Multiparty Quantum Coin Flipping

Andris Ambainis* IAS and U. of Latvia
Harry Buhrman‡ CWI and U. of Amsterdam
Yevgeniy Dodis§ New York University
Hein Röhrig‡ U. of Calgary

Abstract

We investigate coin-flipping protocols for multiple parties in a quantum broadcast setting:

- *We propose and motivate a definition for quantum broadcast. Our model of quantum broadcast channel is new.*
- *We discovered that quantum broadcast is essentially a combination of pairwise quantum channels and a classical broadcast channel. This is a somewhat surprising conclusion, but helps us in both our lower and upper bounds.*
- *We provide* tight *upper and lower bounds on the optimal bias ε of a coin which can be flipped by k parties of which exactly g parties are honest: for* any $1 \leq g \leq k$, $\varepsilon = \frac{1}{2} - \Theta\left(\frac{g}{k}\right)$.

Thus, as long as a constant fraction of the players are honest, they can prevent the coin from being fixed with at least a constant probability. This result stands in sharp contrast with the classical setting, where no non-trivial coin-flipping is possible when $g \leq \frac{k}{2}$.

1. Introduction

1.1. The problem

Consider k parties out of which at least $g \geq 1$ are honest and at most $(k-g)$ are dishonest; which players are dishonest is fixed in advance but unknown to the honest players. The players can communicate over broadcast channels. Initially they do not share randomness, but they can privately flip coins; the probabilities below are with respect to the private random coins. A coin-flipping protocol establishes among the honest players a bit b such that

- if all players are honest, $\Pr[b=0] = \Pr[b=1] = \frac{1}{2}$
- if at least g players are honest, then $\Pr[b=0], \Pr[b=1] \leq \frac{1}{2} + \varepsilon$

ε is called the *bias*; a small bias implies that colluding dishonest players cannot strongly influence the outcome of the protocol. Players may abort the protocol.

1.2. Related work

Classically, if a (weak) majority of the players is bad then no bias $< \frac{1}{2}$ can be achieved and hence no meaningful protocols exist [16]. For example, if we only have two players and one of them is dishonest, then no protocols with bias $< \frac{1}{2}$ exist. For a minority of bad players, quite non-trivial protocols exist. For example, Feige [8] elegantly showed that $(\frac{1}{2}+\delta)$-fraction of good players can achieve bias $\frac{1}{2} - \Omega(\delta^{1.65})$, while achieving bias better than $\frac{1}{2} - \delta$ is impossible.

Allowing quantum bits (qubits) to be sent instead of classical bits changes the situation dramatically. Surprisingly, already in the two-party case coin flipping with bias $< \frac{1}{2}$ is possible, as was first shown by Aharonov et al. [2]. The best known bias is $\frac{1}{4}$ and this is optimal for a special class of three-round protocols [4]; for a bias of ε at least $\Omega(\log\log\frac{1}{\varepsilon})$ rounds of communication are necessary [4]. Kitaev (unpublished, see [12]) showed that in the two-party case no bias smaller than $\frac{1}{\sqrt{2}} - \frac{1}{2}$ is possible.

A weak version of the coin-flipping problem is one in which we know in advance that outcome 0 benefits Alice and outcome 1 benefits Bob. In this case, we only need to bound the probabilities of a dishonest Alice convincing Bob that the outcome is 0 and a dishonest Bob convincing Alice that the outcome is 1. In the classical setting, a standard argument shows that even weak coin flipping with a bias $< \frac{1}{2}$ is impossible when a majority of the players is dishonest. In the quantum setting, this scenario was first studied under the name *quantum gambling* [9]. Subsequently, Spekkens and Rudolph [17] gave a quantum protocol for weak coin flipping with bias $\frac{1}{\sqrt{2}} - \frac{1}{2}$ (i.e., no party can achieve the *desired outcome* with probability greater than $\frac{1}{\sqrt{2}}$). This was recently improved to 0.192 by Mochon [14]. Notice that this

* supported in part by NSF Grant DMS-0111298
‡ supported in part by the EU fifth framework projects QAIP, IST-1999-11234, and RESQ, IST-2001-37559, and a NWO grant
§ supported in part by the National Science Foundation under CAREER Award No. CCR-0133806 and Trusted Computing Grant No. CCR-0311095

is a better bias than in the best strong coin flipping protocol of [4].

We also remark that Kitaev's lower bound for strong coin flipping does not apply to weak coin flipping. Indeed, Mochon's protocol has a better bound than Kitaev's lower bound. Thus, weak protocols with arbitrarily small $\varepsilon > 0$ may be possible. The only known lower bounds for weak coin flipping are that the protocol of [17] is optimal for a restricted class of protocols [3] and that a protocol must use at least $\Omega(\log\log\frac{1}{\varepsilon})$ rounds of communication to achieve bias ε (shown in [4] for strong coin flipping but the proof also applies to weak coin flipping).

1.3. Our contribution

In this paper, we focus on quantum coin flipping for more than two players. However, for our multiparty quantum protocols we will first need a new two-party quantum protocol for *coin flipping with penalty for cheating*. In this problem, players can be heavily penalized for cheating, which will allow us to achieve lower cheating probability as a function of the penalty. This primitive and the quantum protocol for it are presented in Section 2; they may be of independent interest.

One way to classically model communication between more than two parties is by a primitive called *broadcast*. When a player sends a bit to the other players he broadcasts it to all the players at once [6]. However, when we deal with qubits such a broadcast channel is not possible since it requires to clone or copy the qubit to be broadcast and cloning a qubit is not possible [19]. In Section 3 we develop a proper quantum version of the broadcast primitive, which generalizes the classical broadcast. Somewhat surprisingly, we show that our quantum broadcast channel is essentially as powerful as a combination of pairwise quantum channels and a classical broadcast channel. This could also be of independent interest.

Using this broadcast primitive we obtain our main result:

Theorem 1 *For k parties out of which g are honest, the optimal achievable bias is $(\frac{1}{2} - \Theta(\frac{g}{k}))$.*

We prove Theorem 1 by giving an efficient protocol with bias $(\frac{1}{2} - \Omega(\frac{g}{k}))$ in Section 4 and showing a lower bound of $(\frac{1}{2} - \mathrm{O}(\frac{g}{k}))$ in Section 5. Our protocol builds upon our two-party coin-flipping with penalties which we develop in Section 2, and the classical protocol of Feige [8] which allows to reduce the number of participants in the protocol without significantly changing the fraction of good players present. Our lower bound extends the lower bound of Kitaev [12].

To summarize, we show that quantum coin flipping is significantly more powerful than classical coin flipping. Moreover, we give *tight* tradeoffs between the number of cheaters tolerated and the bias of the resulting coin achievable by quantum coin-flipping protocols. We also remark that the fact that we obtain tight bounds in the quantum setting is somewhat surprising. For comparison, such tight bounds are unknown for the classical setting.

In the remainder of the paper, we assume some familiarity with quantum computing. We recommend the book of Nielsen and Chuang [15] for background information on this topic.

1.4. Semidefinite programming

Some of our proofs make use of duality in semidefinite programming. For a review of semidefinite programming, see e.g., [13]. Semidefinite programming is a generalization of linear programming. In addition to linear constraints, semidefinite programs (SDPs) may have constraints that require that a square matrix of variables is positive semidefinite, i.e., that is is symmetric and all its eigenvalues are nonnegative.

We make use of the following basic properties of semidefinite matrices. Let A, B, and C denote square matrices acting on some linear space $\mathcal{V}$ and $\mathcal{W} \subseteq \mathcal{V}$ a subspace. If A is positive semidefinite, we write $A \geq 0$. We define $A \geq B :\Leftrightarrow A - B \geq 0$. Then

$$A \geq B \Leftrightarrow \forall|\psi\rangle : \langle\psi|A|\psi\rangle \geq \langle\psi|B|\psi\rangle$$
$$A = B + C \text{ and } C \geq 0 \Rightarrow A \geq B$$
$$A \geq B \Rightarrow \mathrm{tr}_{\mathcal{W}}(A) \geq \mathrm{tr}_{\mathcal{W}}(B)$$

Here $\mathrm{tr}_{\mathcal{W}}(A)$ denotes the partial trace.

In the Lagrange-multiplier approach, a constrained optimization problem (called the *primal* problem)

$$\max_{x\geq 0} f(x) \text{ subject to } g(x) \leq a \quad \text{for fixed } a > 0$$

is reformulated as an unconstrained optimization problem

$$\max_x \inf_{\lambda\geq 0} f(x) - \lambda\cdot(g(x) - a) \ ,$$

which is bounded from above by the constrained optimization problem (the *dual* problem)

$$\min_{\lambda\geq 0} \lambda\cdot a \text{ subject to } (f - \lambda\cdot g)(x) \leq 0 \text{ for all } x \geq 0 \ .$$

In linear programming, $(f - \lambda\cdot g)(x) \leq 0$ for all $x \geq 0$ if and only if $f - \lambda\cdot g \leq 0$. Therefore the preceding optimization problem can be simplified to

$$\min_{\lambda\geq 0} \lambda\cdot a \text{ subject to } f - \lambda\cdot g \leq 0 \ .$$

The same construction applies to SDPs using matrices as variables and $A\cdot B := \mathrm{tr}(A^*B)$. A feasible solution of the

dual yields an upper bound on the optimal value of the primal problem. Strong duality (i.e., that the optimal values coincide) does not hold in general; however, we will not need this below.

2. Two-party coin flipping with penalty for cheating

We consider the following model for coin flipping. We have two parties: Alice and Bob, among whom at least one is assumed to be honest. If no party is caught cheating, the winner gets 1 coin, the loser gets 0 coins. If honest Alice catches dishonest Bob, Bob loses v coins but Alice wins 0 coins. Similarly, if honest Bob catches dishonest Alice, she loses v coins but Bob wins 0 coins.

Theorem 2 *If Alice (Bob) is honest, the expected win by dishonest Bob (Alice) is at most* $\frac{1}{2} + \frac{1}{\sqrt{v}}$*, for* $v \geq 4$.

Proof. The protocol is as follows. Let $\delta = \frac{2}{\sqrt{v}}$. For $a \in \{0,1\}$, define $|\psi_a\rangle = \sqrt{\delta}|a\rangle|a\rangle + \sqrt{1-\delta}|2\rangle|2\rangle \in \mathbb{C}^3 \otimes \mathbb{C}^3$.

1. Alice picks $a \in \{0,1\}$ uniformly at random, generates the state $|\psi_a\rangle$ and sends the second register to Bob.
2. Bob stores this state in a quantum memory, picks $b \in \{0,1\}$ uniformly at random and sends b to Alice.
3. Alice then sends a and the first register to Bob and Bob verifies if the joint state of the two registers is $|\psi_a\rangle$ by measuring it in a basis consisting of $|\psi_a\rangle$ and everything orthogonal to it. If the test is passed, the result of coin flip is $a \oplus b$, otherwise Bob catches Alice cheating.

Theorem 2 follows immediately from the following two lemmas. □

Lemma 3 *Bob cannot win with probability more than* $\frac{1}{2} + \frac{1}{\sqrt{v}}$*, thus his expected win is at most* $\frac{1}{2} + \frac{1}{\sqrt{v}}$.

Proof. Let ρ_a be the density matrix of the second register of $|\psi_a\rangle$. Then, for the trace distance between ρ_0 and ρ_1 we have $\|\rho_0 - \rho_1\|_t = 2\delta$.

The trace distance is a measure for the distinguishability of quantum states analogous to the total-variation distance of probability distributions; see e.g., [1]. In particular, the probability of Bob winning is at most $\frac{1}{2} + \frac{\|\rho_0-\rho_1\|_t}{4} = \frac{1}{2} + \frac{\delta}{2} = \frac{1}{2} + \frac{1}{\sqrt{v}}$. □

Lemma 4 *Dishonest Alice's expected win is at most* $\frac{1}{2} + \frac{1}{\sqrt{v}}$.

Proof. Alice is trying to achieve $a \oplus b = 0$, which is equivalent to $a = b$. We describe the optimal strategy of Alice as a semidefinite program.

The variables are semidefinite matrices over subspaces of $\mathcal{X} \otimes \mathcal{A} \otimes \mathcal{B}$, where $\mathcal{X}$ is Alice's private storage, $\mathcal{A} \cong \mathbb{C}^3$ holds the first qutrit of the state to be sent in the protocol and $\mathcal{B} \cong \mathbb{C}^3$ holds the second qutrit. For $a, b \in \{0,1\}$, let $\rho_{ba} \in \mathcal{A} \otimes \mathcal{B}$ denote the state that Bob has in the last round, when he has sent b and Alice has sent a (and some qutrit). For $b \in \{0,1\}$, let $\rho_b \in \mathcal{X} \otimes \mathcal{A} \otimes \mathcal{B}$ denote the state before Alice decides on a. Finally, let $\rho_{\text{initial}} \in \mathcal{X} \otimes \mathcal{A} \otimes \mathcal{B}$ denote the state that Alice prepares initially and of which she sends the $\mathcal{B}$ part to Bob.

Then we have the following constraints. The initial state is an arbitrary density matrix:

$$\operatorname{tr}(\rho_{\text{initial}}) = 1 \tag{1}$$

When Alice learns b, she cannot touch $\mathcal{B}$ anymore, but she can apply an arbitrary unitary U_b on $\mathcal{X} \otimes \mathcal{A}$ to store her choice a in $\mathcal{X}$ and to prepare the $\mathcal{A}$ register in the desired state:

$$\operatorname{tr}_{\mathcal{XA}}(\rho_{\text{initial}}) = \operatorname{tr}_{\mathcal{XA}}(\rho_b) \qquad \text{for all } b \in \{0,1\} \tag{2}$$

She will then measure $\mathcal{X}$ register in the computational basis to obtain a. Therefore we have

$$\operatorname{tr}_{\mathcal{X}}(\rho_b) = \rho_{b0} + \rho_{b1} \qquad \text{for all } b \in \{0,1\} \ . \tag{3}$$

Note that this implies $\operatorname{tr}(\rho_{b0}) + \operatorname{tr}(\rho_{b1}) = 1$, so that in general the ρ_{ba} are not density matrices.

Now Bob checks ρ_{ba}. This gives rise to the following objective function for Alice's optimal cheating strategy:

$$\max \sum_{\beta\in\{0,1\}} \sum_{\alpha\in\{0,1\}} \Pr[b=\beta] \Pr[a=\alpha|b=\beta] \cdot \left(\delta_{\alpha\beta} \Pr[\rho_{\beta\alpha} \text{ passes}] - v \Pr[\rho_{\beta\alpha} \text{ fails}]\right) \tag{4}$$

Here the Kronecker-Delta $\delta_{\alpha\beta} = 1$ if and only if $\alpha = \beta$ measures whether Alice managed to get a and b to match. Maximizing (4) is equivalent to maximizing

$$\max \sum_{\beta\in\{0,1\}} \sum_{\alpha\in\{0,1\}} \Pr[b=\beta] \Pr[a=\alpha|b=\beta] \cdot \Pr[\rho_{\beta\alpha} \text{ passes}] \left(\delta_{\alpha\beta} + v\right) \ . \tag{5}$$

Bob plays honestly, therefore $\Pr[b=0] = \Pr[b=1] = \frac{1}{2}$. Moreover, $\Pr[a=\alpha|b=\beta] = \operatorname{tr}(\rho_{\beta,\alpha})$ and

$$\Pr[\rho_{\beta\alpha} \text{ passes}] = \operatorname{tr}\left(|\psi_\alpha\rangle\langle\psi_\alpha| \frac{\rho_{\beta,\alpha}}{\operatorname{tr}(\rho_{\beta,\alpha})}\right) \ .$$

Hence, $\Pr[a=\alpha|b=\beta] \Pr[\rho_{\beta\alpha} \text{ passes}] = \langle\psi_\alpha|\rho_{\beta\alpha}|\psi_\alpha\rangle$. Substituting this into (5) and discarding the constant factor $\frac{1}{2}$ gives the final objective function

$$\max \sum_{\beta\in\{0,1\}} \sum_{\alpha\in\{0,1\}} \langle\psi_\alpha|\rho_{\beta\alpha}|\psi_\alpha\rangle \left(\delta_{\alpha\beta} + v\right) \ . \tag{6}$$

We now proceed to constructing the dual of the SDP formed by the objective function (6) together with the constraints

(1), (2), and (3). The Lagrange ansatz is

$$\max_{\mathcal{P}} \inf_{\mathcal{D}} \sum_{a,b\in\{0,1\}} \mathrm{tr}\left((\delta_{ab}+v)|\psi_a\rangle\langle\psi_a|\rho_{ba}\right) + \sum_{b\in\{0,1\}} \mathrm{tr}\left(L_b\left(\mathrm{tr}_{\mathcal{X}}(\rho_b)-\rho_{b0}-\rho_{b1}\right)\right) - \sum_{b\in\{0,1\}} \mathrm{tr}\left(M_b\,\mathrm{tr}_{\mathcal{XA}}(\rho_b-\rho_{\text{initial}})\right) - \mathrm{tr}(\lambda(\rho_{\text{initial}}-\mathbb{1})) \quad (7)$$

where $\mathcal{P}$ are the primal variables as before, i.e.,

$$\mathcal{P} = \{(\rho_{\text{initial}},\rho_0,\rho_1,\rho_{00},\rho_{01},\rho_{10},\rho_{11}) : \rho_{\text{initial}},\rho_0,\rho_1 \in S(\mathcal{X}\otimes\mathcal{A}\otimes\mathcal{B}), \rho_{00},\rho_{01},\rho_{10},\rho_{11}\in S(\mathcal{A}\otimes\mathcal{B})\}$$

and the dual variables (Lagrange multipliers) are

$$\mathcal{D} = \{(L_0,L_1,M_0,M_1,\lambda) : L_0,L_1\in H(\mathcal{A}\otimes\mathcal{B}), M_0,M_1\in H(\mathcal{B}),\lambda\in\mathbb{R}\} \ .$$

Here $H(\mathcal{V})$ and $S(\mathcal{V})$ denote the Hermitian and semidefinite matrices, respectively, operating on the linear space $\mathcal{V}$. Collecting the primal variables in (7), we get for ρ_{initial}

$$\mathrm{tr}\left((M_0+M_1-\lambda\mathbb{1}_{\mathcal{B}})\,\rho_{\text{initial}}\right) \ .$$

For ρ_b, $b\in\{0,1\}$, we obtain

$$\mathrm{tr}\left((L_b-(\mathbb{1}_{\mathcal{A}}\otimes M_b))\,\mathrm{tr}_{\mathcal{X}}(\rho_b)\right) \ .$$

For ρ_{ba}, $a,b\in\{0,1\}$, we obtain

$$\mathrm{tr}\left((-L_b+(\delta_{ab}+v)|\psi_a\rangle\langle\psi_a|)\,\rho_{ba}\right) \ .$$

The terms in (7) not involving primal variables are just λ. Hence, the following dual SDP will give an upper bound on the optimal value of our primal SDP:

$$\begin{aligned}
&\text{minimize } \lambda \text{ subject to} && (8)\\
&M_0+M_1\le\lambda\mathbb{1}_{\mathcal{B}} && (9)\\
&L_b\le\mathbb{1}_{\mathcal{A}}\otimes M_b \ \text{ for all } b\in\{0,1\} && (10)\\
&(v+\delta_{ab})|\psi_a\rangle\langle\psi_a|\le L_b \ \text{ for all } a,b\in\{0,1\} && (11)\\
&(L_0,L_1,M_0,M_1,\lambda)\in\mathcal{D} && (12)
\end{aligned}$$

We now construct a feasible solution for the dual SDP. We restrict our attention to M_0 and M_1 of the form

$$M_0=\begin{pmatrix} m_0 & & \\ & m_1 & \\ & & m_2\end{pmatrix} \text{ and } M_1=\begin{pmatrix} m_1 & & \\ & m_0 & \\ & & m_2\end{pmatrix}$$

for some $m_0, m_1, m_2\in\mathbb{R}$ with

$$m_0\ge 0\ , \qquad m_1\ge 0\ , \qquad m_2=\frac{1}{2}(m_0+m_1) \ . \quad (13)$$

Moreover, we also impose the restriction $L_b=\mathbb{1}_{\mathcal{A}}\otimes M_b$ for $b\in\{0,1\}$. Since then $\lambda\ge m_0+m_1$, our goal reduces to minimizing m_0 and m_1 subject to

$$\begin{aligned}
L_0-(v+1)|\psi_0\rangle\langle\psi_0|&\ge 0 && (14)\\
L_0-v|\psi_1\rangle\langle\psi_1|&\ge 0 && (15)\\
L_1-v|\psi_0\rangle\langle\psi_0|&\ge 0 && (16)\\
L_1-(v+1)|\psi_1\rangle\langle\psi_1|&\ge 0 \ . && (17)
\end{aligned}$$

Constraints (14) and (17) are satisfied if

$$\begin{aligned}
m_0&\ge(v+1)\delta && (18)\\
m_2&\ge(v+1)(1-\delta) && (19)\\
m_0m_2&\ge(v+1)(1-\delta)m_0+(v+1)\delta m_2 \ . && (20)
\end{aligned}$$

Similarly, Constraints (15) and (16) require that

$$\begin{aligned}
m_1&\ge v\delta && (21)\\
m_2&\ge v(1-\delta) && (22)\\
m_1m_2&\ge v(1-\delta)m_1+v\delta m_2 \ . && (23)
\end{aligned}$$

A solution to the system (13),(18)-(23) is

$$\begin{aligned}
m_0&=\tfrac{1}{2}(1+v)\Big(2-\delta(1+2v) +\sqrt{4-4\delta+(\delta+2\delta v)^2}\Big)\\
m_1&=\tfrac{1}{2}v\Big(2+\delta+2\delta v-\sqrt{4-4\delta+(\delta+2\delta v)^2}\Big) \ .
\end{aligned}$$

From this and the definition of δ, we get that there is feasible solution of the dual SDP with

$$\begin{aligned}
\lambda&=m_0+m_1\\
&=2v+\frac{-1+\sqrt{v}-2v+\sqrt{1-2\sqrt{v}+5v+4v^2}}{\sqrt{v}}\\
&\le 2v+1+\frac{1}{4\sqrt{v}} \ .
\end{aligned}$$

From the earlier transformations of the primal objective function, it follows that the optimal expected payoff of Alice is bounded from above by $\frac{1}{2}\lambda-v\le\frac{1}{2}+\frac{1}{8\sqrt{v}}$. □

3. The multiparty model

3.1. Adversaries

In this work, we assume computationally unbounded adversaries. However, they have to obey quantum mechanics and cannot read the private memory of the honest players (but they can communicate secretly with each other). Moreover, we assume that they can only access the message space in between rounds or when according to the protocol it is their turn to send a message.

3.2. The broadcast channel

A classical broadcast channel allows one party to send a classical bit to all the other players. In the quantum setting this would mean that a qubit would be sent to all the other players. However, when there are more than two players in total we would have to *clone* or *copy* the qubit in order to send it to the other players. Even if the sender knows a classical preparation of the state he wants to send, we cannot allow him to prepare copies because he may be a cheater and send different states to different parties. It is well known that it is impossible to clone a qubit [19], because cloning is not a unitary operation. This means that we will have to take a slightly different approach. Quantum broadcast channels have been studied in an information-theoretic context before [5, 18] but not in the presence of faulty or malicious parties.

Our quantum broadcast channel works as follows. Suppose there are k players in total and that one player wants to broadcast a qubit that is in the state $\alpha|0\rangle + \beta|1\rangle$. What will happen is that the channel will create the k-qubit state $\alpha|0^k\rangle + \beta|1^k\rangle$ and send one of the k qubits to each of the other players. The state $\alpha|0^k\rangle + \beta|1^k\rangle$ can be easily created from $\alpha|0\rangle + \beta|1\rangle$ by taking $k-1$ fresh qubits in the state $|0^{k-1}\rangle$. This joint state can be written as $\alpha|0^k\rangle + \beta|10^{k-1}\rangle$. Next we flip the last $k-1$ bits conditional on the first bit being a 1, thus obtaining the desired state $\alpha|0^k\rangle + \beta|1^k\rangle$. This last operation can be implemented with a series of controlled-not operations. Note that this state is not producing k copies of the original state, which would be the k-fold product state $(\alpha|0\rangle + \beta|1\rangle) \otimes \ldots \otimes (\alpha|0\rangle + \beta|1\rangle)$.

Theorem 5 *In the following sense, a quantum broadcast channel between k parties is comparable to models where the parties have a classical broadcast channel and/or pairwise quantum channels:*

- *If all parties are honest:*
  1. *One use of the quantum broadcast channel can be simulated with $2(k-1)$ uses of pairwise quantum channels.*
  2. *One use of a classical broadcast channel can be simulated with one use of the quantum broadcast channel.*
  3. *One use of a pairwise quantum channel can be simulated by $k+1$ uses of the quantum broadcast channel.*
- *If all but one of the parties are dishonest, using one of the simulations above in place of the original communication primitive does not confer extra cheating power.*

Proof. We first give the simulations and argue that they work in case all players are honest.

1. The sender takes $k-1$ fresh qubits in state $|0^k\rangle$. He applies $k-1$ times CNOT where the subsystem to be broadcast is the control of the CNOT and the fresh qubits are the destination. He then sends each of the $k-1$ qubits via the pairwise quantum channels to the $k-1$ other parties. Each recipient j flips a (private) classical random bit r_j and if $r_j = 1$ performs a σ_z phase flip on the received qubit. Here $\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ is the Pauli matrix that multiplies the relative phase between the $|0\rangle$ and the $|1\rangle$ state by -1. He then sends r_j back to the sender. The sender computes the parity of the r_j and if it is odd, he performs a σ_z phase flip on his part of the broadcast state, thus restoring the correct relative phase. (This randomization is a countermeasure; its utility is explained below.)

2. When the sender wants to broadcast bit $b \in \{0,1\}$, he uses the quantum broadcast channel on qubit $|b\rangle$. The recipients immediately measure their qubit in the computational basis to obtain the classical bit.

3. The quantum broadcast channel can be used to create an EPR pair $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ between two players P_i and P_j with the assistance of the other $(k-2)$ players. i and j are determined by the protocol.

   First one player broadcasts the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, resulting in the k qubit state $|\varphi\rangle = \frac{1}{\sqrt{2}}(|0^k\rangle + |1^k\rangle)$. Now one after the other, the $k-2$ remaining players perform a Hadamard transformation on their qubit, measure it in the computational basis, and broadcast the classical result. Next, if P_i receives a 1 he applies a phase flip σ_z to his part of $|\varphi\rangle$ (P_j does nothing). After this operation, $|\varphi\rangle$ will be an EPR state between P_i and P_j unentangled with the other $k-2$ parties. Using a shared EPR pair, a protocol called *teleportation* [7] can be used to simulate a private quantum channel between P_i and P_j. Teleportation requires the transmission of two bits of classical information.

For the case of all but one party being dishonest:

1. If the sender is honest, the recipients obtain exactly the same subsystems as for the quantum broadcast channel.

   If one of the recipients is honest, he may receive an arbitrary quantum subsystem up to the randomized relative phase. However, exactly the same can be achieved with a quantum broadcast channel with $k-1$ cheating parties, who each perform a Hadamard transformation on their subsystem followed by a measurement in the computational basis.

2. If the sender is honest, all recipients obtain the same computational-basis state.

If one of the recipients is honest, he obtains a classical bit that is possibly randomized in case the dishonest sender does not broadcast a basis state. Since the sender can flip a coin himself, this does not give more cheating power.

3. If the sender is honest, we can assume without loss of generality that all cheating action is done after the EPR pair has been established, because the (merged) cheaters can easily recreate the original broadcast state and also compensate any phase flipping of the honest sender. However, after the EPR pair has been established, the sender unilaterally performs his part of the teleportation circuit and measurements and sends the two bits of classical information. So the most general cheating action is to apply a quantum operation after the reception of the two classical bits. Furthermore, we can even assume that the cheating action is done *after* the correction circuit of teleportation (this is similar to the teleportation of quantum gates [10]) and, hence, amounts to cheating on a pairwise quantum channel.

   If one of the recipients is honest, the best the cheaters can aim for is to give an arbitrary quantum state to the honest recipient. This they can also achieve over a pairwise quantum channel.

□

4. Multiparty quantum protocols

We will first consider the case of only one good player (i.e., $g = 1$) among k players and later extend our results to general g.

One honest player. We need to construct a protocol with bias $\frac{1}{2} - \Omega(\frac{1}{k})$. Before proceeding to our actual protocol, let us consider a simple protocol which trivially extends the previous work in the two-party setting, but does not give us the desired result. The protocols is as follows: player 1 flips a random coin with player 2, player 3 flips a random coin with player 4 and so forth. In each pair, the player with the higher id wins if the coin is 1 and the one with the lower id if the coin is 0. The winners repeat the procedure. With each round of the tournament, half of the remaining players are eliminated (if there is an odd number of players at any moment, the one with the highest id advances to the next round). When there are only two players left, the coin they flip becomes the output of the protocol. (Above we assume we have private point-to-point quantum channels and a classical broadcast channel, which is justified by Theorem 5.)

The elimination step can be implemented using the weak two-party coin-flipping protocol by Spekkens and Rudolph [17] and the last round by the the strong two-party coin-flipping protocol by Ambainis [4]. If there is only one good player, the probability that he makes it to the last round is

$$\left(1 - \frac{1}{\sqrt{2}}\right)^{\lceil -1+\log k \rceil} ;$$

in this case, the probability that the bad players can determine the output coin is $\frac{3}{4}$. In case the good player gets eliminated, the bad players can completely determine the coin. Hence, the overall probability that the bad players can determine the coin is

$$1 - \frac{1}{4}\left(1 - \frac{1}{\sqrt{2}}\right)^{\lceil -1+\log k \rceil} \leq 1 - \frac{1}{4k^{1.78}} ,$$

which corresponds to bias

$$\frac{1}{2} - \Omega\left(\frac{1}{k^{1.78}}\right) .$$

Using the protocol by Mochon [14] improves the exponent slightly to ≈ 1.7 but not all the way to 1. To improve the bound above to the desired value $\frac{1}{2} - \Omega(\frac{1}{k})$, we will use our coin-flipping protocol with penalty from Section 2. The idea is that in normal quantum coin-flipping protocols for two parties, there are three outcomes for a given player: "win," "lose," and "abort." Looking at the elimination tournament above, if an honest player loses a given coin flipping round, he does not complain and the bad player wins the game. However, if the honest player detects cheating, he can and will abort the entire process, which corresponds to the failure of the dishonest players to fix the coin. Of course, if there are few elimination rounds left, bad players might be willing to risk the abort if they gain significant benefits in winning the round. However, if the round number is low, abort becomes prohibitively expensive: a dishonest player might not be willing to risk it given there are plenty more opportunities for the honest player to "lose normally." Thus, instead of regular two-party coin-flipping protocols, which do not differentiate between losing and aborting, we can employ our protocol for coin flipping with penalty, where the penalties are very high at the original rounds, and eventually get lower towards the end of the protocol. Specific penalties are chosen in a way which optimizes the final bias we get, and allows us to achieve the desired bias $\frac{1}{2} - \Omega(\frac{1}{k})$.

Theorem 6 *There is a strong quantum coin-flipping protocol for k parties with bias at most $\frac{1}{2} - \frac{c}{k}$ for some constant c, even with $(k-1)$ bad parties.*

Proof. We assume that $k = 2^n$ for some $n > 0$, as it changes c by at most a constant factor. Let Q_v be the maximum expected win in a two-party protocol with penalty v. Consider the following protocol with n rounds numbered 1 to n.

At the beginning of round i, we have 2^{n+1-i} parties remaining. We divide them into pairs. Each pair performs the two-party coin-flipping protocol with penalty $(2^{n-i}-1)$. The party with the lower id plays Alice and wins if the outcome is 0; the party with the higher id plays Bob and wins if the outcome is 1. The winners proceed to round $(i+1)$.

At the beginning of round $(n-2)$, there are just 8 parties remaining. They perform three rounds of regular coin flipping with no penalty using the protocol of [4, 11], in which no cheater can determine the coin with probability more than $\frac{3}{4}$. This results in maximum probability of $\frac{63}{64}$ for fixing the outcome when there is at least one good player among the 8. The result of the last round is the result of our 2^n-party protocol.

Assume that the honest player has won the first $(n-j)$ coin flips and advanced to round $(j+1)$. Assume that the all other players in round $(j+1)$ are dishonest. Let P_j be the maximum probability with which (2^j-1) dishonest players can fix the outcome in this case.

Lemma 7

$$1-P_j \geq (1-P_{j-1})(1-Q_{2^{j-1}-1}) \qquad (24)$$

Proof. Let p_w, p_l, p_c be the probabilities of the honest player winning, losing and catching the other party cheating in the round $(j+1)$ of the protocol. Notice that $p_w + p_l + p_c = 1$. Then, the probability P_j of 2^j-1 dishonest parties fixing the coin is at most $p_l + p_w P_{j-1}$. (If the honest player loses, they win immediately. If he wins, they can still bias the coin in $j-1$ remaining rounds to probability at most P_{j-1}. If he catches his opponent cheating, he exits the protocol and the dishonest players have no more chances to cheat him.) Using $p_w = 1-p_l-p_c$, we have

$$\begin{aligned} P_j &\leq p_l + p_w P_{j-1} = P_{j-1} + (1-P_{j-1})p_l - P_{j-1}p_c \\ &= P_{j-1} + (1-P_{j-1})\left(p_l - \frac{P_{j-1}}{1-P_{j-1}}p_c\right) \qquad (25) \end{aligned}$$

Next, notice that $P_{j-1} \geq 1 - \frac{1}{2^{j-1}}$. This is because $2^{j-1}-1$ bad players could just play honestly when they face the good player and fix the coin flip if two bad players meet in the last round. Therefore, $\frac{P_{j-1}}{1-P_{j-1}} \geq 2^{j-1}-1$ and (25) becomes

$$P_j \leq P_{j-1} + (1-P_{j-1})(p_l - (2^{j-1}-1)p_c)$$

The term $p_l - (2^{j-1}-1)p_c$ is at most $Q_{2^{j-1}-1}$ because we can interpret it as the expected payoff of the cheater that plays with the honest player. Hence,

$$P_j \leq P_{j-1} + (1-P_{j-1})Q_{2^{j-1}-1} \ ,$$

which is equivalent to the desired (24). □

By applying the lemma inductively, we obtain

$$\begin{aligned} 1-P_n &\geq (1-P_8)\prod_{j=4}^{n}(1-Q_{2^{j-1}-1}) \\ &\geq \frac{1}{64}\prod_{j=4}^{n}(1-Q_{2^{j-1}-1}) \ . \end{aligned}$$

Using the bound in Theorem 2 we get

$$\begin{aligned} 1-P_n &\geq \frac{1}{64}\prod_{j=3}^{n-1}(1-Q_{2^j-1}) \\ &\geq \frac{1}{64}\prod_{j=3}^{n-1}\left(\frac{1}{2} - \frac{1}{\sqrt{2^j-1}}\right) \\ &\geq \frac{1}{8\cdot 2^n}\prod_{j=3}^{\infty}\left(1 - \frac{2}{\sqrt{2^j-1}}\right) \ . \end{aligned}$$

The last product is a positive constant. Therefore, for some constant $c>0$ we have $1-P_n \geq \frac{c}{2^n} = \frac{c}{k}$, which means that the bias is at most $\frac{1}{2} - \Omega(\frac{1}{k})$. □

Extending to many honest players. We can extend Theorem 6 to any number $g \geq 1$ of good players by using the classical lightest-bin protocol of Feige [8]. This protocol allows us to reduce the total number of players until a single good player is left without significantly changing the fraction of good players, after which we can run the quantum protocol of Theorem 6 to get the desired result. Specifically, Lemma 8 from [8] implies that starting from $g=\delta k$ good players out of k players, the players can (classically) select a sub-committee of $\mathrm{O}(\frac{1}{\delta}) = \mathrm{O}(\frac{k}{g})$ players containing at least one good player with probability at least $\frac{1}{2}$. Now this sub-committee can use the quantum protocol of Theorem 6 to flip a coin with bias $\frac{1}{2} - \Omega(\frac{g}{k})$, provided it indeed contains at least one honest player. But since the latter happens with probability at least $\frac{1}{2}$, the final bias is at most $\frac{1}{2} - \frac{1}{2}\cdot\Omega(\frac{g}{k}) = \frac{1}{2} - \Omega(\frac{g}{k})$, as desired.

5. Lower bound

5.1. The two-party bound

For completeness and to facilitate the presentation of our generalization, we reproduce here Kitaev's unpublished proof [12] that any two-party strong quantum coin-flipping protocol must have bias at least $\frac{1}{\sqrt{2}}$. The model here is that the two parties communicate over a quantum channel.

Definition 8 *Let $\mathcal{H} := \mathcal{A}\otimes\mathcal{M}\otimes\mathcal{B}$ denote the Hilbert space of the coin-flipping protocol composed of Alice's private space, the message space, and Bob's private space. A $2N$-round two-party coin-flipping protocol is a tuple*

$$\begin{aligned} &(U_{A,1},\ldots,U_{A,N},U_{B,1},\ldots,U_{B,N}, \\ &\Pi_{A,0},\Pi_{A,1},\Pi_{B,0},\Pi_{B,1}) \end{aligned}$$

where

- $U_{A,j}$ *is a unitary operator on* $\mathcal{A}\otimes\mathcal{M}$ *for* $j = 1, \ldots, N$,
- $U_{B,j}$ *is a unitary operator on* $\mathcal{M}\otimes\mathcal{B}$ *for* $j = 1, \ldots, N$,
- $\Pi_{A,0}$ *and* $\Pi_{A,1}$ *are projections from* $\mathcal{A}$ *onto orthogonal subspaces of* $\mathcal{A}$ *(representing Alice's final measurements for outcome* 0 *and* 1*, respectively),*
- $\Pi_{B,0}$ *and* $\Pi_{B,1}$ *are projections from* $\mathcal{B}$ *onto orthogonal subspaces of* $\mathcal{B}$ *(representing Bob's final measurements for outcome* 0 *and* 1*, respectively),*

so that for

$$|\psi_N\rangle := (1_\mathcal{A}\otimes U_{B,N})(U_{A,N}\otimes 1_\mathcal{B})(1_\mathcal{A}\otimes U_{B,N-1}) (U_{A,N-1}\otimes 1_\mathcal{B})\cdots(1_\mathcal{A}\otimes U_{B,1})(U_{A,1}\otimes 1_\mathcal{B})|0\rangle$$

holds

$$(\Pi_{A,0}\otimes 1_\mathcal{M}\otimes 1_\mathcal{B})|\psi_N\rangle = (1_\mathcal{A}\otimes 1_\mathcal{M}\otimes \Pi_{B,0})|\psi_N\rangle \tag{26}$$

$$(\Pi_{A,1}\otimes 1_\mathcal{M}\otimes 1_\mathcal{B})|\psi_N\rangle = (1_\mathcal{A}\otimes 1_\mathcal{M}\otimes \Pi_{B,1})|\psi_N\rangle \tag{27}$$

$$\|(\Pi_{A,0}\otimes 1_\mathcal{M}\otimes 1_\mathcal{B})|\psi_N\rangle\| = \|(\Pi_{A,1}\otimes 1_\mathcal{M}\otimes 1_\mathcal{B})|\psi_N\rangle\| \tag{28}$$

The first two conditions ensure that when Alice and Bob are honest, they both get the same value for the coin and the third condition guarantees that when Alice and Bob are honest, their coin is not biased. A player aborts if her or his final measurement does not produce outcome 0 or 1; of course, it is no restriction to delay this action to the end of the protocol.

Lemma 9 *Fix an arbitrary two-party quantum coin-flipping protocol. Let* p_{1*} *and* p_{*1} *denote the probability that Alice or Bob, respectively, can force the outcome of the protocol to be* 1 *if the other party follows the protocol. Denote by* p_1 *the probability for outcome* 1 *when there are no cheaters. Then* $p_{1*}p_{*1} \geq p_1$.

Hence, if $p_1 = \frac{1}{2}$, then $\max\{p_{1*}, p_{*1}\} \geq \frac{1}{\sqrt{2}}$. To prove Lemma 9, we construct the view of a run of the protocol from an honest Alice's point of view, with Bob wanting to bias the protocol towards 1. The problem of optimizing Bob's strategy is a semidefinite program.

Lemma 10 *The optimal strategy of Bob trying to force outcome* 1 *is the solution to the following SDP over the semidefinite matrices* $\rho_{A,0}, \ldots, \rho_{A,N}$ *operating on* $\mathcal{A}\otimes\mathcal{M}$*:*

$$\text{maximize} \quad \operatorname{tr}\left((\Pi_{A,1}\otimes 1_\mathcal{M})\rho_{A,N}\right) \text{ subject to} \tag{29}$$

$$\operatorname{tr}_\mathcal{M}(\rho_{A,0}) = |0\rangle\langle 0|_\mathcal{A} \tag{30}$$

$$\operatorname{tr}_\mathcal{M}(\rho_{A,j}) = \operatorname{tr}_\mathcal{M}(U_{A,j}\rho_{A,j-1}U_{A,j}^*) \quad (1 \leq j \leq N) \tag{31}$$

Proof. Alice starts with her private memory in state $|0\rangle_\mathcal{A}$ and we permit Bob to determine the $\mathcal{M}$ part of the initial state. Therefore all Alice knows is that initially, the space accessible to her is in state $\rho_{A,0}$ with $\operatorname{tr}_\mathcal{M}(\rho_{A,0}) = |0\rangle\langle 0|_\mathcal{A}$. Alice sends the first message, transforming the state to $\rho'_{A,0} := U_{A,1}\rho_{A,0}U_{A,1}^*$. Now Bob can do any unitary operation on $\mathcal{M}\otimes\mathcal{B}$ leading to $\rho_{A,1}$, so the only constraint is $\operatorname{tr}_\mathcal{M}(\rho_{A,1}) = \operatorname{tr}_\mathcal{M}(\rho'_{A,0})$. In the next round, honest Alice applies $U_{A,2}$, then Bob can do some operation that preserves the partial trace, and so forth. The probability for Alice outputting 1 is $\operatorname{tr}((\Pi_{A,1}\otimes 1_\mathcal{M})\rho_{A,N})$ because the final state for Alice is $\rho_{A,N}$ and she performs an orthogonal measurement on $\mathcal{A}$ with projections $\Pi_{A,0}$, $\Pi_{A,1}$, and $1_\mathcal{A} - \Pi_{A,0} - \Pi_{A,1}$ (which represents "abort"). □

Lemma 11 *The dual SDP to the primal SDP in Lemma 10 is*

$$\text{minimize} \quad \langle 0|Z_{A,0}|0\rangle \text{ subject to} \tag{32}$$

$$Z_{A,j}\otimes 1_\mathcal{M} \geq U_{A,j+1}^*(Z_{A,j+1}\otimes 1_\mathcal{M})U_{A,j+1} \quad (\text{for all } j : 0 \leq j \leq N-1) \tag{33}$$

$$Z_{A,N} = \Pi_{A,1} \tag{34}$$

over the Hermitian matrices $Z_{A,0}, \ldots Z_{A,N}$ *operating on* $\mathcal{A}$*.*

Proof. We form the dual of the SDP in Lemma 10 as follows: it is equivalent to maximizing over the $\rho_{A,j}$ the minimum of

$$\operatorname{tr}((\Pi_{A,1}\otimes 1_\mathcal{M})\rho_{A,N}) - \operatorname{tr}(Z_{A,0}(\operatorname{tr}_\mathcal{M}(\rho_{A,0}) - |0\rangle\langle 0|_\mathcal{M})) - \sum_{j=1}^{N}\operatorname{tr}(Z_{A,j}\operatorname{tr}_\mathcal{M}(\rho_{A,j} - U_{A,j}\rho_{A,j-1}U_{A,j}^*)) \tag{35}$$

subject to the operators $Z_{A,j}$ on $\mathcal{M}$ being Hermitian (for $0 \leq j \leq N$). In the sum above, the terms containing $\rho_{A,j}$ for $0 \leq j < N$ are

$$-\operatorname{tr}(Z_{A,j}\operatorname{tr}_\mathcal{M}(\rho_{A,j})) + \operatorname{tr}(Z_{A,j+1}\operatorname{tr}_\mathcal{M}(U_{A,j+1}\rho_{A,j}U_{A,j+1}^*)) \ ,$$

which equals

$$\operatorname{tr}\Big(\big(-(Z_{A,j}\otimes 1_\mathcal{M}) + U_{A,j+1}^*(Z_{A,j+1}\otimes 1_\mathcal{M})U_{A,j+1}\big)\rho_{A,j}\Big) \ .$$

Since this term must be non-positive, we arrive at the inequality (33).

For $j = N$, we obtain the dual equality constraint (34) and the dual objective function becomes the only summand of (35) that does not involve any $\rho_{A,j}$. □

Proof of Lemma 9. Let $Z_{A,j}$ and $Z_{B,j}$ $(0 \leq j \leq N)$ denote the optimal solutions for the dual SDPs for a cheating Bob and a cheating Alice, respectively. For each j, $0 \leq j \leq N$, let

$$|\psi_j\rangle := (1_\mathcal{A}\otimes U_{B,j})(U_{A,j}\otimes 1_\mathcal{B})\cdots (1_\mathcal{A}\otimes U_{B,1})(U_{A,1}\otimes 1_\mathcal{B})|0\rangle$$

denote the state of the protocol in round j when both parties are honest. Let $F_j := \langle\psi_j|(Z_{A,j} \otimes 1_{\mathcal{M}} \otimes Z_{B,j})|\psi_j\rangle$. We claim

$$p_{1*}p_{*1} = F_0 \tag{36}$$
$$F_j \geq F_{j+1} \qquad (0 \leq j < N) \tag{37}$$
$$F_N = p_1. \tag{38}$$

Combining (36)–(38), we obtain the desired $p_{1*}p_{*1} \geq p_1$. We now proceed to prove these claims.

Note that the primal SDP from Lemma 10 is strictly feasible: Bob playing honestly yields a feasible solution that is strictly positive. The strong-duality theorem of semidefinite programming states that in this case, the optimal value of the primal and the dual SDPs are the same, and therefore $p_{1*} = \langle 0|_{\mathcal{A}} Z_{A,0}|0\rangle_{\mathcal{A}}$ and $p_{*1} = \langle 0|_{\mathcal{B}} Z_{B,0}|0\rangle_{\mathcal{B}}$ and

$$\begin{aligned} p_{1*}p_{*1} &= \langle 0|_{\mathcal{A}} Z_{A,0}|0\rangle_{\mathcal{A}} \cdot \langle 0|_{\mathcal{M}} 1_{\mathcal{M}}|0\rangle_{\mathcal{M}} \cdot \langle 0|_{\mathcal{B}} Z_{B,0}|0\rangle_{\mathcal{B}} \\ &= \langle 0|(Z_{A,0} \otimes 1_{\mathcal{M}} \otimes Z_{B,0})|0\rangle = F_0. \end{aligned}$$

The inequalities (37) hold because of the constraints (33). Equality (38) holds because by constraint (34) we have

$$\begin{aligned} &\langle\varphi|(Z_{A,N} \otimes 1_{\mathcal{M}} \otimes Z_{B,N})|\varphi\rangle = \\ &\qquad \|(\Pi_{A,1} \otimes 1_{\mathcal{M}} \otimes 1_{\mathcal{B}})(1_{\mathcal{A}} \otimes 1_{\mathcal{M}} \otimes \Pi_{B,1})|\varphi\rangle\|^2 \end{aligned}$$

for any $|\varphi\rangle$; $|\psi_N\rangle$ is the final state of the protocol when both players are honest, so by equation (27),

$$\begin{aligned} &\|(\Pi_{A,1} \otimes 1_{\mathcal{M}} \otimes 1_{\mathcal{B}})(1_{\mathcal{A}} \otimes 1_{\mathcal{M}} \otimes \Pi_{B,1})|\psi_N\rangle\|^2 \\ &\qquad = \|(\Pi_{A,1} \otimes 1_{\mathcal{M}} \otimes 1_{\mathcal{B}})|\psi_N\rangle\|^2 = p_1. \end{aligned}$$

□

5.2. More than two parties

We will now extend Kitaev's lower bound to k parties. As with the upper bounds, we first start with a single honest player ($g = 1$), and then extend the result further to any g.

Theorem 12 *Any strong quantum coin-flipping protocol for k parties has bias at least*

$$\frac{1}{2} - \frac{\ln 2}{k} - \mathrm{O}\left(\frac{1}{k^2}\right)$$

if it has to deal with up to $(k-1)$ bad parties.

We consider the model of private pairwise quantum channels between the parties; by Theorem 5 the results immediately carry over to the quantum broadcast channel.

Definition 13 *Let $\mathcal{H} := \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_k \otimes \mathcal{M}$ denote the Hilbert space composed of the private spaces of k parties and the message space. An N-round k-party coin-flipping protocol is a tuple*

$$(i_1, \ldots, i_N, U_1, \ldots, U_N, \Pi_{1,0}, \Pi_{1,1}, \ldots, \Pi_{k,0}, \Pi_{k,1})$$

where

- *i_j with $1 \leq i_j \leq k$, $1 \leq j \leq N$, indicates whose turn it is to access the message space in round j,*
- *U_j is a unitary operator on $\mathcal{A}_{i_j} \otimes \mathcal{M}$ for $j = 1, \ldots, N$,*
- *for $1 \leq i \leq k$, $\Pi_{i,0}$ and $\Pi_{i,1}$ are projections from $\mathcal{A}_i$ to orthogonal subspaces of $\mathcal{A}_i$ (representing the measurement that party i performs to determine outcome 0 or 1, respectively),*

so that for $|\psi_N\rangle := \tilde{U}_{i_N} \cdots \tilde{U}_{i_1}|0\rangle$ and each pair $1 \leq i < i' \leq k$ and any $b \in \{0,1\}$ holds

$$\tilde{\Pi}_{i,b}|\psi_N\rangle = \tilde{\Pi}_{i',b}|\psi_N\rangle \tag{39}$$
$$\|\tilde{\Pi}_{i,b}|\psi_N\rangle\| = \|\tilde{\Pi}_{i,1-b}|\psi_N\rangle\|. \tag{40}$$

Here $\tilde{U}_j$ denotes the extension of U_j to all of $\mathcal{H}$ that acts as identity on the tensor factors $\mathcal{A}_{i'}$ for $i' \neq i_j$; $\tilde{\Pi}_{i,b} := (1_{\mathcal{A}_1} \otimes \cdots \otimes 1_{\mathcal{A}_{i-1}} \otimes \Pi_{i,b} \otimes 1_{\mathcal{A}_{i+1}} \otimes \cdots \otimes 1_{\mathcal{A}_k})$ is the extension of $\Pi_{i,b}$ to $\mathcal{H}$.

Lemma 14 *Fix an arbitrary quantum coin flipping protocol. For $b \in \{0,1\}$, let p_b be the probability of outcome b in case all players are honest. Let $p_{i,b}$ denote the probability that party i can be convinced by the other parties that the outcome of the protocol is $b \in \{0,1\}$. Then*

$$p_{1,b} \cdot \ldots \cdot p_{k,b} \geq p_b$$

Proof of Lemma 14. The optimal strategy for $k-1$ bad players trying to force outcome 1 is the solution to the SDP from Lemma 10 where all the cheating players are merged into a single cheating player.

Let $(Z_{i,j})_{0 \leq j \leq N}$ denote the optimal solution for the dual SDP for good player i, $1 \leq i \leq k$. For each j, $0 \leq j \leq N$, let $|\psi_j\rangle := \tilde{U}_j \cdots \tilde{U}_1|0\rangle$ denote the state of the protocol in round j when all parties are honest. Let $F_j := \langle\psi_j|(Z_{1,j} \otimes \cdots \otimes Z_{k,j} \otimes 1_{\mathcal{M}})|\psi_j\rangle$. By a similar argument as in the proof of Lemma 9, we have

$$p_{1,1} \cdot \ldots \cdot p_{k,1} = F_0 \tag{41}$$
$$F_j \geq F_{j+1} \qquad (0 \leq j < N) \tag{42}$$
$$F_N = p_1 \tag{43}$$

Hence, $p_{1,1} \cdot \ldots \cdot p_{k,1} \geq p_1$. Repeating the argument with the cheaters aiming for outcome 0 completes the proof. □

Theorem 12 is an immediate consequence.

Proof of Theorem 12. Using the notation of Lemma 14, we have $p_0 = \frac{1}{2}$. Let $q = \max_i p_{i,0}$ denote the maximum probability of any player forcing output 0. By Lemma 14, $q^k \geq p_{1,0} \cdot \ldots \cdot p_{k,0} \geq \frac{1}{2}$, from which follows that

$$q \geq \left(\frac{1}{2}\right)^{1/k} \geq 1 - \frac{\ln 2}{k} - \mathrm{O}\left(\frac{1}{k^2}\right) \ .$$

By Theorem 5 this result applies both to private pairwise quantum channels and the quantum broadcast channel. □

Extending to many honest players. Extension to any number of honest players follows almost immediately from Theorem 12. Indeed, take any protocol Π for k parties tolerating $(k-g)$ cheaters. Arbitrarily partition our players into $k' = \frac{k}{g}$ groups and view each each as one "combined player." We get an induced protocol Π' with k' "super-players" which achieves at least the same bias ε as Π, and can tolerate up to $(k'-1)$ bad players. By Theorem 12, $\varepsilon \geq \frac{1}{2} - \mathrm{O}(\frac{1}{k'}) = \frac{1}{2} - \mathrm{O}(\frac{g}{k})$.

Acknowledgements

We thank L. Fortnow and J.-H. Hoepman for useful discussions.

References

[1] D. Aharonov, A. Y. Kitaev, and N. Nisan. Quantum circuits with mixed states. In *Proceedings of 30th ACM STOC*, pages 10–20, 1998, quant-ph/9806029.

[2] D. Aharonov, A. Ta-Shma, U. Vazirani, and A. Yao. Quantum bit escrow. In *Proceedings of 32nd ACM STOC*, pages 705–714, 2000, quant-ph/0004017.

[3] A. Ambainis. Lower bound for a class of weak quantum coin flipping protocols. 2002, quant-ph/0204063.

[4] A. Ambainis. A new protocol and lower bounds for quantum coin flipping. *Journal of Computer and Systems Sciences*, 68(2):398–416, 2004, quant-ph/0204022. Earlier version in STOC 2001.

[5] H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa, and B. Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76(15):2818–2821, 1996.

[6] M. Ben-Or and N. Linial. Collective coin-flipping. In *Randomness and Computation*, pages 91–115, 1990.

[7] C. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70:1895–1899, 1993.

[8] U. Feige. Noncryptographic selection protocols. In *Proceedings of 40th IEEE FOCS*, pages 142–152, 1999.

[9] L. Goldenberg, L. Vaidman, and S. Wiesner. Quantum gambling. *Physical Review Letters*, 82:3356–3359, 1999.

[10] D. Gottesman and I. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single qubit operations. *Nature*, 402(6760):390–393, 1999.

[11] J. Kerenidis and A. Nayak. Weak coin flipping with small bias. *Information Processing Letters*, 89:131–135, 2004, quant-ph/0206121.

[12] A. Y. Kitaev. Quantum coin-flipping. Talk at QIP 2003 (slides and video at MSRI), December 2002.

[13] M. Laurent and F. Rendl. Semidefinite programming and integer programming. In K. Aardal, G. Nemhauser, and R. Weismantel, editors, *Discrete Optimization*, Handbooks in operations research and management science. Elsevier, 2004. `http://www.optimization-online.org/DB_HTML/2002/12/585.html`.

[14] C. Mochon. Quantum weak coin-flipping with bias of 0.192. Technical Report CALT-68-2486, California Institute of Technology, 2004, quant-ph/0403193.

[15] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[16] M. Saks. A robust noncryptographic protocol for collective coin flipping. *SIAM J. Discrete Math.*, 2(2):240–244, 1989.

[17] R. W. Spekkens and T. Rudolph. A quantum protocol for cheat-sensitive weak coin flipping. *Physical Review Letters*, 89:227901, 2002, quant-ph/0202118.

[18] R. Wilmink. *Quantum Broadcast Channels and Cryptographic Applications* PhD thesis, Universität Bielefeld, 2002.

[19] W. K. Wootters and W. H. Zurek. A single quantum cannot be copied. *Nature*, 299:802–803, 1982.