11] and the n -player minority game [5]. Of more direct relevance for economics are games involving correlated choices without communication [18, 22]. This raises the possibility of creating new economic mechanisms using quantum information processing technologies.

A quantum mechanism to produce public goods is one such example [8, 27]. Provisioning of public goods is one of the most important economic contexts giving rise to the free-rider problem. The free-rider problem cannot be solved with traditional means without either a third party to enforce agreements or a repeated game scenario in which participants can self-police (e.g., tit-for-tat behaviors [3]).

While quantum information processing thus offers the possibility of creating new economics applications, many important questions are still unanswered. The most significant one is how people would, in fact, play quantum games. To date, the work on quantum games only involves game theoretic analysis. Can people learn to play as well as game theory predicts without requiring training in quantum mechanics? Even for conventional games, the predictions of game theory do not always match human behavior because the theory often assumes unrealistic levels of rationality of its players [24, 7]. Furthermore, people are known not to play mixed strategies even when game theory indicates rational players would do so. Proposed quantum games often have no single-strategy Nash equilibrium and multiple mixed-strategy equilibria. In such cases, game theory does not make unique predictions even for rational players. Furthermore, quantum games often have probabilistic outcomes, i.e., the payoffs aren’t necessarily the same each time players make the same choices. These properties make it even more difficult for people to achieve the rational play predicted by game theory [6].

A second major question is how readily quantum games can be physically implemented, as they require the creation, transmission, storage and operation on entangled quantum states. While this issue is beyond the scope of this paper, it is interesting to note that many quantum games do not require long sequences of coherent operations and hence are more likely to realize than large-scale quantum computations such as required to factor numbers large enough to be of cryptographic relevance. Physical implementation of the game reported in this paper should be feasible in the near future.

This paper reports a sequence of laboratory experiments to address the crucial question of whether people can understand and play quantum games as game theory has predicted. We used a quantum version of the prisoner’s dilemma for this study. Our major goal is to test whether actual human behavior results in the predicted higher levels of cooperation compared to zero

cooperation in the classical version. The prisoner’s dilemma, of course, is one of the most studied games in the literature. The rich background of prisoner’s dilemma research allows us to contrast our results, about behavior of human subjects playing quantum games, with behavior in the classical version of the game, which is widely known. Furthermore, the prisoner’s dilemma is also a simple version of the public goods game, for which a quantum mechanics based mechanism performs efficiently for large groups [8]. Specifically, for two players with equal preferences and endowments, the public goods problem reduces to the prisoner’s dilemma. The prisoner’s dilemma illustrates the free-rider problem [17] in the simplest context of a two-person game, in which each player has the choice to “cooperate” or “defect”. Payoffs for both players are higher when both of them choose to cooperate instead of both defecting. However, each individual is better off by defecting. The prisoner’s dilemma involves the possibility of altruistic behaviors in which participants can either select actions that most benefit themselves or those that benefit the group as a whole but at some individual loss [2].

Game theory predicts defection independent of the number of players for the classical n -player prisoner’s dilemma (public goods game). However, the quantum version, when extended to multiple players, results in higher efficiencies when the number of players increases. This difference allows us to further test whether the quantum version can induce behavior substantially different from the classical version.

Instead of waiting for development of a physical implementation of the quantum game, we simulated all the quantum components since the subjects in a laboratory setting would not be able to tell the difference. Specifically, instead of performing operations on physical quantum states, players sent their choices of quantum operators to a computer server which then performed the operations on simulated quantum states. Simulation of quantum games [26] is suitable for laboratory studies which guarantee participants follow the rules of the game without need for the security properties of a physical implementation or developing the legal sanctions of real economic contracts. Moreover, the game involves only a modest number of quantum states, so the exponential increase in time and memory required for simulation on conventional machines is not significant. From a research perspective, a simulation of a quantum game has the additional advantage of allowing a detailed analysis of behavior, since it gives access to the operators and probabilities of each outcome, not just the single observed outcome.

The practical benefit of such quantum mechanisms depends on other issues, such as computational feasibility of the simulation or security requirements of the mechanism, rather than the absolute need of a physical implementation. Simulated quantum mechanics can easily implement any quantum game provided that the size of the relevant quantum systems are small enough so the exponential overhead associated with such simulations remains manageable. The differences in security and communication costs as well as the level of trust assumed for the simulator are important factors determining the value of a true physical quantum implementation. For instance, the quantum version allows

	“Lo” payoffs		“Hi” payoffs	
	C	D	C	D
C	150; 150	75; 175	180; 180	90; 190
D	175; 75	100; 100	190; 90	100; 100

Table 1: Payoffs for the two prisoner’s dilemma games we investigated. First player’s choice to cooperate (C) or defect (D) are shown in the rows, and second player’s choice in the columns. Each entry shows the payoffs to the first and second players for that combination of choices. For example, in the second set of payoffs, if the first player cooperates but the second defects, the first player receives a payoff of 90.

only a single measurement of the outcome. Other properties of the quantum state are destroyed by this measurement, reducing information available to infer individual individual operator choices, and hence providing additional privacy. For example, an auction implemented with a quantum mechanism can guarantee that losing bids will never be revealed. Such privacy can also be achieved via conventional cryptographic methods but with security based on the computational difficulty of deciphering encrypted messages rather than guaranteed by the laws of quantum physics. Moreover, with conventional cryptography, private information could later be revealed either deliberately by one of the parties with the key or via a legal requirement to produce the information.

The paper is organized as follows. Sec. 2 describes the quantum formulation of the prisoner’s dilemma game. We focus on comparing with the standard classical prisoner’s dilemma for simplicity, leaving for future work more complicated comparisons with conventional economic approaches to free riding or allowing use of classically correlated rather than entangled quantum states to enhance the game [21]. Sec. 3 discusses the design of the experiment on human subjects playing this game. Sec. 4 describes resulting behaviors we observed and how they compare to game theory predictions. Sec. 6 concludes with possible extensions to our experiments.

2 The Quantum Prisoner’s Dilemma

Two sets of payoffs, labelled “Lo” and “Hi” respectively, were used in our prisoner’s dilemma experiments, as shown in Table 1. In both cases, defecting is the dominant strategy but cooperating is more efficient, in a Pareto sense, than defecting. The two instances differ in the benefit of individual defection (25 and 10, respectively) and the difference in payoffs between both cooperating, i.e., the efficient outcome, and both defecting, i.e., the Nash equilibrium (50 and 80, respectively). Thus compared to the first set of payoffs, the second set, labelled “Hi”, gives players less temptation to defect and a greater incentive for mutual cooperation.

We presented the prisoner’s dilemma in our experiments as a contribution

game in which participants contribute (cooperate) or not (defect) toward the production of a public good, in the form of additional payoffs received by the two players. Specifically, in a single instance of the game, each of two players is given an initial wealth $W = 100$, and can keep it or contribute all of it to the group. Let c_i denote the amount, 0 or W , player i contributes, so the total contribution from all players is $C = \sum_{j=1}^n c_j$, where n is the number of players (2 in the case of the prisoner's dilemma). The total contribution is multiplied by a and distributed equally among the players. This two-player contribution game is identical to a public goods game where the size of the group is $n = 2$ and the production function of the payoff from the public good is aC/n where $1 < a < n$. Thus the payoff to player i is

$$W - c_i + \frac{a}{n}C \tag{1}$$

Since $a < n$, each player obtains a higher payoff by not contributing no matter what choices other players make. If all players make this dominant choice, each receives a payoff of W . If they all contribute, each payoff would be the larger value aW . Thus the group is better off if all contribute, but each person prefers not to contribute and free-ride on the public good produced by others' contributions.

The two sets of payoffs in Table 1 correspond to $a = 1.5$ and 1.8 , respectively. Thus the players' choices are framed in terms of a decision to contribute or not, which correspond to the cooperate and defect choices of the prisoner's dilemma. We chose this framework to place the prisoner's dilemma in an economically relevant setting, namely a two-player version of the public goods game. This setting naturally extends to public goods games with larger groups. This mapping to a public goods game is only relevant in the larger context of using quantum mechanics to solve economics problems. In this paper, we primarily focus narrowly on behavior observed in the two prisoner's dilemma games described by Table 1.

2.1 Playing the Quantum Game

Quantum games [20] can be described in the standard mathematical framework for quantum mechanics [14], which we briefly describe here in the context of the quantum version of the prisoner's dilemma game. Players manipulate physical systems with two physical states (such as a photon with vertical or horizontal polarization). In analogy with two-state devices used to implement bits, the quantum version of such systems are called *qubits*. A set of n classical bits can exist in one of 2^n different configurations, and various logic operations on these bits change them from one configuration to another. The quantum system is more complicated: a *superposition* of the 2^n states. A superposition corresponds to a complex-valued vector, with one component (called an *amplitude*) for each of the 2^n configurations. Quantum operations alter the superposition by multiplying the original vector by a complex-valued matrix. The configuration of the quantum system can also be observed or measured. Doing so produces a

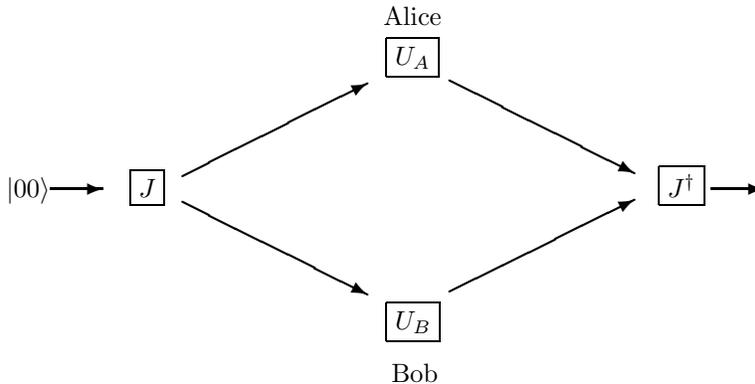


Figure 1: Two-player quantum game.

definite state of each bit, i.e., a single configuration. The configurations are produced probabilistically with probabilities equal to the squared magnitudes of the amplitudes associated with the configurations in the superposition. The superposition vector itself is not observable, only its consequence in determining the probability of observing the various possible configurations of states.

The quantum prisoner's dilemma is set up as follows: first create an initial pair of qubits (with 0 and 1 representing cooperate and defect, respectively), physically transmit one qubit to each player, allow each to operate on their individual qubit before sending them back to the source, then combine the result (by undoing the initial entangling pair-creating operation). To preserve the correspondence with the original game, the game is designed so two particular quantum operations correspond to the classical choices. This allows the players, if they wish, to play the original prisoner's dilemma by restricting their choices to these particular operations. The final measurement gives a definite value for each qubit, which then corresponds to the individuals' choices. Fig. 1 schematically illustrates this game scenario.

This game involves two qubits, one for each player, so the quantum system is described by a vector of $2^2 = 4$ amplitudes, with one amplitude for each possible combination of values for the two bits, i.e., $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$. Initially, each of the two bits is set to 0, i.e., they are in the configuration $|00\rangle$, corresponding to the initial vector $v = (1, 0, 0, 0)$. After completing the steps of the game, the final vector for the pair is

$$\psi = J^\dagger(U_A \otimes U_B)Jv \tag{2}$$

where J is the entanglement operator acting on the two bits, J^\dagger , its adjoint, is the disentanglement operator, and U_A , U_B are the single-bit operators selected by the two players. Aside from the players' choices of U_A , U_B the rest of the

game is fixed and public knowledge. Specifically, J is the 4×4 matrix

$$J = \frac{1}{\sqrt{2}}(I + i\sigma_x \otimes \sigma_x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & i \\ 0 & 1 & i & 0 \\ 0 & i & 1 & 0 \\ i & 0 & 0 & 1 \end{pmatrix} \quad (3)$$

where σ_x is the 2×2 Pauli matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. The players choose from among the general single-qubit operators, given by

$$U(\theta, \phi, \alpha) = \begin{pmatrix} e^{-i\phi} \cos \frac{\theta}{2} & e^{i\alpha} \sin \frac{\theta}{2} \\ -e^{-i\alpha} \sin \frac{\theta}{2} & e^{i\phi} \cos \frac{\theta}{2} \end{pmatrix} \quad (4)$$

up to an irrelevant overall phase factor. That is, each player is given physical possession of one of the qubits, and can operate on that bit with any quantum operator. Hence the players' choices correspond to the general single-qubit operators. By contrast, the initialization operator J operates on both qubits together to produce a superposition of the two configurations $|00\rangle$ and $|11\rangle$, described by the vector $(1/\sqrt{2}, 0, 0, i/\sqrt{2})$. The inverse operator J^\dagger , applied after the players make their choices, also operates on both qubits.

Given the choices $\theta_A, \phi_A, \alpha_A$ and $\theta_B, \phi_B, \alpha_B$ of the players, matrix multiplication gives the resulting probabilities for the four outcomes, $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$ as [8]

$$\begin{aligned} & \left(\cos \left(\frac{\theta_A}{2} \right) \cos \left(\frac{\theta_B}{2} \right) \cos(\phi_A + \phi_B) - \sin \left(\frac{\theta_A}{2} \right) \sin \left(\frac{\theta_B}{2} \right) \sin(\alpha_A + \alpha_B) \right)^2 \\ & \left(\sin \left(\frac{\theta_A}{2} \right) \cos \left(\frac{\theta_B}{2} \right) \cos(\alpha_A - \phi_B) - \cos \left(\frac{\theta_A}{2} \right) \sin \left(\frac{\theta_B}{2} \right) \sin(\phi_A - \alpha_B) \right)^2 \\ & \left(\sin \left(\frac{\theta_A}{2} \right) \cos \left(\frac{\theta_B}{2} \right) \sin(\alpha_A - \phi_B) + \cos \left(\frac{\theta_A}{2} \right) \sin \left(\frac{\theta_B}{2} \right) \cos(\phi_A - \alpha_B) \right)^2 \\ & \left(\cos \left(\frac{\theta_A}{2} \right) \cos \left(\frac{\theta_B}{2} \right) \sin(\phi_A + \phi_B) + \sin \left(\frac{\theta_A}{2} \right) \sin \left(\frac{\theta_B}{2} \right) \cos(\alpha_A + \alpha_B) \right)^2 \end{aligned}$$

respectively. These expressions capture the full behavior of this 2-player quantum game, and hence allow implementing the game via simulation with minimal computational cost.

This quantum game contains the classical prisoner's dilemma. Specifically the operators $U(0, 0, 0) = I$ and $U(\pi, 0, \pi/2)$ correspond to cooperate and defect, respectively. That is, if both players restrict their choices to these two operators, the outcomes are always the same as that of the classical game. For example, if both players pick $U(0, 0, 0)$, then the outcome $|00\rangle$ (i.e., both players cooperate) has probability one.

2.2 Nash Equilibria for the Quantum Game

In the classical game, defect is the dominant strategy. By contrast, in the quantum game, no such dominant strategy exists. Specifically, for given choices

$\theta_A, \phi_A, \alpha_A$ made by Alice, Bob can always arrange for the outcome in which Alice cooperates but Bob does not, i.e., the highest possible payoff for Bob, by selecting

$$\begin{aligned}\theta_B &= \theta_A + \pi \\ \phi_B &= \alpha_A \\ \alpha_B &= \phi_A - \pi/2\end{aligned}\tag{5}$$

Hence the game strategy space is analogous to the game *stone, paper, scissors* rather than conventional prisoner's dilemma (where a single choice, i.e., defect, is always the best response). Thus there is no single strategy equilibrium. Mixed strategies provide a Nash equilibrium but there are many such equilibria, all with the same expected payoff. We use the Bayesian Nash equilibrium as the rational solution concept for the quantum game. Each individual will play a strategy (pure or mixed) such that they are mutually maximizing their expected payoff. None have an incentive to make a unilateral change to their strategy. Rational players could pick among these equilibria arbitrarily, or mix among them.

As an example of a mixed strategy, suppose Alice is considering the choice $\theta_A, \phi_A, \alpha_A$. If she always makes this choice, Bob could respond with Eq. (5), forcing Alice to cooperate while Bob does not. Realizing this response on the part of Bob, Eq. (5) gives Alice an alternate set of values $\theta_{A'}, \phi_{A'}, \alpha_{A'}$ equal to $\theta_{A'} = \theta_A, \phi_{A'} = \phi_A - \pi/2, \alpha_{A'} = \alpha_A - \pi/2$, that will instead force Bob to cooperate while Alice does not. In response to these, Bob has the choices $\theta'_B = \theta_A + \pi, \phi'_B = \alpha_A - \pi/2, \alpha'_B = \phi_A - \pi$, and finally Alice's original values are a best response to these. Thus if Alice picks randomly among $\theta_A, \phi_A, \alpha_A$ and $\theta_{A'}, \phi_{A'}, \alpha_{A'}$ while Bob picks randomly among $\theta_B, \phi_B, \alpha_B$ and $\theta_{B'}, \phi_{B'}, \alpha_{B'}$, then half the time Alice will have a best response against Bob, receiving payoff $W(1 + a/2)$. But the other half, roles will be reversed, giving Alice a payoff of $Wa/2$, and hence an overall expected payoff of $W(1 + a)/2$. Note that such a mixed strategy will be apparent in the outcome probabilities: half the time giving probability 1 of outcome 01, and the other half a probability 1 of outcome 10. Averaging over many runs of such choices, we would see outcomes 01 and 10 have average probabilities 1/2.

Since there is such a mixed strategy equilibrium for any choice of $\theta_A, \phi_A, \alpha_A$, the extreme case is a uniform mix of these possible strategies, amounting to random selection of the angles. If players use a random selection strategy, the expected values of the probabilities for each outcome are 1/4, again leading each player to cooperate half the time, but with different average probabilities than the case with a best response mixed strategy.

For any of these equilibria, we expect cooperation with probability 1/2. Thus it is more efficient than the classical prisoner's dilemma where game theory predicts no cooperation although it does not achieve the full efficient outcome.

Thus, game theory provides only weak predictions of the strategies rational players will pick. However, it supplies a strong prediction on the average rate of co-operation independent of which mixed strategy equilibrium the players will

play. Notice that game theory is predicting that the quantum version of the prisoner’s dilemma will result in substantially more (50% vs 0%) cooperation than the classical version. This potential differentiation, even when altruism is a significant factor, allows us to measure experimentally whether people behave differently in the quantum and classical versions of the prisoner’s dilemma.

3 Experimental Design

The primary goal of the experiments is to determine whether the quantum version of the prisoner’s dilemma is more efficient than the classical version. Both versions were implemented in the HP Experimental Economics software platform [9]. A second goal is to ascertain how well subjects play quantum games without specific training in quantum mechanics. While difficult to obtain a precise answer to this question even in an experimental setting, we approach it in two ways: quantitatively by comparing how close people come to the game theory predicted outcome, and qualitatively by asking participants to describe their strategy in a written paragraph after completing the experiment.

Before the experiments, subjects were directed to a web site¹ where the instructions were posted. They had to pass a web-based quiz correctly before they were allowed to participate.

3.1 Experimental Procedure

Each experiment consisted of a series of periods, during which each person played a single instance of the game. To induce behavior corresponding to randomly-matched game situations, players were randomly paired at the beginning of each period. This randomization is designed to reduce repeated game effects in which players can encourage cooperation through the expectation of future rewards or punishment, e.g., as results in the tit-for-tat behavior seen in the repeated version of the classical prisoner’s dilemma [3].

The classical version of the game is easily implemented: each player makes a binary decision of whether to cooperate or not.

In the quantum game [8], players manipulate a shared entangled state. For the quantum game, except when the choices of both players give probability one to a single outcome, the payoffs of the quantum game aren’t necessarily the same each time players make the same choices due to the probabilistic selection of outcomes. This property makes it difficult for people to achieve the rational play predicted by game theory because of uncertain feedback even if players want to co-ordinate equilibrium play. To facilitate co-ordination between players, in each period we allowed subjects two rounds of communication before they played the game for real. Here is a chronological order of the events in a period:

1. The computer randomly grouped the players into groups of two. (The experiments were conducted with even number of players.) Players were

¹[http://www.hpl.hp.com/econexperiment/Quantum Public Goods/instructions_overview.htm](http://www.hpl.hp.com/econexperiment/Quantum%20Public%20Goods/instructions_overview.htm)

not told who their opponents were.

2. Each player had two chances to send a message to his opponent in the group. The message consisted only of a set of three numbers, in the range 0 to 12.
3. They played the game by each entering three numbers.
4. The outcome was revealed.

All quantum interactions were handled through software simulation. No real quantum device was used. That is, between steps 3 and 4 the simulation used the numbers from each player to specify single-bit quantum operators, performed the operations and produced the outcome corresponding to observing the final state of the system. From the perspective of the users, this simulation provided the same behavior as would an actual physical implementation. In particular, the participants learned the outcome but not the actual choices made by their opponent.

We used the two sets of payoffs shown in Table 1. Nash equilibrium analysis indicates that both sets should give the same rates of cooperation in either the quantum or the classical version of the game. Previous work in the classical prisoner’s dilemma has shown that even in randomly-matched games, people cooperated to some extent, usually explained by non-Nash behavioral effects such as altruism [19]. Thus we varied the payoffs to see if the quantum games shows similar behavior effects. As an additional comparison, we also ran a “repeated game” version of the experiment in which players were randomly grouped into pairs just once, at the start of the experiment, and played repeated with the same opponent for all periods.

3.2 Interface and Training

In the quantum version of the game, each player picks three numbers, representing the angles of corresponding quantum operators. We chose to specify angles as numbers between 0 and 12, telling participants that the numbers corresponded to hours on a clock so 12 was equivalent to 0. Since it is not practical to teach subjects quantum mechanics or game theory, we developed tools to present the game without requiring such training. Specifically, we developed “what-if” scenario tools to show the players consequences of possible choices. Quantum mechanics was not mentioned to the subjects. All the quantum processes were described as black boxes and the “what-if” tools allowed the players to learn the behavior of these processes.

Two “what-if” tools were provided. The first one allows a subject to input 6 numbers; three for his own choices and three as guesses to his opponent’s input. Using these six numbers, the tool shows the probability of the four possible outcomes (cooperate/cooperate, cooperate/defect, defect/cooperate, defect/defect).

For this particular quantum game, for any quantum operator one person chooses, there is a corresponding operator the other can choose to force the

experiment	multiplier	number of participants	number of periods	
			classical	quantum
Lo-pilot	$a = 1.5$	6	–	21
Lo-1	$a = 1.5$	10	20	30
Hi-1	$a = 1.8$	12	30	30
Hi-2	$a = 1.8$	10	30	30
Repeated	$a = 1.5$	10	30	32

Table 2: Summary of the experiments, giving the multiplier a , number of participants and number of periods for each game. Except for the pilot experiment, each experiment involved both the classical and quantum versions of the games with the same group of people during a single afternoon.

probability of any outcome to one. The second “what-if” tool allows a subject to input three numbers as guesses of what his opponent might do. It then shows the corresponding operators (3 numbers representing each operator) that will result in certainty for each of the four possible outcomes.

Fig. 2 shows a portion of the user interface for the experiments. In the *Decision Support* section, the player examines the consequence of the choices 5, 0, 8 (i.e., θ, ϕ, α equal to $5\pi/6, 0, 4\pi/3$, respectively) when the other player makes choices 0, 0, 0. For instance, with these choices, with probability of about 70%, the player will not contribute while the opponent does, a situation giving the player a payoff of 175, as seen in the payoff table in the middle section of the interface. The lower portion of the *Decision Support* shows the player could obtain this outcome with 100% probability using the choices 6, 0, 9 (again under the assumption that the opponent chooses 0, 0, 0).

4 Results

Table 2 summarizes our experiments. In the Repeated experiment, players were randomly paired just once, at the start of the experiment, while in the others players were randomly paired at the start of each period. We used two sets of payoffs, labelled as “Lo” and “Hi”, giving lower and higher incentives for cooperation, respectively, as indicated in Table 1.

We used bootstrapping methods [23] for many statistical tests reported in this paper. Many comparisons, of cooperation rates and other measurements, were evaluated with permutation tests. Furthermore, all calculations were performed using the standard bootstrapping library of S-plus.

4.1 Cooperation Rates: Quantum vs. Classical

In all experiments, the levels of cooperation in the quantum games exceeded those of the corresponding classical games, as shown in Table 3 and as pre-

experiment	cooperation rate				p -value
	classical		quantum		
Lo-pilot	–	–	41%	52/126	–
Lo-1	16%	32/200	52%	155/300	0.002
Hi-1	50%	179/360	57%	205/360	0.08
Hi-2	25%	76/300	46%	137/300	0.002
ALL-single	33%	287/860	51%	549/1086	0.002
Repeated	66%	236/360	86%	330/384	0.002

Table 3: Average cooperation frequencies in the classical and quantum games for each experiment. In addition to the percentage of cooperation outcomes, we also show, as a fraction, the number of times a player cooperated and the total number of opportunities for cooperation (i.e., twice the number of games played during the experiment, which equals the product of the numbers of players and periods). The p -values give the probability cooperation rates would be at least as different as we observed by chance, assuming there was no difference between the classical and quantum games.

dicted by game theory [8]. This is strong evidence that the quantum game outperformed its classical counterpart even in an environment where full rationality seemed unlikely. One interpretation of this result is that while players did not solve complicated quantum physics mathematics, they found the solution through a learning process. It is analogous to any child catching a ball without consciously solving the complex equations governing the motion of the ball.

Two-population permutation tests were used to determine whether the cooperation rate averaged over the classical games was different from that of the quantum games. The typical number of re-samples is 1000. While we expect some variations in the p -values that we reported, because of bootstrapping, we do not expect it will affect our conclusions. Furthermore, the cooperation rates in the quantum games were similar for the two different sets of payoffs of Table 1. This indicates that our results are robust with variations in the payoffs.

The cooperation rates for the classical game are nonzero, even for the randomly-matched games, indicating a degree of altruistic behavior or influence of the possibility of some repeated play due to the small group size. On the other hand, the cooperation rates for the randomly-matched quantum experiments in Table 3 are consistent with the game theory prediction of 50%, as indicated in Table 4. Specifically, the table also reports the two-sided binomial test of whether the cooperation rates are significantly different from what would be expected if the players have 50% probability to cooperate each time. As one can see, we cannot reject the hypothesis of a 50% cooperation rate in all but one of the quantum games (at 5% significance). Furthermore, the null hypothesis could not be rejected when the test was performed on the pooled data of all quantum experiments. Thus, we concluded that the aggregate behavior is consistent with that predicted for the mixed strategy Nash equilibria in the case

experiment	p -value
Lo-pilot	0.06
Lo-1	0.60
Hi-1	0.01
Hi-2	0.15
ALL-single	0.73
Repeated	0.00

Table 4: Binomial test for significant difference from 50% cooperation for each quantum experiment individually as well as all randomly-matched experiments combined. The p -values are the probabilities to find cooperation rates at least as different from 50% as we observed if each case independently had 50% probability to cooperate.

experiment	p -value	
	classical	quantum
Lo-pilot	–	0.068
Lo-1	0.004	0.146
Hi-1	0.002	0.616
Hi-2	0.216	0.788
ALL-single	0.002	0.138
Repeated	0.060	0.176

Table 5: Binomial test for significant difference in cooperation between the first and second halves of the periods in each experiment. The p -values are the probabilities the difference in cooperation rates between the two halves would be at least as large as we observed if games in both halves had the same cooperation rate.

of the quantum game.

4.2 Time Dependence

Another measurement of rationality is in the time dependency of cooperation rates. Game theory predicts the cooperation rates to be zero in prisoner’s dilemma experiments, whether the games are finitely repeated or matched with a random opponent every period. Similarly, game theory predicts that the quantum version will result in 50% cooperation rate independent of whether the games are repeated with the same opponent or not. However, it is a commonly observed phenomenon in classical prisoner’s dilemma experiments that cooperation rates were more than zero and tended to decrease over time. It is interesting to see if the the quantum version will have the same result.

No significant time dependence was observed in the randomly-matched quantum games. Two-population permutation tests of cooperation rates in the first

and second halves of all experiments reveal no statistical differences, as shown in Table 5.

On the other hand, cooperation rates in the randomly matched classical games tend to decrease toward the end, perhaps indicating our randomly-matched setup is not viewed completely as a randomly-matched game by participants. Since the focus of this research is to determine whether people can play quantum games effectively, we were less concerned about strictly eliminating the end-game effect. Methods such as using a larger group of subjects who only play single games or using a random stopping time were procedurally inefficient and were not employed.

Our results on cooperation rates and time dependence illustrate an interesting paradox. While game theory requirement of rationality is *more* stringent in the case of the quantum game than the classical version, we observe aggregate behavior in the quantum game *closer* to the game theory predictions!

4.3 Individual Strategy

The results described above indicate game theory is a good predictor of aggregate behavior of the quantum mechanism. Not only does the quantum game give cooperation rates close to the 50% value of the mixed-strategy Nash equilibria, no time dependency of the cooperation rate was observed, which is also consistent with game theory. This contrasts with the case of the classical game in which a significant amount of cooperation is seen in spite of the dominant strategy for the randomly-matched game being complete defection.

Beyond these aggregate predictions, game theory also indicates rational players of the quantum game will adopt a mixed-strategy equilibrium (but does not specify which of the many equilibria will be selected since they all have the same expected payoff). However, the surveys conducted after each experiment indicate that subjects were often trying to guess what the other player would do, and then select a best response based on that guess using the provided what-if decision support tools. This suggests players were not deliberately playing a mixed strategy, contrary to the prediction of game theory. For example, one subject wrote, “I tried to anticipate what the other player was going to enter in, and I would choose a set of numbers that would result in the other player investing and me not investing”. While this is anecdotal evidence that players understood the game and tried to optimize their payoffs, they were not consciously considering mixed-strategy Nash equilibria.

To examine the question of strategy choice quantitatively, we examine more detailed game theory predictions than just the overall cooperation rate. A particular complication of the quantum prisoner’s dilemma is that random choices of the angles in Eq. (4) is one of the mixed-strategy equilibria. Thus if players chose angles randomly we could not distinguish between a sophisticated choice on their part to play a Nash equilibrium or simply uninformed random guesses as to a suitable set of angles.

The Kolmogorov-Smirnov test was used to determine whether participants were making random choices. The null hypothesis is that participants chose

θ, ϕ, α randomly with a uniform distribution. Each game gives 6 observations, 3 per participant, of these choices. The number of observations for each experiment is then six times the number of periods. The number of observations ranged from 378 to 1152. In all quantum experiments, the hypothesis of independent random choices is rejected with p -values of zero, or too small to be precisely calculated by the statistical package that we used. This is overwhelming evidence, consistent with our survey results, that participants did not play randomly.

Furthermore, because the average outcome probabilities are close to $1/4$, we can also exclude the possibility players are using a simple best-response Nash equilibrium mixed-strategy based on each player using a pair of best response choices to those of the other [8]. Such choices would result in half the time one player cooperating while the other defects, and vice versa the other half of the time, i.e., average outcome probabilities of 0.5 for $|01\rangle$ and $|10\rangle$, and zero for $|00\rangle$ and $|11\rangle$. These values do not match the observed values close to $1/4$.

In summary, both the players' descriptions of their strategies and the observed outcome probabilities indicate players are not using the mixed-strategy Nash equilibria predicted by game theory. This observation is analogous to human players of the stone-paper-scissors game trying to outguess each other rather than generate random choices.

The observed outcome probabilities give a further consistency check on our data, the probability for outcomes $|01\rangle$ and $|10\rangle$ should be the same, because in each game the choice of who is considered to be the first player and who is the second is made randomly by the system. This choice is irrelevant for the strategic properties of the game, and does not appear in the user interface. Instead, results from the "what-if" tools are always presented in terms of "my" and "other's" choices. Confirming this expectation, a two-sided test of the observed outcome probabilities gave a 99.6% probability our observations could have arisen from equal values when averaged over games.

4.4 Repeated Game

The experiment with the repeated version of the quantum prisoner's dilemma provides additional evidence of the player's understanding of the quantum game.

The analysis of repeated games is difficult for game theory since the folk theorem [1] shows that, with a suitable punishment strategy, any outcome can be a Nash equilibrium given that the game is repeated infinitely and the discount rate is high enough. While this theorem does not strictly apply to our experiments, due to the finite number of repetitions during an afternoon, it does indicate how repeated play can lead to different equilibria than single-shot games. In particular, it makes mutual cooperation a possible equilibrium in the classical prisoner's dilemma, e.g., using a punishment strategy such as tit-for-tat [3]. In practice, people do cooperate more if they are playing with the same opponent.

Our experiment also showed increased cooperation in repeated games. Cooperation rates in Table 3 show the repeated games, both classical and quantum, give higher cooperation rates than the corresponding randomly-matched games

(with p -value of 0.002). Moreover, subjects cooperated at a higher rate in the quantum version of the game than in the classical one (with p -value 0.002). This is in spite of the fact that the quantum version offers a less effective punishment strategy because a player cannot guarantee his or her defection 100% of the time.

Furthermore, this behavior is strong evidence that when the subjects wanted to, they could co-ordinate effectively at the both-cooperate outcome, indicating subjects developed a good understanding of the quantum game.

4.5 Variation Among Individuals

We also examined individual cooperation rates. In every experiment where we ran both classical and quantum versions (i.e., all except the pilot), the standard deviation of cooperation rates amongst individuals was substantially higher in the classical game than in the quantum version of the prisoner’s dilemma. A paired two-population t-test of standard deviations of individual cooperation rates gives a p -value of 0.0042, rejecting the hypothesis that the two versions have the same variation. This is strong evidence of less variation among individuals in the quantum game. In classical game, a few individuals cooperate significantly more than others. This suggests the quantum game is not only more efficient in aggregate, but also has an equalizing effect on cooperation rates. The quantum mechanism can be construed as a more “fair” system than the classical game since people cooperate at more similar rates.

5 Multiple-Player Prisoner’s Dilemma

In theory, the quantum mechanism described above can be extended to address the public goods game among groups of n players [8], with payoffs given by Eq. (1). Game theory predicts that cooperation rates will go up when the size of the group increases when playing under this quantum mechanism, with the expected payoff to each player of $W(a - (a - 1)2^{-(n-1)})$. This is better than the classical game’s equilibrium payoff W , and only slightly below the efficient outcome, aW , for large n . This increasing efficiency contrasts with the conventional observation that, in practice, free riding becomes *more* likely as group size increases due to the increased difficulty of monitoring behavior of group members [15]. Our experimental setup allows examining these assertions empirically.

In the game we study [8], for n players, a source creates $n(n - 1)/2$ entangled pairs of qubits and sends them to the players so that each pair of players in the group shares an entangled pair of qubits. The game consists of simultaneously played “mini-games” between each pair of players, using their shared entangled qubits. Thus each player participates simultaneously in $n - 1$ mini-games, each of which is identical to the 2-player quantum game described in Sec. 2. For simplicity, players are constrained to make the same choices for all of their mini-games. In the case of homogeneous preferences and randomly selected groups

that we study, relaxing this constraint does not change the strategic aspects of the game [27]. With multiple bits for each player, we also need a rule relating final observed bit values to the player’s contribution. In our case [8], for each player, if any of their $n - 1$ bits is observed to be a 1, they contribute to the public good, i.e., in effect they preauthorize a charge to their account based on this outcome.

Game theory indicates the value of a (between 1 and n) has no effect on equilibrium contributions in either the quantum or the classical version of the game. However, a higher a means subjects gain more if they cooperate and could affect actual behavior. To address this possibility, we examined two cases. First, we used $a = 0.75n$, so the difference in payoff to each person when everyone contributes and when no one does, $(a - 1)W$, grows with n . Thus as group size increases, people have more to gain by finding a way to create the public good rather than not producing any. However, the temptation to defect remains constant: no matter what contributions others make, an individual gains $(1 - a/n)W = 0.25W$ by not contributing. This choice of payoffs corresponds to a public goods scenario, such as building a local park, in which the quality of the good increases with the size of the group involved. This situation provides a stringent test for whether the quantum game can outperform the classical one, especially in the small groups feasible to test in the laboratory, because some altruistic behavior is often observed. Thus the increasing benefit of the public good with group size may lead the classical game to have significant contribution rather than the game theory prediction of none.

The second choice for a was the fixed value $a = 1.5$, independent of group size. In this case, the difference in payoffs between everyone contributing and nobody contributing is independent of n while the temptation to defect, $(1 - a/n)W$, *increases* with n . With relatively less to gain from the efficient outcome of all contributing, and less influence on that gain as the group size increases, we expect the classical game contribution to be smaller than when a increases with group size. Thus this scenario provides a difficult case for public goods provisioning. In this case, our main interest is not just whether the quantum version is more efficient than the classical one, but rather whether the quantum version can provide a significant level of contribution at all.

For $n = 2$, both choices for a give the same value, i.e., $a = 1.5$, corresponding to the “Lo”-payoff experiments listed in Table 2.

Table 6 summarizes the experiments with larger groups.

In all experiments, the levels of contributions in the quantum games exceeded those of the corresponding classical games, as shown in Table 3 and as predicted [8]. Fig. 3 shows the effect of group size on contributions. Moreover, contribution rates in the quantum games showed no significant dependence on time (i.e., number of periods).

The contribution rates for the quantum games in Table 7 are consistent with the game theory prediction of $1 - 2^{-(n-1)}$, as indicated in Table 8, which shows the two-sided binomial test of whether the contribution rates are significantly different from what would be expected if the players have $1 - 2^{-(n-1)}$ probability to contribute each time.

experiment	multiplier	group size	number of participants	number of periods	
				classical	quantum
E3a	$a = 2.25$	3	12	29	12
E3b	$a = 2.25$	3	12	28	23
E4a	$a = 3$	4	16	15	13
E4Low	$a = 1.5$	4	16	31	29

Table 6: Summary of the experiments, giving the multiplier a , number of participants and number of periods for each game.

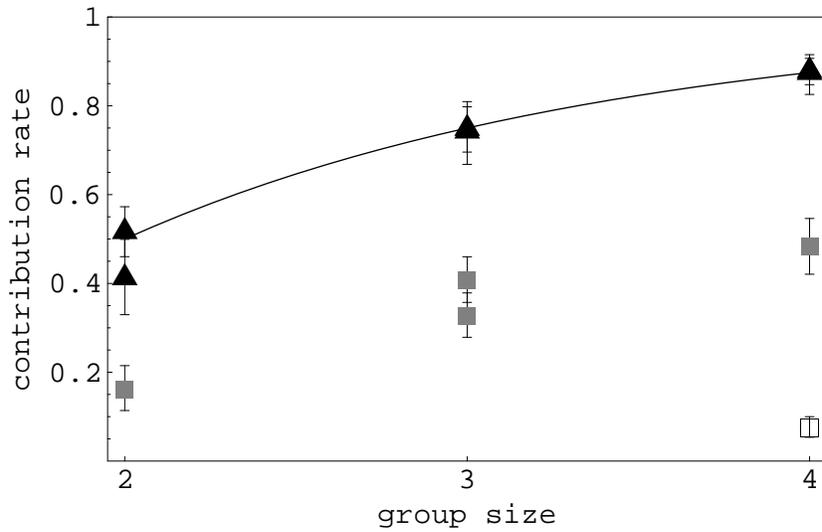


Figure 3: Contribution rates for classical (box) and quantum (black triangle) games as a function of group size n . Error bars show the 95% confidence intervals for the contribution rates based on the sample size. The curve shows the contribution rates for the quantum game predicted by game theory, $1 - 2^{-(n-1)}$. The open box is classical contribution rate for experiment E4Low. The two experiments with each of $n = 3$ and 4 have indistinguishable contribution rates for the quantum games.

experiment	contribution			
	classical		quantum	
E3a	41%	142/348	74%	107/144
E3b	33%	110/336	75%	206/276
E4a	48%	116/240	88%	182/208
E4Low	7%	37/496	88%	408/464

Table 7: Average contribution frequencies in the classical and quantum games shown as a percentage and as a fraction of the number of times a player contributed and the total number of opportunities for contribution (i.e., the product of the numbers of players and periods). In all experiments with both classical and quantum games, the p -value for the probability contribution rates would be at least as different as we observed by chance, assuming there was no difference between the classical and quantum games, was less than 0.002.

experiment	p -value
E3a	0.85
E3b	0.89
E4a	1
E4Low	0.83

Table 8: Binomial test for significant difference from the game theory predicted $1 - 2^{-(n-1)}$ contribution rates. The p -values are the probabilities to find contribution rates at least as different from predicted values as we observed if each case independently had the predicted probability to contribute.

6 Conclusion

We reported the first experimental evidence that people, without training in quantum physics, can play a simple quantum game effectively. Furthermore, our observations are consistent with the game theory predictions that the quantum version of the prisoner’s dilemma is more efficient, in a Pareto sense, than the classical version. In addition, we found that although the aggregate predictions (e.g., about cooperation rate) of game theory were accurate, there was substantial deviation when individual strategies were analyzed.

Statistical tests with significance based on a p -value of less than 5% of the difference arising from chance leads to the following conclusions. The quantum game gives higher cooperation than classical. Changing the temptation to defect and the benefit of mutual cooperation did not affect cooperation rates for the quantum game (as predicted by game theory). The quantum randomly-matched game gives less cooperation than repeated classical play. Hence the quantum mechanism, while improving over the randomly-matched classical game, is not as effective as reputation effects from repeated play. The quantum game with repeated play gives more cooperation than repeated classical play. There is no time-dependence in the randomly-matched quantum games, but classical games do show time dependence.

When extended to larger groups, the level of cooperation in the quantum game increased as predicted. This is strong evidence that subjects were responding to the strategic considerations of the quantum game.

In addition, we found significantly larger variation in cooperation rates among individuals in the classical games than in the quantum games. In this sense, the quantum games are more “fair” to participants, in spite of the quantum game being significantly more complicated to analyze.

As for the strategies players used to determine their choices, we found they did not simply make random choices, nor use a best-response mixed-strategy among two choices. Either of these are Nash equilibria. We found players often selected multiples of $\pi/2$, which arise in the best-response values displayed by our “what-if” tools under the assumption that the other player uses 0,0,0 or best-response iterates of that choice. But even in these cases, they did not select randomly among these iterates of best responses.

There are many future directions for this line of experimental research. While the quantum mechanism solved the free-rider problem in the experiments, it is not yet practical as an economic mechanism in the real world. Two issues need to be addressed: application to more complex economic situations and comparison with conventional economic methods.

For the first issue, the game we considered used the same set of payoffs for every player. In more realistic scenarios, players could have different preferences which further complicate the mechanism design. For example, some players may get such a low benefit from cooperating that the efficient outcome is for them to defect. In this case the mechanism should give an outcome in which the low-benefit players do not contribute, while ensuring those with high benefits contribute (e.g., by not having incentives to pretend to have low benefits). It

would therefore be interesting to experiment with heterogeneous preferences among the players.

The second issue is that there are plenty of classical economic mechanisms that solve the free-rider problem in different environments. For example, taxation is a common method, with thousands of years of history, to provide public goods for a community. The dominant issue in such a scenario is to determine preferences, which can be heterogeneous, in the community so the taxing authority can determine the optimal amount of public goods to procure and hence set the corresponding taxation rate. A seminal classical solution is the Groves-Ledyard mechanism [16, 4], which can solicit truthful revelation of preferences in an incentive compatible manner. Another example is the provision point mechanism, which is popular amongst charities. In the provision point mechanism, contributions are gathered and if the total passes a certain threshold, the public goods will be produced. If not, the contributions will be refunded by a trusted third party. While these approaches can be effective, they have limitations, e.g., in requiring a third party enforcer who may not be able to elicit truthful preferences, or destroying some contributions to maintain correct incentives, or provide no incentives to participate. Moreover, if the players expect to repeat the game in the future, reputation-based mechanisms can help encourage contribution. Thus there is a need to experimentally evaluate the trade-off between quantum and classical mechanisms to identify situations where the quantum mechanism may be superior in practice. The quantum mechanism will only be worthwhile to implement, given the high costs of developing new hardware, if we can find significant benefits over existing classical solutions.

Game theory suggests an efficiency gain with multiple players even if some pairs of qubits are not in fact implemented correctly [27] (e.g., due to noise when creating the pairs, or deliberate choice by the provider of pairs). This gain occurs even if the fraction of incorrect pairs is public knowledge provided participants do not know *which* particular pairs these are. Since people may not treat risks and probabilities in a rational manner [25], it would be interesting to evaluate how such noise affects behavior. In particular, whether the game theory prediction that efficiency decreases smoothly with noise is consistent with human behavior.

This research establishes potential benefits of quantum economics from the perspective of behavioral issues. The fact that participants untrained in quantum mechanics were able to achieve higher performance with the quantum game suggests that future efforts on physical implementation will have the benefits predicted by game theory, as well as the security and trust guarantees provided by a physical implementation. Prior to our study, it was not clear whether the substantial effort required for physical implementation would achieve the economic payoffs predicted by game theory. This also opens possibilities of quantum economics applications in various economic areas that involve some need for coordination, including public goods, correlated-value auctions, coordination games and digital rights management. Many research institutions, including HP Labs, are in the process of physically implementing the quantum devices required for these mechanisms. In the foreseeable future, the benefits

seen in these experiments could be realized as feasible quantum economic mechanisms.

Acknowledgments

We thank Raymond Beausoleil, Issac Chuang, Phil Kuekes and Li Zhang for helpful discussions.

References

- [1] Robert Aumann and Sergiu Hart, editors. *Handbook of Game Theory with Economics Applications*, volume 1. Elsevier, 1992.
- [2] R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, March 1981.
- [3] Robert Axelrod. The evolution of strategies in the iterated prisoner’s dilemma. In Lawrence Davis, editor, *Genetic Algorithms and Simulated Annealing*, chapter 3, pages 32–41. Morgan Kaufmann, Los Altos, CA, 1987.
- [4] Mark Bagnoli and Barton L. Lipman. Provision of public goods: Fully implementing the core through private contributions. *Review of Economic Studies*, 56:583–601, 1989.
- [5] Simon C. Benjamin and Patrick M. Hayden. Multiplayer quantum games. *Physical Review A*, 64:030301, 2001.
- [6] Colin Camerer. Individual decision making. In John Kagel and Alvin E. Roth, editors, *The Handbook of Experimental Economics*, chapter 8, pages 587–703. Princeton Univ. Press, 1995.
- [7] Colin Camerer, Teck Ho, and Juin-Kuan Chong. A cognitive hierarchy model of one-shot games. *Quarterly Journal of Economics*, 119(3):861–898, 2004.
- [8] Kay-Yut Chen, Tad Hogg, and Raymond Beausoleil. A quantum treatment of public goods economics. *Quantum Information Processing*, 1:449–469, 2002. arxiv.org preprint quant-ph/0301013.
- [9] Kay-Yut Chen and Ren Wu. Computer games and experimental economics. In *Proc. of the 5th Intl. Conf. on Enterprise Information Systems (ICEIS)*, April 2003.
- [10] Jiangfeng Du et al. Entanglement enhanced multiplayer quantum games. *Physics Letters A*, 302:229–233, 2002. arxiv.org preprint quant-ph/0110122.

- [11] Jiangfeng Du et al. Experimental realization of quantum games on a quantum computer. *Physical Review Letters*, 88:137902, 2002. arxiv.org preprint quant-ph/0104087.
- [12] J. Eisert, M. Wilkens, and M. Lewenstein. Quantum games and quantum strategies. *Physical Review Letters*, 83:3077–3080, 1999. arxiv.org preprint quant-ph/9806088.
- [13] Jens Eisert and Martin Wilkens. Quantum games. *J. Modern Optics*, 47:2543–2556, 2000. arxiv.org preprint quant-ph/0004076.
- [14] Richard P. Feynman. *QED: The Strange Theory of Light and Matter*. Princeton Univ. Press, NJ, 1985.
- [15] Natalie S. Glance and Bernardo A. Huberman. Dynamics of social dilemmas. *Scientific American*, 270(3):76–81, March 1994.
- [16] Theodore Groves and John O. Ledyard. Optimal allocation of public goods: A solution to the ‘free rider’ problem. *Econometrica*, 45:783–809, 1977.
- [17] G. Hardin. The tragedy of the commons. *Science*, 162:1243, 1968.
- [18] Bernardo A. Huberman and Tad Hogg. Quantum solution of coordination problems. *Quantum Information Processing*, 2:421–432, 2003. arxiv.org preprint quant-ph/0306112.
- [19] John Ledyard. Public goods: A survey of experimental research. In John Kagel and Alvin Roth, editors, *Handbook of Experimental Economics*, chapter 2, pages 111–181. Princeton Univ. Press, 1995.
- [20] David A. Meyer. Quantum strategies. *Physical Review Letters*, 82:1052–1055, 1999. arxiv.org preprint quant-ph/9804010.
- [21] David A. Meyer. Quantum communication in games. In *Quantum Communication, Measurement and Computing*, volume 734, pages 36–39. AIP Conference Proceedings, 2004.
- [22] Pierfrancesco La Mura. Correlated equilibria of classical strategic games with quantum signals. arxiv.org preprint quant-ph/0309033, Sept. 2003.
- [23] Eric Noreen. *Computer-Intensive Methods for Testing Hypotheses*. John Wiley, NY, 1989.
- [24] Thomas R. Palfrey and Howard Rosenthal. Testing game-theoretic models of free riding: New evidence on probability bias and learning. In Thomas R. Palfrey, editor, *Laboratory Research in Political Economy*, pages 239–267. Univ. of Michigan Press, 1991.
- [25] Richard H. Thaler. Anomalies: The ultimatum game. *J. of Economic Perspectives*, 2(4):195–206, 1988.

- [26] S. J. van Enk and R. Pike. Classical rules in quantum games. *Physical Review A*, 66:024306, 2002.
- [27] Li Zhang and Tad Hogg. Reduced entanglement for quantum games. *Intl. J. of Quantum Information*, 1(3):321–335, 2003.