

A Quantum Protocol for Sampling Correlated Equilibria Unconditionally and without a Mediator^{*}

Jordanis Kerenidis¹ and Shengyu Zhang²

¹ Laboratoire d'Informatique Algorithmique: Fondements et Applications,
Univ. Paris Diderot 7, and CNRS, Centre for Quantum Technologies, Singapore
jkeren@liafa.jussieu.fr

² Department of Computer Science and Engineering and The Institute of Theoretical
Computer Science and Communications, The Chinese University of Hong Kong
syzhang@cse.cuhk.edu.hk

Abstract. A correlated equilibrium is a fundamental solution concept in game theory that enjoys many desirable mathematical and algorithmic properties: it can achieve more fair and higher payoffs than a Nash equilibrium and it can be efficiently computed for a vast class of games. However, it requires a trusted mediator to assist the players in sampling their moves, which is a major drawback in many practical applications.

A computational solution to this problem was proposed by Dodis, Halevi and Rabin [DHR00]. They extended the original game by adding a preamble stage, where the players communicate with each other and then they perform the original game. For this extended game, they show that the players can achieve payoffs at least as high as in any correlated equilibrium, provided that the players are computationally bounded and can communicate before the game.

The introduction of cryptography with computational security in game theory is of great interest both from a theoretical and more importantly from a practical point of view. However, the main game-theoretic question remained open: can we achieve any correlated equilibrium for 2-player games without a trusted mediator and also unconditionally?

In this paper, we provide a positive answer to this question. We show that if the players can communicate via a *quantum* channel before the game, then for 2-player games, payoffs at least as high as in any correlated equilibrium can be achieved, without a trusted mediator and unconditionally. This provides another example of a major advantage of

^{*} Most of the work was done when the authors visited Centre of Quantum Technologies (CQT), Singapore in early January, 2011, under the support of CQT. I.K.'s research was also supported by French projects ANR-09-JCJC-0067-01, ANR-08-EMER-012 and the project QCS (grant 255961) of the E.U. S.Z.'s research was supported by China Basic Research Grant 2011CBA00300 (sub-project 2011CBA00301), Research Grants Council of Hong Kong (Project no. CUHK418710, CUHK419011), and benefited from research trips under the support of China Basic Research Grant 2007CB807900 (sub-project 2007CB807901).

quantum information processing: quantum communication enables players to achieve a real correlated equilibrium unconditionally, a task which is impossible in the classical world.

More precisely, we prove that for any correlated equilibrium p of a strategic game G , there exists an extended game (with a quantum communication initial stage) Q with an efficiently computable approximate Nash equilibrium σ , such that the expected payoff for both players in σ is at least as high as in p .

The main cryptographic tool used in the construction is the quantum weak coin flipping protocol of Mochon [Moc07].

1 Introduction

Game theory is a research area of great importance that studies the behavior of two or more players, when interacting with each other in order to achieve individual goals. It has found far reaching applications in the fields of economics, biology, computer science, sociology, political sciences, the study of Internet and stock markets, among others.

Most games fall into two broad categories: 1) Strategic games, where all players choose their strategies simultaneously or without knowing the other players' moves. The payoffs depend on the joint strategy that is performed by all players, and the game is usually described in a matrix form when there are only two players. 2) The extensive games, where the players take turns in making moves. Examples of strategic games include the Battle of the Sexes, Prisoner's Dilemma, Vickrey auction, etc. Examples of extensive games include chess, the eBay auction system, etc.

In order to study stable behaviors in games, the concept of an equilibrium has been put forward [vNM44]. A Nash equilibrium, the most fundamental notion of an equilibrium, is a joint strategy of all players, such that no player has any incentive to change her own strategy given that all other players retain theirs. One of the seminal results in this area is that every game with finite players and finite strategies has a mixed Nash equilibrium [vNM44, Nas51], *i.e.* one where the strategy of each player is a distribution over deterministic strategies. Note that these distributions are uncorrelated across different players and hence, each player can sample independently her strategy.

Even though the importance of Nash equilibria is undisputed, there are some drawbacks. First, the recent breakthrough results by [DGP09, CDT09] have shown that finding a Nash equilibrium is a computationally hard problem. To make matters worse, in many games there are more than one Nash equilibrium and it is really unclear whether the players will end up in one of them, and if yes, which one and how. Note that in many cases these equilibria are not fair, and thus different players have a preference for a different equilibrium.

Let us see a simple example, the Battle of the Sexes. A couple needs to decide where to go for holidays. Partner A prefers Amsterdam to Barcelona, and Partner B prefers Barcelona to Amsterdam. But both players prefer going to the same place than ending up in different places; see the following payoff Table, where

the pair of numbers in each entry represents the payoffs of the two partners in order.

	Amsterdam	Barcelona
Amsterdam	(4,2)	(0,0)
Barcelona	(0,0)	(2,4)

So where should they go? There are two pure Nash equilibria in the above game. They both go to Amsterdam, and hence have payoffs 4 and 2 respectively, or both go to Barcelona and have payoffs 2 and 4 respectively. Even though these are Nash equilibria, none of them is fair, causing the *battle* of the sexes. There is actually a third Nash equilibrium, a mixed one, where each player independently flips a coin and decides to go to their preferred place with probability $2/3$ and to the preferred place of the other player with probability $1/3$. In this case, the expected payoff is the same for both players and equal to $4/3$. Even though this is a fair equilibrium, it is pretty inefficient, since now both players have payoff even less than in the case of the unfair pure equilibrium. Moreover, there is a $5/9$ chance the couple goes in different places, which they really do not prefer.

One simple way to rectify all of these problems is the introduction of the notion of a correlated (Nash) equilibrium [Aum74]. In such an equilibrium, we allow the strategies of the players to be drawn from a correlated distribution p , and same as for a Nash equilibrium, we require that each player has no incentive to deviate given the current sample of his strategy and the information of the distribution p (but not the sampled strategies of the other players). There are many nice properties of these equilibria. First, they form a superset of Nash equilibria and hence they always exist. Moreover, it is not hard to exhibit games with a correlated equilibrium which enjoys fairness and whose social welfare (*i.e.* the total payoff of the players) is arbitrarily better than that of any Nash equilibrium. Second, unlike Nash equilibria, it is computationally easy to compute an optimal correlated equilibrium by solving an LP, for many types of games, including constant-player, polymatrix, graphical, hypergraphical, congestion, local effect, scheduling, facility location, network design and symmetric games [PR08, VNRT07]. In our previous example, a correlated equilibrium is the strategy where with probability $1/2$ the couple goes to Amsterdam or to Barcelona. The expected payoff for each player is then 3 and the couple is equally happy.

However, in general it is hard to sample from a correlated distribution. In fact, even for the case of two players and the distribution of the correlated equilibrium being just one fair coin, it is well known that without any computational assumptions, it is impossible to achieve just that; actually in any classical protocol one player has a strategy to get his/her desired outcome with probability 1. A canonical solution to this problem is to introduce a trusted mediator, who in this case flips the fair coin. However, for many real life scenarios, trusted mediators are simply not available.

A computational solution to this problem was proposed by Dodis, Halevi and Rabin [DHR00], who showed that classical cryptographic protocols can provide an elegant way to achieve a correlated equilibrium under standard computational

hardness assumptions. Both in their paper and in ours, by achieving a correlated equilibrium, we mean that the players achieve payoffs which are at least as high as the ones in the correlated equilibrium, but not that they necessarily sample the joint strategies with the distribution according to the correlated equilibrium. It is an open question if this stronger property is achievable.

More specifically, for any strategic game where the correlated equilibrium can be efficiently computed, Dodis *et al.* do the following: Before playing the game, the players communicate in order to sample a joint strategy from the equilibrium distribution, in such a way that each player at the end of the protocol only knows her strategy and has no information about the other players' moves apart from the fact that they come from the equilibrium distribution. The privacy and correctness of this procedure are guaranteed by the fact that the players are computationally bounded and the assumption that a primitive, equivalent to Oblivious Transfer, exists. Then, the players play the original game. Since they have no information about the other players' strategies and the joint strategy is sampled from a correlated equilibrium of the original game, they have no incentive to deviate. In other words, being honest during the communication phase and playing the move that resulted from the communication phase is a Nash equilibrium of the new extended game that achieves payoff equal to the correlated equilibrium of the original game. The introduction of cryptography in game theory is a very promising idea that nevertheless needs to be used with caution due to the many nuances in the two models. Note, last, that the use of cryptography by Dodis *et al.* provides a solution only when one is willing to accept the notion of computational equilibria, which are very different than the equilibria used by game theorists. Since then, a series of works have studied the relation between cryptography and game theory [FS02, LMPS04, IML05, ADGH06].

In our paper we show that, in fact, one need not resort to computational equilibria, if we allow the players to communicate via a quantum channel instead of a classical one. This provides another example of a major advantage of quantum information processing: quantum communication enables players to achieve a real correlated equilibrium. Note that we only make the communication before the game quantum but the game itself remains a classical one.

A priori, it is not clear that quantum communication can provide any significant advantage, since we know that Oblivious Transfer, the primitive that Dodis, Halevi and Rabin need for their construction, is impossible even in the quantum world [Lo97]. We overcome this problem by providing a new way to extend any game with an efficient correlated equilibrium into a new game that has an efficient Nash equilibrium achieving equal or even better payoffs (up to an arbitrarily small ϵ). The construction is based on the existence of a weaker primitive, called Weak Coin Flipping. This primitive is impossible classically without any computational assumptions. In the quantum world, however, Mochon [Moc07] has showed in a powerful result that there exists a quantum coin flipping protocol, where player A prefers Head and player B prefers Tail (which is exactly the case in the Battle of the Sexes), such that if one player plays the

honest strategy, then no matter how the other player plays, the bias of the coin cannot exceed an arbitrarily small ε .

Let us first focus on a subset of all correlated equilibria, that we call Nash-support correlated equilibria, where the distribution is over joint strategies that are pure Nash equilibria. Such correlated equilibria exist whenever the game has pure Nash equilibria and in many cases are optimal. For example, in the Battle of the Sexes, and more generally in all coordination games, the optimal correlated equilibrium is a uniform distribution over the two pure Nash equilibria.

As in Dodis *et al.*, we construct an extended game, in which the players first exchange messages, then play the original game by choosing strategies. A Nash equilibrium in the extended game is a sequence of moves of all players such that no unilateral deviation by one player can increase her payoff. At a high level, the new game we construct has the following stages: 1) *Communication stage*: the players use as a subroutine the quantum weak coin flipping protocol in order to sample a joint strategy from the distribution of the original Nash-support correlated equilibrium. 2) *Game stage*: the players play the original game and their payoff is the same as in the original game.

It is not hard to see that being honest during the communication stage and playing the strategy that corresponds to the sampled Nash equilibrium is an approximate Nash equilibrium and it achieves payoff (almost) equal to the correlated equilibrium of the original game. Let us assume that one of the players is dishonest while the other is playing the honest strategy. The cheating player can deviate during the coin flipping process but this only increases his payoff by at most ε by the security of the coin flipping protocol. Then, he can deviate by not playing the suggested strategy, but since the suggested strategy is a Nash equilibrium, he cannot increase his payoff.

Theorem 1. *For any Nash-support correlated equilibrium p of a game G with at most n strategies for each player, there exists an extended game Q with an ε -Nash equilibrium σ computable in time $\text{poly}(n, 1/\delta, 1/\varepsilon)$, such that the expected payoffs for both players in σ is at least as high as in p minus δ .*

For general correlated equilibria, we further extend our game as follows. Since we do not preserve privacy of the moves, it may be to someone's advantage to change their strategy instead of following the suggestion. We remedy this situation by adding a final stage to the game and by using the usual "Punishment for Deviation" method. Therefore, in 3) *Checking stage*: the players submit an Accept/Reject move, where a player plays Reject if the strategy of the other player during the second stage is not equal to the suggested one. The payoff of the players is equal to the one in the original game if they both play Accept in the last phase, and 0 otherwise. Note that we do not need the Accept/Reject moves to be simultaneous and that without loss of generality we assume that all payoffs are in $[0, 1]$.

Again, it is not hard to see that being honest during the communication stage and playing the suggested move is an approximate Nash equilibrium for this game and it achieves payoff equal to the correlated equilibrium of the original game. Let us assume that one of the players is dishonest while the other is

playing the honest strategy. The cheating player can deviate during the coin flipping process but this will only increase his payoff by at most an ϵ fraction by the security of the coin flipping protocol. Then, in the second stage, he can deviate by not playing the suggested strategy, but then his payoff will be 0 since the honest player will play Reject in the Checking stage. Hence, there is no significant advantage for any player to deviate from the honest strategy.

Theorem 2. *For any correlated equilibrium p of a game G with at most n strategies for each player, there exists an extended game Q with an ϵ -Nash equilibrium σ computable in time $\text{poly}(n, 1/\delta, 1/\epsilon)$, such that the expected payoffs for both players in σ is at least as high as in p minus δ .*

Let us make a more detailed comparison with the results of Dodis, Halevi and Rabin [DHR00]. They describe an extended game, first introduced by Barany [Bár92], that involves a communication stage and then the game stage. In the communication stage, they securely compute a functionality that they call Correlated Element Selection. This consists of two players sampling a joint strategy from a correlated distribution, with the extra privacy property that at the end each player knows only his/her own move. Then, in the second stage, the players play the original game. If a player catches the other one cheating during the communication stage, then he plays his minmax move in the second stage, i.e. the move that minimizes the other players' payoff.

On one hand, in our protocol, the communication stage achieves something weaker. We sample from the correlated distribution in a way that at the end, both players know the joint strategy. By removing the privacy constraint we are able to achieve the sampling using the weaker primitive of Weak Coin Flipping. A nice property of our procedure is that the honest player is guaranteed to have an output, regardless of the dishonest player's strategy. For the case of the Nash-support correlated equilibria, we do not have to resort to the minmax punishment, since even if the honest player catches the other player cheating, he can still play the suggested move. In the case of general correlated equilibria, we need to be more explicit in our punishment by adding the Accept/Reject stage, in order to dissuade the players from deviating from the suggested move.

On the other hand, we achieve something much stronger than before, which is that we do not make any assumptions about the computational power of the players. Hence we are able to use quantum communication to achieve a real correlated equilibrium for a large array of different types of games with unconditionally powerful players and without a trusted mediator.

A few remarks are in order for this extra checking stage that we add in the case where the correlated equilibrium has support on joint strategies that are not Nash equilibria. First, note that all the equilibria remain unchanged, since we specified the payoffs of any joint strategy with a Reject move as 0. Hence, sampling a correlated equilibrium in the new game is equivalent to sampling a correlated equilibrium in the original game. This means that the quantum advantage comes from the sampling part and not due to the checking part. For a fair comparison, we can also augment the classical game with the choice

of Accept/Reject. It is not hard to see that the players still cannot sample a correlated equilibrium in this new game; otherwise they would have been able to do a strong coin flipping which is impossible.

Second, in many practical situations, breaking preagreed rules is considered losing (and thus given the least payoff) automatically. Many games in sports are of this nature. For example, when the referee tosses a coin to decide the side of the court for each team, both teams know the outcome of this random process and are not allowed to disagree no matter the outcome; otherwise the team will be claimed to lose by the referee immediately. Moreover, in extensive games, the checking phase is already implicitly present. In the middle of a chess game, only a subset of moves is compatible with the stage of the game and hence if a player decides to play some other move, then the other player will Reject either immediately or at the end of the game. Hence, adding an Accept/Reject stage only makes explicit what is implicitly present in any game, that if a player breaks the rules then the other one rejects the outcome of the game.

Third, our Accept/Reject stage is not simultaneous. One has to be very careful with adding simultaneous moves to a game, since two players can flip a fair coin with a simultaneous move where each plays one of two possible moves at random. If the two moves are the same then the coin is Head and if different the coin is Tail. Here, we do not add the ability to play simultaneously.

Fourth, one may wonder why the honest player would prefer to reject and receive 0 payoff — she could instead choose to accept even though the other player cheated and receive a possibly positive payoff. Note that this is not a defect of our protocol, rather, it is an inherent property of Nash equilibria in extensive games. As explained in Dodis *et al.* where there was again a punishment step, the Nash equilibrium property requires merely local optimality by considering the scenario where at most one player deviates from the protocol; nothing is guaranteed if both players cheat. Moreover, the insistence of the honest player to punish the cheater forces the other player not to cheat in the first place (or to stop cheating if the game is repeated). One possible way to remedy the situation would be to consider subgame perfect equilibria, however neither our protocol nor the one in Dodis *et al.* has this property.

Note that our protocol does not provide a quantum algorithm to compute a Nash equilibrium. However, it almost renders this question moot. Instead of a quantum algorithm to compute a Nash equilibrium, there is a quantum protocol where the players can generate a correlated equilibrium, which enjoys desirable properties such as fairness and higher payoff.

Since our protocol uses quantum channels, one may wonder whether the power of two-way quantum communication enables us to achieve any *quantum* equilibrium with payoff higher than any classical correlated equilibrium. This is actually not possible: Any quantum protocol eventually generates a joint strategy s according to some correlated distribution p . If the players' behaviors in the protocol form a Nash equilibrium, then the resulting distribution p is a quantum correlated equilibrium of the quantized game, because otherwise the players would

like to change their behaviors in the last step. By an observation in [Zha12], p is also a (classical) correlated equilibrium of the original (classical) game, which the present paper already gives a way to generate.

2 Preliminaries

Game Theory. In a classical strategic game with n players, labeled by $\{1, 2, \dots, n\}$, each player i has a set S_i of strategies. We use $s = (s_1, \dots, s_n)$ to denote the *joint strategy* selected by the players and $S = S_1 \times \dots \times S_n$ to denote the set of all possible joint strategies. Each player i has a utility function $u_i : S \rightarrow \mathbb{R}$, specifying the *payoff* or *utility* $u_i(s)$ to player i on the joint strategy s . For simplicity of notation, we use subscript $-i$ to denote the set $[n] - \{i\}$, so s_{-i} is $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, and similarly for S_{-i} , p_{-i} , etc.

In a classical extensive game with perfect information, the players take moves in turns and all players know the entire history of all players' moves. An extensive game can be transformed into strategic form by tabulating all deterministic strategies of the players, which usually results in an exponential increase in size.

A game is $[0, 1]$ -*normalized* if all utility functions are in $[0, 1]$. Any game can be scaled to a normalized one. For a fair comparison, we assume that all games in this paper are normalized.

A Nash equilibrium is a fundamental solution concept in game theory. Roughly, it says that in a joint strategy, no player can gain more by changing her strategy, provided that all other players keep their current strategies unchanged.

Definition 1. A pure Nash equilibrium is a joint strategy $s = (s_i, \dots, s_n) \in S$ satisfying

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}), \quad \forall i \in [n], \forall s'_i \in S_i.$$

Pure Nash equilibria can be generalized by allowing each player to independently select her strategy according to some distribution, leading to the following concept of *mixed Nash equilibrium*.

Definition 2. A (mixed) Nash equilibrium (NE) is a product probability distribution $p = p_1 \times \dots \times p_n$, where each p_i is a probability distributions over S_i , satisfying

$$\sum_{s_{-i}} p_{-i}(s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i}} p_{-i}(s_{-i}) u_i(s'_i, s_{-i}), \quad \forall i \in [n], \forall s_i, s'_i \in S_i \text{ with } p_i(s_i) > 0.$$

Definition 3. A correlated equilibrium (CE) is a probability distribution p over S satisfying

$$\sum_{s_{-i}} p(s_i, s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i}} p(s_i, s_{-i}) u_i(s'_i, s_{-i}), \quad \forall i \in [n], \forall s_i, s'_i \in S_i.$$

If the correlated equilibrium is a distribution over pure Nash equilibria, then we call it a *Nash-support* correlated equilibrium.

We will also need an approximate version of equilibrium, which basically says that no Player i can gain much by changing the suggested strategy s_i . Depending on whether we require the limit of the gain for each possible s_i in the support of p or on average of p , one can define worst-case and average-case approximate equilibrium. It turns out that the average-case one, as defined below, has many nice properties, such as being the limit of a natural dynamics of minimum regrets ([VNRT07], Chapter 4) and hence it is the one we will use.

Definition 4. *An ε -correlated equilibrium is a probability distribution p over S satisfying*

$$\mathbf{E}_{s \leftarrow p}[u_i(s'_i(s_i), s_{-i})] \leq \mathbf{E}_{s \leftarrow p}[u_i(s_i, s_{-i})] + \varepsilon,$$

for any i and any function $s'_i : S_i \rightarrow S_i$. An ε -correlated equilibrium p is an ε -Nash equilibrium if it is a product distribution $p = p_1 \times \cdots \times p_n$.

We can also define equilibria for extensive games by defining the corresponding equilibria on their strategic form.

Cryptography. We provide the formal definition of a weak coin flipping protocol.

Definition 5. *A weak coin flipping protocol between Alice and Bob is a protocol where Alice and Bob interact and at the end, Alice outputs a value $c_A \in \{0, 1\}$ and Bob outputs a value $c_B \in \{0, 1\}$. If $c_A = c_B$, we say that the protocol outputs $c = c_A$. If $c_A \neq c_B$ then the protocol outputs $c = \perp$.*

An (a, ε) -weak coin flipping protocol ($WCF(a, \varepsilon)$) has the following properties:

- *If $c = a$, we say that Alice wins. If $c = 1 - a$, we say that Bob wins.*
- *If Alice and Bob are honest then $\Pr[\text{Alice wins}] = \Pr[\text{Bob wins}] = 1/2$.*
- *If Alice cheats and Bob is honest then $P_A^* = \Pr[\text{Alice wins}] \leq 1/2 + \varepsilon$.*
- *If Bob cheats and Alice is honest then $P_B^* = \Pr[\text{Bob wins}] \leq 1/2 + \varepsilon$.*

P_A^ and P_B^* are the cheating probabilities of Alice and Bob. The cheating probability of the protocol is defined as $\max\{P_A^*, P_B^*\}$.*

Note that in the definition the players do not abort, since a player that wants to abort can always declare victory rather than aborting without reducing the security of the protocol.

We will use the following result by Mochon.

Proposition 1. [Moc07] *For every $\varepsilon > 0$ and $a \in \{0, 1\}$, there exists a quantum $WCF(a, \varepsilon)$ protocol P .*

Moreover, the protocol uses a number of qubits and rounds which is polynomial in $\frac{1}{\varepsilon}$. Note that this is a weaker definition of a usual coin flip, since here, we assign a winning value for each player. Even though each player cannot bias the coin towards this winning value, he or she can bias the coin towards the losing value with probability 1. Weak coin flipping is possible using quantum communication, though for the strong coin flipping the optimal cheating probability for any protocol is $1/\sqrt{2}$ [Kit03, CK09].

In the following section we will use weak coin flipping as a subroutine for the following cryptographic primitive, that enables two players to jointly sample from a correlated distribution, in a way that no dishonest player can force a distribution which is far from the honest one.

Definition 6. A Correlated Strategy Sampling protocol between two players P_1 and P_2 is an interactive protocol where the players receive as input a game G with an efficiently computable correlated equilibrium¹ p and at the end, P_1 outputs a joint strategy $s = (s_1, s_2) \in S_1 \times S_2$ and P_2 outputs a joint strategy $s' = (s'_1, s'_2) \in S_1 \times S_2$. If $s = s'$, we say the protocol outputs $s = (s_1, s_2)$. If $s \neq s'$ then we say the protocol outputs \perp .

An (ε, δ) -Correlated Strategy Sampling procedure satisfies the following properties:

1. If both players follow the honest strategy, then they both output the same joint strategy $s = (s_1, s_2)$, where $s \leftarrow p_h$ for some distribution p_h , s.t.

$$\text{for both } i \in \{1, 2\}, \quad \mathbf{E}_{s \leftarrow p_h}[u_i(s)] \geq \mathbf{E}_{s \leftarrow p}[u_i(s)] - \delta$$

2. If Player 1 is dishonest and Player 2 is honest (similarly for the other case), then Player 2 outputs a joint strategy s distributed according to some q , s.t.

$$\mathbf{E}_{s \leftarrow q}[u_2(s)] \geq \mathbf{E}_{s \leftarrow p_h}[u_2(s)] - \varepsilon, \quad \mathbf{E}_{s \leftarrow q}[u_1(s)] \leq \mathbf{E}_{s \leftarrow p_h}[u_1(s)] + \varepsilon.$$

Note again, that similar to the case of the weak coin flip, the players do not abort, since a player that wants to abort can always choose the joint strategy that is best for him rather than aborting without reducing the security of the protocol.

3 The Extended Game

For simplicity, we consider a two-player strategic game G with at most n strategies for each player, but our results easily extend to more players. We describe how to derive an extended game Q from any such G .

Similar to the DHR extended game, we assume that the players can communicate with each other before they start playing the game, but now via a quantum channel. In this preamble stage they perform a quantum protocol that we call Correlated Strategy Sampling.

In the following section we show how to implement this procedure unconditionally, using a Weak Coin Flipping subroutine with bias $\varepsilon' = O(\varepsilon / \log n)$.

Then, we extend the original game G to a 2-stage game, where the first stage is identical to the game G and for the second stage, which we call the Checking stage, the available moves for each player are Accept or Reject. We define the payoff for any joint strategy where some player outputs Reject in the second stage to be 0.

¹ A correlated equilibrium p is efficiently computable if there is a Turing machine that, on an input game, outputs p in time polynomial in the input size.

<p>Extended Game Q</p> <ol style="list-style-type: none"> 1. Communication Stage: The two players perform the Correlated Strategy Sampling procedure for the game G and correlated equilibrium p. 2. Game Stage: The two players play the original game G. 3. Checking Stage: The two players each play a move from the set $\{A, R\}$.

We can now restate and prove our main theorem.

Theorem 2. *For any correlated equilibrium p of the game G with at most n strategies for each player, and for any $\varepsilon, \delta > 0$, there exists an extended game Q with an ε -Nash equilibrium σ that can be computed in time $\text{poly}(n, 1/\delta, 1/\varepsilon)$ and such that the expected payoff for both players in σ is at least as high as the one in p minus δ .*

Proof. We describe Player 1's strategy in the ε -Nash equilibrium σ as follows (Player 2's strategy is symmetric): In the Communication Stage, Player 1 is honest and obtains an output $s = (s_1, s_2)$. In the Game Stage, he plays the move s_1 . In the Checking Stage, he plays A if Player 2's move in the Game Stage was s_2 and R otherwise.

Let us show that this is indeed an ε -Nash equilibrium. A dishonest player (assume Player 1) can try to increase his payoff by first deviating from the protocol in the Communication Stage. If Player 2 outputs a joint strategy $s = (s_1, s_2)$ then we know from the security of the Correlated Strategy Sampling procedure that this is a sample from a distribution q s.t.

$$\mathbf{E}_{s \leftarrow q}[u_1(s)] \leq \mathbf{E}_{s \leftarrow p_h}[u_1(s)] + \varepsilon. \quad (1)$$

Hence, if Player 1 is dishonest during Stage 1 and then plays s_1 in Stage 2, then his gain is at most ε . If he decides to change his move, then the honest player would play R in Stage 3, so his payoff would be 0. Overall, no matter what strategy the dishonest player follows he cannot increase his payoff more than ε from the honest strategy mentioned above, and hence this strategy is an ε -approximate Nash equilibrium.

Note that from the security of the Correlated Strategy Sampling procedure we also have

$$\mathbf{E}_{s \leftarrow q}[u_2(s)] \geq \mathbf{E}_{s \leftarrow p_h}[u_2(s)] - \varepsilon. \quad (2)$$

Hence, we have the following interesting corollary

Corollary 1. *In the extended game Q , the expected payoff of the honest player will not decrease by more than ε , no matter how the dishonest player deviates, unless the dishonest player makes both players' payoff equal to 0.*

In other words, the honest strategy remains an equilibrium even if the objective of a player is not to maximize his own payoff but rather maximize the difference between the players' payoffs.

In the special case of a Nash-support correlated equilibrium, the extended game consists only of the first two stages. Similarly to the general case, we can prove that the honest strategy is an approximate Nash equilibrium. Note that now, the reason the dishonest player cannot increase his payoff by deviating from the suggested joint strategy is because the joint strategy is a pure Nash equilibrium, which, by definition, leaves each player no incentive to deviate even if she knows the other players' pure strategies to be played. Hence we do not need the checking stage.

4 The Correlated Strategy Sampling Procedure

Let us first describe the intuition behind our procedure. First, we think of all distributions as uniform distributions over a multiset of joint strategies of size K , and hence sampling a joint strategy is equivalent to uniformly sampling a number in $[K]$. Then, we sample sequentially $\log K$ bits. At each step, the players declare the value of the bit they prefer, by calculating their expected payoffs conditioned on the already sampled bits. If the players agree on the value, then this is the sampled bit. If they disagree, then they perform a weak coin flip in order to sample the bit. The fact that the weak coin flip is almost perfect, implies that no dishonest player can bias the distribution by a lot.

We now provide the technical details. In a two-player game G with at most n strategies for each player, let p be an efficiently computable correlated equilibrium that the players know and aim to generate. A typical scenario is that p is the lexicographically first correlated equilibrium that maximizes the total payoff. If the distribution is not uniform we can emulate it by a uniform distribution \bar{p} on a multiset S of joint strategies of size $K = 2^k \in [n/\delta, 2n/\delta]$ for some integer k , such that the distance between the two distributions is at most δ . We can equivalently think of the distribution \bar{p} as a distribution on $\{0, 1\}^k$ by associating each element $\bar{s} \in \{0, 1\}^k$ with an element $s \in S$.

Let \bar{p}_h be the distribution on $\{0, 1\}^k$ that arises from our procedure when both players are honest and \bar{q} the distribution of the honest player's output when the other player is dishonest. The distributions \bar{p}_h and \bar{q} naturally give rise to the distributions p_h and q on the set of joint strategies (where the probability of s in p_h is the sum of the probabilities according to \bar{p}_h of the elements in $\{0, 1\}^k$ that correspond to s).

For a probability distribution μ over $X = X_1 \times \dots \times X_k$, we use the standard notation $\mu(\cdot | x_{m+1} \dots x_k)$ to denote the conditional distribution on $X_1 \times \dots \times X_m$, i.e. μ conditioned on the last $k - m$ variables being $x_{m+1} \dots x_k$. We also use $\mu(x_1 \dots x_m)$ to denote the probability in the marginal distribution on the first m variables. By $x \leftarrow \mu$, we mean to draw a sample x from distribution μ . Let $\text{sign}(a)$ be the function which is 1 if $a \geq 0$ and -1 if $a < 0$. The protocol appears in the following figure.

(ε, δ) -Correlated Strategy Sampling Protocol

Input: A game G with at most n strategies for each player, and an efficiently computable correlated equilibrium p .

1. Each Player i computes locally the equilibrium p and emulates p by a uniform distribution \bar{p} on a multiset of joint strategies (i.e. on $\{0, 1\}^k$), with $k = O(\log n)$.
2. **for** $j = 1$ to k
 - (a) Each Player i computes and announces his preference $a_i^j = \text{sign}(\mathbf{E}_{\bar{s}^{j+1} \dots \bar{s}^k \leftarrow \bar{p}_h(\cdot | \bar{s}^1 \dots \bar{s}^{j-1} 0)}[u_i(s)] - \mathbf{E}_{\bar{s}^{j+1} \dots \bar{s}^k \leftarrow \bar{p}_h(\cdot | \bar{s}^1 \dots \bar{s}^{j-1} 1)}[u_i(s)])$.
 - (b) **if** $a_1^j a_2^j = -1$,
Run $WCF(a_1^j, \varepsilon/k)$. Let the outcome of Player i be $\bar{s}_i^j \in \{0, 1\}$.
else
Set $\bar{s}_1^j = \bar{s}_2^j$ to be their commonly desirable value.
3. Each Player i outputs s according to the jointly flipped coins $\bar{s} = \bar{s}_i^1 \dots \bar{s}_i^k$.

Analysis. First, if both players are honest then their expected utility is at least as high as in the original CE, up to an additive error δ due to the precision of using k bits to emulate p . If in all rounds they flip a fair coin then their expected utility is exactly the same as in p . If at some round they both agree on a preferred value then this increases both players expected utility.

We now prove that no dishonest player can increase his utility by much. Let us assume without loss of generality that Player 1 is dishonest and Player 2 is honest. We prove that after round m ,

Claim. For any $m = 1, \dots, k$, we have

$$\begin{aligned} & \sum_{\bar{s}^1 \dots \bar{s}^m} (\bar{q}(\bar{s}^1 \dots \bar{s}^m) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^m)) \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(s) \\ \leq & \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} (\bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1})) \sum_{\bar{s}^m \dots \bar{s}^k} \bar{p}_h(\bar{s}^m \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^{m-1}) u_1(s) + \frac{\varepsilon}{k}. \end{aligned} \quad (3)$$

The proof is in Appendix. Adding the inequalities in the claim for all m , we have

$$\sum_{\bar{s}^1 \dots \bar{s}^k} (\bar{q}(\bar{s}^1 \dots \bar{s}^k) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^k)) u_1(s) \leq \varepsilon.$$

By going back to the space of joint strategies we have $\mathbf{E}_{s \leftarrow q}[u_1(s)] \leq \mathbf{E}_{s \leftarrow p_h}[u_1(s)] + \varepsilon$.

Moreover, for the honest player we have, by a similar argument (changing u_1 to u_2 , ε/k to $-\varepsilon/k$, and changing the direction of the inequality in Claim 4), it is also easy to show the claimed Eq.(2).

The same analysis holds when Player 2 is dishonest. Also, it is easy to see that the complexity of the protocol is polynomial in n/δ and $1/\varepsilon$. This completes the proof of our main theorem.

A final remark is that the same protocol can be used for general k -player games. In each round, some players prefer \bar{s}^m to be 0 and some players prefer 1. We can then let two representatives, one from each group, to do the weak coin flipping, at the end of which the representatives announce the bits. If one representative lies, then the other reject in the third stage. The previous analysis then easily applies to this scenario as well.

References

- [ADGH06] Abraham, I., Dolev, D., Gonen, R., Halpern, J.: Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In: Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing, pp. 53–62 (2006)
- [Aum74] Aumann, R.: Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96 (1974)
- [Bár92] Bárány, I.: Fair distribution protocols or how the players replace fortune. *Mathematics of Operations Research* 17, 327–340 (1992)
- [CDT09] Chen, X., Deng, X., Teng, S.: Settling the complexity of computing two-player nash equilibria. *Journal of the ACM* 56(3) (2009)
- [CK09] Chailloux, A., Kerenidis, I.: Optimal quantum strong coin flipping. In: The 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 527–533 (2009)
- [DGP09] Daskalakis, C., Goldberg, P., Papadimitriou, C.: Computing a nash equilibrium is PPAD-complete. *SIAM Journal on Computing* 39(1), 195–259 (2009)
- [DHR00] Dodis, Y., Halevi, S., Rabin, T.: A Cryptographic Solution to a Game Theoretic Problem. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 112–130. Springer, Heidelberg (2000)
- [FS02] Feigenbaum, J., Shenker, S.: Distributed algorithmic mechanism design: recent results and future directions. In: Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, pp. 1–13 (2002)
- [IML05] Izmalkov, S., Micali, S., Lepinski, M.: Rational secure computation and ideal mechanism design. In: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, pp. 585–595 (2005)
- [Kit03] Kitaev, A.: Quantum coin-flipping. Presentation at The 6th Workshop on Quantum Information Processing, QIP 2003 (2003)
- [LMPS04] Lepinski, M., Micali, S., Peikert, C., Shelat, A.: Completely fair SFE and coalition-safe cheap talk. In: Proceedings of the Twenty-third Annual ACM Symposium on Principles of Distributed Computing, pp. 1–10 (2004)
- [Lo97] Lo, H.-K.: Insecurity of quantum secure computations. *Physical Review A* 56(2) (1997)
- [Moc07] Mochon, C.: Quantum weak coin flipping with arbitrarily small bias. arXiv:0711.4114 (2007)
- [Nas51] Nash, J.: Non-cooperative games. *The Annals of Mathematics* 54(2), 286–295 (1951)
- [PR08] Papadimitriou, C.H., Roughgarden, T.: Computing correlated equilibria in multi-player games. *Journal of the ACM* 55(3) (2008)
- [vNM44] von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press (1944)

[VNRT07] Vazirani, V., Nisan, N., Roughgarden, T., Tardos, É.: Algorithmic Game Theory. Cambridge University Press (2007)

[Zha12] Zhang, S.: Quantum strategic game theory. In: Proceedings of the 3rd Innovations in Theoretical Computer Science, pp. 39–59 (2012); earlier at arXiv:1012.5141 and QIP 2011

A Proof of Claim 4

For the convenience of notation, we sometimes write $u_1(\bar{s})$ to mean $u_1(s)$ where s corresponds \bar{s} . First by expanding the probability to marginal times conditional probabilities, we have

$$\begin{aligned} & \sum_{\bar{s}^1 \dots \bar{s}^m} (\bar{q}(\bar{s}^1 \dots \bar{s}^m) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^m)) \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(s) \\ = & \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} \left(\bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) \sum_{\bar{s}^m} \bar{q}(\bar{s}^m | \bar{s}^1 \dots \bar{s}^{m-1}) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1}) \sum_{\bar{s}^m} \bar{p}_h(\bar{s}^m | \bar{s}^1 \dots \bar{s}^{m-1}) \right) \\ & \cdot \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(s). \end{aligned}$$

For those $\bar{s}^1 \dots \bar{s}^{m-1}$ that the two players have the same preference on \bar{s}^m , the best for Player 1 is then just to follow the honest protocol. Thus the corresponding part in the inequality in Claim 4 is true even without the ϵ/k term. For the rest $\bar{s}^1 \dots \bar{s}^{m-1}$, the two players have different preferences; without loss of generality, assume that Player 1 prefers \bar{s}^m to be 0. Then the best for Player 1 raising her utility is to try to bias \bar{s}^m in the coin flipping to 0 as much as possible. By the security of the weak coin flipping (which holds against a dishonest player that may possess a quantum auxiliary input, hence includes the situation where the dishonest player may try to entangle the different executions of the coin flips), the above quantity is at most

$$\begin{aligned} & \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} \left(\bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) \left(\frac{1}{2} + \frac{\epsilon}{k} \right) \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(\bar{s}^1 \dots \bar{s}^{m-1} 0 \bar{s}^{m+1} \dots \bar{s}^k) \right. \\ & \quad \left. + \bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) \left(\frac{1}{2} - \frac{\epsilon}{k} \right) \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(\bar{s}^1 \dots \bar{s}^{m-1} 1 \bar{s}^{m+1} \dots \bar{s}^k) \right) \\ & \quad - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1}) \frac{1}{2} \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(\bar{s}^1 \dots \bar{s}^{m-1} 0 \bar{s}^{m+1} \dots \bar{s}^k) \\ & \quad - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1}) \frac{1}{2} \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(\bar{s}^1 \dots \bar{s}^{m-1} 1 \bar{s}^{m+1} \dots \bar{s}^k) \end{aligned}$$

$$\begin{aligned}
&= \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} (\bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1})) \sum_{\bar{s}^m} \bar{p}_h(\bar{s}^m | \bar{s}^1 \dots \bar{s}^{m-1}) \\
&\quad \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) u_1(s) + \frac{\epsilon}{k} \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} \bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) \\
&\quad \sum_{\bar{s}^{m+1} \dots \bar{s}^k} \bar{p}_h(\bar{s}^{m+1} \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^m) \\
&\quad \cdot [u_1(\bar{s}^1 \dots \bar{s}^{m-1} 0 \bar{s}^{m+1} \dots \bar{s}^k) - u_1(\bar{s}^1 \dots \bar{s}^{m-1} 1 \bar{s}^{m+1} \dots \bar{s}^k)]. \\
&\leq \sum_{\bar{s}^1 \dots \bar{s}^{m-1}} (\bar{q}(\bar{s}^1 \dots \bar{s}^{m-1}) - \bar{p}_h(\bar{s}^1 \dots \bar{s}^{m-1})) \sum_{\bar{s}^m \dots \bar{s}^k} \bar{p}_h(\bar{s}^m \dots \bar{s}^k | \bar{s}^1 \dots \bar{s}^{m-1}) u_1(s) + \frac{\epsilon}{k}.
\end{aligned}$$

where we used the fact that $\bar{p}_h(\bar{s}^m | \bar{s}^1 \dots \bar{s}^{m-1}) = 1/2$ in the equality, and the fact that the game is $[0, 1]$ -normalized in the inequality.