

Behavior of Multi-Agent Protocols Using Quantum Entanglement

Kay-Yut Chen and Tad Hogg and Bernardo A. Huberman

HP Labs
Palo Alto, CA

Abstract

We describe how entangled quantum states can aid in coordination, cooperation and resource allocation in multi-agent systems. These protocols provide alternatives to conventional methods, with different trade-offs of capabilities and information privacy. We also present results of human-subject experiments with simulated versions of some of these methods, showing people can learn to use entangled states effectively without training in quantum mechanics. Thus these quantum protocols are suitable for mixed systems consisting of human and software agents. These techniques are beneficial with even a few bits and operations, making their physical implementation much easier than quantum applications to hard computational problems such as factoring or search.

Introduction

Proposed applications of quantum information processing (Nielsen & Chuang 2000) mainly focus on hard computational problems, such as factoring and search, and cryptographic key exchange. However, the ability to distribute and manipulate entangled particles also offers a new resource to aid multiparty strategic decision-making, giving rise to mechanisms based on quantum games (Meyer 1999; Eisert, Wilkens, & Lewenstein 1999; Eisert & Wilkens 2000; Du & others 2002a; 2002b; Benjamin & Hayden 2001). By altering the strategic incentives, such games offer novel mechanisms for multi-agent systems, consisting of software agents, robots, people or combinations of these. For example, the alteration of the quantum state upon observation can prevent agents from learning details of the state created by other agents. This information hiding can simplify the design of agents since there is no need to consider strategic consequences of information that is used by never available to other agents. More generally, quantum mechanisms change the range of choices within a game, giving the possibility of more favorable strategic outcomes than the corresponding classical mechanism.

Quantum mechanisms for multi-agent systems could be useful with even a modest number of quantum bits, and hence aid in developing commercial applications of early

quantum information technology (Spiller & Munro 2006). The mechanisms rely on distributed access to a source of entangled particles (Lloyd, Shahriar, & Hemmer 2004), analogous to the current reliance on utilities providing telecommunications. The ability to exchange entangled states over distances of tens of kilometers and perform coherent operations on a few states has already been demonstrated (Stucki & others 2002; Vandersypen *et al.* 2000; Steffen *et al.* 2003). Thus the technological basis for implementing economic mechanisms is likely to arise well before applications, such as factoring large numbers, requiring long coherence times on many bits.

In this paper, we describe using entangled states for three key interactions among self-interested agents: coordination, cooperation and resource allocation. Specifically, we discuss quantum protocols for 1) helping agents coordinate behavior in the presence of an adversary (Huberman & Hogg 2003; Mura 2003), 2) inducing cooperation in a public-goods scenario, in which each agent is tempted to free-ride on the efforts of others, without relying on reputation or a trusted third party enforcer (Chen, Hogg, & Beausoleil 2002; Zhang & Hogg 2003), and 3) allocating resources with allocative externalities via auctions that maintain privacy of agent preferences (Harsha, Hogg, & Chen 2006 in preparation). These three protocols have increasing requirements for physical implementation: first, distributing, storing and measuring entangled states; second, performing a single operation on such states; and, third, performing a series of operations with repeated communication of entangled states among the agents. Before describing these protocols, we discuss evaluating their behavior, including their use by people as part of a multi-agent system.

Evaluating Quantum Protocols

Determining how realistic agents may use quantum mechanisms is important before investing considerable effort in physically implementing the mechanisms. Moreover, understanding actual behavior could help improve mechanism design by indicating how to structure a quantum mechanism to facilitate learning and performance. We consider both theoretical and empirical approaches to this question.

Game theory provides a theoretical framework for evaluating the behavior of idealized rational agents facing strategic choices. This theory applies to protocols involving quan-

tum entanglement, though generally involving a larger range of choices than in conventional games.

For example, a two-state quantum system (a *qubit*) can be placed in arbitrary superpositions of those two physical states. Such superpositions are linear combinations of the two states, denoted as $a|0\rangle + b|1\rangle$ where a and b are complex numbers and $|0\rangle$ and $|1\rangle$ denote the two states (e.g., horizontal and vertical polarization of a photon, or spin up and spin down of an electron). Upon measurement of such a superposition, one observes either 0 or 1 with probabilities $|a|^2$ and $|b|^2$, respectively. When manipulating such superpositions, agents can choose to apply any single-qubit operator, whose general form is defined by three angles θ, ϕ, α as the matrix

$$U(\theta, \phi, \alpha) = \begin{pmatrix} e^{-i\phi} \cos \frac{\theta}{2} & e^{i\alpha} \sin \frac{\theta}{2} \\ -e^{-i\alpha} \sin \frac{\theta}{2} & e^{i\phi} \cos \frac{\theta}{2} \end{pmatrix} \quad (1)$$

up to an irrelevant overall phase factor. Multiplying a superposition, viewed as a 2-element vector, by this matrix produces another superposition. By contrast, in a classical game with a single bit, the only choices are the identity and inversion operations.

Game theory usually expresses the strategic consequences of a set of choices and specified payoffs for the possible outcomes in terms of equilibrium concepts (Eisert & Wilkens 2000), with the presumption that players' choices will correspond to an equilibrium. In particular, a Nash equilibrium is a set of choices for the players for which no single player can increase their payoff by making a different choice. A game may have zero, one or many such equilibria. A more complicated situation involves mixed-strategy equilibria, in which players select among various choices randomly according to a particular probability distribution with no gain possible for a single player choosing differently.

Unfortunately, agent behavior does not always correspond to the predictions of game theory, even when there is a unique equilibrium. This is true for both human and software agents, due to various computational and behavioral limitations on full rationality assumed by game theory. In the case of humans, game theory often assumes unrealistic levels of player rationality (Palfrey & Rosenthal 1991; Camerer, Ho, & Chong 2004). Furthermore, people are assumed to understand fully the consequences of manipulating quantum states. In fact, people have particular difficulty achieving rational play predictions for games involving multiple mixed-strategy Nash equilibria and probabilistic outcomes (Camerer 1995). Such properties are often found in quantum games. Similarly, full rationality for software agents may require they solve NP-hard combinatorial problems and simulate other agents doing the same, so their design may instead involve approximate heuristic methods to determine their choices, leading to systems whose behavior does not match predictions of game theory.

Thus the suitability of quantum mechanisms depends not only on how they may function in theory but also on how people would use them or program software agents to use them. Because people have limited intuition for the behavior of quantum entanglement and little experience with games involving such states, people may face significant challenges in choosing appropriate actions in such games or designing

agents to do so on their behalf. As a first step to address the question of actual behavior with quantum protocols, we studied human behavior in a laboratory setting.

Instead of waiting for development of a physical implementation of distribution and storage of entangled states and quantum operations on those states, we simulated all the quantum components since people 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 (van Enk & Pike 2002) 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 games we consider involve only a modest number of qubits, so the exponential increase in time and memory required for simulating quantum operations 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 possible outcome, not just the single observed outcome.

We implemented both classical and quantum versions of the games using the HP Experimental Economics software platform (Chen & Wu 2003). These economic experiments, in which people were paid according to how well they performed in the game, each consisted of a series of periods played over a few hours. In each period, players were randomly assigned into smaller groups, e.g., groups of three people, and the people in each group played a game independently of the other groups. Players received experimental dollars based on the payoffs from the choices they and others in their group made in each period of the experiment. The exchange rate of experimental dollars into real dollars (paid at the end of the experiment) was announced in advance. In the quantum games, players specify a choice of quantum operator defined in Eq. (1). To do so they specify the angles corresponding to their choice. In the experiments with the auction mechanism described below, only the first column of the operator is relevant for the outcomes, in which case players specify only two angles: θ and ϕ . Before the experiments, instructions and a required quiz were available. The experimental sessions started with a review of these instructions and a brief practice session with the game software.

Coordination

Coordination challenges for multi-agent systems arise when agents must make correlated choices to achieve high payoff. These situations may involve sharing a common resource that only one agent at a time can access, or agents helping each other perform a task beyond the ability of a single agent (Lerman & others 2001). From a game-theory perspective, a coordination problem (Shelling 1960; Fudenberg & Tirole 2000; Camerer 2003) involves a payoff matrix which yields several Nash equilibria. These equilibria may give the same payoff to all players, in which case

the problem is for them to agree on which one to use, or different payoffs, leading to a competitive coordination game in which players prefer different equilibria. More generally, games with multiple equilibria may require players to coordinate their choices to achieve a particular equilibrium (Aumann 1974). Without such coordination the players can spend inordinate amounts of time trying to settle on an equilibrium, with consequent loss of the opportunity for high-payoff from coordinated choices.

A simple example of a cooperative coordination game is two people choosing whether to drive on the left or the right side of the road. The payoff matrix for this game involves a benefit if players make the same choices and a large penalty for opposite choices. Thus this game has two Nash equilibria, with equal payoffs, corresponding to both drivers choosing the same side of the road. The coordination problem consists in both drivers finding a way to agree on which side of the road to drive.

Coordination becomes more difficult in competitive contexts, e.g., with an adversary whose payoff depends on preventing other agents from agreeing on their choices. Interesting competitive coordination games arise, for example, in the case of two players trying to coordinate on a mixed strategy against a third one without resorting to previous agreement or communication (due to concerns of the adversary learning the coordinated choice in time to act to reduce its payoff). For example, consider two military allies, on opposite sides of a field, who want to get a target held by an adversary on either the left or right sides of the field. The first of the allies can create a distraction. The other has the personnel and equipment needed to find the target provided they can do so undisturbed. The allies need to decide whether to send their forces to the left or right sides of the field. For any chance of success, the distraction from ally 1 must be on the opposite side of the field from where ally 2 goes. The resulting payoffs express the fact that ally 2 must choose the same place as where the adversary hid the target (which will happen 1/2 the time in case the adversary uses a mixed strategy of each choice randomly selected with equal probability), but also be sure to do the opposite of ally 1. The allies wish to avoid communication and want to delay their decision as long as possible since the adversary could move the target if he knows which choice the allies make well in advance of their action, for example by spies or eavesdropping. This game has a mixed strategy equilibrium and the requirement for anti-correlation among the allies. Specifically, the adversary could select each choice with probability 1/2. If the allies similarly make random choices without coordination, their expected payoff is only 1/4. If instead they always coordinate their choices, i.e., never have both picking left or both picking right, their expected payoff increases to 1/2.

These and many other instances of coordination problems can be solved in several ways. One solution resorts to a third party who knows the preferences of the participants and has the authority to pick an equilibrium which is then broadcast to the players. In the case of competitive coordination problems the third party, e.g., a government, may also have enforcement powers, since some players may wish to move the group to another equilibrium with higher personal utility.

Another solution to coordination problems involves communication among players so that they can negotiate a choice. In the case of cooperative games even one player flipping a coin and broadcasting the result as the corresponding choice provides an effective solution. In a competitive setting, the negotiation is more complicated as different players may prefer different equilibria. A third mechanism for solving coordination problems invokes social norms, in which common knowledge of the participants' preferences can distinguish one equilibrium from the others, as in the case of choosing the largest river as a boundary between two countries. Such distinguished equilibria are often called focal points (Shelling 1960; Huyck, Battalio, & Rankin 1997).

While these mechanisms can solve coordination problems, there are times when none of these options are available, either because they are too expensive, slow or difficult to implement, or because privacy concerns about revealing preferences to third parties. Furthermore, a constraint from a larger context, such as the need to use a mixed strategy, might make it disadvantageous for players to have their choices revealed in advance. For instance, relying on communication could fail if jammed by an adversary, or used to identify the fact that players are communicating from specific locations – information that in itself could be damaging to reveal. In particular, if one player awaits an acknowledgement from the other, uncertainty over whether lack of response is due to jamming the response or lack of reception of the original communication could prevent establishing common knowledge needed for coordination.

The requirement for random but correlated choices without communication at the time the choices are made could be achieved by flipping a coin in advance, setting another coin to match the first one, hiding each coin in a separate box, and giving one box to each player. Opening a box at some later time gives an outcome correlated with that of the other player. A corresponding algorithmic approach would be prior agreement on a secret seed for use with a pseudorandom number generator. This amounts to prior agreement on the *method* for determining a choice, rather than the choice itself. Still, an adversary could learn this method of how the players will choose long before they actually do, use it to determine the choice to be made, and thus adjust its strategy accordingly.

If these considerations lead the players to prefer not using these conventional techniques, it would appear that the only choice left for the participants is to choose at random which strategy to pursue, which would lead to many instances of coordination failures and a consequent reduction in their respective payoffs, especially with a large number of agents. Nevertheless, a protocol using quantum entanglement provides an alternative solution to coordination problems without communication, trusted third parties or pre-arranged choices. It thus provides an additional option for addressing coordination problems, with a different set of strengths and limitations from those of the conventional approaches.

The quantum protocol involves creating entangled states among particles which are held separately by the participants (Huberman & Hogg 2003). Players can measure their

particles independently, to achieve a random choice correlated with all other players. The randomness avoids the risk of revealing a prearranged choice and the correlation ensures coordination without communication.

In the simplest case, where players face two choices, A , and B , they can use entangled particles with two physically observable states, such as their spin or polarization. These states are, by prior agreement, mapped to the two choices (e.g., spin up corresponds to choice A). The particles are created in the entangled state $(|AA\rangle + |BB\rangle)/\sqrt{2}$, and each agent takes physical possession of one of the particles. At a time of their choosing, each participant observes their particle, resulting in either A or B as the outcome, and makes the corresponding choice. The key aspect that makes this technique different from random choices is that entanglement implies a definite correlation between the two measurements, i.e. both players get either outcome A or B , irrespective of the spatial separation between them, and without communication. Entanglement thus allows players to get correlated random bits, without communication or prior agreement on the specific choices they will make. Consequently, they can use these bits to coordinate their choices.

Inducing Cooperation

Self-interested agents can face social dilemmas in which the group as a whole benefits if agents incur some individual costs, but each agent is individually better off by free-riding on the efforts of others in the group (Hardin 1968). A prototypical example of the challenge for inducing cooperation is a generalization of the Prisoner's Dilemma to n players. Specifically, we consider a contribution game in which each of the n players is given an initial wealth W and must decide whether to keep it or contribute all of it to the group. The "public good" each player receives is equal to the fraction a/n of total contributions from the group, where a , between 1 and n , is a multiplier giving the increased value of the public good to the group as a whole. Thus if c_i denotes the amount, 0 or W , player i contributes, the payoff to player i is $W - c_i + \frac{a}{n} \sum_{j=1}^n c_j$. Since $a < n$, each player obtains a higher payoff by not contributing no matter what choice other players make. If all players make this dominant choice, each receives a payoff of W . On the other hand, if they all contribute, each payoff would be the larger value aW . Thus, the dominant strategy equilibrium is not efficient: 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. While this simple game ignores real world issues such as heterogeneous preferences and the ability to contribute a fraction of one's wealth, it captures the essence of the social dilemma posed by public goods economics. The two-player version of this game is the well-known Prisoner's Dilemma.

We compare behavior of the conventional ("classical") and quantum versions of the game.

Quantum Protocol

The quantum version of the contribution game only requires manipulation of pairwise entangled quantum states, which

are far easier to implement than those using higher-order entangled states. Specifically, in the game we study (Chen, Hogg, & Beausoleil 2002), 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. 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 (Zhang & Hogg 2003). A player's choice in the game consists in applying a single one-bit quantum operator, given by Eq. (1), to each of their $n-1$ qubits. The players then return the qubits to the source which applies the inverse of the original entanglement operation to each pair and measures the bits. 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. This quantum game has no dominant strategy, and multiple mixed-strategy Nash equilibria, each 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 quantum game generalizes the classical one in that a restricted set of choices for the quantum operators exactly reproduce the choices and payoffs of the classical public goods game.

Game theory predicts contribution rates of the quantum version will be closer to the economically efficient level than those of the classical game (Chen, Hogg, & Beausoleil 2002). Moreover, the economic efficiency of the quantum game is predicted to improve with group size. This predicted improvement 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 (Glance & Huberman 1994).

Human Behavior

We examined human performance with a simulated version of the public goods game (Chen & Hogg 2006). Each experiment consisted of a series of periods. Players were randomized into groups of size n at the beginning of each period. This randomization reduced reputation or repeated game effects (Axelrod 1987) by increasing player anonymity to exacerbate the free-rider problem, thereby providing a challenging test case for an economic mechanism.

In the quantum game, each player picks three angles specifying a one-bit quantum operator given by Eq. (1). We provided two what-if scenario tools showing the consequences of players' choices in the mini-game. The first tool allows a subject to specify his own choices and a guess of the opponent's choice. The tool then shows the probability of the four possible outcomes of a mini-game (i.e., measured bit values 00, 01, 10, or 11). The second tool takes a guess of the opponent's choices and shows corresponding choices the player

could use to produce each of the four possible outcomes with probability one. To facilitate co-ordination between players, in each period we allowed subjects two rounds of communication. This communication tool allowed each player to broadcast a set of three numbers to other members of the group. Players could use this communication as they wished, with no suggestions provided in the instructions. After the communication rounds, players entered their actual choices. Participants then learned the outcome but not the actual choices made by their opponents.

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 in relatively small groups whereas larger groups tend to be more anonymous and each person's contribution is a smaller fraction of the whole, leading to greater likelihood of free riding. 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.

We ran six experiments (Chen & Hogg 2006), with group sizes of 2 to 4 and 6 to 16 total participants. In all but the first experiment, the participants played both the classical and quantum versions of the game during a single afternoon.

In all experiments, the levels of contributions in the quantum games exceeded those of the corresponding classical games, as predicted (Chen, Hogg, & Beausoleil 2002). Fig. 1 shows the effect of group size on contributions.

The contribution rates for the classical game are nonzero,

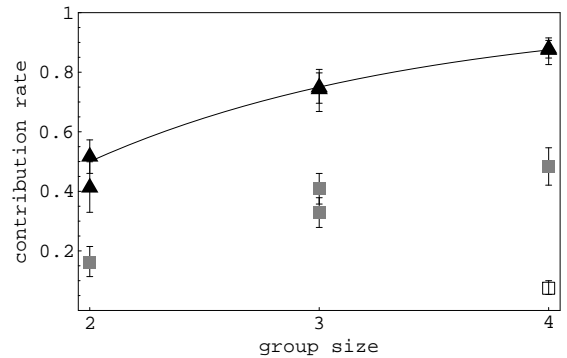


Figure 1: Contribution rates for classical (box) and quantum (black triangle) games as a function of group size n . Error bars show the 95% confidence intervals for the contribution rates based on the sample size. The curve shows the contribution rates for the quantum game predicted by game theory, $1 - 2^{-(n-1)}$. The open box is classical contribution rate for experiment with $n = 4$ and $a = 1.5$. The other experiments used $a = 0.75n$. The two experiments with each of $n = 3$ and 4 have indistinguishable contribution rates for the quantum games.

indicating a degree of altruistic behavior or influence of a reputation effect in spite of the random group selection for each period. On the other hand, the contribution rates for the randomly-matched quantum experiments are consistent with the game theory prediction of $1 - 2^{-(n-1)}$. i.e., the aggregate behavior is consistent with that predicted for the mixed-strategy Nash equilibria.

These experiments illustrate people are able to understand their choices to play well in the quantum public-goods protocol. Thus this protocol is particularly well-suited for use in multi-agent systems containing both human and software agents. Thus its is relatively easy to design software agents to effectively use the quantum protocol.

Resource Allocation via Auctions

The coordination and cooperation scenarios described above involve games where all relevant information on payoffs is known to all participants. In contrast, many multi-agent interactions involve scenarios with incomplete information. For instance, many economic games involve private information, e.g., knowledge about valuations, business cost structures and technical capabilities. Similarly, agents representing physical robots have limited access to environmental information, and the available information can vary among agents. These situations give rise to incomplete information games as participants are not aware of some information others use to determine their payoffs. Moreover, participants with private information may not want that information disclosed to others, fearing loss of competitive advantage in subsequent interactions. Nevertheless, they could find it advantageous to participate in mechanisms allowing them to jointly exploit their information.

Resource allocation problems addressed by auctions pro-

vide one such economic scenario. In this case, the private information arises from the valuations of the bidders.

Quantum Protocol

Quantum information processing may be useful in auctions with privacy concerns because when bids are encoded in quantum states, auction processes can be designed to reveal nothing but the winning bid and allocation. Cryptographic protocols (Naor, Pinkas, & Sumner 1999) can also have this property, but with differing assumptions about collusion and the security. In particular, with quantum auctions, the intermediate states containing the bids is irrevocably destroyed by the measurement at the end of the protocol. Thus, unlike cryptographic protocols where cryptographic keys may accidentally or intentionally be revealed after the auction, the quantum protocol ensures this is not possible. In addition, the quantum protocol can compactly express complex interdependencies among items and bidders via entanglement in the bidding superpositions.

In our auction protocol (Harsha, Hogg, & Chen 2006 in preparation), each bidder selects an operator that produces the desired bid from a prespecified initial state (e.g., all bits set to zero). This protocol allows general superpositions to represent bids for multiple items, or bundles of items (as in combinatorial auctions). Moreover, two or more bidders can entangle their superpositions to specify correlations in outcomes, e.g., two bidders want either they both win their respective items or neither wins. Such entanglement allows expressing allocative externalities within the context of a single auction protocol. In the simplest case, the bidders would need only the single-qubit operator of Eq. (1) to create their superposition two states: one representing not winning the auction and the other representing the amount the bidder wishes to bid for a single item.

To find the winning bidder(s), the auctioneer repeatedly asks the bidders to apply their individual operators in a distributed implementation of a quantum search algorithm. The distributed nature of the search ensures neither the auctioneer nor the other bidders gain any information about losing bids when following the protocol. We use a distributed version of the adiabatic search method (Farhi & others 2001) for finding optimal values.

The use of quantum superpositions and search introduces a broad range of incentive issues beyond those examined in prior quantum games. For instance, quantum search methods are probabilistic in nature so are not guaranteed to always return the highest bid. In this respect, they are similar to using approximate or heuristic methods in combinatorial auctions to avoid the general NP-hard problem of identifying winners. However, in the quantum case, the outcome is probabilistic even in the simple case of finding the highest bid for a single-item auction. Provided bidders create superpositions with uniform amplitudes, the probability of finding the correct answer can be made as high as one wishes by running the search with sufficiently many steps, the small residue probability of awarding the item not to the highest bidder may change bidding behavior. In addition to the probabilistic selection of a winner, there is also a chance that no one wins, analogous to auctions with a reservation price in

which no bidder wins if their bids do not exceed that price.

A second game theory issue is that encoding bids in quantum states extends the strategic choices for the bidders. For instance, a bidder can alter the amplitudes associated with different bids in the superposition. In the standard formulation of adiabatic search (Farhi & others 2001), such nonuniform amplitudes can adversely affect the outcome of the auction by allowing agents with low bids to win with high probability. Fortunately, a simple permutation of values the auctioneer assigns to the states in superposition during the adiabatic search removes the incentive for any single bidder to deviate from uniform amplitudes in their superposition (Harsha, Hogg, & Chen 2006 in preparation). Thus, game theory predicts the permuted search method is more efficient than the standard search method because the permuted ordering eliminates incentives to create superpositions of bids with nonuniform amplitudes. From a quantum algorithm perspective, this construction of the search method in the auction protocol illustrates how incentive issues affect algorithm design, in contrast to the more common concern with computational issues.

Human Behavior

For experiments with auctions, we considered a simple first-price sealed-bid auction in which each bidder submits a secret bid for a single item and the highest bid wins the auction. This type of auction has been well studied both theoretically and experimentally, providing a simple case for studying the use of quantum mechanisms with incomplete information.

The main experimental question is whether an appropriate design of the quantum auction protocol induces bidders to exhibit similar behavior as in classical auction in spite of the ability to manipulate amplitudes in superposition. While game theoretic analysis seems to indicate that may be the case, there are intuitive reasons, such as the fact that people may play non-equilibrium strategies that are more Pareto efficient than the Nash equilibrium, to cast doubt on whether these assertions will be true if real human bidders are involved.

We ran a set of experiments, each consisting of multiple periods, to address this question. In each period, participants were grouped randomly into groups of three, i.e., we considered auctions with three bidders, and a value between 0 and 100 was generated randomly for each person. If a person won the auction, he or she received the difference between his or her value and bid. Participants were not told who the other members of their group were.

In the quantum auction, participants specified their choice of operator (in Eq. (1)) to create their bid superposition through two angles (θ and ϕ – for specifying their initial state, only the first column of the operator matrix is significant, so there is no need to specify the third angle α in this experiment) in addition to their desired bid. This contrasts with the classical auction where participants just specify a bid. The operator choices made by the three participants in a quantum auction determine the probability for each bidder to win – it is not the case that the highest bid always wins. Three treatments were used: a classical first price auction

used as a benchmark, and two different quantum protocols, using the standard and permuted search methods.

We provided a decision support tool to help participants understand consequences of their choices. Specifically, a person could enter values for the choices of the three members of the group and see the resulting probability of each bidder winning, and the probability of having no winner.

For resource allocation in multi-agent systems, an important criterion is allocative efficiency. It measures the extent to which the item(s) are allocated to high value uses, in this case a person with highest value for the item. Specifically, we examined the ratio of the achieved value to the maximum possible value. In the case of auctions, this is simply the ratio of the value of the winning bidder to the highest value amongst the bidders. As expected, the classical auction achieved high allocative efficiency (96%+) in all experiments. The efficiency dropped to about 33% for the quantum auctions, with no significant difference between the standard and permuted search protocols. A significant percentage of the quantum auctions resulted in no winner. For the auctions that did result in a winner, the allocative efficiency was higher for permuted search (81%) than for the standard search (70%).

Furthermore, we observed that the permuted search protocol resulted in higher revenue just as the theory has predicted. These results are strong evidence that subjects, while having little knowledge of quantum physics, learned to use the quantum protocol effectively.

The quantum auction protocol is significantly more complicated than the public-goods protocol. In our experiments, we found numerous cases with no winners, lowering the overall efficiency. Although this might be improved somewhat with further training, it illustrates a trade-off of the quantum protocol as providing more privacy than classical auctions but at a cost of lowered economic efficiency. An open question is how people would perform in a mixed multi-agent system consisting of both people and software agents competing for the resources via the quantum auction protocol. Moreover, it would be interesting to examine extended scenarios in which privacy concerns would arise endogenously, e.g., via a repeated allocation scenario where revelation of losing bids at one time may reveal a bidders cost structure or other private information, which may place those bidders at a disadvantage for subsequent allocation competitions.

Discussion

We presented protocols for using entangled states for three types of interactions arising with self-interested agents, both human and software: coordination in adversarial situations, cooperation in spite of temptation to free-ride on the efforts of others, and resource allocation with allocative externalities and information hiding preferences.

These applications of quantum entanglement to multi-agent interactions contrast with the bulk of research in quantum algorithms for NP-hard search problems (Grover 1997; Farhi & others 2001; van Dam, Mosca, & Vazirani 2001; Hogg 2003) and genetic algorithms (Spector *et al.* 1999).

Our laboratory experiments show people can play the quantum public goods game effectively without specialized training in quantum physics or game theory. Instead, simple what-if scenario tools allow participants to understand consequences of their choices well enough to play effectively. Thus scenario tools are a useful way to present quantum games to people. The results suggest it is possible to achieve the benefits predicted by game theory, as well as the security and trust guarantees provided by physics, if suitable quantum information processing system can be implemented. Our experimental work involved human subjects, and software agents were only used as a debugging tool. An obvious extension to our experimental work is designing software agents capable of participating in these games, and identifying suitable strategies for their decision-making. Analysis of experimental observations of human behavior in quantum games suggests ideas for agent designs. The software system used in the experiments allows software agents to participate in the same games that human players are playing. Thus, it will be easy to switch among experiments with people, software agents and combining both people and software agents for a general multi-agent system.

An interesting extension to this work is to identify aspects of quantum games that give benefits for agents without requiring agents to solve NP-hard problems and also allow effective performance by people. On a more speculative note, with future development of quantum computers able to manipulate many qubits over long periods of time, quantum mechanisms could extend to use by software agents running on such quantum computers. Such “quantum agents” would be able to directly participate in the quantum evolution of the protocols, unlike the situations presented in this paper involving software agents on conventional computers and people where final actions derive from measured outcomes from the quantum states. The direct participation of quantum agents would allow for a broader range of protocols, including decisions based on counterfactual events (Kwiat & others 1999). While it remains to determine the extent to which such protocols would be useful, the development of such agents could greatly extend the scope of quantum interactions within multi-agent systems.

Acknowledgments

We thank Raymond Beausoleil, Saikat Guha, Pavithra Harsha, Phil Kuekes, Li Zhang and Cecilia Zenteno for helpful discussions.

References

- Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1:67–96.
- Axelrod, R. 1987. The evolution of strategies in the iterated prisoner’s dilemma. In Davis, L., ed., *Genetic Algorithms and Simulated Annealing*. Los Altos, CA: Morgan Kaufmann. chapter 3, 32–41.
- Benjamin, S. C., and Hayden, P. M. 2001. Multiplayer quantum games. *Physical Review A* 64:030301.

- Camerer, C.; Ho, T.; and Chong, J.-K. 2004. A cognitive hierarchy model of one-shot games. *Quarterly Journal of Economics* 119(3):861–898.
- Camerer, C. 1995. Individual decision making. In Kagel, J., and Roth, A. E., eds., *The Handbook of Experimental Economics*. Princeton Univ. Press. chapter 8, 587–703.
- Camerer, C. 2003. *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton University Press.
- Chen, K.-Y., and Hogg, T. 2006. How well do people play a quantum prisoner’s dilemma? *Quantum Information Processing* 5:43–67.
- Chen, K.-Y., and Wu, R. 2003. Computer games and experimental economics. In *Proc. of the 5th Intl. Conf. on Enterprise Information Systems (ICEIS)*.
- Chen, K.-Y.; Hogg, T.; and Beausoleil, R. 2002. A quantum treatment of public goods economics. *Quantum Information Processing* 1:449–469. arxiv.org preprint quant-ph/0301013.
- Du, J., et al. 2002a. Entanglement enhanced multiplayer quantum games. *Physics Letters A* 302:229–233. arxiv.org preprint quant-ph/0110122.
- Du, J., et al. 2002b. Experimental realization of quantum games on a quantum computer. *Physical Review Letters* 88:137902. arxiv.org preprint quant-ph/0104087.
- Eisert, J., and Wilkens, M. 2000. Quantum games. *J. Modern Optics* 47:2543–2556. arxiv.org preprint quant-ph/0004076.
- Eisert, J.; Wilkens, M.; and Lewenstein, M. 1999. Quantum games and quantum strategies. *Physical Review Letters* 83:3077–3080. arxiv.org preprint quant-ph/9806088.
- Farhi, E., et al. 2001. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* 292:472–476.
- Fudenberg, D., and Tirole, J. 2000. *Game Theory*. Cambridge, MA: MIT Press.
- Glance, N. S., and Huberman, B. A. 1994. Dynamics of social dilemmas. *Scientific American* 270(3):76–81.
- Grover, L. K. 1997. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters* 79:325–328. arxiv.org preprint quant-ph/9706033.
- Hardin, G. 1968. The tragedy of the commons. *Science* 162:1243.
- Harsha, P.; Hogg, T.; and Chen, K.-Y. 2006 (in preparation). Quantum auctions. Technical report, HP Labs.
- Hogg, T. 2003. Adiabatic quantum computing for random satisfiability problems. *Physical Review A* 67:022314. arxiv.org preprint quant-ph/0206059.
- Huberman, B. A., and Hogg, T. 2003. Quantum solution of coordination problems. *Quantum Information Processing* 2:421–432. arxiv.org preprint quant-ph/0306112.
- Huyck, J. V.; Battalio, R.; and Rankin, F. 1997. On the origin of convention: Evidence from coordination games. *Economic Journal* 107:576–597.
- Kwiat, P. G., et al. 1999. High-efficiency quantum interrogation measurements via the quantum Zeno effect. *Physical Review Letters* 83:4725–4728.
- Lerman, K., et al. 2001. A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life* 7:375–393.
- Lloyd, S.; Shahriar, M. S.; and Hemmer, P. R. 2004. Teleportation and the quantum internet. Technical report, MIT.
- Meyer, D. A. 1999. Quantum strategies. *Physical Review Letters* 82:1052–1055. arxiv.org preprint quant-ph/9804010.
- Mura, P. L. 2003. Correlated equilibria of classical strategic games with quantum signals. arxiv.org preprint quant-ph/0309033.
- Naor, M.; Pinkas, B.; and Sumner, R. 1999. Privacy preserving auctions and mechanism design. In *Proc. of the ACM Conference on Electronic Commerce*, 129–139. NY: ACM Press.
- Nielsen, M. A., and Chuang, I. L. 2000. *Quantum Computation and Quantum Information*. Cambridge Univ. Press.
- Palfrey, T. R., and Rosenthal, H. 1991. Testing game-theoretic models of free riding: New evidence on probability bias and learning. In Palfrey, T. R., ed., *Laboratory Research in Political Economy*. Univ. of Michigan Press. 239–267.
- Shelling, T. C. 1960. *The Strategy of Conflict*. Oxford University Press.
- Spector, L.; Barnum, H.; Bernstein, H. J.; and Swamy, N. 1999. Finding a better-than-classical quantum AND/OR algorithm using genetic programming. In Angeline, P., ed., *Proc. of the 1999 Congress on Evolutionary Computing*. Washington, DC: IEEE.
- Spiller, T. P., and Munro, W. J. 2006. Towards a quantum information technology industry. *Journal of Physics: Condensed Matter* 18:V1–V10.
- Steffen, M.; van Dam, W.; Hogg, T.; Breyta, G.; and Chuang, I. 2003. Experimental implementation of an adiabatic quantum optimization algorithm. *Physical Review Letters* 90:067903.
- Stucki, D., et al. 2002. Quantum key distribution over 67km with a plug and play system. *New J. of Physics* 4:41.
- van Dam, W.; Mosca, M.; and Vazirani, U. 2001. How powerful is adiabatic quantum computation? In *Proc. of the 42nd Annual Symposium on Foundations of Computer Science (FOCS2001)*, 279–287. IEEE.
- van Enk, S. J., and Pike, R. 2002. Classical rules in quantum games. *Physical Review A* 66:024306.
- Vandersypen, L. M. K.; Steffen, M.; Sherwood, M. H.; Yannoni, C. S.; Breyta, G.; and Chuang, I. L. 2000. Implementation of a three-quantum-bit search algorithm. *Applied Physics Letters* 76:646–648.
- Zhang, L., and Hogg, T. 2003. Reduced entanglement for quantum games. *Intl. J. of Quantum Information* 1(3):321–335.