Quantum protocols for Rabin oblivious transfer

Erika Andersson*[3], Akshay Bansal†[1], James T. Peat‡[3], Jamie Sikora§[1], and Jiawei Wu¶[2]

[1]Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA
[2]Centre for Quantum Technologies, National University of Singapore, Singapore
[3]SUPA, Institute of Photonics and Quantum Sciences, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, Scotland, UK

July 8, 2025

Abstract

Rabin oblivious transfer is the cryptographic task where Alice wishes to receive a bit from Bob but it may get lost with probability 1/2. In this work, we provide protocol designs which yield quantum protocols with improved security. Moreover, we provide a constant lower bound on any quantum protocol for Rabin oblivious transfer. To quantify the security of this task with asymmetric cheating definitions, we introduce the notion of cheating advantage which may be of independent interest in the study of other asymmetric cryptographic primitives.

1 Introduction

Cryptographic tasks where an adversary remains oblivious to the input or the data of other parties are popularly studied under the banner of *oblivious transfer*. Several variants of oblivious transfer tasks exist in the literature today with varying goals and security definitions. A commonly studied variant of oblivious transfer, known as 1-out-of-2 oblivious transfer (OT), is a cryptographic primitive involving two parties, Alice and Bob, where Bob has two bits of information, and Alice wishes to learn one of them based on her choice of an index bit. The security of OT protocols is usually a two-fold evaluation where one needs to protect honest Bob from a cheating Alice who may wish to obtain more information that the protocol allows (perhaps by learning both bits), meanwhile Bob may wish to learn Alice's input, that is, which bit she wishes to learn (see Definition A.1 for a formal definition of 1-out-of-2 OT and its underlying security notions).

Oblivious transfer, its variants, and other simple cryptographic tasks, are referred to as primitives as they could be used as building blocks for other useful tasks such as coin flipping [Kit02,

*E.Andersson@hw.ac.uk
†akshaybansal14@gmail.com
‡jtp2000@hw.ac.uk
§sikora@vt.edu
¶constchar0212@gmail.com

Moc07, ARV21], bit commitment [CK11], oblivious circuit evaluation [BBCS92, Cré94], and even general multiparty computation [Kil88, KOS16].

In this work, we also consider the task of coin flipping where Alice and Bob wish to flip a coin but have a different preferred choices for the outcome. It has been shown that quantum protocols exist for coin flipping up to arbitrary levels of security [Moc07] with explicit constructions given in [ARV21, ARVW24]. This has been generalized to construct secure z-balanced coin flipping in [CK09] where the two parties output heads with probability z and tails with the remaining probability of $(1-z)$ securely. Moreover, optimal protocols for the *semi-honest* version of 1-out-of-2 OT were developed in [CGS16], which also provides an the best known lower bound on the security of 1-out-of-2 oblivious transfer as a corollary.

The focus of this work is to examine quantum protocols for a variant of oblivious transfer known as Rabin oblivious transfer. It has been shown that Rabin oblivious transfer is equivalent to 1-out-of-2 oblivious transfer [Cré88], and that 1-out-of-2 oblivious transfer is universal [Kil88]. However, there are only a few papers that studied Rabin oblivious transfer in a quantum setting [BS25, SPK$^+$24] where the aim is to study its information-theoretic security. The goal of our work is to provide improved bounds on its overall security by finding better protocols as well as lower bounds on how secure they can be.

1.1 Rabin oblivious transfer

In this work, we consider a variant of oblivious transfer, known as *Rabin oblivious transfer*, denoted by ROT. This task was first proposed by Rabin in [Rab79, Rab05], whence the name. The task was later revisited in [FMR96] to give a protocol that guarantees security under certain computational hardness assumptions. ROT is the cryptographic task which accomplishes the following:

- At the beginning of the protocol, Bob outputs a uniformly random bit $y \in \{0,1\}$.
- At the end of the protocol, Alice receives y with probability $1/2$, else the bit is lost.

The security of ROT protocols is defined in a two-fold manner. Dishonest Alice wishes to maximize her chances of learning the bit y while dishonest Bob wishes to maximize his chances of successfully guessing whether Alice *asserts* the receipt of the data or the data loss event. Note that if Bob cheats, Alice might think that she learned the bit y, but the learned bit could very well be different from y. Thus, a cheating Bob only cares if Alice *thinks* she correctly learned y. Note that in contrast to the analysis described in [FMR96], we work under the regime of information-theoretic security where the adversaries are not bounded under any computational assumptions.

We now provide a more formal definition.

Definition 1.1 (Rabin OT)**.** Rabin OT (ROT) is the cryptographic task between two parties (Alice and Bob) such that:

- Bob chooses $y \in \{0,1\}$ uniformly at random.
- At the end of the protocol, Alice outputs a bit $g \in \{0,1\}$ with probability $1/2$ (which corresponds to her asserting that she learned the bit y) or NULL with probability $1/2$ (which corresponds to her losing the bit).

The completeness and security notions for an ROT protocol are below.

- *Completeness:* If both Alice and Bob are honest, neither party aborts. Alice outputs $g = y$ with probability $1/2$ or NULL with probability $1/2$.

- *Cheating Alice:* If Bob is honest, then dishonest Alice's cheating probability is defined as

$$P_A^* = \max\{\Pr[\text{Alice successfully learns } y]\},$$

where the maximum is taken over all cheating strategies of Alice. Note that $P_A^* \geq 3/4$ as Alice can simply choose to follow the protocol to learn y with probability $1/2$ and randomly guess y in the event of NULL.

- *Cheating Bob*: If Alice is honest, then dishonest Bob's cheating probability is defined as

$$P_B^* = \max\{\Pr[\text{Bob correctly guesses whether Alice } \textit{asserts} \text{ the receipt of } y \text{ or NULL}]\},$$

where the maximum is taken over all cheating strategies of Bob. More formally, if we denote Alice's assertion by bit a, then $a = 0$ denotes that Alice asserted the receipt of y, while $a = 1$ denote that she asserted the receipt of NULL. Here, P_B^* could be alternatively expressed as the maximum probability with which Bob correctly learns a.

We emphasize that Alice asserting (possibly privately) the receipt of y (by outputting g) does not necessarily imply the receipt of the actual y, i.e., it could be the case that $g \neq y$, or even that y does not even exist. Note that $P_B^* \geq 1/2$ as Bob can simply choose to follow the protocol honestly and make a uniformly random guess to whether Alice asserts y or NULL.

As the cheating notions of Alice and Bob in the ROT are asymmetric, especially with respect to the lower bounds of $3/4$ and $1/2$ on their respective cheating probabilities, we define a single security measure of how good an ROT protocol is with respect to its pair of cheating probabilities (P_A^*, P_B^*). To achieve this, we introduce the concept of the *cheating advantage* of a protocol, denoted as γ, defined as:

$$\begin{aligned} \gamma &:= \max\{\gamma_A, \gamma_B\}, \\ \gamma_A &:= P_A^*/P_A^{ideal} \quad \text{where} \quad P_A^{ideal} = 3/4, \quad \text{and} \\ \gamma_B &:= P_B^*/P_B^{ideal} \quad \text{where} \quad P_B^{ideal} = 1/2. \end{aligned} \tag{1}$$

The ultimate objective of designing ROT protocols is thus to find one which *minimizes the cheating advantage*, with the hopeful goal of making γ as close to 1 as possible. (We note that $\gamma \leq 2$ since both P_A^* and P_B^* are no greater than 1.)

We remark that this notion of *cheating advantage* is similar to what is done in [CKS13] for the case of k-out-of-n (forcing) oblivious transfer which also has vastly different ideal cheating probabilities for Alice and Bob. We believe that this is a natural security measure that can be adopted by other primitives, even symmetric ones. However, since we are dealing with two cheating parameters, it could be the case that some applications would desire to minimize a different function of these two probabilities, different than the cheating advantage. While we do look at cheating advantage in this work, we note that our focus is on the design of protocols that minimize both the cheating probabilities, and the cheating advantage is just a way to assign a quality measure to a given protocol.

1.2 Prior work on quantum protocols for Rabin oblivious transfer and related cryptographic tasks

A first quantum ROT protocol with non-trivial security guarantees was given in [BS25] where a couple different variants in the security definition were considered including the one described in this work. Specifically, the authors first developed a framework of stochastic switching and then

later used it to combine a couple of different insecure Rabin OT protocols to yield one with better security. This protocol yields $\gamma_A = 1.29$ and $\gamma_B = 1.87$, implying $\gamma = 1.87 < 2$ for the security definition considered in Section 1.1.

In [SPK+24], a quantum ROT protocol using two pure states is investigated, and shown to outperform classical protocols in some parameter regimes. Also, it was previously claimed that there exists a perfectly secure quantum protocol for Rabin OT [HW06]. However, this protocol turned out to be vulnerable to cheating [PA24]. In fact, a constant lower bound on the overall security of a variant of ROT in the setting of secure function evaluation (with additional inputs) can be shown from the general lower bound in [OS22] and is given explicitly in the thesis [Osb22, Section 4.7], hinting at the existence of general lower bounds for ROT.

If one makes no assumptions on the cheating powers of Alice and Bob, for example, that they are computationally bounded, the limits on the unconditional security for 1-out-of-2 OT was studied in [CKS13] which provided a constant lower bound on the worst-case security for all complete[1] protocols. This result extends the no-go result in [Lo97] that shows that general quantum computations cannot be done securely, which includes the special case of OT. The lower bound was later improved in [CGS16] which showed that a variant of OT, known as *semi-honest* oblivious transfer, where dishonest Alice is restricted to the set of strategies that allows her to learn the chosen bit perfectly.

More recently, a lower bound on the security of 1-out-of-2 OT [ASR+21] has been used to deduce fundamental trade-offs between circuit privacy, data privacy and correctness for a broad family of quantum homomorphic encryption protocols [HOT23].

In terms of finding protocols, it was shown in [CKS13] that partial security for the 1-out-of-2 OT task can be obtained, that is, the maximum probability with which Alice or Bob could cheat successfully is strictly below 1, a strict (quantum) improvement over classical protocols where it is known that either Alice or Bob can always cheat perfectly. Finding optimal protocols for 1-out-of-2 OT is still an interesting, and important, open problem.

1.3 A few remarks about the security setting considered in this work

Here we discuss how our study of Rabin oblivious transfer fits into the bigger picture of two-party quantum cryptography.

First and foremost, we are considering *information-theoretic* security, where we do not place any restriction on a cheating party. For example, we do not make any assumptions on computational hardness [Yao86, GMW87], space limitations [Sch07, DFSS08], noisy hardware [WST08, STW09], etc. While these are interesting in their own right, and can be well-motivated, we are studying this task in the most demanding setting possible, that is, when Alice and Bob are only bounded by the laws of quantum mechanics.

Secondly, we are assuming that all communication channels are noiseless as the problem could become easier in some noisy settings. Consider a binary erasure channel which transmits a bit x and, with probability p, the bit is replaced with a flag $\perp$. Then for $p = 1/2$, this can accomplish exactly Rabin oblivious transfer. Moreover, Alice and Bob cannot cheat in this setting. *In this sense, Rabin oblivious transfer seeks to emulate a binary erasure channel where the noise is controlled by Alice or Bob, not the environment.* Indeed, it is possible to efficiently implement other tasks such as bit commitment and oblivious transfer in the noisy setting [Cré97].

Lastly, we would like to emphasize that unlike some of the previous works that studied the composability of various two-party tasks such as coin flipping [VPdR19, WHBT25], here we only

[1] Complete protocols are defined such that honest participants achieve the desired objective.

consider the standalone security of the primitives relevant to this work, including Rabin oblivious transfer, 1-out-of-2 oblivious transfer and coin flipping. Therefore, the protocols we propose in the later sections of this work may not be necessarily secure under arbitrary compositions, and we leave this research direction for future work.

1.4 New protocols for Rabin oblivious transfer

In this section, we present novel constructions of Rabin oblivious transfer protocols that offer a strict improvement in overall security compared to those described in previous works [BS25, SPK+24]. Here, we adopt the security notion discussed in Section 1.1.

1.4.1 A quantum protocol via weak coin flipping

One way to reduce the cheating advantages of many two-party cryptographic protocols is to balance them using *weak coin flipping*. Suppose that you are given two protocols *for the same task* $\mathcal{P}_1$ with $\mathcal{P}_2$ and suppose that Alice can cheat more in $\mathcal{P}_1$ than in $\mathcal{P}_2$ and Bob can cheat more in $\mathcal{P}_2$ than in $\mathcal{P}_1$. In other words, if Alice and Bob were faced with deciding on whether to use $\mathcal{P}_1$ or $\mathcal{P}_2$, then they would want opposing decisions. One way to deal with this is to have them *flip a coin*, and if they output HEADS then use $\mathcal{P}_1$, else use $\mathcal{P}_2$. This also has the potential benefit of reducing the overall cheating, as described in the following lemma.

Lemma 1.2 (Balancing Lemma)**.** Suppose we are given two quantum ROT protocols with $\mathcal{P}_1$ with cheating advantages $(\gamma_A^{(1)}, \gamma_B^{(1)})$ and $\mathcal{P}_2$ with cheating advantages $(\gamma_A^{(2)}, \gamma_B^{(2)})$. If Alice prefers $\mathcal{P}_1$ (i.e., $\gamma_A^{(1)} \geq \gamma_A^{(2)}$), Bob prefers $\mathcal{P}_2$ (i.e., $\gamma_B^{(1)} \leq \gamma_B^{(2)}$), and $0 < \frac{\gamma_B^{(2)} - \gamma_A^{(2)}}{\gamma_B^{(2)} - \gamma_A^{(2)} + \gamma_A^{(1)} - \gamma_B^{(1)}} < 1$ then there exists a family of quantum ROT protocols approaching

$$\gamma = \gamma_A = \gamma_B = \frac{\gamma_A^{(1)}\gamma_B^{(2)} - \gamma_B^{(1)}\gamma_A^{(2)}}{\gamma_B^{(2)} - \gamma_A^{(2)} + \gamma_A^{(1)} - \gamma_B^{(1)}} \leq \min\{\gamma_1, \gamma_2\}, \tag{2}$$

where $\gamma_1 = \max\{\gamma_A^{(1)}, \gamma_B^{(1)}\}$ and $\gamma_2 = \max\{\gamma_A^{(2)}, \gamma_B^{(2)}\}$. The improvement is strict, i.e., we have $\gamma < \min\{\gamma_1, \gamma_2\}$ under the previous conditions if and only if Alice has a strict preference for $\mathcal{P}_1$ ($\gamma_A^{(1)} > \gamma_A^{(2)}$) and Bob strictly prefers $\mathcal{P}_2$ ($\gamma_B^{(1)} < \gamma_B^{(2)}$).

Remark 1.3. This balancing lemma above can in general be used for any two protocols for a given two-party task (coin flipping, bit commitment, etc.) to potentially yield another protocol with improved security. We note that this trick uses a quantum protocol for coin flipping for the information-theoretic security setting. In other settings, one must be mindful of the coin flipping protocol security when balancing in this way.

We note that if both Alice and Bob prefer the same protocol over the other, e.g., $\gamma_A^{(1)} \leq \gamma_A^{(2)}$ and $\gamma_B^{(1)} \leq \gamma_B^{(2)}$, then this balancing lemma does not offer any advantages. However, its power comes from the ability to combine protocols for which Alice and Bob have different preferences.

Proof Sketch. The proof follows immediately using the task of *unbalanced weak coin flipping* for which optimal quantum protocols are known [CK09], based on optimal weak (balanced) coin flipping protocols [Moc07, ARV21]. Specifically, for the z-unbalanced weak coin flipping task between Alice and Bob, for all choices of $\epsilon > 0$, there exist a sequence of protocols which in the limit have the following features:

- The probability that the outcome is HEADS is z when Alice and Bob are honest, and they both output the same coin flip;
- Alice cannot force the outcome HEADS with probability greater than z;
- Bob cannot force the outcome TAILS with probability greater than $1 - z$.

Suppose Alice and Bob use a z-unbalanced weak coin flipping protocol with $z = \frac{\gamma_B^{(2)} - \gamma_A^{(2)}}{\gamma_B^{(2)} - \gamma_A^{(2)} + \gamma_A^{(1)} - \gamma_B^{(1)}}$, then on observing the outcome HEADS, they perform $\mathcal{P}_1$ and on outcome TAILS then perform $\mathcal{P}_2$. The calculation of γ_A or γ_B follows directly by calculating $z\gamma_A^{(1)} + (1-z)\gamma_A^{(2)}$ or $z\gamma_B^{(1)} + (1-z)\gamma_B^{(2)}$ for $z = \frac{\gamma_B^{(2)} - \gamma_A^{(2)}}{\gamma_B^{(2)} - \gamma_A^{(2)} + \gamma_A^{(1)} - \gamma_B^{(1)}}$. □

This lemma turns out to be valuable for minimizing the cheating advantage in the protocols in this paper, mostly due to the fact that Alice and Bob's cheating probabilities are rather imbalanced. But before that, we demonstrate that this simple act of balancing can turn bad protocols into decent protocols.

Consider the following two *bad* protocols for ROT.

Protocol 1. A bad Rabin oblivious transfer protocol $\mathcal{P}_1$.

- Bob chooses $y \in \{0, 1\}$ uniformly at random.
- Bob flips a fair coin.
  - If HEADS, Bob announces y to Alice.
  - If TAILS, Bob announces "NULL" to Alice.

Clearly this protocol satisfies $P_A^* = 3/4$ since Alice either learns y or not, and if not she has to guess. Moreover, $P_B^* = 1$ as Bob knows exactly what Alice would assert. Thus, for this protocol we have

$$\gamma_A = \frac{3/4}{3/4} = 1, \quad \gamma_B = \frac{1}{1/2} = 2, \quad \text{and} \quad \gamma = \max\{\gamma_A, \gamma_B\} = 2, \tag{3}$$

which is the worst value of γ possible.

Protocol 2. Another bad Rabin oblivious transfer protocol $\mathcal{P}_2$.

- Bob chooses $y \in \{0, 1\}$ uniformly at random.
- Bob prepares $|y\rangle$ and sends it to Alice.
- Alice flips a fair coin.
  - If HEADS, Alice measures $|y\rangle$ to learn y.
  - If TAILS, Alice throws the qubit away and accepts NULL as her outcome.

Clearly this protocol satisfies $P_A^* = 1$ since she can simply measure the qubit to learn it perfectly. Moreover, $P_B^* = 1/2$, since Alice has full control over what she asserts, and no information is sent to Bob after Alice decides to measure or not. Thus, for this protocol we have

$$\gamma_A = \frac{1}{3/4} = \frac{4}{3}, \quad \gamma_B = \frac{1/2}{1/2} = 1, \quad \text{and} \quad \gamma = \max\{\gamma_A, \gamma_B\} = \frac{4}{3}, \tag{4}$$

which is not the worst value of γ possible, but still not that good.

Even though protocols $\mathcal{P}_1$ and $\mathcal{P}_2$ have poor security, we can apply the Balancing Lemma to improve their overall security resulting in the following theorem.

Theorem 1.4. There exists a quantum protocol for ROT with cheating advantage $\gamma = 5/4$.

A few remarks are in order. Note that even though this protocol is easy to describe, just use optimal unbalanced weak coin flipping with the two bad ROT protocols, it is important to note that optimal unbalanced weak coin flipping protocols are very complicated. The only way we know how to construct such protocols is via optimal weak coin flipping [Moc07]. Despite much work done to simplify such protocols [ACG$^+$16, ARV21], they are still very technical. Moreover, it has been shown that such protocols cannot be efficient [Mil20]. Thus, this begs the question:

Does there exist ROT protocols that are simple and offer decent security?

In the following two sections, we describe how to achieve this.

1.4.2 A quantum protocol via a reduction from oblivious transfer

Typically, reductions between tasks are useful to establish relationships between them, which may include non-trivial security bounds [CKS13]. In some cases, reductions can also be used to yield protocols with desirable security. However, some of the fundamental security notions, such as *completeness* might be compromised while reducing one task to another. Here, it is crucial to devise appropriate reductions for tasks where no such violations are permitted.

We next depict a folklore (classical) reduction from OT to ROT. Alice and Bob communicate only classically outside the OT subroutine to perform the ROT task. We characterize this as a classical reduction given that the communication outside of the OT subroutine is classical.

Reduction 1. Rabin oblivious transfer from 1-out-of-2-oblivious transfer.

- Bob chooses data $y \in \{0,1\}$ and a dummy bit $z \in \{0,1\}$ uniformly at random.
- Bob chooses a permutation value $p \in \{0,1\}$ uniformly at random and assigns the bits $(x_0, x_1) = (y, z)$ if $p = 0$, or $(x_0, x_1) = (z, y)$, otherwise.
- **Alice and Bob perform 1-out-of-2 oblivious transfer** so that Alice learns either x_0 or x_1 (see details on 1-out-of-2 OT in Definition A.1).
- Bob sends p to Alice. Alice now knows whether she learned y or z.
- If she learned z, then she outputs NULL. Otherwise, she outputs y.

It is straightforward to observe that an ROT implementation using Reduction 1 is *complete* whenever the underlying OT protocol is *complete*. Additionally, the reduction is tight in the sense that given a perfectly secure OT protocol the resultant ROT protocol also enjoys perfect security. Of course, perfectly secure OT is impossible as discussed earlier, thus this reduction does not give us perfectly secure ROT protocols. However, it does give partial security, as discussed below.

Using this reduction, our idea to create ROT protocols is now clear, we can simply choose an existing (or develop a new) quantum OT protocol and slot it into Reduction 1. However, since this is a quantum protocol, it is *not clear* how Alice and/or Bob may cheat with regards to composing various cheating strategies. That is, an optimal cheating strategy in the OT protocol may not

be the optimal cheating strategy in the overall protocol. This optimal cheating behavior (and thus the worst case cheating probabilities) needs to be proven for this reduction, since we do not know much about the security aspects of 1-out-of 2 oblivious transfer protocols when used as a subroutine in a larger composition. Thus, we provide a rigorous self-contained security analysis of the Rabin OT protocols resulting from the previous reduction.

We now exhibit a complete quantum protocol for ROT, below.

Protocol 3. Using Reduction 1 with the 1-out-of-2 OT protocol from [CKS13].

- Alice chooses a bit $a \in \{0,1\}$ uniformly at random and creates the two-qutrit state

$$|\phi_a\rangle = \frac{1}{\sqrt{2}}|aa\rangle + \frac{1}{\sqrt{2}}|22\rangle \in \mathcal{A} \otimes \mathcal{B}.$$

- Alice sends the qutrit $\mathcal{B}$ to Bob.
- Bob chooses data bit $y \in \{0,1\}$ and $z \in \{0,1\}$ uniformly at random.
- Bob chooses a permutation bit $p \in \{0,1\}$ uniformly at random and assigns $(x_0, x_1) = (y, z)$ if $p = 0$, or $(x_0, x_1) = (z, y)$, otherwise.
- Bob applies the unitary

$$U_{x_0 x_1} = \begin{bmatrix} (-1)^{x_0} & 0 & 0 \\ 0 & (-1)^{x_1} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

  to $\mathcal{B}$ and sends it back to Alice.

- Alice determines the value of x_a using the two-outcome measurement:

$$\{\Pi_0 \coloneqq |\phi_a\rangle\langle\phi_a|_{\mathcal{A}\otimes\mathcal{B}}, \Pi_1 \coloneqq \mathbb{1}_{\mathcal{A}\otimes\mathcal{B}} - |\phi_a\rangle\langle\phi_a|\}.$$

- Bob sends p to Alice.
- If $a = p$, Alice asserts the receipt of $y = x_a$, else she asserts NULL.

The completeness of the 1-out-of-2 protocol in [CKS13] ensures that this ROT protocol is also *complete*. However, it is not clear whether the optimal cheating probability for dishonest Alice where she tries to learn Bob's data bit y is equivalent to Alice learning both x_0 and x_1 in the reduced 1-out-of-2 OT protocol. Although learning just these two bits (x_0 and x_1) qualifies as a valid cheating strategy, she might deploy a more clever strategy to improve her chances of successfully learning y, i.e., we have $P_A^* \geq P_A^{OT}$. Note that quantum communication in the reduced 1-out-of-2 OT protocol attributes to a non-trivial overall analysis as dishonest Alice could now correlate her messages with Bob's input data. We give an argument of Alice's possible strategies below, and provide the missing details in the appendix.

A general strategy for a cheating Alice has her creating a state of the form below

$$|\psi'\rangle = \alpha|\psi_\alpha\rangle_{\mathcal{A}'}|0\rangle_{\mathcal{B}} + \beta|\psi_\beta\rangle_{\mathcal{A}'}|1\rangle_{\mathcal{B}} + \gamma|\psi_\gamma\rangle_{\mathcal{A}'}|2\rangle_{\mathcal{B}},$$

where $\alpha, \beta, \gamma \geq 0$ satisfy $\alpha^2 + \beta^2 + \gamma^2 = 1$, and sends the subsystem $\mathcal{B}$ to Bob. Once Bob applies the unitary $U_{x_0 x_1}$ on the subsystem $\mathcal{B}$ and sends it back, Alice holds the state

$$|\psi'_{x_0 x_1}\rangle = \alpha(-1)^{x_0}|\psi_\alpha\rangle_{\mathcal{A}'}|0\rangle_{\mathcal{B}} + \beta(-1)^{x_1}|\psi_\beta\rangle_{\mathcal{A}'}|1\rangle_{\mathcal{B}} + \gamma|\psi_\gamma\rangle_{\mathcal{A}'}|2\rangle_{\mathcal{B}}.$$

It is important to observe that Alice's choice of the states $\{|\psi_\alpha\rangle, |\psi_\beta\rangle, |\psi_\gamma\rangle\}$ is not relevant in maximizing her chances of successfully cheating in Protocol 3 as for any choice of $|\psi_\alpha\rangle$, $|\psi_\beta\rangle$ and $|\psi_\gamma\rangle$, there exists a unitary U such that

$$U|\psi'\rangle = |\psi\rangle,$$

where $|\psi\rangle = \alpha|0\rangle_{\mathcal{A}'}|0\rangle_{\mathcal{B}} + \beta|0\rangle_{\mathcal{A}'}|1\rangle_{\mathcal{B}} + \gamma|0\rangle_{\mathcal{A}'}|2\rangle_{\mathcal{B}}$.
Thus, Alice could restrict to creating states of the form $|\psi\rangle$, or equivalently just the states of the form

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle \in \mathcal{B}.$$

Subsequently, her task reduces to extracting x_0 or x_1 (depending on p) from the reduced state

$$|\psi_{x_0x_1}\rangle = \alpha(-1)^{x_0}|0\rangle + \beta(-1)^{x_1}|1\rangle + \gamma|2\rangle \in \mathcal{B}. \tag{5}$$

If $p = 0$, the probability of Alice successfully guessing y (or x_0 in this case) is given by

$$\Pr[\text{Alice successfully learns } y|p=0] = \frac{1}{2} + \frac{1}{4}\left\|\rho_0(0) - \rho_1(0)\right\|_1,$$

where $\rho_0(0) = \frac{1}{2}|\psi_{00}\rangle\langle\psi_{00}| + \frac{1}{2}|\psi_{01}\rangle\langle\psi_{01}|$ and $\rho_1(0) = \frac{1}{2}|\psi_{10}\rangle\langle\psi_{10}| + \frac{1}{2}|\psi_{11}\rangle\langle\psi_{11}|$.
Similarly, if $p = 1$, the probability of Alice successfully guessing y (or x_1 in this case) is given by

$$\Pr[\text{Alice successfully learns } y|p=1] = \frac{1}{2} + \frac{1}{4}\left\|\rho_0(1) - \rho_1(1)\right\|_1,$$

where $\rho_0(1) = \frac{1}{2}|\psi_{00}\rangle\langle\psi_{00}| + \frac{1}{2}|\psi_{10}\rangle\langle\psi_{10}|$ and $\rho_1(1) = \frac{1}{2}|\psi_{01}\rangle\langle\psi_{01}| + \frac{1}{2}|\psi_{11}\rangle\langle\psi_{11}|$.
Optimizing the overall success probability over all choices of α, β and γ, we get $P_A^* = \cos^2(\pi/8)$ when $\alpha = \beta = 1/2$ and $\gamma = 1/\sqrt{2}$.

We now state the overall security of this protocol and defer the rest of the proof to Appendix B.

Lemma 1.5. The ROT Protocol 3 has cheating probabilities $P_A^* = \cos^2(\pi/8)$ and $P_B^* = 3/4$ implying $\gamma_A = 1.138$ and $\gamma_B = 1.5$, resulting in a cheating advantage of $\gamma = 1.5$.

By combining Protocol 3 with Protocol 2 using the Balancing Lemma 1.2, we get an ROT protocol with security strictly better than the protocol discussed in Section 1.4.1 and is stated in the following theorem.

Theorem 1.6. There exists a quantum protocol for ROT with cheating advantage $\gamma = 1.239$.

We now describe a way to execute many ROT protocols in succession using a simpler protocol which exhibits the same cheating probabilities as derived previously for Protocol 3.

1.5 A simple sequence of Rabin oblivious transfer protocols

We now propose another way of performing Rabin oblivious transfer based on the OT protocols described in [ASR+21]. We describe the protocol below.

Protocol 4. A sequence of Rabin oblivious transfers.

- For each $i \in \{1, \ldots, N\}$, Bob picks a pair of bits (x_0^i, x_1^i), all independently and uniformly at random.

- For each $i \in \{1, \ldots, N\}$, Bob prepares the two-qubit states $|\Psi_{x_0^i x_1^i}\rangle$, where we define:

$$|\Psi_{00}\rangle = |00\rangle\,, \quad |\Psi_{01}\rangle = |{+}{+}\rangle\,, \quad |\Psi_{10}\rangle = |{-}{-}\rangle\,, \quad |\Psi_{11}\rangle = |11\rangle\,, \tag{6}$$

  and $|\pm\rangle := (|0\rangle \pm |1\rangle)/\sqrt{2}$. Bob sends the states to Alice.

- Alice selects a subset S of the set $\{1, \ldots, N\}$ of size $\left\lfloor \sqrt{N} \right\rfloor$ uniformly at random. For each $j \in S$, Alice asks Bob to announce the state $|\Psi_{x_0^j, x_1^j}\rangle$.

- If Bob announces $|{+}{+}\rangle$ or $|{-}{-}\rangle$, then Alice measures both qubits in the $\{|+\rangle\,, |-\rangle\}$ basis, otherwise Alice measures in the $\{|0\rangle\,, |1\rangle\}$ basis. Alice aborts if any of her outcomes are not consistent with Bob's purported states. If Alice does not abort, then the $\left\lfloor \sqrt{N} \right\rfloor$ states used for these tests are discarded.

- For each $j \notin S$, Alice measures the first qubit of $|\Psi_{x_0^j x_1^j}\rangle$ in the $\{|0\rangle\,, |1\rangle\}$ basis and the second qubit of $|\Psi_{x_0^j x_1^j}\rangle$ in the $\{|+\rangle\,, |-\rangle\}$ basis. This way Alice learns one of Bob's bits with certainty while gaining no knowledge of the other (each of her measurement outcomes allows her to rule out one of Bob's possible states). For example, if her measurement outcome is

  - 00: then we have $(x_0^j, x_1^j) = (0,0)$ or $(0,1)$, i.e., the first bit is 0;
  - 01: then we have $(x_0^j, x_1^j) = (0,0)$ or $(1,0)$, i.e., the second bit is 0;
  - 10: then we have $(x_0^j, x_1^j) = (0,1)$ or $(1,1)$, i.e., the second bit is 1;
  - 11: then we have $(x_0^j, x_1^j) = (1,0)$ or $(1,1)$, i.e., the first bit is 1.

- For each $j \notin S$, Bob declares which of the two bits will be used for Rabin oblivious transfer. This is chosen uniformly at random for each j.

As before, the completeness follows from the completeness of the underlying OT protocol. Again we have to carefully consider what a cheating Alice can do to further her cheating probability over what she can achieve if she uses the cheating strategy which is optimal in the OT protocol.

The optimal cheating probabilities for Alice and Bob are given in the following lemma.

Lemma 1.7. The ROT Protocol 4 has cheating probabilities $P_A^* = \cos^2(\pi/8)$ for all N and $P_B^* = 3/4$ in the limit $N \to \infty$ implying $\gamma_A = 1.138$ and $\gamma_B = 1.5$, resulting in a cheating advantage of $\gamma = 1.5$ in the limit.

The full security analysis for the above protocol can be found in Appendix C.

1.6 A constant lower bound on γ for any quantum protocol for Rabin oblivious transfer

The two main objectives for understanding the limitations of the security for any cryptographic task is to (a) find good protocols and the corresponding security measures, and (b) find lower bounds on such security measures. By improving either of these objectives brings us that much closer to finding the *best* protocol.

So far, we have been concentrating on objective (a), finding good quantum protocols. Now, we consider objective (b), finding lower bounds.

Why care about lower bounds? Suppose one were to show a constant lower bound on γ, that is, suppose $\gamma \geq c$ for some constant $c > 1$. This implies two things. First, this means that obviously one cannot find protocols with $\gamma < c$, which could suggest that any existing protocols are better than expected (by being now closer to the optimal γ). Secondly, and importantly, this *rules out security amplification attempts*. One may wonder if ROT can be repeated many times and then somehow by putting all the pieces together reduce the cheating probabilities to their ideal goals. Having a constant lower bound rules out such attempts, since near-perfect protocols cannot exist, even in the limit of a family of protocols. Thus, just by having a constant lower bound, regardless of the constant, provides insight towards how one may design optimal protocols. We therefore ask the question:

Does there exist a constant lower bound for Rabin oblivious transfer?

At first glance, it is not always obvious how to guess this. For example, quantum protocols for strong coin flipping (defined shortly) have a constant lower bound [Kit02] while quantum protocols for weak coin flipping do not [Moc07]. What about Rabin oblivious transfer? In this section, we answer this question in the affirmative.

1.6.1 Strong coin flipping

We have discussed (unbalanced) weak coin flipping between Alice and Bob, the task of strong coin flipping is similar. Roughly, a strong coin flipping protocol is defined such that:

- The probability that the outcome is HEADS is $1/2$ when Alice and Bob are honest, and they both output the same coin flip;
- Alice can force the outcome HEADS with probability $P_{A,0}$;
- Alice can force the outcome TAILS with probability $P_{A,1}$;
- Bob can force the outcome HEADS with probability $P_{B,0}$;
- Bob can force the outcome TAILS with probability $P_{B,1}$.

A formal definition is given in Definition A.2. While much is now known about the limits of the achievable values of $P_{A,0}$, $P_{A,1}$, $P_{B,0}$, $P_{B,1}$, in this discussion we only require one bound, proven by Kitaev [Kit02] (see Proposition A.3):

$$P_{A,0} \cdot P_{B,0} \geq 1/2. \tag{7}$$

In other words, in *any* coin flipping protocol, either Alice or Bob can force the outcome 0 with probability at least $1/\sqrt{2} > 1/2$.

1.6.2 A reduction from Rabin oblivious transfer to strong coin flipping

We present our reduction below.

Protocol 5. Coin flipping from Rabin oblivious transfer

- Alice and Bob perform an ROT protocol with Bob's data bit denoted by y. We denote Alice's ROT output by the trit $\hat{g} \in \{0, 1, NULL\}$. If Alice asserts the receipt of data, then $\hat{g} = y$, else $\hat{g} = NULL$.

- Alice defines the assert bit $a \in \{0,1\}$ to be 0 if $\hat{g} \in \{0,1\}$, and to be 1 if $\hat{g} = NULL$.
- Bob chooses b uniformly at random and sends b to Alice.
- Alice sends a and $\hat{g}$ to Bob. Bob checks if $(a, y, \hat{g})$ are consistent, i.e., he checks if $\hat{g} = y$ when $a = 0$. If $\hat{g} \neq y$, Bob aborts.
- If Alice and Bob do not abort, then they both output HEADS if $a = b$, and TAILS otherwise.

We first remark that this protocol is complete, i.e., both Alice and Bob output the same coin flip whose value is generated uniformly at random. Below, we investigate the security of this coin flipping protocol with respect to the security of the ROT protocol.

Lemma 1.8. In Protocol 5, we have

$$P_{A,0} = \frac{1}{2} + \frac{1}{2}P_A^* \quad \text{and} \quad P_{B,0} = P_B^*, \tag{8}$$

where P_A^* and P_B^* are the cheating probabilities in the ROT subroutine.

The proof of the above lemma can be found in Appendix D.

By combining Equations (1) and (7), i.e., that

$$\gamma = \max\{\gamma_A, \gamma_B\}, \quad \gamma_A = \frac{P_A^*}{3/4}, \quad \gamma_B = \frac{P_B^*}{1/2}, \quad \text{and} \quad P_{A,0} \cdot P_{B,0} \geq 1/2 \tag{9}$$

together with Lemma 1.8, one gets the quadratic inequality

$$\frac{1}{2}\gamma\left(\frac{1}{2} + \frac{1}{2}\cdot\frac{3}{4}\gamma\right) \geq \frac{1}{2}, \tag{10}$$

implying that $\gamma \geq \frac{2}{3}\left(\sqrt{7} - 1\right) > 1$.

Therefore, we have the following theorem.

Theorem 1.9. Any quantum protocol for Rabin oblivious transfer has a cheating advantage satisfying $\gamma \geq \frac{2}{3}\left(\sqrt{7} - 1\right) \approx 1.097 > 1$.

By combining Theorems 1.6 and 1.9, we get the following corollary.

Corollary 1.10. The optimal quantum protocol for Rabin oblivious transfer has cheating advantage $\gamma \in [1.097, 1.239]$.

Acknowledgments

JS thanks Dominique Unruh for pointing out the OT to ROT reduction. AB thanks Or Sattath and Atul Mantri for useful discussions and pointing out the references for protocols under noisy communication. JW is supported by the National Research Foundation, Singapore and A*STAR under its Quantum Engineering Programme (NRF2021-QEP2-01-P06). EA was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under Grants No. EP/T001011/1 and EP/Z533208/1.

References

[ACG+16] Dorit Aharonov, André Chailloux, Maor Ganz, Iordanis Kerenidis, and Loïck Magnin. A simpler proof of the existence of quantum weak coin flipping with arbitrarily small bias. *SIAM Journal on Computing*, 45(3):633–679, 2016.

[ARV21] Atul Singh Arora, Jérémie Roland, and Chrysoula Vlachou. Analytic quantum weak coin flipping protocols with arbitrarily small bias. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 919–938, 2021.

[ARVW24] Atul Singh Arora, Jérémie Roland, Chrysoula Vlachou, and Stephan Weis. Protocols for quantum weak coin flipping. arXiv preprint arXiv:2402.15855, 2024.

[ASR+21] Ryan Amiri, Robert Stárek, David Reichmuth, Ittoop V. Puthoor, Michal Mičuda, Ladislav Mišta, Jr., Miloslav Dušek, Petros Wallden, and Erika Andersson. Imperfect 1-Out-of-2 Quantum Oblivious Transfer: Bounds, a Protocol, and its Experimental Implementation. *PRX Quantum*, 2:010335, 2021.

[BBCS92] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. Practical quantum oblivious transfer. In *Advances in Cryptology — CRYPTO '91*, pages 351–366. Springer Berlin Heidelberg, 1992.

[BCR87] Gilles Brassard, Claude Crépeau, and Jean-Marc Robert. All-or-nothing disclosure of secrets. In *Proceedings on Advances in Cryptology—CRYPTO '86*, page 234–238, Berlin, Heidelberg, 1987.

[BS25] Akshay Bansal and Jamie Sikora. Breaking barriers in two-party quantum cryptography via stochastic semidefinite programming. *Quantum*, 9:1602, 2025.

[CGS16] André Chailloux, Gus Gutoski, and Jamie Sikora. Optimal bounds for semi-honest quantum oblivious transfer. *Chicago Journal of Theoretical Computer Science*, 2016(13):1–17, 2016.

[CK09] André Chailloux and Iordanis Kerenidis. Optimal quantum strong coin flipping. In *2009 IEEE 50nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 527–533, 2009.

[CK11] André Chailloux and Iordanis Kerenidis. Optimal bounds for quantum bit commitment. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 354–362, 2011.

[CKS13] André Chailloux, Iordanis Kerenidis, and Jamie Sikora. Lower bounds for quantum oblivious transfer. *Quantum Info. Comput.*, 13(1–2):158–177, 2013.

[Cré88] Claude Crépeau. Equivalence between two flavours of oblivious transfers. In *Advances in Cryptology — CRYPTO '87*, pages 350–354. Springer Berlin Heidelberg, 1988.

[Cré94] Claude Crépeau. Quantum oblivious transfer. *Journal of Modern Optics*, 41(12):2445–2454, 1994.

[Cré97] Claude Crépeau. Efficient cryptographic protocols based on noisy channels. In *Advances in Cryptology — EUROCRYPT '97*, pages 306–317. Springer Berlin Heidelberg, 1997.

[DFSS08] Ivan B. Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner. Cryptography in the bounded-quantum-storage model. *SIAM Journal on Computing*, 37(6):1865–1890, 2008.

[FMR96] Michael J Fischer, Silvio Micali, and Charles Rackoff. A secure protocol for the oblivious transfer. *Journal of Cryptology*, 9(3):191–196, 1996.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play ANY mental game. In *Proceedings of the 19th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 218–229, 1987.

[HOT23] Yanglin Hu, Yingkai Ouyang, and Marco Tomamichel. Privacy and correctness tradeoffs for information-theoretically secure quantum homomorphic encryption. *Quantum*, 7:976, 2023.

[HW06] Guang Ping He and ZD Wang. Oblivious transfer using quantum entanglement. *Physical Review A—Atomic, Molecular, and Optical Physics*, 73(1):012331, 2006.

[Kil88] Joe Kilian. Founding crytpography on oblivious transfer. In *Proceedings of the 20th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 20–31, 1988.

[Kit02] Alexei Kitaev. Quantum coin flipping. Unpublished result. *Talk at the 6th Annual workshop on Quantum Information Processing (QIP 2003)*, 2002.

[KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 830–842. Association for Computing Machinery, 2016.

[Lo97] Hoi-Kwong Lo. Insecurity of quantum secure computations. *Phys. Rev. A*, 56:1154–1162, 1997.

[Mil20] Carl A Miller. The impossibility of efficient quantum weak coin flipping. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 916–929, 2020.

[Moc07] Carlos Mochon. Quantum weak coin flipping with arbitrarily small bias. arXiv preprint arXiv:0711.4114, 2007.

[OS22] Sarah Anne Osborn and Jamie Sikora. A constant lower bound for any quantum protocol for secure function evaluation. In *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, volume 232 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:14, 2022.

[Osb22] Sarah Anne Osborn. Constant lower bounds on the cryptographic security of quantum two-party computations. Master's thesis, Virginia Tech, 2022.

[PA24] James T. Peat and Erika Andersson. Cheating in quantum Rabin oblivious transfer using delayed measurements. arXiv preprint arXiv:2408.12388, 2024.

[Rab79] Michael O. Rabin. Digitalized signatures and public-key function as intractable as factorization. *Technical Report, Computer Science, MIT/LCS/TR-212*, 1:100–123, 1979.

[Rab05] Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Paper 2005/187, 2005.

[Sch07] Christian Schaffner. Cryptography in the bounded-quantum-storage model. arXiv preprint arXiv:0709.0289, 2007.

[SPK+24] Lara Stroh, James T. Peat, Mats Kroneberg, Ittoop V. Puthoor, and Erika Andersson. Quantum Rabin oblivious transfer using two pure states. *Phys. Rev. Res.*, 6:043004, 2024.

[STW09] Christian Schaffner, Barbara Terhal, and Stephanie Wehner. Robust cryptography in the noisy-quantum-storage model. *Quantum Info. Comput.*, 9(11):963–996, 2009.

[VPdR19] V Vilasini, Christopher Portmann, and Lídia del Rio. Composable security in relativistic quantum cryptography. *New Journal of Physics*, 21(4):043057, 2019.

[WHBT25] Jiawei Wu, Yanglin Hu, Akshay Bansal, and Marco Tomamichel. On the composable security of weak coin flipping. *Quantum*, 9:1780, 2025.

[WST08] Stephanie Wehner, Christian Schaffner, and Barbara M. Terhal. Cryptography from noisy storage. *Phys. Rev. Lett.*, 100:220502, 2008.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *1986 IEEE 27th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 162–167, 1986.

A Definition of other cryptographic tasks

We now define a variant of oblivious transfer, namely 1-out-of-2 OT, and coin flipping with their underlying security notions.

Definition A.1 (1-out-of-2 OT)**.** 1-out-of-2 OT is the cryptographic task between two parties (Alice and Bob) such that

- Alice chooses $a \in \{0,1\}$ uniformly at random and Bob chooses $(x_0, x_1) \in \{0,1\}^2$ uniformly at random.
- If both Alice and Bob are honest, then neither Alice nor Bob aborts and Alice outputs $z = x_a$.

We define the following notions of security for protocols that implement 1-out-of-2 OT.

- *Completeness:* If both Alice and Bob are honest, Alice learns the bit x_a unambiguously.
- *Cheating Alice:* A dishonest Alice attempts to learn both x_0 and x_1. The cheating probability of Alice is given by

$$P_A^{OT} = \max \Pr[\text{Alice correctly guesses both } x_0 \text{ and } x_1],$$

  where the maximum is taken over all cheating strategies of Alice. Note that $P_A^{OT} \geq 1/2$ since she can simply follow the protocol to learn x_a perfectly (completeness) and can correctly guess $x_{\bar{a}}$ by making a guess uniformly at random.

- *Cheating Bob:* A dishonest Bob attempts to learn the bit a. The cheating probability of Bob is given by
$$P_B^{OT} = \max \Pr[\text{Bob correctly guesses } a],$$
where the maximum is taken over all cheating strategies of Bob. Note that $P_B^{OT} \geq 1/2$ since he can correctly guess Alice's bit a by guessing one of the two values uniformly at random.

We note that there are alternative definitions of this task where the inputs are fixed beforehand or possibly generated randomly within the protocol itself. Many of these definitions are equivalent to the one presented above and we refer the reader to [BCR87, Cré88, CKS13, ASR+21] for proofs and discussions.

We next discuss coin flipping and provide its security definitions.

Definition A.2 (Coin flipping)**.** Coin flipping is the cryptographic task between two parties (Alice and Bob) where they communicate to agree on a common binary outcome (0 or 1). A weak version of this task is a variant where Alice *wins* if the outcome is 1 while Bob *wins* if the outcome is 0.

We consider the following notions of security for a given weak coin flipping protocol.

- *Completeness:* If both Alice and Bob are honest, then neither party aborts and the shared outcome is generated uniformly at random.
- *Cheating Alice:* If Bob is honest, then Alice's cheating probability is defined with respect to the outcome $d \in \{0, 1\}$ as
$$P_{A,d} = \max \Pr[\text{Alice successfully forces Bob to accept the outcome } d],$$
where the maximum is taken over all possible cheating strategies of Alice. Note that here $P_{A,d} \geq 1/2$ since Alice can simply choose to follow the protocol honestly to observe the outcome d uniformly at random.
- *Cheating Bob:* If Alice is honest, then Bob's cheating probability is defined with respect to the outcome $d \in \{0, 1\}$ as
$$P_{B,d} = \max \Pr[\text{Bob successfully forces Alice to accept the outcome } d],$$
where the maximum is taken over all possible cheating strategies of Bob. As before, here $P_{B,d} \geq 1/2$ since Bob can simply choose to follow the protocol honestly to observe the outcome d uniformly at random.

The variant of the coin flipping task where Alice and Bob do not have any preferred choice of outcome is known as strong coin flipping. We now state Kitaev's lower bound for strong coin flipping which is useful to deduce lower bounds for other two-party primitives.

Proposition A.3 (Kitaev's lower bound [Kit02])**.** For all quantum coin flipping protocols, we have
$$P_{A,d} P_{B,d} \geq 1/2 \tag{11}$$
where $d \in \{0, 1\}$ and $P_{A,d}$ and $P_{B,d}$ are as defined above.

B Security proof of Protocol 3

We now provide a proof of Lemma 1.5. We reproduce the protocol below, for convenience.

- Alice chooses a bit $a \in \{0,1\}$ uniformly at random and creates the two-qutrit state
$$|\phi_a\rangle = \frac{1}{\sqrt{2}}|aa\rangle + \frac{1}{\sqrt{2}}|22\rangle \in \mathcal{A}\otimes\mathcal{B}.$$
- Alice sends the qutrit $\mathcal{B}$ to Bob.
- Bob chooses data bit $y \in \{0,1\}$ and $z \in \{0,1\}$ uniformly at random.
- Bob chooses a permutation bit $p \in \{0,1\}$ uniformly at random and assigns $(x_0, x_1) = (y, z)$ if $p = 0$, or $(x_0, x_1) = (z, y)$, otherwise.
- Bob applies the unitary
$$U_{x_0x_1} = \begin{bmatrix} (-1)^{x_0} & 0 & 0 \\ 0 & (-1)^{x_1} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
and sends back $\mathcal{B}$ to Alice.
- Alice determines the value of x_a using the two-outcome measurement:
$$\{\Pi_0 \coloneqq |\phi_a\rangle\langle\phi_a|_{\mathcal{A}\otimes\mathcal{B}}, \Pi_1 \coloneqq \mathbb{1}_{\mathcal{A}\otimes\mathcal{B}} - |\phi_a\rangle\langle\phi_a|\}.$$
- Bob sends p to Alice.
- If $a = p$, Alice asserts the receipt of $y = x_a$, else she asserts NULL.

Lemma B.1. The optimal success probability (P_A^*) with which dishonest Alice can cheat in Protocol 3 is $\cos^2(\pi/8)$.

Proof. We have already argued in the main text that the most general strategy for Alice has her creating the state
$$\alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle, \tag{12}$$
where α, β, γ are nonnegative and satisfy $\alpha^2 + \beta^2 + \gamma^2 = 1$, and sending this qutrit to Bob. After Bob's unitary, Alice now has the state
$$|\psi_{x_0,x_1}\rangle = \alpha(-1)^{x_0}|0\rangle + \beta(-1)^{x_1}|1\rangle + \gamma|2\rangle, \tag{13}$$
which encodes his two bits. Since Alice wants to learn y, she should wait until p is revealed which tells here whether she wants to learn x_0 or x_1 (that is, which bit of (x_0, x_1) is really y). If $p = 0$, that is, Alice wants to learn x_0, and she can do this with maximum probability given by
$$\frac{1}{2} + \frac{1}{4}\|\rho_0(0) - \rho_1(0)\|_1, \tag{14}$$
where the norm is the trace norm and $\rho_0(0) = \frac{1}{2}|\psi_{00}\rangle\langle\psi_{00}| + \frac{1}{2}|\psi_{01}\rangle\langle\psi_{01}|$ is the encoding with $x_0 = 0$ and $\rho_1(0) = \frac{1}{2}|\psi_{10}\rangle\langle\psi_{10}| + \frac{1}{2}|\psi_{11}\rangle\langle\psi_{11}|$ is the encoding with $x_0 = 1$. A quick calculation shows that
$$\rho_0(0) - \rho_1(0) = \begin{pmatrix} 0 & 0 & 2\alpha\gamma \\ 0 & 0 & 0 \\ 2\alpha\gamma & 0 & 0 \end{pmatrix} \tag{15}$$

which has at most two nonzero eigenvalues, $\pm 2\alpha\gamma$. Therefore, Alice can cheat with maximum probability

$$\frac{1}{2} + \alpha\gamma \tag{16}$$

in this case.

If $p = 1$, that is, Alice wants to learn x_1, and she can do this with maximum probability given by

$$\frac{1}{2} + \frac{1}{4}\|\rho_0(1) - \rho_1(1)\|_1, \tag{17}$$

where $\rho_0(1) = \frac{1}{2}|\psi_{00}\rangle\langle\psi_{00}| + \frac{1}{2}|\psi_{10}\rangle\langle\psi_{10}|$ is the encoding with $x_1 = 0$ and $\rho_1(1) = \frac{1}{2}|\psi_{01}\rangle\langle\psi_{01}| + \frac{1}{2}|\psi_{11}\rangle\langle\psi_{11}|$ is the encoding with $x_1 = 1$. A quick calculation shows that

$$\rho_0(1) - \rho_1(1) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 2\beta\gamma \\ 0 & 2\beta\gamma & 0 \end{pmatrix} \tag{18}$$

which has at most two nonzero eigenvalues, $\pm 2\beta\gamma$. Therefore, Alice can cheat with maximum probability

$$\frac{1}{2} + \beta\gamma \tag{19}$$

in this case.

Since the choice of p is uniformly random, Alice's overall cheating probability is given as

$$\frac{1}{2} + \frac{1}{2}\beta\gamma + \frac{1}{2}\alpha\gamma. \tag{20}$$

A simple optimization over the parameters $\alpha, \beta, \gamma \geq 0$ satisfying $\alpha^2 + \beta^2 + \gamma^2 = 1$ yeilds that $\alpha = \beta = 1/2$ and $\gamma = 1/\sqrt{2}$ is the optimal solution from which we can calculate the optimal cheating probability as $\frac{1}{2} + \frac{1}{2\sqrt{2}} = \cos^2(\pi/8)$, as desired.

□

Lemma B.2. The optimal success probability (P_B^*) with which dishonest Bob can cheat in Protocol 3 is $3/4$.

Proof. It is necessary and sufficient for dishonest Bob in Protocol 3 to simply learn Alice's bit a as Bob can then set his choice for p accordingly, thereby successfully guessing Alice's assertion of receiving y or NULL. Bob can maximally distinguish the two states $\mathrm{Tr}_{\mathcal{A}}(|\phi_0\rangle\langle\phi_0|)$ and $\mathrm{Tr}_{\mathcal{A}}(|\phi_1\rangle\langle\phi_1|)$ in the Holevo-Helstrom basis resulting in $P_B^* = 3/4$. □

C Security proof of Protocol 4

We next provide a proof of Lemma 1.7.

C.1 Cheating Alice

Theorem C.1. The optimal success probability (P_A^*) with which dishonest Alice can cheat in Protocol 4 is $\frac{1}{4}(2 + \sqrt{2})$ for all N.

Proof. In this protocol Alice only receives states, and sends nothing to Bob. This means that the only cheating method available to Alice is to use a measurement which is different from what she would measure if honest. As Bob decides which of the two bits that counts for each of the remaining $N - \left\lfloor \sqrt{N} \right\rfloor \approx N$ states, Alice's best strategy is to store these states until Bob declares which bit value counts. Alice then makes a minimum-error measurement individually on each state, aiming to learn the bit that Bob chose. Due to symmetry, Alice's probability of correctly obtaining the first bit is the same as for the second. When Bob declares that the first bit should be used, the two states Alice has to distinguish between are given by

$$\rho_0 = \frac{1}{2}(|00\rangle\langle 00| + |{+}{+}\rangle\langle {+}{+}|), \quad \rho_1 = \frac{1}{2}(|11\rangle\langle 11| + |{-}{-}\rangle\langle {-}{-}|). \tag{21}$$

Here Alice's best strategy is to make a minimum-error measurement on each instance individually. Since there are only two states to distinguish between, we use the standard result for the probability of success, the so-called Holevo-Helstrom measurement, yielding

$$p_s = \frac{1}{2}(1 + \mathrm{Tr}(|p_0\rho_0 - p_1\rho_1|)), \tag{22}$$

where $|A| = \sqrt{A^\dagger A}$. This gives Alice's cheating probability as

$$P_A^* = \frac{1}{4}(2 + \sqrt{2}) \approx 0.853 \tag{23}$$

for all choices of N. □

C.2 Cheating Bob

Lemma C.2. The optimal success probability (P_B^*) with which dishonest Bob can cheat in Protocol 4 is $\frac{3}{4}$ as $N \to \infty$.

Proof. The proof follows the same argument as before. A cheating Bob in the 1-out-of-2 protocol wants to know which bit Alice received. This is exactly what he needs to know whether or not Alice has received the bit value in the Rabin protocol we have constructed as well. Therefore, Bob's cheating probability is

$$P_B^* = \frac{3}{4}, \tag{24}$$

which follows directly from [ASR+21], and is only valid in the limit as $N \to \infty$. This limit assures the fraction tested becomes negligible disallowing further cheating possibilities for Bob. □

D Security proof of Protocol 5

We now prove the security of the coin flipping protocol with respect to the security of the ROT subroutine. We reproduce the protocol below, for convenience.

- Alice and Bob perform an ROT protocol with Bob's data bit denoted by y. We denote Alice's ROT output by the trit $\hat{g} \in \{0, 1, NULL\}$.
- Alice defines the assert bit $a \in \{0, 1\}$ to be 0 if $\hat{g} \in \{0, 1\}$, and to be 1 if $\hat{g} = NULL$.
- Bob chooses b uniformly at random and sends b to Alice.

- Alice sends a and $\hat{g}$ to Bob. Bob checks if $(a, y, \hat{g})$ are consistent, i.e., he checks if $\hat{g} = y$ when $a = 0$. If $\hat{g} \neq y$, Bob aborts.
- If Alice and Bob do not abort, then they both output HEADS if $a = b$, and TAILS otherwise.

D.1 Cheating Alice

Since Alice wants the outcome HEADS, she must send back $a = b$ in the last message. However, this is only accepted if $(a, y, \hat{g})$ are consistent.

If $b = 0$, then Alice must return $a = 0$, in which case Bob accepts the value of a if $\hat{g} = y$. The maximum probability that Alice can pass the test in this case is the maximum probability that she can learn y in the ROT protocol, which is precisely P_A^*.

If $b = 1$, then Alice must return $a = 1$, but in this case, Bob is expecting $\hat{g} = NULL$, and thus if she sends this, Bob has nothing to check, and will simply accept the value of a.

Since b is chosen independently from the rest of the protocol, we have that $P_{A,0} = \frac{1}{2} + \frac{1}{2}P_A^*$, as desired.

D.2 Cheating Bob

Since Bob want HEADS, he must infer the value of a ahead of time when he sends b. In other words, he must send $b = a$. Therefore, he is successful in forcing Alice to output HEADS with the same probability with which he can learn a in the protocol. Since the value of a corresponds to exactly whether or not Alice asserts the value of Bob's bit y, we have that the maximum probability with which Bob learns a in the coin flipping protocol is simply P_B^*.