

Quantum Science and Technology



PAPER

Demonstration of a Bayesian quantum game on an ion-trap quantum computer

RECEIVED
29 April 2018

REVISED
19 June 2018

ACCEPTED FOR PUBLICATION
26 June 2018

PUBLISHED
10 July 2018

Neal Solmeyer^{1,7,8} , Norbert M Linke², Caroline Figgatt², Kevin A Landsman², Radhakrishnan Balu^{3,4} , George Siopsis⁵  and C Monroe⁶

¹ The MITRE Corporation 7596 Colshire Drive, McLean, VA, 22102-7539, United States of America

² Joint Quantum Institute, Department of Physics, and Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, United States of America

³ Computer and Information Sciences Directorate, Army Research Laboratory, Adelphi, MD, 21005-5069, United States of America

⁴ Computer Science and Electrical Engineering, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250, United States of America

⁵ Department of Physics and Astronomy, The University of Tennessee, Knoxville, TN 37996-1200, United States of America

⁶ Joint Quantum Institute, Department of Physics, and Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, USA, and IonQ Inc., College Park, MD, 20740, United States of America

⁷ This work was conducted while N Solmeyer was at the Army Research Laboratory. The author's affiliation with the MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or the support for, the positions, opinions, or viewpoints expressed by the author. Approved for Public Release; Distribution Unlimited. Case Number 17-4775.

⁸ Author to whom any correspondence should be addressed.

E-mail: meyer.neal@gmail.com, radhakrishnan.balu.civ@mail.mil, radbalu1@umbc.edu and siopsis@tennessee.edu

Keywords: quantum games, quantum computation, ion traps

Abstract

We demonstrate a Bayesian quantum game on an ion-trap quantum computer with five qubits. The players share an entangled pair of qubits and perform rotations on their qubit as the strategy choice. Two five-qubit circuits are sufficient to run all 16 possible strategy choice sets in a game with four possible strategies. The data are then parsed into player types randomly in order to combine them classically into a Bayesian framework. We exhaustively compute the possible strategies of the game so that the experimental data can be used to solve for the Nash equilibria of the game directly. Then we compare the payoff at the Nash equilibria and location of discontinuous changes in the payoff obtained from the experimental data to the theory, and study how it changes as a function of the amount of entanglement.

1. Introduction

Game theory, originally developed in the 1940s and 1950s [1–3], has since been the source of important contributions in fields such as economics [4], political science [5], biology [6], and computer science [7]. The addition of quantum information concepts into games led to the formalization of quantum games [8, 9]. Since their introduction, quantum games have been studied in a variety of contexts. With the growing prevalence of quantum computers and quantum networks, quantum games emerge as strong candidates for real world applications in quantum security protocols [10], distributed quantum computing algorithms [11], or improving the efficiency of classical network routing algorithms [12]. Because game theory is required to describe any system that has multiple independent agents, quantum games may be useful and even necessary to the rapidly growing field of quantum machine learning where quantum systems will behave autonomously and independently.

In contrast with many interesting quantum computing algorithms, quantum games can be demonstrated with small numbers of qubits, making them an attractive application for early demonstrations on quantum computers. There have been several experimental demonstrations of quantum games using NMR quantum computers [13, 14] and various linear optical quantum computing schemes [15–21]. In these demonstrations, the circuit equivalent of the game was executed with the strategy choices of the Nash equilibrium, as determined

by theoretical analysis, and compared the expected payoff at Nash equilibrium. This amounts to a type of benchmarking of the performance of the quantum computer under the framework of game theory.

The most common use of quantum computing is to produce an algorithm with a speedup over the classical version, in terms of resources used or gates employed. A natural question is whether or not quantum information or quantum games could be used to speedup the calculation of the Nash equilibria in classical games, where computational speed is limiting. Optimization of finding the solutions in the quantum context is not explored and is beyond the scope of this paper. For our purposes, quantum information is introduced into the games in order to produce equilibria that do not exist in the classical version of the game.

One might argue that a true demonstration of a quantum game would actually involve real players, either humans or computers, playing the game on true quantum hardware. With real players playing on a classically simulated quantum computer, people may actually play quantum games more rationally than they play classical games, even if they have no prior training in quantum mechanics [22, 23]. This may be due to the fact that people may have fewer preconceived notions about the quantum strategies, and are thus more likely to simply play for the highest payoff. Though this may have interesting implications for potential real world applications of quantum games, there is a gap between the potential uses of quantum games and the availability of quantum hardware.

This work aims to partially address that gap by performing a more complete demonstration of a game on a scalable quantum architecture that can be applied to more complex game scenarios in the future. Ion-trap-based systems are promising candidates for quantum computers, which is a prerequisite for quantum games. Further, ion-trap architectures are also promising for quantum networking [24], where nodes are remotely located and entangled, which may also be a requirement for quantum game applications that require the players to be remotely located.

We realize a game with incomplete information, i.e., a Bayesian game [25]. The amount of incomplete information is determined by a probability distribution of different player types. Bayesian games are of interest because of their deep connection to Bell's inequalities [26]. The Bayesian game we analyze is not directly derived from a Bell's inequality, but is rather a Bayesian game formed by incorporating incomplete information classically into a quantum game. The motivation behind this approach is that incomplete information may often be a feature of any potential application of quantum games, not just games specifically designed to violate Bell's inequalities, and the interplay of classical probability and quantum statistics can lead to a rich structure [27].

This is the first experimental demonstration of a quantum game using a scalable architecture. Because of the sophistication of current ion-trap quantum computers compared to previously used quantum computers, a more extensive implementation is possible. Using a novel parallelization scheme, we perform an exhaustive computation of all possible strategy choice combinations. This allows us to solve for the Nash equilibria of the game based only on the experimental outcomes. Therefore, we can look not only at the payoff at Nash equilibria, but also study where the equilibria occur and observe the effect of experimental noise on the location of the discontinuous changes in the payoff that can occur when one set of equilibria changes to another as the entanglement or amount of incomplete information is changed [28]. We find that the deviation of the expected to measured payoff grows with the degree of entanglement in the game. We also find that the deviation between the experimental and theoretical location of the discontinuity grows with entanglement.


2. Quantum game implementation

We experimentally demonstrate a Bayesian game that displays several features worth exploring such as multiple simultaneous Nash equilibria and discontinuous behavior as functions of both the Bayesian probability and the amount of entanglement. A detailed theoretical analysis of this game can be found in [27]. Being composed of two player games, the demonstration is relatively straightforward to implement experimentally, yet has a rich enough structure to observe several features of Bayesian games that are interesting from a game theoretical perspective.

First, we describe the two player game, which is used to construct the Bayesian game, see figure 1. An entangling gate is applied to two qubits, initialized in the state $|00\rangle$, each qubit corresponding to one of two players. The entangling operation is given by the general XX gate:

$$J(\chi) = e^{i\chi X \otimes X} = \cos \chi + iX \otimes X \sin \chi$$

$$= \begin{pmatrix} \cos \chi & 0 & 0 & -i \sin \chi \\ 0 & \cos \chi & -i \sin \chi & 0 \\ 0 & -i \sin \chi & \cos \chi & 0 \\ -i \sin \chi & 0 & 0 & \cos \chi \end{pmatrix}, \quad (1)$$



where $\chi \in [0, \pi/4]$. For $\chi = 0$, J is the identity matrix, where the resulting state, $|00\rangle$ has a concurrence of 1. For $\chi = \pi/4$, a Bell state with maximal entanglement is produced, $1/\sqrt{2}(|00\rangle + |11\rangle)$, with a concurrence of 1. After the entangling operation, the players apply their strategy choices U_A and U_B . Finally we will need the conjugate transpose of $J(\chi)$, $(J(\chi))^\dagger = J(-\chi)$. The final state for the two player game is given, in the basis $|A, B\rangle$, by:

$$|\psi_f\rangle = J^\dagger(\chi) \cdot (U_A \otimes U_{B1}) \cdot J(\chi)|00\rangle. \tag{2}$$

The payoff is then calculated based on the measurements of the qubits. If the outcome of the measurement is $|0\rangle$, this corresponds to one strategy choice in the classical game (the analogue of cooperation (C) in the prisoner’s dilemma), whereas if the outcome is $|1\rangle$, this corresponds to the analogue of defection (D). The payoff for a given strategy choice set is determined by the payoff matrix for the game:

$$\begin{matrix} A|B_1 & |0\rangle(C) & |1\rangle(D) \\ |0\rangle(C) & (11, 9) & (1, 10) \\ |1\rangle(D) & (10, 1) & (6, 6) \end{matrix} \quad \begin{matrix} A|B_2 & |0\rangle(C) & |1\rangle(D) \\ |0\rangle(C) & (11, 9) & (1, 6) \\ |1\rangle(D) & (10, 1) & (6, 0). \end{matrix}$$

The two payoff matrices are for the games versus the two different player types for player B . We will call the type of player B that uses the first payoff matrix B_1 , where player B_2 uses the second payoff matrix. The rows and columns represent the outcomes (which are the strategy choices in classical games) of player A and B respectively, and the numbers $(\$^A, \$^B)$ are the payoffs for player A and B .

The game between A and B_1 is the standard prisoner’s dilemma, while in the game between A and B_2 , player B_2 believes player A is the DA’s brother, which gives player A an advantage resulting in an asymmetric payoff between the players.

The payoff for player A is given by the expectation value of the final state weighted by the elements of the payoff matrix:

$$\langle \$^A \rangle = \sum_i |\langle i|\psi_f\rangle|^2 \$^A_i, \tag{3}$$

where the sum is over all four possible measurement outcomes of the two-qubit system, and the $\A_i are the corresponding elements of the payoff matrix.

A game is also defined by the allowable strategy choices. We implement the game using four possible strategy choices, compared to the most commonly used set of three. The four-choice single player strategy set, which we label U , is given by the three Pauli matrices (X , Y , and Z) plus the identity (I). This choice of possible strategies bounds the results of a game with arbitrary continuous strategy choices [29, 30]. With no entanglement, the strategies I and X correspond to the classical strategy choices of cooperation (C) and defection (D) respectively. There are no classical analogues for the strategy choices Y and Z . The classical analogue of entanglement is a correlated equilibria in classical game theory, which requires communication between the players.

The Bayesian game consists of two players A and B . Player A is of one type, and player B can be one of two types, B_1 or B_2 . Player A plays with either B_1 or B_2 , with some probability p , which parametrizes the incomplete information held by player A which is maximal when $p = 0.5$, and minimal when $p = 0$ or 1 . This game uses three qubits, one for each player type, and proceeds as follows. All three qubits are initialized to state $|0\rangle$. The entangling operation $J(\chi)$ is performed between either A and B_1 or A and B_2 probabilistically. Each player chooses a strategy U_i from the set U and applies it to their qubit. Then the appropriate unentangling operation $J^\dagger(\chi)$ is performed. Finally, the three qubits are measured and the expected payoff for player A is given by the weighted average of playing with B_1 and B_2 :

$$\langle \$^A \rangle = \langle \$^A(A, B_1) \rangle(p) + \langle \$^A(A, B_2) \rangle(1 - p) \tag{4}$$




Figure 2. Circuit used for parallel implementation of two player quantum games. The top two qubits correspond to the two players, and the bottom three qubits are auxiliary qubits used to run various strategy choices in parallel. The two player qubits are entangled, $J(\chi)$ and a Hadamard is applied to each of the auxiliary qubits. The CNOT and $C-Z$ gates followed by either an I or X gate, entangle the player qubits with the auxiliary qubits so that all 16 strategy choice combinations can be run with two 5-qubit circuits. Finally, an unentangling gate is applied and a measurement is performed on all qubits.

and the payoff for the B players is given by $\langle \$^{B1}(A, B_1) \rangle$ and $\langle \$^{B2}(A, B_2) \rangle$. There is structure in the game as the amount of entanglement, χ , is varied, and as the probability to play with either player, p , is varied. We demonstrate this by varying χ in the experiment, and varying p in the analysis. To implement a Bayesian game between players A , B_1 , and B_2 , we run many versions of the two player circuit shown in figure 1, and then combine them into a Bayesian game in the data analysis. This is similar to what has been done to produce initial states for quantum games which are a mixture of entangled states that exhibit quantum discord [19].

The demonstration requires running $4 \times 4 = 16$ two-qubit circuits for different strategy combinations, and to make efficient use of the hardware available, we employ a novel parallelization scheme. We compute the circuit for multiple strategy choices simultaneously by using auxiliary qubits in superposition. If the parallelization scheme were not used, the computer could only compute two strategy choice combinations at once which would therefore require eight different circuits to be run. Switching the circuit is the primary contribution to overhead time, and thus, the parallelization scheme decreases the run time by a factor of four.

Two of the qubits must be assigned to the players, while the three remaining qubits can be used to simultaneously run eight of the 16 strategy choice combinations. The circuit for the parallel implementation is given in figure 2. Starting with the three auxiliary qubits in the state $|000\rangle_{123}$, we apply a Hadamard gate H on each of them. Together with the other 2 qubits, we form $|00+++ \rangle_{AB123}$, where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

We perform $J(\chi)$ on qubits A and B , resulting in the state

$$J_{AB}(\chi)|00+++ \rangle_{AB123}. \quad (5)$$

We use the first auxiliary qubit as control to perform a controlled-NOT (CNOT) gate on qubit A .

$$\frac{1}{\sqrt{2}} \sum_{x=0}^1 X_A^x J_{AB}(\chi) |00x++ \rangle_{AB123}. \quad (6)$$

We use the second auxiliary qubit as control to perform a controlled- Z ($C-Z$) gate on qubit A .

$$\frac{1}{2} \sum_{x,y=0}^1 Z_A^y X_A^x J_{AB}(\chi) |00xy+ \rangle_{AB123}. \quad (7)$$

Then we use the third auxiliary qubit as control to perform $C-Z$ on qubit B .

$$\frac{1}{2^{3/2}} \sum_{x,y,z=0}^1 Z_B^z Z_A^y X_A^x J_{AB}(\chi) |00xyz \rangle_{AB123}. \quad (8)$$

Finally, we apply $J(-\chi)$ on qubits A and B , to obtain

$$\frac{1}{2^{3/2}} \sum_{x,y,z=0}^1 J_{AB}(-\chi) Z_B^z Z_A^y X_A^x J_{AB}(\chi) |00xyz \rangle_{AB123} \quad (9)$$

and then measure all qubits. The final state contains 8 terms corresponding to different strategies, $U_A \in \{I, X, Y, Z\}$ and $U_B \in \{I, Z\}$.

For the remaining strategies, we need to apply X on qubit B before the final step. In this case, the final state is

$$\frac{1}{2^{3/2}} \sum_{x,y,z=0}^1 J_{AB}(-\chi) X_B Z_B^z Z_A^y X_A^x J_{AB}(\chi) |00xyz \rangle_{AB123}. \quad (10)$$

Its terms correspond to the remaining eight strategies, $U_A \in \{I, X, Y, Z\}$ and $U_B \in \{X, Y\}$.

3. Experimental procedure

The experimental setup constitutes a programmable quantum computer. It consists of five $^{171}\text{Yb}^+$ ions which are trapped in a linear arrangement using a radiofrequency (RF) Paul trap, and laser cooled close to their motional ground state. Two states in the $^2S_{1/2}$ ground level are used as the qubit states ($|0\rangle = |F = 0, m_F = 0\rangle$, and $|1\rangle = |F = 1, m_F = 0\rangle$). They differ in energy by the 12.642 821 GHz hyperfine splitting which is insensitive to the magnetic field to first order. This so-called ‘atomic clock’ qubit has a typical coherence time of 0.5 s which can be further extended by suppressing magnetic field noise. Optical pumping initializes the entire five-qubit register, and readout is performed collectively by detecting state-dependent fluorescence [31].

Each ion is imaged onto its own channel of a multi-channel photomultiplier tube which allows its state to be determined with 99.4(1)% average fidelity, while the 5-qubit state detection is limited by channel-to-channel crosstalk to 95.7(1)% average fidelity. For averaged state probabilities, this state-preparation and measurement (SPAM) error can be straightforwardly re-normalized by applying an independently determined crosstalk-matrix. We drive qubit operations by applying a pair of Raman beams that are configured to form beat notes near the qubit frequency to the ions. Both beams are derived from a single 355 nm mode-locked laser. One beam is applied globally to the entire chain while the second is split into a linear array of individual addressing beams, each of which is focused onto a single ion using a multi-channel acousto-optic modulator (AOM) [32].

Single qubit gates (so-called R-gates) are applied by driving resonant Rabi flopping on any individual ion with the duration, phase, and amplitude defined by the RF signals on the multi-channel AOM. We achieve two-qubit gates (so-called XX-gates) by applying bichromatic Raman beat notes near the motional sideband frequencies. They create an effective XX-Ising interaction between the spin degrees of freedom containing the qubit mediated by all of the collective motional modes in the ion chain [33–35]. In order to leave spin and motion disentangled at the end of the operation, we employ a pulse-shaping scheme during the gates [36, 37].

Any pair of qubits can be entangled in this way, which makes this a fully-connected system of qubits [38]. A classical compiler breaks down a library of computational gates, such as Hadamard, CNOT, or controlled-Phase, into the native R- and XX-operations. Since any context-dependence (such as calibration parameters) of the native operations is handled by the compiler, the high-level gates become modular. Arbitrary circuits can then be implemented from a user interface by specifying a sequence of computational and/or native gates which makes the system programmable. Our compiler breaks down computational gates into basic gates by using a look-up table which contains pre-optimized circuits for standard gates such as CNOT, as well as a list of calibration parameters which are periodically updated. The theoretical concept behind the compiler is described in [39]. Native single- and two-qubit gate fidelities are typically 99.5(2)% and 98.5(5)%, respectively. Gate times are about 10 μs for single- and 210 μs for two-qubit gates.

4. Results

We perform the simulation for nine values of χ between zero and maximal in steps of $\frac{\pi}{40}$. The nominal values of χ were, in units of π : (0, 0.025, 0.05, 0.075, 0.1, 0.125, 0.15, 0.175, 0.2, 0.225, 0.25). For each of the two five-qubit circuits depicted in figure 1, we have $\sim 30\,000$ runs. Before and after each data run (except for $\chi = (0, 0.125, 0.25)$), a run was taken to measure the value of χ to take any deviation from the nominal values into account. The procedure for $\chi = (0, 0.125, 0.25)$ were taken at an earlier date before the calibration of χ was implemented. Such deviations are the result of calibration inaccuracy in the experiment primarily due to laser power and alignment over the course of the experimental data run. The average of the measured χ s are used for the analysis and were found to be: (0, 0.027(2), 0.054(3), 0.080(4), 0.108(4), 0.125, 0.151(5), 0.178(5), 0.201(5), 0.224(6), 0.25). For each experimental run, the quantum computer outputs the measured value (i.e., either $|0\rangle$ or $|1\rangle$) for all five qubits.

The data are parsed into two groups representing games of A versus B_1 and A versus B_2 respectively. The value of p , which determines the probability that player A plays with either B_1 or B_2 , is chosen for a given analysis, and each data point is sorted randomly into the two categories with probability p . Next, the expectation value of all 32 possible outcomes is computed to form the output population vector for the two data sets. SPAM correction is then applied to the population vector to correct for readout errors.

Next, in order to compute the payoff, equation (4), for the players for a given set of strategy choices, we must determine the outcomes of the qubits for player A and B depending on the 8 possible outcomes for the three auxiliary qubits. For example, if the ‘ P ’ circuit in figure 2 outputs $|000\rangle$ for the three auxiliary qubits, this corresponds to players A and B having applied the strategies $\{I, I\}$. The components with the auxiliary output of $|000\rangle$ are summed to form the 4-component vector representing the expectation values of the A and B qubits for each of the 8 auxiliary output combinations. These population vectors are each then re-normalized and the experimental payoff is computed with equation (3). This is done for both the I -circuit and the X -circuit of




figure 2 in order to compute the experimental payoff for all 16 possible strategy choice combinations of the two player payoff matrix. The data for both types of player B receive the same treatment.

From the two 4×4 two player payoff matrices, we can compute the 64 element $4 \times 4 \times 4$ payoff matrix of the Bayesian game according to equation (4). The same data are analyzed for p values ranging from 0 to 1 in steps of 0.01 in order to observe the structure of the game as a function of the incomplete information.

In order to compute the Nash equilibria, the best response curves must be constructed. The best response curve for player A is given by the list $\beta_A = \{i^*, j, k\}$, where j and k run over all of the possible strategy choices of B_1 and B_2 , and i^* is the strategy choice that gives A the highest payoff for the choices of j and k for players B_1 and B_2 . The best response curves for the B players are similarly calculated, $\beta_{B1} = \{i, j^*, k\}$, and $\beta_{B2} = \{i, j, k^*\}$. The Nash equilibria are given by the intersection of the best response curves, $\beta_A \cup \beta_{B1} \cup \beta_{B2}$. In other words, any strategy choice combination where each player is playing their best response to the other player's choices is a Nash equilibrium. In general, this can result in one or more Nash equilibria or none.


For the computation of the best response curves from the experimental data, we consider a payoff to be the best response if it is within an amount δ of the maximum. If this is not included, the best response for each player will only be one particular strategy choice set, that will be determined by the experimental noise, and we almost never get a Nash equilibrium. The data presented use $\delta = 0.1$.

We plot the payoffs of the three players, as a function of the entanglement χ , both theoretically and experimentally, for the nine different values of entanglement in figure 3. These data are analyzed for the case where $p = 0.5$. The error in the horizontal represents the error in entanglement calibration data taken before and after each set of data. The error in the payoff is the statistical error determined by the error in a binomial distribution of ion populations assuming that, on average, there were 3000 shots contributing to each point, with ~ 300 of them in $|1\rangle$, which was found to be the case for the dominant equilibria.

The data show that for no entanglement, $\chi = 0$, the experimental data very closely match the theoretical data. With growing χ there is a systematic shift by which the experimental results fall predominantly below the theory values, probably a result of the growing error introduced by the gate. However this effect is smaller than our estimated statistical error. The theoretical and experimental data have no Nash equilibrium above a critical value of entanglement, $\chi \sim 0.175\pi$. Note also that the experimental data for $\chi = 0.025\pi$ show a Nash equilibrium even though none is predicted theoretically. This is indicative of the fact that experimental errors can alter the critical values of the discontinuous equilibria changes.

The addition of experimental noise to quantum games is known to not affect the existence of a Nash equilibrium, but it tends to lower the payoff at Nash equilibria [40]. This is due to the fact that the quantum games that are chosen for demonstrations are those in which the payoff at Nash equilibrium is larger than the payoff for other strategy choices. This is also the case for our implementation, see figure 3.

To characterize the size of the systematic shift of the expected payoff for theory versus experiment, we plot the root-mean-square deviation (RMSD) of the experimental data to the theoretical calculation as a function of entanglement. The results are plotted in figure 4. Because not all equilibria exist for all values of entanglement, the data are analyzed for $p = 0$, instead of because at $p = 0$ there is Nash equilibrium that is constant for all values of entanglement, and has the theoretical payoffs for players (A, B_1, B_2) equal to $(11, 10, 9)$ with the strategy



choices of the three players given by $\{I, X, I\}$ as well as $\{Z, Y, Z\}$. This payoff could be verified by plugging the strategy choices for A and B_2 into figure 1, and computing the payoff from the resulting state and the payoff matrices. The trend is that the deviation from the theoretically calculated payoff grows as entanglement, and hence experimental error, increases. Notably, for large values of entanglement, the shift is larger in the case of $p = 0$ than it is for the case of $p = 0.5$ plotted in figure 3. This is because there is a larger deviation in the game of A versus B_2 than there is in the game of A versus B_1 , which stems from the differences in the particular gates applied by the players in the equilibria in different regions of the Bayesian game.

The data for B players are scaled to the probability of playing with A, to retain the same scale as the payoff matrix. The data for B_1 at $p = 0$, where he does not actually play the game with A, are calculated as the asymptotic value as p approaches zero from an infinitesimally small non-zero probability. This was done in order to be able to analyze the data, whereas in an actual Bayesian game, the payoff of B_1 in this case is nothing, which is determined by the definition of the Bayesian game.

In addition, we can see from the data how the discontinuous behavior of the Nash equilibria change with experimental noise. Our step size of $\chi = \pi/40$ does not permit a detailed study of the threshold for discontinuity as the entanglement is varied. But, because we incorporate the probability in the analysis, we can see the change in the threshold probability with fine steps allowing a systematic study of the threshold probability at which the equilibria change. In the top part of figure 5 we plot the data for one example value of entanglement, $\pi/20$, as a function of p both theoretically and experimentally as a black line and red triangles, respectively.

First, there are two discontinuities, one near $p = 0.16$ and the other near $p = 0.55$. It can be seen that the critical value of p for the discontinuity is different from theory and experiment. If we increase the value of the best response thresholding parameter δ , the experimental Nash equilibria will extend further in either direction, also, additional equilibria may arise.

In the region of the equilibrium between $p = 0.6$ and $p = 1$, another feature of note occurs. The main equilibrium in this region is given by $\{X, Y, Z\}$. Though for p between 0.81 and 1, a second Nash equilibrium appears. This equilibrium is given by $\{Y, X, I\}$ and has a slightly lower payoff. The presence of this equilibrium is due to the finite δ parameter in the analysis of the best response. If δ were smaller, the region of the second equilibria would shrink, but so would the region with the real equilibrium. For this, and other similar observed secondary equilibria, the transition is usually blurred, meaning the second equilibrium can appear and vanish several times with increasing p before it is reliably present.

We analyze the data for each value of entanglement for all values of p in order to see how the deviation of the threshold p changes as the entanglement grows. For consistency, we analyze the location of the discontinuity near $p = 0.16$. Theoretically this discontinuity occurs at $p = 0.16$ for all values of entanglement. The deviation of the theoretical to experimental threshold p also grows slightly with the entanglement as seen in the bottom of figure 5. For the values of $\chi > .175\pi$, another equilibrium appeared in the analysis (not shown in any graph) with the strategy choices of $\{X, Z, Y\}$. This equilibrium closely resembles the main one in this region with strategy choices $\{I, X, I\}$ and $\{Z, Y, Z\}$ with nearly the same payoffs. However, this equilibrium has a threshold p which is slightly larger, $p \sim 0.29$. As χ increases, the payoffs, and threshold p of this equilibrium converge to the main equilibrium.




Figure 5. Top: example of an analysis with varying p for $\chi = \pi/20$. There is a region in the center where the equilibrium disappears, and a region for high p where a second equilibrium appears. Bottom: the deviation of the experimentally determined threshold p as a function of entanglement.

5. Discussion

Deviations in the payoff for the experimentally determined equilibria are expected due to the finite fidelity of the circuit, figure 2. If there is an error in one of the player qubits, the experimentally determined payoff will be incorrect, while if there is an error in one of the auxiliary qubits, the outputs for the player qubits in that run will be misidentified. In either case this results in a different output than that predicted by the circuit. Because the game implemented is one where the payoffs at the equilibria are typically some of the higher payoffs in the game, deviation from the theoretical behavior will typically result in a lower payoff. If a game were chosen where the equilibrium was not the highest payoff in the game, i.e., if it were not Pareto optimal, then deviation from the theoretical payoff at equilibrium would tend to increase the payoff, which we see if we change the payoff matrices for the analysis in our game. Gate errors tend to be larger for larger values of entanglement because they depend on the coherence of relative phases of multiple qubits, which suffers from experimental error such as differential Stark shifts, whereas unentangled states only rely on their individual coherences. Thus, the deviation from theoretical behavior increases with χ .

When the experimentally determined payoff deviates significantly from the theoretical value, this can result in the disappearance or appearance of Nash equilibria. This is seen both at the boundaries between regions with different equilibria and in the appearance of new equilibria as in the bottom and top of figure 5. The value of $\chi = \pi/20$ was simply chosen as a representative example, the results for other values of entanglement are similar. These transitions can be blurred, so that they are not as precisely defined as they are theoretically. This could have impact in the applications of quantum games such as mechanism design where the game is structured in order to steer the players towards certain Nash equilibria, so that the play is self-reinforcing and thus stable. If the experimental errors remove some expected equilibrium, the players would not converge their play on the strategies of the equilibrium as expected and the game could become unstable.

When an alternate competing equilibrium appears, such as in the top of figure 5, the players could become stuck on the ‘wrong’ equilibrium. If the game dynamically changes, for instance, if p changes, the players may continue to follow the lower payoff equilibrium.

We have performed what we believe is a unique demonstration of a quantum game for several reasons. It is the first experimental demonstration performed on an ion-trap-based quantum computer, and we employ a novel parallelization scheme. The sophistication of the ion-trap quantum computer has enabled a much more

extensive demonstration that has allowed us to generate enough data to demonstrate a more complicated Bayesian game than the more commonly demonstrated two player games. This is also the first time the experimental data have been solely used to solve for the Nash equilibria of the game, which allowed us to experimentally observe the behavior of the payoff as well as the equilibria discontinuities as the amount of entanglement changes. We are also the first to show explicitly how the deviation of the theory to experiment varies as a function of the entanglement.

The value of entanglement is interpreted as being set by a referee of the game. In the case of four allowed strategy choices, maximal entanglement is not always desired from the point of view of an optimal equilibrium. In fact, the behavior of the game becomes less predictable for larger entanglement because of the increased gate errors. These considerations would have to come into any design of a game or choices by a referee in order to promote the desired behavior.

All the ingredients to move trapped ion quantum computing devices into the noisy intermediate-scale quantum era [41] have been demonstrated. Stable chains of over 100 ions [42], and the technology to manipulate them exists. The best entangling gates achieved to date were realized with trapped ions, and their errors are one order of magnitude lower [43] than the gates used in this work. The combination of these capabilities will produce medium size systems able to surpass classical computers in the next few years. This means quantum games which can only be run on a quantum computer will be implemented and perhaps used to test quantum computers and communication protocols. Even with these low error rates, the system size and number of operations needed will require error mitigation techniques such as [44] to achieve useful performance before error corrected operations eventually come into reach.

Advances in ion-trap quantum computers underscore the reason that they are a promising platform for quantum games in particular. In addition, ion-trap-based quantum computers are a promising candidate for quantum networking, which is crucial for some quantum game applications that require the agents to be remotely located in order to be useful. The potentially long coherence time of trapped ions would also enable the quantum hardware to interact with other systems, such as humans, classical computers, sensors, etc, as may be required for quantum game applications.

Quantum games may play an important role in the applications of quantum computers and quantum networks as they begin to become more available. We believe this demonstration brings us one step closer to that reality.

Acknowledgments

This work was supported by the IARPA LogiQ program, the ARO MURI on Modular Quantum Circuits, the AFOSR MURI on Optimal Quantum Measurements, the ARL Center for Distributed Quantum Information, and the NSF Physics Frontier Center at JQI.

ORCID iDs

Neal Solmeyer  <https://orcid.org/0000-0002-0366-7842>

Radhakrishnan Balu  <https://orcid.org/0000-0003-1494-5681>

George Siopsis  <https://orcid.org/0000-0002-1466-2772>

References

- [1] von Neumann J and Morgenstern O 1944 *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press)
- [2] Nash J 1950 Equilibrium points in n-person games *Proc. Natl Acad. Sci.* **36** 48
- [3] Nash J 1951 Non-cooperative games *Ann. Math.* **54** 286–95
- [4] Shubik M 1981 Game theory models and methods in political economy *Handbook of Mathematical Economics* ed K J Arrow and M D Intriligator vol 1 (New York: Elsevier) 285–330
- [5] Levy G and Razin R 2004 It takes two: an explanation for the democratic peace *J. Eur. Econ. Assoc.* **5** 1–29
- [6] Axelrod R M and Dion D 1988 The further evolution of cooperation *Science* **242** 1385
- [7] Shoham Y 2008 Computer science and game theory *Commun. ACM* **51** 74
- [8] Meyer D 1999 Quantum strategies *Phys. Rev. Lett.* **82** 1052–5
- [9] Eisert J, Wilkens M and Lewenstein M 1999 Quantum games and quantum strategies *Phys. Rev. Lett.* **83** 3077–80
- [10] Maitra A et al 2015 Proposal for quantum rational secret sharing *Phys. Rev. A* **92** 022305
- [11] Li Q, He Y and Jiang J-P 2009 A novel clustering algorithm based on quantum games *J. Phys. A: Math. Theor.* **42** 445303
- [12] Zableta O G, Barrangú J P and Arizmendi C M 2017 Quantum game application to spectrum scarcity problems *Physica A* **466** 455–61
- [13] Du J, Li H, Xu X, Shi M, Wu J, Zhou X and Han R 2002 Experimental realization of quantum games on a quantum computer *Phys. Rev. Lett.* **88** 137902
- [14] Mitra A, Sivapriya K and Kumar A 2007 Experimental implementation of a three qubit quantum game with corrupt source using nuclear magnetic resonance quantum information processor *J. Magn. Reson.* **187** 306–13

- [15] Prevedel R, Andre S, Walther P and Zeilinger A 2007 Experimental realization of a quantum game on a one-way quantum computer *New J. Phys.* **9** 205
- [16] Zhang P, Zhang Y-S, Huang Y-F, Peng L, Li C-F and Guo G-C 2008 Optical realization of quantum gambling machine *Europhys. Lett.* **82** 30002
- [17] Altepeter J, Hall M, Medic M, Patel M, Meyer D and Kumar P 2009 Experimental realization of a multi-player quantum game *Advances in Optical Sciences Congress* (Washington, DC: Optical Society of America) [PDNTuA2](#)
- [18] Schmid C, Flitney A, Wieczorek W, Kiesel N, Weinfurter H and Hollenberg L 2010 Experiental implementation of a four-player quantum game *New J. Phys.* **12** 063031
- [19] Zu C, Wang Y-X, Chang X-Y, Wei Z-H, Zhang S-Y and Duan L-M 2012 Experimental demonstration of quantum gain in a zero-sum game *New J. Phys.* **14** 033002
- [20] Pinheiro A R C, Souza C, Caetano D, Juguenin J, Schmidt A and Khoury A 2013 Vector vortex implementaion of a quantum game *J. Opt. Soc. Am. B* **30** 3210–4
- [21] Balthazar W, Passos M, Schmidt A, Caetano D and Huguenin J 2015 Experiemntal realization of the quantum duel game using linear optical circuits *J. Phys. B: At. Mol. Opt. Phys.* **48** 165505
- [22] Chen K-Y and Hogg T 2006 How well do people play a quantum prisoner’s dilemma? *Quantum Inf. Process.* **5** 43–67
- [23] Chen K-Y and Hogg T 2008 Experiments with probabilistic quantum auctions *Quantum Inf. Process.* **7** 139–52
- [24] Blinov B B, Moehring D L, Duan L-M and Monroe C 2004 Observation of entanglement between a single trapped atom and a single photon *Nature* **428** 153–7
- [25] Harsanyi J C 1967 Games with incomplete information played by Bayesian players *Manage. Sci.* **14** 159
- [26] Brunner N and Linden N 2013 Connection between Bell nonlocality and Bayesian game theory *Nat. Commun.* **4** 2057
- [27] Solmeyer N, Dixon R and Balu R 2017 Characterizing the Nash equilibria of a three-player Bayesian quantum game *Quantum Inf. Process.* **16** 146
- [28] Du J, Li H, Xu X, Zhou X and Han R 2003 Phase-transition-like behaviour of quantum games *J. Phys. A: Math. Gen.* **36** 6551–62
- [29] Solmeyer N and Balu R 2017 Characterizing the Nash equilibria of three-player Bayesian quantum games *Proc. SPIE* **10212** 102120T
- [30] Landsburg E S 2011 Nash equilibria in quantum games *Proc. Am. Math. Soc.* **139** 4423
- [31] Olmschenk S, Younge K C, Moehring D L, Matsukevich D N, Maunz P and Monroe C 2007 Manipulation and detection of a trapped Yb^+ hyperfine qubit *Phys. Rev. A* **76** 052314
- [32] Debnath S, Linke N M, Figgatt C, Landsman K A, Wright K and Monroe C 2016 Demonstration of a small programmable quantum computer module using atomic qubits *Nature* **536** 63–6
- [33] Mølmer K and Sørensen A 1999 Multiparticle entanglement of hot trapped ions *Phys. Rev. Lett.* **82** 1835–8
- [34] Milburn G J, Schneider S and James D F V 2000 Ion trap quantum computing with warm ions *Fortschr. Phys.* **48** 801–10
- [35] Solano E, de Matos Filho R L and Zagury N 1999 Deterministic Bell states and measurement of the motional state of two trapped ions *Phys. Rev. A* **59** R2539–2543
- [36] Zhu S-L, Monroe C and Duan L-M 2006 Trapped ion quantum computation with transverse phonon modes *Phys. Rev. Lett.* **97** 050505
- [37] Choi T, Debnath S, Manning T A, Figgatt C, Gong Z-X, Duan L-M and Monroe C 2014 Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement *Phys. Rev. Lett.* **112** 190502
- [38] Linke N M, Maslov D, Roetteler M, Debnath S, Figgatt C, Landsman K A, Wright K and Monroe C 2016 Experimental comparison of two quantum computing architectures *Proc. Natl Acad. Sci.* **114** 3305–10
- [39] Maslov D 2017 Basic circuit compilation techniques for an ion-trap quantum machine *New J. Phys.* **19** 023035
- [40] Buluta I M, Fujiwara S and Hasegawa S 2006 Quantum games in ion traps *Phys. Lett. A* **358** 100
- [41] Preskill J 2018 Quantum Computing in the NISQ era and beyond arXiv:1801.00862v2
- [42] Pagano G et al 2018 *Cryogenic Trapped-Ion System for Large Scale Quantum Simulation* arXiv:1802.03118
- [43] Ballance C J, Harty T P, Linke N M, Sepiol M A and Lucas D M 2016 High-fidelity quantum logic gates using trapped-ion hyperfine qubits *Phys. Rev. Lett.* **117** 060504
- [44] Li Y and Benjamin S C 2017 Efficient variational quantum simulator incorporating active error minimization *Phys. Rev. X* **7** 021050