CountCrypt: Quantum Cryptography between QCMA and PP

Eli Goldin[1], Tomoyuki Morimae[2], Saachi Mutreja[3], and Takashi Yamakawa[4]

[1]New York University (eli.goldin@nyu.edu)
[2]Yukawa Institute for Theoretical Physics, Kyoto University (tomoyuki.morimae@yukawa.kyoto-u.ac.jp)
[3]Columbia University (saachi@berkeley.edu)
[4]NTT Social Informatics Laboratories (takashi.yamakawa@ntt.com)

October 24, 2024

Abstract

We construct a quantum oracle relative to which **BQP** = **QCMA** but quantum-computation-classical-communication (QCCC) key exchange, QCCC commitments, and two-round quantum key distribution exist. We also construct an oracle relative to which **BQP** = **QMA**, but quantum lightning (a stronger variant of quantum money) exists. This extends previous work by Kretschmer [Kretschmer, TQC22], which showed that there is a quantum oracle relative to which **BQP** = **QMA** but pseudorandom state generators (a quantum variant of pseudorandom generators) exist.

We also show that QCCC key exchange, QCCC commitments, and two-round quantum key distribution can all be used to build one-way puzzles. One-way puzzles are a version of "quantum samplable" one-wayness and are an intermediate primitive between pseudorandom state generators and EFI pairs, the minimal quantum primitive. In particular, one-way puzzles cannot exist if **BQP** = **PP**.

Our results together imply that aside from pseudorandom state generators, there is a large class of quantum cryptographic primitives which can exist even if **BQP** = **QCMA**, but are broken if **BQP** = **PP**. Furthermore, one-way puzzles are a minimal primitive for this class. We denote this class "CountCrypt".

Contents

1 Introduction

Nearly all cryptographic primitives require some form of computational assumption [IL89]. For example, the existence of one-way functions (OWFs) implies that $\mathbf{P} \neq \mathbf{NP}$. However, some primitives seemingly require stronger computational assumptions than others. For instance, it appears that key exchange protocols require stronger computational assumptions than OWFs. In particular, OWFs can be built from key exchange protocols [BCG90], but there exist strong barriers to building key exchange from OWFs [IR89]. In general, the relationships between classical cryptographic primitives are well-understood. Notably, nearly all primitives can be used to build OWFs, and so the existence of any sort of cryptography necessitates that $\mathbf{P} \neq \mathbf{NP}$.

Everything changes once quantum computers enter the picture. It is not hard to define quantum versions of classical cryptographic primitives, where we allow either the protocol description and/or protocol inputs/outputs to be quantum. For instance, [JLS18] defines pseudorandom state generators (PRSGs), a quantum analogue of pseudorandom generators where the output is a quantum state, and it is indistinguishable from a Haar random quantum state. There exists a quantum oracle under which PRSGs exist but $\mathbf{P} = \mathbf{NP}$ [KQST23][1] or $\mathbf{BQP} = \mathbf{QMA}$ [Kre21]. This gives evidence that some quantum primitives require weaker assumptions even than OWFs.

Since [Kre21, KQST23], a slew of recent work has begun to map out the relationships between different forms of cryptographic hardness. To quickly summarize recent progress, it seems that most quantum cryptographic primitives fall into one of several buckets:

1. "QuantuMania": quantum cryptographic primitives which are broken if $\mathbf{BQP} = \mathbf{QCMA}$. This class includes three types of primitives. The first type is just quantumly-secure classical primitives, such as post-quantum OWFs or public-key encryption (PKE). The second type is primitives with "genuinely quantum" additional functionalities such as PKE with certified deletion [KNY23] and digital signatures with revocable signing keys [MPY23]. The third type is most non-interactive quantum-computation-classical-communication (QCCC) primitives, such as PKE or digital signatures with classical public keys, ciphertexts, or signatures [CGG24].

2. "CountCrypt": quantum cryptographic primitives which may exist if $\mathbf{BQP} = \mathbf{QCMA}$, but are broken by a $\#\mathbf{P}$ oracle. That is, CountCrypt primitives do not exist if $\mathbf{BQP} = \mathbf{PP}$. Thus, this category includes primitives such as PRSGs [JLS18, Kre21], (pure output) one-way state generators (OWSGs) [MY22, MY24], and one-way puzzles (OWPuzzs) [KT24a, CGG+23]. There is evidence that these primitives may exist even if $\mathbf{BQP} = \mathbf{QCMA}$ [Kre21]. On the other hand, (post-quantum) classical cryptography does not exist if $\mathbf{BQP} = \mathbf{QCMA}$. Thus, it seems plausible that these primitives could be built from weaker assumptions than those necessary for classical cryptography.[2]

3. "NanoCrypt": quantum cryptographic primitives which can exist if $\mathbf{BQP} = \mathbf{PP}$, and which we do not know how to break using *any* classical oracle. This category contains exclusively (single-copy) "quantum output" primitives. Examples of this category are EFI pairs [BCQ23], quantum bit commitments [Yan22], multiparty computation, and 1-copy secure PRSGs (1-PRSGs) [MY22]. There exists an oracle under which these primitives are secure against one-query to an arbitrary classical

[1] More precisely, this paper shows the existence of 1-PRSGs.

[2] Unlike the classical MicroCrypt, CountCrypt does not collapse to a single primitive. There exists a black-box separation between OWSG and OWPuzz, two CountCrypt primitives [BMM+24, BCN24].

oracle [LMW24]. Nevertheless, it is possible that these primitives all can indeed be broken by (multiple queries to) a #**P** oracle, and it is just that no attack has been discovered. There are also some black-box separations between some primitives in NanoCrypt and those in CountCrypt. In particular, there is an oracle under which 1-PRSGs (and therefore EFI pairs) exist, but (multi-copy secure) PRSGs do not [CCS24, AGL24].

Note that each of these categories also has an associated minimal primitive. Nearly every primitive in QuantuMania implies the existence of efficiently verifiable one-way puzzles (EV-OWPuzzs) [CGG24]. Like-wise, (inefficiently verifiable) one-way puzzles (OWPuzzs) appear to be minimal for CountCrypt [KT24a, KT24b], while EFI pairs seem minimal for NanoCrypt [MY24, BCQ23]. From this perspective, we can define QuantuMania as the class of primitives which can be used to build EV-OWPuzzs. We can define CountCrypt as the class of primitives which can be used to build OWPuzzs but not EV-OWPuzzs. And we can define NanoCrypt as the class of primitives which can be used to build EFI pairs but not OWPuzzs. For a figure illustrating the relationships between primitives in these three classes, see Figure 1.

While a great amount of progress has been made in recent years, there are a few primitives which we still do not know how to place within these broad categories.

Interactive protocols with classical communication. One interesting class of quantum primitives are those where communication is required to be done over classical channels. This class contains variants of most classical cryptographic primitives, where the only difference is that we allow the protocol itself to perform quantum operations. Examples include classical-communication key exchange (KE), classical-communication digital signatures, and EV-OWPuzz. We will denote such primitives as “QCCC” primitives [CLM23], referring to the fact that they use quantum computation and classical communication.[3] [KT24a, CGG24] showed that many QCCC primitives can be used to build EV-OWPuzzs. In particular, QCCC versions of encryption, signatures, and pseudorandomness all imply EV-OWPuzzs and so are broken if $\mathbf{BQP} = \mathbf{QCMA}$.

However, there are a few notable exceptions to this result. In particular, it is not clear how to build EV-OWPuzzs from QCCC primitives whose security game includes multiple rounds of communication. Key examples of such primitives, which we consider in this work, are QCCC KE and QCCC commitments. It is easy to see that QCCC KE and QCCC commitments can be broken with a #**P** oracle. Thus, QCCC KE and QCCC commitments should lie in either QuantuMania or CountCrypt. If we look from the perspective of minimal primitives, even less is known. It is unknown whether or not QCCC KE and QCCC commitments can be used to build any of EV-OWPuzzs, OWPuzzs, or EFI pairs[4].

2-round quantum key distribution. One of the most surprising initial results in quantum cryptography is that, with quantum communication, KE can exist *unconditionally* [Wie83, BB14][5]. Such protocols, termed quantum key distribution or QKD, are

[3] In the field of entanglement theory, such a setup has been called LOCC (which stands for local operation and classical communication). However, LOCC is usually used for statistical settings, i.e., local parties can do any unbounded computation.

[4] Note that [KT24a] constructs OWPuzzs from QCCC commitments with a *non-interactive* opening phase. However, this is not enough for our purposes. In particular, our construction of QCCC commitments in the oracle setting has an *interactive* opening phase.

[5] Assuming the existence of classical authenticated channel.

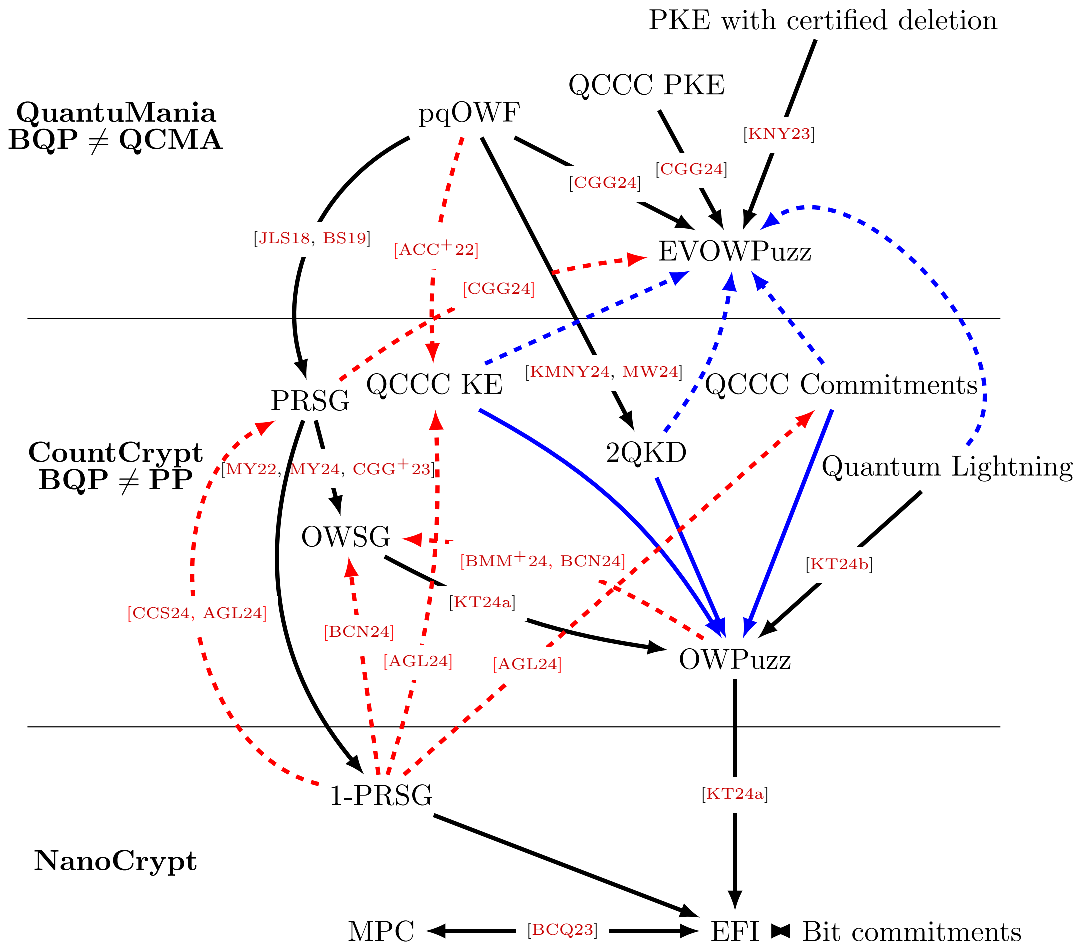

Figure 1: A graph of some known implications between primitives in QuantuMania, CountCrypt, and NanoCrypt. Dashed lines represent black-box separations. Blue lines are new in our work. pqOWF means post-quantum (i.e., quantumly-secure) OWFs. OWSG refer to pure state one-way state generators [MY22].

secure when performed over an authenticated classical channel and an *arbitrary, corrupted* quantum channel.

Interestingly, QKD can be performed in three rounds and still maintain information-theoretic security [GYZ17][6]. On the other hand, it is known that 2-round QKD (2QKD) is impossible information-theoretically [MW24]. However, there is no other known lower bound on the computational hardness of 2QKD. The best known upper bound is that 2QKD can be constructed from OWFs [KMNY24, MW24]. It is entirely an open question whether 2QKD lies in QuantuMania, CountCrypt, or NanoCrypt.

Quantum money/quantum lightning. Quantum money [Wie83, JLS18, AC12] is one of the most fundamentally "quantum" cryptographic primitives. Quantum money essentially models digital cash. In a quantum money scheme, a bank has the ability to mint quantum states which represent some currency. Money states should be verifiable: either the bank or other users should be able to tell if a money state is valid. Money states should also be uncloneable: it should not be possible for users to take several copies of a money state and turn it into more copies. Note that quantum money is fundamentally impossible if we require that states are classical. This is because it is easy to copy any classical string.

For the purposes of this work, we will consider a stronger variant of quantum money, known as quantum lightning [Zha19]. A quantum lightning scheme consists of a minter and a verifier. The minter produces a money state along with a corresponding serial number. The verifier, given a money state and a serial number, checks if the serial number matches. Security says that no adversary can construct two money states which verify under the same serial number. This is a stronger definition than standard quantum money, since adversary can choose which state they are trying to clone.

It is again easy to see that quantum lightning is broken if $\mathbf{BQP} = \mathbf{PP}$, since given the verification algorithm a $\#\mathbf{P}$ algorithm can simply find an accepting state ($\#\mathbf{P}$ oracles can be used to synthesize many implicitly defined states [KT24b, CGG$^+$23, INN$^+$22]). Furthermore, recent work has shown that OWPuzzs can be built from quantum lightning [KT24b], but given a quantum lightning scheme, it is not known how to build an EV-OWPuzz. Thus, it seems that quantum lightning lies either in CountCrypt or QuantuMania, but it is not clear which is the case. Note that the same reasoning applies to quantum money schemes as well.

Our results. The major question of our work is the following.

Do these primitives (QCCC KE, QCCC commitments, 2QKD, quantum lightning) lie in CountCrypt or QuantuMania?

and more specifically

Is there a way to build EV-OWPuzzs from any of these primitives?

It turns out that the answer to this latter question seems to be no. In particular, we show

1. There exists a quantum oracle relative to which $\mathbf{BQP} = \mathbf{QCMA}$, but QCCC KE, QCCC commitments, and 2QKD exist.

2. There exists a (different) quantum oracle relative to which $\mathbf{BQP} = \mathbf{QCMA}$, but quantum lightning exists.

[6]Note that in this protocol, only one party learns if the execution succeeds.

3. If any of QCCC KE, QCCC commitments, or 2QKD exist, then OWPuzzs exist.

Together, these results imply that quantum lightning, quantum money, QCCC KE and QCCC commitments, as well as 2QKD are all CountCrypt primitives.

2 Technical Overview

An oracle relative to which QCCC KE exist and BQP = QCMA. The key idea behind this oracle is that it functions as idealized KE, but where the internal state is forced to be quantum. As a toy example, consider the simple idealized KE protocol defined as follows.

1. Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be a random function accessible only to Alice and Bob.
2. Alice chooses x and sends x to Bob.
3. Bob chooses y and sends y to Alice.
4. Alice and Bob both output $F(x, y)$.

Note that this protocol is secure against any adversary who does not have access to F. This is because just given x and y, the adversary has no way to compute $F(x, y)$.

In order to define an oracle under which this KE protocol is secure, we need an oracle which allows Alice and Bob to query $F(x, y)$, while other parties cannot. To do this, we associate to each x and y a quantum key $|\phi_x\rangle$ and $|\phi_y\rangle$. We then guarantee that the only way for a party to compute $|\phi_x\rangle$ is if it chose x itself. We mandate that in order to query $F(x, y)$, a party must have access to either $|\phi_x\rangle$ or $|\phi_y\rangle$.

Our oracle will be defined respective to a Haar random unitary U and a random function $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. Our oracle consists of the following three oracles, SG, Mix, and $PSPACE$.

1. SG: samples a bit string x uniformly at random and outputs $x, |\phi_x\rangle$. Here $|\phi_x\rangle = U|x\rangle$.
2. $Mix(x, y, |\psi\rangle)$: Mix applies $U^\dagger$ on $|\psi\rangle$, measures the state in the computational basis, and outputs $F(x, y)$ if and only if the result is x or y.
3. $PSPACE$: an oracle for a **PSPACE** complete problem.

This leads to the following natural KE protocol.

1. Alice runs $SG \to (x, |\phi_x\rangle)$ and sends x to Bob.
2. Bob runs $SG \to (y, |\phi_y\rangle)$ and sends y to Alice.
3. Alice outputs $Mix(x, y, |\phi_x\rangle) \to F(x, y)$.
4. Bob outputs $Mix(x, y, |\phi_y\rangle) \to F(x, y)$.

Security follows from the fact that given x and y, it is impossible for an eavesdropper to compute $|\phi_x\rangle$ or $|\phi_y\rangle$, since with high probability queries to SG will never return these values. But without $|\phi_x\rangle$ or $|\phi_y\rangle$, it is impossible to query $F(x, y)$, and so the agreed key appears uniformly random. Note that since this security holds information-theoretically, it also holds against adversaries with access to the $PSPACE$ oracle.

Theorem 1 (Theorem 13 restated)**.** *QCCC KE exists relative to* $(SG, Mix, PSPACE)$*.*

BQP = QCMA relative to $(SG, Mix, PSPACE)$**.** We show that for any input x and oracle verifier $\mathsf{Ver}^{SG,Mix,PSPACE}$, there exists a verifier V such that for all w,

$$\Pr[V^{PSPACE}(x,w) \to 1] \approx \Pr[\mathsf{Ver}^{SG,Mix,PSPACE}(x,w) \to 1].$$

Thus, given x and $\mathsf{Ver}^{SG,Mix,PSPACE}$, we can use the $PSPACE$ oracle to check if V^{PSPACE} has a witness w accepting x. Since V^{PSPACE} acts similarly to $\mathsf{Ver}^{SG,Mix,PSPACE}$, this also reveals whether $\mathsf{Ver}^{SG,Mix,PSPACE}$ has a witness accepting x. And so we can use the $PSPACE$ oracle to decide languages in $\mathbf{QCMA}^{SG,Mix,PSPACE}$.

We take inspiration from [Kre21], which observes that with all but doubly exponential probability over U, any efficient algorithm querying U acts identically if the algorithm instead samples U itself. Thus, by a union bound, $\mathsf{Ver}(x,\cdot)$ acts the same on all witnesses when we replace U with a freshly sampled Haar random unitary. We observe that a similar property holds for random oracles, as long as the algorithm querying the random oracle can only query at random points.

Thus, we define V^{PSPACE} to simulate a Haar random unitary U' and random oracle F' using a t-design and t-wise independent function respectively. V^{PSPACE} will act the same as $\mathsf{Ver}^{SG,Mix,PSPACE}$, but whenever Ver would query either SG or Mix, V will simulate the query replacing U with U' and F with F' respectively.

Theorem 2 (Theorem 15 restated)**.** $\mathbf{BQP}^{SG,Mix,PSPACE} = \mathbf{QCMA}^{SG,Mix,PSPACE}$.

A concentration bound for random oracles. The key observation that underlies our proof that **BQP** = **QCMA** is that with very high probability over the choice of a fixed random function F, F is indistinguishable from a fresh random function F' under random query access. The intuition here stems from the similarity between a fixed random oracle and the auxiliary-input random oracle model (AI-ROM) [Unr07].

The AI-ROM is a variant of the random oracle model which allows adversaries to have non-uniform advice *depending on the chosen random oracle*. One of the most useful ways to prove things in the AI-ROM is the so-called presampling technique [Unr07, CDGS18, GLLZ21]. A bit-fixing, or presampling, oracle is a random function where its evaluation on polynomially many fixed points is chosen adversarially. The presampling technique states that any protocol secure with bit-fixing oracles is also secure in the AI-ROM.

Intuitively, the only way to distinguish a fixed random function F from a fresh F' is by hardcoding some information about F into the distinguisher. This is morally equivalent to giving the distinguisher non-uniform advice depending on F, i.e. the AI-ROM. In fact, we show that a variant of the presampling lemma also holds for fixed random functions.

Lemma 1 (Presampling lemma for fixed random functions, Theorem 12 restated)**.** *Let $\mathcal{D}$ be the uniform distribution over functions $\{0,1\}^n \to \{0,1\}^{n'}$. Let $\mathcal{A}^{\mathcal{O}}$ be any PPT algorithm making at most T queries to its oracle. With probability $\geq 1 - 2^{-2^{n/8}}$ over $f \leftarrow \mathcal{D}$, there exists a bit-fixing source $\mathcal{D}_f$ fixing at most $\frac{2^{n/8}(2^{n/8}+2n)}{n'}$ points such that* $\left|\Pr[\mathcal{A}^f \to 1] - \Pr_{f' \leftarrow \mathcal{D}_f}[\mathcal{A}^{f'} \to 1]\right| \leq \mathsf{negl}(n)$.

But note that under random query access, since the fixed-points take up a negligible fraction of query inputs, the probability that a fixed-point is queried is negligible. Thus, bit-fixing oracles are indistinguishable from truly random functions under random query access. And so, under random query access, with extremely high probability a fixed random function is indistinguishable from a fresh random function.

QCCC commitments exist under this oracle. We can observe that under the same oracle, we can also build QCCC commitments. Our protocol goes as follows.

1. The committer runs $SG \to x, |\phi_x\rangle$ until the first bit of x is the bit it is trying to commit to. The committer then sends all but the first bit of x to the receiver, which we will denote $x_{>1}$.

2. To open to a bit b, the receiver runs $SG \to y, |\phi_y\rangle$ and sends y to the committer.[7]

3. Upon receiving y, the committer computes $Mix(x, y, |\phi_x\rangle) \to c$ and sends c to the receiver.

4. The receiver sets $x = (b, x_{>1})$ and accepts if $Mix(y, x, |\phi_y\rangle) = c$.

Hiding of this protocol is trivial, since the committer just sends a random string in the commit phase. Binding is slightly more difficult, but the proof follows very similar lines to the security of KE. In particular, the only way that the committer can open to a bit b is if it knows $|\phi_{b,x_{>1}}\rangle$. However, it is impossible for the receiver to know both $|\phi_{0,x_{>1}}\rangle$ and $|\phi_{1,x_{>1}}\rangle$, since it would have needed to receive both of these separately from SG. Thus, once the committer knows $|\phi_{b,x_{>1}}\rangle$, it can not change its mind and decide to open to $\bar{b}$.

Theorem 3 (Theorem 14 restated)**.** *QCCC commitments exist relative to* $(SG, Mix, PSPACE)$.

2QKD exists under this oracle. We also observe that under the same oracle, 2QKD exists. This follows trivially from our construction of KE. In particular, any QCCC KE protocol is also a QKD protocol which does not make use of the quantum channel. Thus, since our QCCC KE protocol is two-round, it is also a 2QKD protocol.

Corollary 1. *2QKD exists relative to* $(SG, Mix, PSPACE)$.

An oracle relative to which quantum lightning exists and BQP = QMA. Using similar ideas, we also describe an oracle under which quantum lightning exists but **BQP** = **QMA** (note that this is strictly stronger than **BQP** = **QCMA**, since **QCMA** $\subseteq$ **QMA**). The oracle will essentially define the quantum lightning scheme. In particular, for a Haar random unitary U, we define our oracle (SG, V, C) as follows.

1. SG: samples a bit string x uniformly at random and outputs $x, |\phi_x\rangle$. Here $|\phi_x\rangle = U|x\rangle$.

2. $V(x, |\psi\rangle)$: V applies $U^\dagger$ on $|\psi\rangle$, measures the state in the computational basis, and outputs 1 if and only if the result is x.

3. C: a non-trivial oracle independent of U serving as a "**QMA** breaker". In particular, this will be the same oracle as the one used in [Kre21].

We can observe that (SG, V) is exactly a quantum lightning scheme. (SG is the minter, where x is the serial number and $|\phi_x\rangle$ is the money state. V is the verifier.) In particular, money states will always verify with their corresponding serial number. However, it is not possible to generate $x, |\phi_x\rangle \otimes |\phi_x\rangle$. Security goes by a reduction to the complexity-theoretic no-cloning theorem defined by [AC12]. In particular, this theorem states that with only

[7] This step can be done in the commit phase, not the opening phase. One advantage of doing it in the opening phase is that the receiver does not need to store state long term and the commit phase becomes non-interactive.

polynomially many queries to a verification oracle, it is impossible to clone a Haar random state over $10n$ qubits with success probability $\Omega(2^{-5n})$. We will then set the length of $|\phi_x\rangle$ to be $10n$ qubits long. To clone a Haar random state $|\psi\rangle$ using a quantum lightning breaker, we will simply replace a random $|\phi_k\rangle$ with $|\psi\rangle$. With probability approximately 2^{-n}, the quantum lightning breaker will clone $|\psi\rangle$, violating complexity-theoretic no-cloning.

Theorem 4 (Theorem 18 restated)**.** *Quantum lightning exists relative to* (SG, V, C).

Showing that **BQP** = **QMA** is easier here, since our oracles are simpler. In particular, the argument goes through similar lines to [Kre21].

Theorem 5. $\mathbf{BQP}^{SG,V,C} = \mathbf{QMA}^{SG,V,C}$.

Building OWPuzzs from QCCC KE. To construct OWPuzzs from QCCC KE, we look to the construction of OWFs from classical KE [IL89]. In particular, this paper constructs distributional OWFs from KE, and shows that distributional OWFs and OWFs are equivalent. A distributional OWF is a OWF which is hard to *distributionally* invert. That is, given $y = f(x)$, it is hard to sample from the set $f^{-1}(y)$.

Recent work [CGG24] analogously defines a distributional variant of OWPuzzs. In particular, a distributional OWPuzzs consists of a sampler producing puzzle-key pairs (s, k) such that given a puzzle s, it is hard to sample from the conditional distribution over keys k. It turns out that distributional OWPuzzs and OWPuzzs are equivalent [CGG24].

It is then not hard to build a distributional OWPuzzs from a QCCC KE. In particular, the puzzle will be the transcript (i.e., the sequence of messages exchanged) between Alice and Bob while the OWPuzz key will be the corresponding shared key of the KE. By correctness, the distribution over shared keys will be close to constant. Distributionally inverting this puzzle is enough to recover the shared key and thus break the KE.

Theorem 6 (Theorem 20 restated)**.** *If there exist QCCC KE, then there exist OWPuzzs.*

Building OWPuzzs from QCCC commitments. To build OWPuzzs from QCCC commitments, we again take advantage of distributional OWPuzzs. In particular, we observe that if distributional OWPuzzs do not exist, then it is possible to do conditional sampling for any sampleable distribution. This gives an approach to break hiding for QCCC commitment schemes.

In particular, given the transcript τ of the commit phase for a commitment to a bit b, a corrupted receiver can sample the committer's first message in the opening phase conditioned on τ.[8] The receiver can then repeat this process to sample the committer's second and third messages, up until the entire opening phase has been simulated. Since the receiver can run the entire opening phase on its own, it can learn the output b. By binding, this should be the same as the committed bit, and so this attack breaks hiding.

Theorem 7 (Theorem 21 restated)**.** *If there exist QCCC commitments, then there exist OWPuzzs.*

Building OWPuzzs from 2QKD. Formally, a QKD protocol is a protocol where Alice and Bob agree on a shared key over a public authenticated classical channel and a public arbitrary quantum channel. In particular, a good QKD protocol should satisfy three properties.

[8] Remember that we are now considering general commitments where the opening phase is not necessarily non-interactive.

1. (Correctness): If the protocol is executed honestly, then Alice and Bob output the same key.
2. (Validity): An adversarial party Eve, who controls the quantum channel, cannot make Alice and Bob output different keys without them noticing and outputting $\perp$.
3. (Security): If either Alice or Bob does not output $\perp$, then Eve can not learn their final key.

A two-round QKD protocol, or 2QKD is simply a QKD protocol where Alice and Bob each send at most one message. It is known that 2QKD protocols can be built from post-quantum OWFs [KMNY24, MW24].

A concurrent work defines a quantum version of OWPuzzs, a state puzzle, and shows that state puzzles are equivalent to OWPuzzs [KT24b]. A state puzzle is essentially a one-way puzzle where the key is a quantum state. Since the key is quantum, there is a canonical verification algorithm which just projects onto the target state.[9]

We observe that the existence of a 2QKD protocol immediately gives a state puzzle. In particular, the puzzle will consist of Alice's first classical message to Bob, and the key will be Alice's internal state as well as her first quantum message.[10] In particular, if this is not a state puzzle, then we can attack the 2QKD protocol by replacing Alice's quantum message with the result from the state puzzle breaker. This will allow an adversary to impersonate Alice and obtain a corresponding private state. Since Alice does not send any more messages beyond this, Bob will not be able to detect this impersonation and so the adversary will be able to predict Bob's output.

Theorem 8 (Theorem 22 restated)**.** *If there exist 2QKD, then there exist OWPuzzs.*

On our separating oracles. The oracles that we consider in this work are arbitrary completely-positive-trace-preserving (CPTP) maps. This is a slightly weaker model than previous quantum oracle separations [Kre21, KQST23, CCS24, AGL24]. In particular, a separation of this form rules out black-box reductions where the reduction does not make inverse queries to the oracles and does not "purify" the oracles, making use of ancilla qubits.

3 Preliminaries

3.1 Notation

Throughout the paper, $[n]$ denotes the set of integers $\{1, 2, \ldots, n\}$. If X is a set, we use $x \leftarrow X$ to denote that x is sampled uniformly at random from X. Similarly, if A is an algorithm and x is an input, we use $y \leftarrow A(x)$ to say that y is the output of A on input x. For a distribution D, $x \leftarrow D$ means that x is sampled according to D. A function f is *negligible* if for every constant $c > 0$, $f(n) \leq \frac{1}{n^c}$ for all sufficiently large $n \in \mathbb{N}$. We will use $\mathsf{negl}(n)$ to denote an arbitrary negligible function, and $\mathsf{poly}(n)$ to denote an arbitrary polynomially bounded function. For two interactive algorithms A and B that interact over a classical or quantum channel, $A \rightleftarrows B$ means the execution of the interactive protocol, or the distribution of the transcript, i.e. the sequence of messages exchanged between A

[9] Note that we can always assume that the key is a pure state, since adding the extra qubits from its purification into the key only improves security.

[10] Remember that in QKD we have authenticated classical channel and non-authenticated quantum channel. We therefore assume that Alice's first message consists of a quantum part and a classical part.

and B. $A \leftrightarrows B \to y$ means that the execution of the interactive protocol produces output y. We use the abbreviation QPT for a uniform quantum polynomial time algorithm. We will also use non-uniform QPT to refer to quantum polynomial time algorithms with polynomial *classical* advice. We use the notation $A^{(\cdot)}$ to refer to a (classical or quantum) algorithm that makes queries to an oracle. We use $\mathsf{TD}(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ to denote the trace distance between density matrices ρ and σ. For two distributions D_1, D_2, we use $\Delta(D_1, D_2) = \frac{1}{2}\sum_x |\Pr[x \leftarrow D_1] - \Pr[x \leftarrow D_2]|$ to denote the total variation distance (aka statistical distance) between D_1 and D_2. We write $D_1 \approx_\epsilon D_2$ to say that D_1 and D_2 are ϵ indistinguishable. That is, for all QPT $\mathcal{A}$, $|\Pr_{x\leftarrow D_1}[1 \leftarrow \mathcal{A}(x)] - \Pr_{x\leftarrow D_2}[1 \leftarrow \mathcal{A}(x)]| \le \epsilon$. We use $D_1 \approx D_2$ to mean $D_1 \approx_{\mathsf{negl}(n)} D_2$. For algorithms which take in a unary security parameter 1^n, we will sometimes omit 1^n as input when the value is implicit.

3.2 Quantum Information

A quantum channel $\mathcal{A}$ is any completely positive trace preserving (CPTP) map. Intuitively, quantum channels represent all "physical" processes. That is, the set of CPTP maps is the quantum analogue of the set of all functions.

For a quantum channel $\mathcal{A}$, we let $\|A\|_\diamond$ denote its *diamond norm*. The diamond norm and trace distance satisfy the following relation:

Property 9 ([NC11]). *Let $\mathcal{A}$ and $\mathcal{B}$ be quantum channels and ρ be a density matrix. Then, $\mathsf{TD}(\mathcal{A}(\rho), \mathcal{B}(\rho)) \le ||\mathcal{A} - \mathcal{B}||_\diamond$.*

We note the following relationship between trace distance and fidelity of pure states.

Lemma 2. *Let $|\phi\rangle, |\psi\rangle$ be two pure states. Then,*

$$\mathsf{TD}(|\phi\rangle\langle\phi|, |\psi\rangle\langle\psi|) = \sqrt{1 - |\langle\phi|\psi\rangle|^2}.$$

In this work, we will often consider quantum channels corresponding to performing some projective measurement, which we call measurement channels.

Definition 1 (Measurement Channels). *Let $\mathcal{M} = \{\Pi, I - \Pi\}$ be a binary measurement corresponding to a projector Π. The measurement channel corresponding to $\mathcal{M}$ is the following channel:*

$$\rho \mapsto \Pi\rho\Pi + (I - \Pi)\rho(I - \Pi).$$

We make use of the following fact about measurement channels.

Lemma 3 ([Kuk18]). *Let $\mathcal{A} = \{\Pi_A, I - \Pi_A\}$, $\mathcal{B} = \{\Pi_B, I - \Pi_B\}$ be two measurement channels. Then*

$$||\mathcal{A} - \mathcal{B}||_\diamond = 2\max_\rho |\mathrm{Tr}(\Pi_A\rho) - \mathrm{Tr}(\Pi_B\rho)|.$$

3.3 Haar Measure And Its Concentration

We use $\mathbb{U}(N)$ to denote the group of $N \times N$ unitary matrices, and μ_N to denote the Haar measure on $\mathbb{U}(N)$. We use $\mathbb{C}^N$ to denote the set of quantum pure states of dimension N, and μ_N^s to denote the Haar measure on $\mathbb{C}^N$. Given a metric space $(\mathcal{M}, d)$ where d denotes the metric on the set $\mathcal{M}$, a function $f : \mathcal{M} \to \mathbb{R}$ is L-lipschitz if for all $x, y \in \mathcal{M}, |f(x) - f(y)| \le L \cdot d(x, y)$. We will in particular be concerned about the Frobenius norm. For a complex matrix M, its Frobenius norm is $\sqrt{\mathrm{Tr}(M^\dagger M)}$.

The following inequality involving Lipschitz continuous functions captures the strong concentration of Haar measure.

Theorem 10 ([Mec19]). *Let $k \geq 1$ be an integer. Given $N_1, N_2, \ldots, N_k \in \mathbb{N}$, let $X = \mathbb{U}(N_1) \bigoplus \cdots \bigoplus \mathbb{U}(N_k)$ be the space of block diagonal unitary matrices with blocks of size $N_1, N_2, \ldots, N_k$. Let $\mu = \mu_{N_1} \times \cdots \times \mu_{N_k}$ be the product of Haar measures on X. Suppose that $f : X \to \mathbb{R}$ is L-Lipshitz with respect to the Frobenius norm. Then for every $t > 0$,*

$$\Pr_{U \leftarrow \mu}[|f(U) - \mathbb{E}_{V \leftarrow \nu}[f(V)]| \geq t] \leq 2 \exp\left(-\frac{(N-2)t^2}{24L^2}\right),$$

where $N = \min(N_1, \ldots, N_k)$.

Lemma 4 ([Kre21]). *Let $A^{(\cdot)}$ be a quantum algorithm that makes T queries to an oracle operating on D qubits, and let $|\psi\rangle$ be any input to $A^{(\cdot)}$. For $U \in \mathbb{U}(D)$, define $f(U) = \Pr[A^U(|\psi\rangle) = 1]$. Then, f is $2T$-Lipshitz with respect to the Frobenius norm.*

3.4 Quantum Cryptography Definitions

Definition 2 (QCCC Key Exchange). *A QCCC key exchange (KE) protocol is an efficient interactive two-party protocol between two parties A, B. At the end of the protocol, $A \leftrightarrows B$ produces a classical transcript τ, A's output bit a, and B's output bit b. A KE protocol must satisfy the following properties.*

1. *Correctness: At the conclusion of the protocol, both parties agree on the output. Formally,*
$$\Pr[a = b : A \leftrightarrows B \to (\tau, a, b)] \geq 1 - \mathsf{negl}(n).$$

2. *Security: An adversary with the ability to see all communication should not be able to intercept the output. That is, for all non-uniform QPT adversaries $\mathcal{A}$,*
$$\Pr[a \leftarrow \mathcal{A}(\tau) : A \leftrightarrows B \to (\tau, a, b)] \leq \mathsf{negl}(n).$$

Definition 3 (Quantum Lightning [Zha19]). *A quantum lightning protocol is a tuple $(Mint, Ver)$ of uniform QPT algorithms with the following syntax.*

1. *$Mint(1^n) \to (\sigma, |\$\rangle)$: takes in the security parameter 1^n and outputs a serial number σ and a money state $|\$\rangle$.*

2. *$Ver(1^n, \sigma, |\$\rangle) \to 0/1$: takes in the security parameter 1^n as well as a serial number σ and money state $|\$\rangle$ and outputs accept (1) or reject (0).*

We require the following properties.

1. *Correctness: Honestly generated money states verify under their serial number. That is,*
$$\Pr[0 \leftarrow Ver(1^n, Mint(1^n))] \leq \mathsf{negl}(n).$$

2. *Security: No attacker can generate two money states for the same serial number. That is, for all non-uniform QPT attackers $\mathcal{A}$,*
$$\Pr\left[\begin{array}{c} 1 \leftarrow Ver(1^n, \sigma, \rho_A) \\ 1 \leftarrow Ver(1^n, \sigma, \rho_B) \end{array} : (\sigma, \rho_{AB}) \leftarrow \mathcal{A}(1^n)\right] \leq \mathsf{negl}(n).$$

   Here, ρ_{AB} is a state over two registers A an B, and ρ_A (ρ_B) is the state on the register A (B).

Definition 4 (Distributional one-way puzzles from [CGG24])**.** *A β distributional one-way puzzle (OWPuzz) is a uniform QPT algorithm $\mathsf{Samp}(1^n) \to (k, s)$ which takes in a security parameter n and produces two classical bit strings, a key k and a puzzle s, such that given a puzzle s, it is computationally infeasible to sample from the conditional distribution over keys. More formally, we require that for all non-uniform QPT algorithms $\mathcal{A}$, for all sufficiently large $n \in \mathbb{N}$,*

$$\Delta((k, s), (\mathcal{A}(1^n, s), s)) \geq \beta(n),$$

where $(k, s) \leftarrow \mathsf{Samp}(1^n)$.

Remark 1. *Distributional OWPuzzs are equivalent to OWPuzzs [CGG24], defined in [KT24a]. A OWPuzz is a QPT sampler paired with an inefficient verifier such that given a puzzle s, it is hard to find a key k which verifies. Since we will not work directly with OWPuzzs, we omit the formal definition.*

Definition 5 (QCCC Commitments from [KT24a])**.** *A QCCC commitment scheme is an efficient two-party protocol between a QPT committer Com and a QPT receiver Rec consisting of a commit stage and an opening stage operating on a private input m described as follows.*

1. *Commit stage: both parties receive a unary security parameter 1^n. The committer Com receives a private input m. It interacts with the receiver Rec using only classical messages, and together they produce a transcript z. At the end of the stage, both parties hold a private quantum state ρ_{Com} and ρ_{Rec} respectively.*

2. *Opening stage: both parties receive the transcript z as well as their private quantum states ρ_{Com} and ρ_{Rec} respectively. The committer Com interacts with the receiver Rec using only classical messages. At the end of the stage, the receiver either outputs a message or the reject symbol $\perp$.*

We require the following two properties.

1. *Correctness: For all messages m, when Com and Rec interact honestly, the probability that Rec outputs m at the end of the opening stage is at least $1 - \mathsf{negl}(n)$.*

2. *(Computational) hiding: For all $m \neq m'$ and for all non-uniform QPT adversarial receivers Rec′, the transcript of the interaction between Rec′ and the committer with input m is indistinguishable from the transcript of the interaction between Rec′ and the committer with input m'. That is,*

$$Com(m) \rightleftarrows Rec' \approx Com(m') \rightleftarrows Rec'.$$

3. *(Computational weak honest) binding: For all m and for any non-uniform QPT adversarial committer Com′, the probability that Com′ wins the following game is $\leq \mathsf{negl}(n)$.*

   (a) In the first stage, an honest receiver Rec interacts with the honest committer Com to produce a transcript z and receiver state ρ_{Rec}.

   (b) In the second stage, the honest receiver Rec is given ρ_{Rec} and z, while Com′ is given z (but not ρ_{Com}). They then proceed to run the opening stage with the committer replaced by Com′, and Rec produces a final output m'. Com′ wins if $m' \neq m$ and $m' \neq \perp$.

Remark 2. *Here we use a very weak security definition for QCCC commitments. This only serves to make our construction of OWPuzz from QCCC commitments stronger. For our separation, we use a stronger security notion detailed in Section 5.3.*

4 Simulating Oracles

Before describing our oracle separations, we need to introduce a crucial procedure that allows an adversary to "simulate" general access to a family of fixed random unitaries $U = \{U_n\}_{n\in\mathbb{N}}$ and a family of fixed random functions $F = \{F_n\}_{n\in\mathbb{N}}$ using only a few queries to U and F.

4.1 Haar Random Unitary Simulation Procedure – High Level Overview

As observed by [Kre21], the main property of Haar random unitaries which allows them to be simulated is their strong concentration. Let V be a quantum algorithm that makes T queries to a Haar random unitary U on n qubits. Then, with high probability over sampling a pair of Haar random unitaries U, U', the output distributions of V^U and $V^{U'}$ are within small TV distance. Concretely, for every $\delta > 0$,

$$\Pr_{U,U'\leftarrow\mu_{2^n}}\left[\left|\Pr[V^U = 1] - \Pr[V^{U'} = 1]\right| \le \delta\right] \ge 1 - \exp\left(-\Omega\left(\frac{2^n\delta^2}{T^2}\right)\right).$$

This inequality tells us that we can *replace* all queries to U with a freshly sampled Haar random unitary U', by having the TV increase only negligibly.
However, the concentration inequality only holds when the dimension 2^n is large enough with respect to the number of queries T. In order to get around this issue (and simulate U_ℓ, $\forall \ell = o(\log(T))$), we can "learn" the unitaries with $2^{\log(poly(T))} = poly(T)$ queries to the unitaries. To summarize, let V be a quantum algorithm that takes inputs of size λ. Suppose V makes $poly(\lambda)$ queries to U. With high probability over U, the output distribution of V on *all* standard basis inputs can be simulated to within a small constant in TV distance with only $poly(\lambda)$ queries to U by:

- Using T-designs to replace the oracle calls to U_ℓ, for large enough ℓ (such as $\ell \ge c \cdot \log\lambda$ for some large enough constant c). This perfectly simulates T queries to a freshly sampled family U_ℓ, and, by the concentration property, this in turn approximates the original output distribution of each run to within a small δ, where we can take δ to be a small constant[11]. This step does not require any queries to U.
- Efficiently learning all of the unitaries in U_ℓ, for all small enough ℓ (e.g. $\ell \le \Theta(\log\lambda)$). This step requires $poly(\lambda)$ time and queries to U.

4.2 Lemmas for Simulating Random Unitaries

We here define the necessary lemmas to make the simulation procedure defined in the previous subsection explicit.

State Tomography.

Theorem 11 (Quantum State Tomography, Proposition 2.2 from [HKOT23])**.** *There is a pure state tomography procedure with the following behavior. Let $d \in \mathbb{N}$. For any $n > d, \epsilon > 0$, given $O(n/\epsilon)$ copies of a pure state $|\phi\rangle \in \mathbb{C}^d$, it outputs a classical description of an estimate pure state $|\phi'\rangle$ such that with probability $\ge 1 - e^{-5n}$,*

$$\left|\langle\phi|\phi'\rangle\right|^2 \ge 1 - \epsilon.$$

[11]Note that we cannot take δ to be exponentially small in λ here because ℓ could be as small as $c \cdot \log\lambda$.

Remark 3. *Note that the statement from [HKOT23] specifies an error depending on the dimension of the state. However, this stronger statement is immediate upon observing the proof, and indeed the predecessor of this result from [GKKT20] makes the dependence on n explicit.*

Approximate t-Designs. An ϵ-approximate quantum unitary t-design is a distribution over unitaries that "ϵ-approximates" a Haar random unitary, when considering their action via a t-copy parallel repetition.

Definition 6 (Approximate Unitary Design [Kre21])**.** *Let $\epsilon \in [0,1]$ and $t, N \in \mathbb{N}$. A probability distribution S over $\mathbb{U}(N)$ is an ϵ-approximate unitary t-design if:*

$$(1-\epsilon)\mathbb{E}_{U \leftarrow \mu_N}[(U(.)U^\dagger)^{\otimes t}] \preceq \mathbb{E}_{U \leftarrow S}[(U(.)U^\dagger)^{\otimes t}] \preceq (1+\epsilon)\mathbb{E}_{U \leftarrow \mu_N}[(U(.)U^\dagger)^{\otimes t}],$$

where $B \preceq A$ means that $A - B$ is positive semidefinite.

It is well-known that there are efficient constructions of such unitary t-designs.

Lemma 5 ([Kre21])**.** *There exists $m : \mathbb{N} \to \mathbb{N}$, such that the following holds. For each $n, t \in \mathbb{N}$, and $\epsilon > 0$, there is a $poly(n, t, \log(\frac{1}{\epsilon}))$-time classical algorithm A that takes $m(n)$ bits of randomness as input, and outputs a description of a unitary quantum circuit on n qubits such that the output distribution of A is a phase-invariant ϵ-approximate unitary t-design (over $\mathbb{U}(2^n)$).*

An ϵ-approximate unitary t-design S is said to be *phase-invariant* if, for any unitary U in the support of S, U and ωU are sampled with the same probability, where ω is the $(t+1)$-th root of unity. The following lemma says that replacing the Haar measure with a phase-invariant ϵ-approximate unitary $2t$-design is undetectable to any algorithm that makes only t queries to the Haar random unitary.

Lemma 6 ([Kre21] for the result without inverse queries, [Yam] for the result with inverse queries)**.** *Let S be a phase-invariant $\frac{\epsilon}{2^t N^{9t}}$-approximate unitary $2t$-design over $\mathbb{U}(N)$, and let $D^{(\cdot)}$ be any t-query quantum algorithm. Then,*

$$\left| \Pr_{U \leftarrow S}[D^{U,U^\dagger} = 1] - \Pr_{U \leftarrow \mu_N}[D^{U,U^\dagger} = 1] \right| \leq \epsilon.$$

4.3 Simulating Fixed Random Functions

We will argue that a **QCMA** verifier Ver is not substantially more powerful than a **BQP** machine at learning non trivial properties about the fixed random function f. Analogously to the unitary case, we will show a "concentration" bound for random functions. In particular, we prove that with extremely high probability, a random function is indistinguishable from a bit-fixing oracle.

Definition 7 (Bit Fixing Oracle, [Unr07, CDGS18])**.** *A k-bit-fixing oracle X is a distribution over functions $\{0,1\}^m \to \{0,1\}^n$ with the following property:*
there exists some distribution Y over tuples of k elements of $\{0,1\}^m \times \{0,1\}^n$ such that X is defined as

1. *Sample $(x_1, y_1), \ldots, (x_k, y_k) \leftarrow Y$.*
2. *Output a random function f conditioned on $f(x_i) = y_i$ for each $i \in [k]$.*

The key observation we make is that the proof from [CDGS18] showing the equivalence between the AI-ROM and bit-fixing oracles actually shows that any flat distribution over random functions with sufficiently high entropy is indistinguishable from a bit-fixing oracle.

For any fixed distinguisher $\mathcal{A}$, we define $\mathcal{D}'_f$ to be the uniform distribution over functions which $\mathcal{A}$ cannot distinguish from f. Our key observation is that with extremely high probability, $\mathcal{D}'_f$ has very high entropy. Thus, by [CDGS18], since $\mathcal{A}$ cannot distinguish $\mathcal{D}'_f$ from f, there is some bit-fixing source which $\mathcal{A}$ cannot distinguish from f. Note that this means that the corresponding bit-fixing source is dependent on the adversary, but this is sufficient for our purposes.

Lemma 7. *Fix any oracle algorithm $\mathsf{Ver}^{(\cdot)}$. Let $\alpha, \beta \in \mathbb{R}$, $N, M \in \mathbb{N}$. For a function $f : [N] \to [M]$, define $acc(f) = \Pr[\mathsf{Ver}^f \to 1]$. Define D'_f to be the uniform distribution over functions $f' : [N] \to [M]$ such that $acc(f) - \alpha \le acc(f') \le acc(f) + \alpha$. Sample f uniformly from functions $[N] \to [M]$. With probability $\ge 1 - \frac{\beta}{\alpha}$,*

$$H_\infty(D'_f) \ge N \log M - \log \frac{1}{\beta}.$$

Proof. Let $val(f) = i$ if and only if $acc(f) \in [i\alpha, (i+1)\alpha)$. Define $w(i) = \Pr_f[val(f) = i]$. Define D'_f to be the uniform distribution over f' such that $val(f') = val(f)$. Note that D'_f and D_f are both uniform, but D'_f samples from a subset of the set D_f samples from. Thus, $H_\infty(D'_f) \le H_\infty(D_f)$. In particular, for any fixed f, D'_f is the uniform distribution over $M^N \cdot w(val(f))$ different functions. Thus,

$$H_\infty(D'_f) = N \log M + \log w(val(f)).$$

We also have

$$\begin{aligned}\Pr_f[w(val(f)) \le \beta] &= \sum_{i:w(i)\le\beta} \Pr_f[val(f) = i] = \sum_{i:w(i)\le\beta} w(i) \\ &\le \sum_{i:w(i)\le\beta} \beta \le \frac{\beta}{\alpha}\end{aligned}$$

and so with probability $\ge 1 - \frac{\beta}{\alpha}$, $w(val(f)) \ge \beta$ and so

$$H_\infty(D'_f) = N \log M + \log w(val(f)) \ge N \log M - \log \frac{1}{\beta}.$$

$\square$

Applying the techniques of [CDGS18] then gives

Theorem 12. *Let $\mathcal{D}$ be the uniform distribution over functions $\{0,1\}^n \to \{0,1\}^{n'}$. Let $\mathcal{A}^{\mathcal{O}}$ be any algorithm making at most T classical queries to its oracle. With probability $\ge 1 - 2^{-2^{n/8}}$ over $f \leftarrow \mathcal{D}$, there exists a k-bit-fixing oracle $\mathcal{D}_f$ such that*

$$\left| \Pr[\mathcal{A}^f \to 1] - \Pr_{f' \leftarrow \mathcal{D}_f}[\mathcal{A}^{f'} \to 1] \right| \le \mathsf{negl}(n)$$

with

$$k = \frac{2^{n/8}(2^{n/8} + 2n)}{n'}.$$

4.4 Proof of Random Function Concentration (Theorem 12)

In this section, we use the techniques of [CDGS18] to formally prove Theorem 12.

Definition 8 (Dense Source [CDGS18])**.** *An (N, M) source is a random variable X with range $[M]^{[N]}$, in other words a distribution over functions $[N] \to [M]$. A source is called $(k, 1-\delta)$ dense if it is fixed (i.e. constant) in at most k coordinates and, for every subset $I \subseteq [N]$,*

$$H_\infty(X_I) \geq (1-\delta) \cdot |I| \cdot \log M = (1-\delta) \cdot \log M^{|I|}.$$

where X_I is the distribution X restricted to coordinates $x_i \in I$.

Definition 9 (Bit Fixing Source)**.** *A (N, M) source X is called k bit-fixing if it is fixed (i.e. constant) on at most k coordinates and is uniform on the rest.*

Let X be distributed uniformly over $[M]^{[N]}$, and $Z := f(X)$, where $f : [M]^{[N]} \to \{0,1\}^S$ is an arbitrary function. Let X_z be the distribution of X conditioned on $f(X) = z$. Let $S_z = N \log M - H_\infty(X_z)$ be the minimum entropy deficiency of X_z. Let $\gamma > 0$ be arbitrary. Then, the following lemma holds:

Lemma 8 ([CDGS18])**.** *For every $\delta > 0$, X_z is γ close to a convex combination of finitely many $(P', (1-\delta))$ dense sources for $P' = \frac{S_z + \log(1/\gamma)}{\delta \cdot \log M}$.*

For every $(P, 1-\delta)$ source X' in the convex combination stated in the above lemma, let Y' be the corresponding bit fixing source. Then, we have the following lemma:

Lemma 9 ([CDGS18])**.** *For any $(P', 1-\delta)$ dense source X' and its corresponding bit fixing source Y', it holds that for any distinguisher D that makes at most T classical queries to its oracle,*

$$\left| \Pr[D^{X'} = 1] - \Pr[D^{Y'} = 1] \right| \leq T\delta \cdot \log M.$$

With these lemmas in hand, in order to simulate queries to the fixed random function f, we will show that:

1. There exists a distribution $\mathcal{D}'_f$ over functions with sufficient minimum entropy such that, with high probability over f, for any input x, and any witness w,

$$\left| \Pr[\mathsf{Ver}^f(w, x) = 1] - \Pr_{f' \leftarrow \mathcal{D}'_f}[\mathsf{Ver}^{f'}(w, x) = 1] \right| \leq \mathsf{negl}(n).$$

2. The distribution $\mathcal{D}'_f$ is "close" to a convex combination of dense sources.

3. Each dense source in the combination is close to a bit fixing source.

Collectively, this implies that, the QCMA verifier Ver making queries to a fixed but randomly sampled function f can be replaced with a new QCMA verifier Ver' which is the same as Ver, but with every query to f replaced with a query to a bit fixing source which is fixed on randomly chosen points.

We now have the tools to prove Theorem 12.

Proof. Set $\alpha = \frac{1}{2^n}$, $\beta = \frac{2^{-2^{n/8}}}{2^n}$. Let $\mathcal{D}'_f$ be the distribution from Lemma 7. We have that with probability $\geq 1 - 2^{-2^{n/8}}$

$$\left| \Pr[\mathcal{A}^f \to 1] - \Pr_{f' \leftarrow \mathcal{D}'_f}[\mathcal{A}^{f'} \to 1] \right| \leq \frac{1}{2^n}, \tag{1}$$

and $H_\infty(\mathcal{D}'_f) \geq 2^n \cdot n' - \log \frac{1}{\beta}$. But note that $\mathcal{D}'_f$ fits the conditions for Lemma 8. In particular, we get that $\mathcal{D}'_f$ is 2^{-n} close to some distribution $\mathcal{D}''_f$, where $\mathcal{D}''_f$ is a convex combination of $\left(P', 1-2^{-n/8}\right)$ dense sources with $P' = \frac{2^{n/8}(2^{n/8}+2n)}{n'}$. We can define $\mathcal{D}'''_f$ to be the convex combination of the corresponding bit-fixing sources from $\mathcal{D}''_f$. Then

$$\left| \Pr_{f'' \leftarrow \mathcal{D}''_f}[\mathcal{A}^{f''} \to 1] - \Pr_{f''' \leftarrow \mathcal{D}'''_f}[\mathcal{A}^{f'''} \to 1] \right| \leq T \cdot 2^{-n/8} \cdot n' \leq \mathsf{negl}(n) \tag{2}$$

And so setting $\mathcal{D}_f = \mathcal{D}'''_f$, triangle inequality gives us

$$\left| \Pr[\mathcal{A}^f \to 1] - \Pr_{f' \leftarrow \mathcal{D}_f}[\mathcal{A}^{f'} \to 1] \right| \leq \mathsf{negl}(n). \tag{3}$$

□

5 An Oracle World Where QCCC Primitives Exists But BQP = QCMA

We formally define the oracle $SG, Mix, PSPACE$. Our oracle will be parameterized by a sequence of unitaries $\{U_n\}_{n\in\mathbb{N}}$ and a family of functions $\{f_n\}_{n\in\mathbb{N}}$. Each U_n will be sampled from the Haar distribution over $2n$ qubits, and each $f_n : \{0,1\}^{2n} \to \{0,1\}^n$ will be sampled randomly from the space of functions from $\{0,1\}^{2n}$ to $\{0,1\}^n$.

We define the oracle $(SG, Mix, PSPACE)$ as follows:

1. SG_n: Sample $k \leftarrow \{0,1\}^n$. Output $k, |\phi_k\rangle$. Here $|\phi_k\rangle := U_n |k, 0^n\rangle$ for $k \in \{0,1\}^n$.
2. $Mix_n(\rho_{ABC})$: On input ρ_{ABC}, measure ρ_A and ρ_B in the standard basis, producing a measurement result (a, b). Measure $U_n^\dagger \rho_C U_n \to z$ (that is, apply $U_n^\dagger$ to the register C and measure in the standard basis, getting a result z). If $z = (a, 0^n)$ or $z = (b, 0^n)$, output $f_n(a, b)$. Otherwise, output $\perp$. For ease of notation, we will sometimes write $Mix(a, b, \psi)$ to denote that the first two registers will be measured in the computational basis.
3. $PSPACE$: An oracle for some **PSPACE**-complete problem.

5.1 Key Lemma

For this lemma, we will formalize the intuition that given only forward query access to a unitary U, you can't generate the state $(x, U|x\rangle)$ without having queried the unitary on $|x\rangle$.

We define the following set of oracles

1. Sample a Haar random unitary U over $2n$ qubits.
2. *Apply*(ρ): measure ρ in the standard basis $\to x \in \{0,1\}^n$, output $U|x, 0^n\rangle$.
3. $Test(\rho_{AB})$: measure ρ_A in the standard basis $\to x \in \{0,1\}^n$. Measure $U^\dagger \rho_B U$ (i.e. apply $U^\dagger$ to register B) in the standard basis $\to x'$. Output 1 if and only if $x' = (x, 0^n)$.
4. $Test'(\rho_{AB})$: measure ρ_A in the standard basis $\to x \in \{0,1\}^n$. Measure $U^\dagger \rho_B U$ (i.e. apply $U^\dagger$ to register B) in the standard basis $\to x'$. Output 1 if and only if $x' = (x, 0^n)$ and x was previously measured in some query to *Apply*.

Formally, our lemma will say that $Test$ and $Test'$ are indistinguishable.

Lemma 10. *Let $\mathcal{A}^{Apply,Test}$ be any inefficient adversary making at most $T = \mathsf{poly}(n)$ queries to its oracles. Then*

$$\left|\Pr[1 \leftarrow \mathcal{A}^{Apply,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply,Test'}]\right| \leq \frac{T}{2^n}.$$

Proof. We will prove this via a hybrid argument. In particular, we will define $\mathbf{G}_i(1^n)$ as follows

1. Sample a Haar random unitary U over $2n$ qubits.
2. Simulate $\mathcal{A}^{\mathcal{O}_1,\mathcal{O}_2}$.
3. When $\mathcal{A}$ makes a query to its first oracle $\mathcal{O}_1$, forward the query to *Apply*.
4. For the first i queries $\mathcal{A}$ makes to $\mathcal{O}_2$, forward the query to $Test'$.
5. For every query after i which $\mathcal{A}$ makes to $\mathcal{O}_2$, forward the query to $Test$.

We then claim that for each i,

$$|\Pr[1 \leftarrow \mathbf{G}_i] - \Pr[1 \leftarrow \mathbf{G}_{i-1}]| \leq \frac{1}{2^n}.$$

In particular, let us define $|\phi_k\rangle = U|k, 0^n\rangle$. We define $\pi_k = |\phi_k\rangle\langle\phi_k|$. We can equivalently redefine $Test$ as

1. Measure ρ_A in the standard basis $\to x \in \{0,1\}^n$.
2. Apply the measurement $\Pi = \{\{\pi_k\}_{k\in\{0,1\}^n}, \mathbb{I} - \sum_{k\in\{0,1\}^n} \pi_k\}$.
3. Output 1 if and only if the measurement result is π_x.

Similarly, we can redefine $Test'$ as

1. Measure ρ_A in the standard basis $\to x \in \{0,1\}^n$.
2. Apply the measurement $\widetilde{\Pi} = \{\{\pi_k\}_{k\in App}, \mathbb{I} - \sum_{k\in App} \pi_k\}$. Here *App* is the set of k that was measured in *Apply*.
3. Output 1 if and only if the measurement result is π_x with $x \in App$.

But note that to sample a unitary over $2n$ qubits is equivalent to sampling each $|\phi_k\rangle$ in sequence from the Haar distribution over the orthogonal space to all previously sampled states.

Let σ_i^j be the mixed state representing the adversary's state just before query j in game G_i. We can observe that $\sigma_i^i = \sigma_{i-1}^i$ by definition. Note that σ_i^i and σ_{i-1}^i only depend on the values of $|\phi_k\rangle$ for k queried to *Apply*. Thus, we can redefine $\mathbf{G}_{i-1}$ to sample each value $|\phi_k\rangle$ lazily as follows

1. Create a database of descriptions of quantum states D.
2. Simulate $\mathcal{A}^{\mathcal{O}_1,\mathcal{O}_2}$.
3. When $\mathcal{A}$ makes a query x to its first oracle $\mathcal{O}_1$, if $x \in D$, output $D[x]$. Otherwise, sample $|\phi_x\rangle$ from the Haar measure on the space orthogonal to D and set $D[x] = |\phi_x\rangle$.

4. For the first $i-1$ queries $\mathcal{A}$ makes to $\mathcal{O}_2$, forward the query to $Test'$.
5. As soon as the ith query to $\mathcal{O}_2$ is made, for all $x \in \{0,1\}^n$ such that $x \notin D$, sample orthogonal $|\phi_x\rangle$ from the Haar measure on the space orthogonal to D and set $D[x] = |\phi_x\rangle$.
6. For every query after $i-1$ which $\mathcal{A}$ makes to $\mathcal{O}_2$, forward the query to $Test$.

Observe that this is an equivalent formulation of $\mathbf{G}_{i-1}$, since $|\phi_k\rangle$ are all sampled according to the same distribution, and are never used before they are sampled.

In particular, during the ith query, $2^n - (i-1)$ orthogonal states are sampled from a space of dimension $2^{2n} - (i-1)$. Note that the only way to distinguish between $\mathbf{G}_{i-1}$ and $\mathbf{G}_i$ is if the ith query to $Test$ in $\mathbf{G}_{i-1}$ produces a measurement result π_x such that $|\phi_x\rangle$ was not previously returned by *Apply*. Let ρ be the mixed state corresponding to σ_{i-1}^i on the second input register of $Test$. We can bound the probability of distinguishing $\mathbf{G}_i$ and $\mathbf{G}_{i-1}$ by the probability that ρ measures to π_x for one of the $2^n - (i-1)$ newly sampled projectors π_x. We have that

$$|\Pr[1 \leftarrow \mathbf{G}_i] - \Pr[1 \leftarrow \mathbf{G}_{i-1}]| \leq \mathbb{E}\left[\mathrm{Tr}\left(\left(\sum_{k \notin D} \pi_k\right)\rho\right)\right] = \frac{2^n - (i-1)}{2^{2n} - (i-1)}. \tag{4}$$

And so we get that

$$\left|\Pr[1 \leftarrow \mathcal{A}^{Apply,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply,Test'}]\right| \leq \sum_{i=1}^{T} \frac{2^n - (i-1)}{2^{2n} - (i-1)} \leq \frac{T}{2^n}. \tag{5}$$

□

Corollary 2. *Let $\mathcal{A}^{\mathcal{O}_1,\mathcal{O}_2}$ be any inefficient adversary making at most $T = \mathsf{poly}(n)$ queries to its oracles. Let $S \subseteq \{0,1\}^n$ be any set such that*

$$\Pr[\mathcal{A}^{Apply,Test} \textit{ makes an Apply query on some } s \in S] \leq \epsilon.$$

We define $Test^S(\rho_{AB})$: measure ρ_A in the standard basis $\to x \in \{0,1\}^n$. Measure $U^\dagger \rho_B U$ in the standard basis $\to x'$. It outputs 1 if and only if $x' = (x, 0^n)$ and $x \notin S$.

Then

$$\left|1 \leftarrow \Pr[\mathcal{A}^{Apply,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply,Test^S}]\right| \leq \epsilon + \frac{T}{2^n}.$$

Proof. We can define $Apply^S$ to be the same as *Apply*, but it rejects whenever queried on any $s \in S$. By the definition of S,

$$\left|\Pr[1 \leftarrow \mathcal{A}^{Apply,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply^S,Test}]\right| \leq \epsilon. \tag{6}$$

But by Lemma 10,

$$\left|\Pr[1 \leftarrow \mathcal{A}^{Apply^S,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply^S,Test'}]\right| \leq \frac{T}{2^n}. \tag{7}$$

But the only way to distinguish $(Apply^S, Test)$ from $(Apply^S, Test^S)$ is to query $Test^S$ on some state measuring to $s \in S$. But Equation (7) shows that this can occur with at most negligible probability. And so

$$\left|\Pr[1 \leftarrow \mathcal{A}^{Apply^S,Test}] - \Pr[1 \leftarrow \mathcal{A}^{Apply^S,Test^S}]\right| \leq \frac{T}{2^n}. \tag{8}$$

Combining Equations (6) and (8) using triangle inequality finishes the proof. □

5.2 QCCC KE Exists Relative To $(SG, Mix, PSPACE)$

Theorem 13. *With probability* 1 *over* $\{U_n\}_{n\in\mathbb{N}}$, $\{f_n\}_{n\in\mathbb{N}}$, *relative to the oracles* $(SG, Mix, PSPACE)$, *QCCC KE exist.*

Our KE protocol operates as follows.

1. Alice runs $SG_n \to x, |\phi_x\rangle$. Alice sends x to Bob.
2. Bob runs $SG_n \to y, |\phi_y\rangle$. Bob sends y to Alice.
3. Alice outputs $Mix_n(x, y, |\phi_x\rangle) \to f_n(x, y)$.
4. Bob outputs $Mix_n(x, y, |\phi_y\rangle) \to f_n(x, y)$.

Agreement follows by construction. It only remains to show security. Without loss of generality we assume the adversary is a uniform quantum algorithm $\mathcal{A}^{SG,Mix,PSPACE}(1^n, x, y, c)$ where x, y is the transcript and $c \in \{0,1\}^{\mathsf{poly}(n)}$ is the advice string.

We will use several hybrids to switch our KE game to one where we can more easily argue that the adversary has negligible advantage. In particular, we can describe our original KE security game $\mathbf{G}_1(1^n)$ as follows.

1. Sample $x, y \leftarrow \{0,1\}^n$.
2. Run $guess \leftarrow \mathcal{A}^{SG,Mix,PSPACE}(x, y, c)$.
3. Output 1 if and only if $guess = f_n(x, y)$.

We also define our second hybrid $\mathbf{G}_2(1^n)$, where we replace the fixed unitary with a fresh random unitary:

1. Sample $x, y \leftarrow \{0,1\}^n$.
2. Sample a Haar random unitary $\widetilde{U}_n$ over $2n$ qubits.
3. Define $\widetilde{SG}_n, \widetilde{Mix}_n$ to be the same as SG_n, Mix_n but with all calls to U_n replaced by calls to $\widetilde{U}_n$.
4. For $n' \neq n$, define $\widetilde{SG}_{n'} = SG_{n'}$ and $\widetilde{Mix}_{n'} = Mix_{n'}$.
5. Run $guess \leftarrow \mathcal{A}^{\widetilde{SG},\widetilde{Mix},PSPACE}(x, y, c)$.
6. Output 1 if and only if $guess = f_n(x, y)$.

Claim 1. *With probability* $\geq 1 - \exp(-2^n)$ *over* $\{U_n\}_{n\in\mathbb{N}}$,

$$|\Pr[1 \leftarrow \mathbf{G}_2(1^n)] - \Pr[1 \leftarrow \mathbf{G}_1(1^n)]| \leq \mathsf{negl}(n).$$

Proof. Lemma 4 says that $\Pr[1 \leftarrow \mathbf{G_1}(1^n)]$ is $2T$-Lipschitz in the Frobenius norm as a function of U_n. But it is clear that

$$\Pr[1 \leftarrow \mathbf{G}_2(1^n)] = \mathop{\mathbb{E}}_{U_n}[\Pr[1 \leftarrow \Pr[\mathbf{G}_1(1^n)]]].$$

Let $\delta = 2^{-n/8}$. Invoking the strong concentration of the Haar measure from Theorem 10 with $t = \delta$ and $L = 2T$ gives us

$$\Pr_{U_n, \widetilde{U}_n}[|\Pr[1 \leftarrow \mathbf{G}_2(1^n)] - \Pr[1 \leftarrow \mathbf{G}_1(1^n)]| \geq \delta] \leq 2\exp\left(-\frac{(2^{2n}-2)\delta^2}{96T^2}\right) \leq \exp(-2^n)$$

for all sufficiently large n. $\square$

We next replace the fixed random function with a bit-fixing oracle, using our concentration theorem for random oracles. In particular, let $\mathcal{D}_{f_n}$ be the bit-fixing oracle from Theorem 12. In particular, we know that such a $\mathcal{D}_{f_n}$ exists with probability $\geq 1 - 2^{-2^{n/4}}$. Note furthermore that $\mathcal{D}_{f_n}$ fixes at most $\frac{2^{n/4}(2^{n/4}+4n)}{n} \leq 2^{3n/4}$ points. We then define our next hybrid $\mathbf{G}_3(1^n)$ by replacing f_n with $\mathcal{D}_{f_n}$:

1. Sample $x, y \leftarrow \{0,1\}^n$.
2. Sample a Haar random unitary $\widetilde{U}_n$ over $2n$ qubits.
3. Sample $f'_n \xleftarrow{\$} \mathcal{D}_{f_n}$.
4. Define $\widetilde{SG}'_n, \widetilde{Mix}'_n$ to be the same as SG_n, Mix_n but with all calls to U_n replaced by calls to $\widetilde{U}_n$ and all calls to f_n replaced by calls to f'_n.
5. For $n' \neq n$, define $\widetilde{SG}'_{n'} = SG_{n'}$ and $\widetilde{Mix}'_{n'} = Mix_{n'}$.
6. Run $guess \leftarrow \mathcal{A}^{\widetilde{SG}', \widetilde{Mix}', PSPACE}(x, y, c)$.
7. Output 1 if and only if $guess = f'_n(x, y)$.

Claim 2. *With probability $\geq 1 - 2^{-2^{n/4}}$ over the choice of $\{f_n\}_{n\in\mathbb{N}}$,*

$$|\Pr[1 \leftarrow \mathbf{G}_3(1^n)] - \Pr[1 \leftarrow \mathbf{G}_2(1^n)]| \leq \mathsf{negl}(n).$$

Proof. This follows immediately from Theorem 12 and observing that f_n is a function from $\{0,1\}^{2n} \to \{0,1\}^n$. □

In the next hybrid, we replace the bit-fixing oracle with a truly random oracle. That is, we define $\mathbf{G}_4(1^n)$ as follows:

1. Sample $x, y \leftarrow \{0,1\}^n$.
2. Sample a Haar random unitary $\widetilde{U}_n$ over $2n$ qubits.
3. Sample a random function $\widetilde{f}_n : \{0,1\}^{2n} \to \{0,1\}^n$.
4. Define $\widetilde{\widetilde{SG}}_n, \widetilde{\widetilde{Mix}}_n$ to be the same as SG_n, Mix_n but with all calls to U_n replaced by calls to $\widetilde{U}_n$ and all calls to f_n replaced by calls to $\widetilde{f}_n$.
5. For $n' \neq n$, define $\widetilde{\widetilde{SG}}_{n'} = SG_{n'}$ and $\widetilde{\widetilde{Mix}}_{n'} = Mix_{n'}$.
6. Run $guess \leftarrow \mathcal{A}^{\widetilde{\widetilde{SG}}, \widetilde{\widetilde{Mix}}, PSPACE}(x, y, c)$.
7. Output 1 if and only if $guess = \widetilde{f}_n(x, y)$.

Claim 3.

$$|\Pr[1 \leftarrow \mathbf{G}_4(1^n)] - \Pr[1 \leftarrow \mathbf{G}_3(1^n)]| \leq \mathsf{negl}(n).$$

Proof. Let us consider the probability that any $k, |\phi_k\rangle$ returned from SG_n satisfies the property that there is some point $f'_n(k, \cdot)$ or $f'_n(\cdot, k)$ fixed by the bit-fixing oracle in $\mathbf{G}_3(1^n)$. Since there are $\leq 2^{3n/4}$ fixed points, the probability this ever occurs is, by union bound, $\leq T\frac{2 \cdot 2^{3n/4}}{2^n} \leq \mathsf{negl}(n)$. Thus, Corollary 2 gives that with all but negligible probability, in $\mathbf{G}_3(1^n)$ the adversary never queries f'_n on any point fixed by the bit fixing oracle. And so the claim follows. □

In our final hybrid, we guarantee that the oracle never queries the random oracle on (x, y). In particular, Mix will simply never accept states $|\phi_x\rangle$ or $|\phi_y\rangle$. In particular, we define $\mathbf{G}_5(1^n)$ as follows.

1. Sample $x, y \leftarrow \{0,1\}^n$.
2. Sample a Haar random unitary $\widetilde{U}_n$ over $2n$ qubits.
3. Sample a random function $\widetilde{f}_n : \{0,1\}^{2n} \rightarrow \{0,1\}^n$.
4. Define Mix^*_n as follows:
   (a) On input ρ_{ABC}, measure ρ_{AB} in the standard basis, getting (a, b).
   (b) Apply $\widetilde{U}_n$ on ρ_C and measure in the standard basis, getting a result z.
   (c) If $z = (x, 0^n)$ or $z = (y, 0^n)$, output $\perp$.
   (d) Otherwise if $z = (a, 0^n)$ or $z = (b, 0^n)$, output $f(a, b)$.
   (e) Otherwise, output $\perp$.
5. For $n' \neq n$, define $Mix^*_{n'} = Mix_{n'}$.
6. Run $guess \leftarrow \mathcal{A}^{\widetilde{\widetilde{\widetilde{SG}}},Mix^*,PSPACE}(x, y, c)$.
7. Output 1 if and only if $guess = \widetilde{f}_n(x, y)$.

Claim 4.

$$|\Pr[1 \leftarrow \mathbf{G}_5(1^n)] - \Pr[1 \leftarrow \mathbf{G}_4(1^n)]| \leq \mathsf{negl}(n).$$

Proof. The probability that SG_n ever returns $x, |\phi_x\rangle$ or $y, |\phi_y\rangle$ is $\leq \frac{2T}{2^n} \leq \mathsf{negl}(n)$. The claim immediately follows from Corollary 2. □

Note that $\mathbf{G}_5(1^n)$ never queries $\widetilde{f}_n(x, y)$ until the very last test. The probability that $\mathcal{A}$ guesses $\widetilde{f}_n(x, y)$ without querying it is $\frac{1}{2^n}$. And so,

$$\Pr[1 \leftarrow \mathbf{G}_5(1^n)] \leq \frac{1}{2^n} = \mathsf{negl}(n). \tag{9}$$

Applying the triangle inequality and union bound gives us that

$$\Pr_{U_n, f_n}[\Pr[1 \leftarrow \Pr[\mathbf{G}_1(1^n)]] \leq \mathsf{negl}(n)] \geq 1 - \exp(-2^n) - 2^{-2^{n/4}}.$$

And so taking the union bound over all c gives us that for any $p = \mathsf{poly}(n)$, $\mathcal{A}^{SG,Mix,PSPACE}$ achieves advantage larger than $1/p$ with probability at most $2^{\mathsf{poly}(n)}\left(\exp(-2^n) + 2^{-2^{n/4}}\right) = \mathsf{negl}(n)$. But note that $\sum_{n=1}^{\infty} \mathsf{negl}(n)$ converges, and so by the Borel-Cantelli lemma [ÉB09, Can17], $\mathcal{A}$ achieves negligible advantage for all but finitely many input lengths with probability 1 over $\{U_n\}_{n\in\mathbb{N}}$ and $\{f_n\}_{n\in\mathbb{N}}$. Taking a union bound over adversaries completes the proof.

5.3 QCCC Commitments Exist Relative To $(SG, Mix, PSPACE)$

Since this section details a construction, we will show that our construction satisfies stronger security requirements than given in Definition 5. Note that if a bit commitment scheme satisfies statistical hiding and sum-binding, then it trivially satisfies computational hiding and weak computational binding.

Definition 10 (Statistical hiding from [HR07])**.** *Let (Com, Rec) be a bit-commitment scheme. We say that (Com, Rec) satisfies statistical hiding if for all (inefficient) adversarial receivers Rec', the transcript of the interaction between Rec' and the committer with input m is statistically close to the transcript of the interaction between Rec' and the committer with input m'. That is,*

$$\Delta(Com(m) \rightleftarrows Rec', Com(m') \rightleftarrows Rec') \leq \mathsf{negl}(n).$$

Definition 11 (Sum-binding from [Unr16])**.** *Let (Com, Rec) be a bit-commitment scheme. We say that (Com, Rec) satisfies sum-binding if for all QPT adversarial senders Com', the probability that Com' wins the following game is $\leq \frac{1}{2} + \mathsf{negl}(n)$.*

1. *In the first stage, an honest receiver Rec interacts with the adversarial committer Com' to produce a transcript z and receiver state ρ_{Rec}. Com' produces the adversarial committer state $\rho_{Com'}$.*
2. *The challenger samples a bit $b \leftarrow \{0, 1\}$.*
3. *In the second stage, the honest receiver Rec is given ρ_{Rec} and z, while Com' is given $z, \rho_{Com'}$, and b. They then proceed to run the opening stage with the committer replaced by Com', and Rec produces a final output b'. Com' wins if $b' = b$.*

Theorem 14. *With probability 1 over $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$ there exists a bit commitment scheme satisfying statistical hiding and sum-binding relative to $SG, Mix, PSPACE$.*

Proof. We construct a commitment scheme $\mathsf{Com}^{SG,Mix,PSPACE}$,$\mathsf{Rec}^{SG,Mix,PSPACE}$ as follows. Note that our commitment scheme will have a non-interactive commitment phase and an interactive opening phase.

1. To commit to b, Com runs $SG_n \rightarrow x, |\phi_x\rangle$. Let x_1 be the first bit of x and let $x_{>1}$ be the last $n-1$ bits of x. Com repeats this process until $x_1 = b$. Com sends $x_{>1}$ to Rec.
2. To open to a bit b, Rec runs $SG_n \rightarrow y, |\phi_y\rangle$. Rec sends y to Com.
3. Upon receiving y, Com computes $Mix_n(x, y, |\phi_x\rangle) \rightarrow c$ and sends c to Rec.
4. Rec accepts if $Mix_n(y, (b, x_{>1}), |\phi_y\rangle) = c$.

Correctness is clear by construction. Note that the commitment $x_{>1}$ is a uniformly random string, and so $(\mathsf{Com}, \mathsf{Rec})$ trivially satisfies statistical hiding.

Thus, it remains to show computational binding. We will define a sequence of hybrids $\mathbf{G}_1, \ldots, \mathbf{G}_6$ such that $\mathbf{G}_1$ is the original sum-binding game and for each i,

$$|\Pr[1 \leftarrow \mathbf{G}_i] - \Pr[1 \leftarrow \mathbf{G}_{i-1}]| \leq \mathsf{negl}(n).$$

Let $\mathcal{A}$ be an adversary against the sum-binding game. We describe the original sum-binding security game $\mathbf{G}_1(1^n)$ as follows

1. Sample $b \leftarrow \{0,1\}$.
2. Sample $y \leftarrow \{0,1\}^n$.
3. Run $x \leftarrow \mathcal{A}^{SG,Mix,PSPACE}$.
4. Run $z \leftarrow \mathcal{A}^{SG,Mix,PSPACE}(b, x, y, \rho_{\mathcal{A}})$.
5. Output 1 if and only if $z = f_n(x, y)$.

In $\mathbf{G}_2(1^n)$, we replace all calls to U_n with calls to a fresh Haar-random $\widetilde{U}_n$.

Claim 5. *With probability* $\geq 1 - \exp(-2^n)$ *over* $\{U_n\}_{n\in\mathbb{N}}$*,*

$$|\Pr[1 \leftarrow \mathbf{G}_2(1^n)] - \Pr[1 \leftarrow \mathbf{G}_1(1^n)]| \leq \mathsf{negl}(n).$$

Proof. Lemma 4 says that $\Pr[1 \leftarrow \mathbf{G_1}(1^n)]$ is $2T$-Lipschitz in the Frobenius norm as a function of U_n. But it is clear that

$$\Pr[1 \leftarrow \mathbf{G}_2(1^n)] = \mathop{\mathbb{E}}_{U_n}[\Pr[1 \leftarrow \mathbf{G}_1(1^n)]].$$

Let $\delta = 2^{-n/8}$. Invoking the strong concentration of the Haar measure from Theorem 10 with $t = \delta$ and $L = 2T$ gives us

$$\Pr_{U_n}\left[\left|\Pr[1 \leftarrow \mathbf{G}_1(1^n)] - \Pr_{\widetilde{U}_n}[1 \leftarrow \mathbf{G}_2(1^n)]\right| \geq \delta\right] \leq 2\exp\left(-\frac{(2^{2n}-2)\delta^2}{96T^2}\right) \leq \exp(-2^n) \tag{10}$$

for all sufficiently large n. □

.

Define $\mathcal{D}_{f_n}$ to be the distribution from Theorem 12. In particular, we know that such a $\mathcal{D}_{f_n}$ exists with probability $\geq 1 - 2^{-2^{n/4}}$. Note furthermore that $\mathcal{D}_{f_n}$ fixes at most $\frac{2^{n/4}(2^{n/4}+4n)}{n} \leq 2^{3n/4}$ points. We define $\mathbf{G}_3(1^n)$ to be the same as $\mathbf{G}_2(1^n)$, but with all calls to f_n replaced with a freshly sampled $f'_n \leftarrow \mathcal{D}_{f_n}$.

Claim 6. *With probability* $\geq 1 - 2^{-2^n}$*,*

$$|\Pr[1 \leftarrow \mathbf{G}_3(1^n)] - \Pr[1 \leftarrow \mathbf{G}_2(1^n)]| \leq \mathsf{negl}(n).$$

Proof. This follows immediately from Theorem 12 and observing that f_n is a function from $\{0,1\}^{2n} \to \{0,1\}^n$. □

$\mathbf{G}_4(1^n)$ will be the same as $\mathbf{G}_3(1^n)$, except we replace all calls to f'_n with a uniformly random $\widetilde{f}_n$

Claim 7. $|\Pr[1 \leftarrow \mathbf{G}_4(1^n)] - \Pr[1 \leftarrow \mathbf{G}_3(1^n)]| \leq \mathsf{negl}(n)$.

Proof. Let us consider the probability that any $k, |\phi_k\rangle$ returned from SG_n satisfies the property that there is some point $f'_n(k, \cdot)$ or $f'_n(\cdot, k)$ fixed by the bit-fixing oracle in $\mathbf{G}_3$. Since there are $\leq 2^{3n/4}$ fixed points, the probability this ever occurs is, by union bound,

$$\leq T\frac{2 \cdot 2^{3n/4}}{2^n} \leq \mathsf{negl}(n).$$

Thus, Corollary 2 gives that with all but negligible probability, in $\mathbf{G}_3(1^n)$ the adversary never queries f'_n on any point fixed by the bit fixing oracle. And so the claim follows. □

$\mathbf{G}_5(1^n)$ will be the same as $\mathbf{G}_4(1^n)$, except that we guarantee SG_n never outputs the same suffix twice. We also guarantee that SG_n never outputs the committer message after it has been sent and that SG_n never outputs the receiver's challenge. That is, we replace SG_n with SG'_n defined as follows

1. Sample $k \leftarrow \{0,1\}^n$.
2. If there is a k' previously sampled such that $k_{>1} = k'_{>1}$, output $\perp$.
3. If the committer message x has already been sent and $k_{>1} = x$, output $\perp$.
4. If $k = y$, output $\perp$.
5. Output $k, \widetilde{U}_n|k, 0^n\rangle$.

Claim 8. $|\Pr[1 \leftarrow \mathbf{G}_5(1^n)] - \Pr[1 \leftarrow \mathbf{G}_4(1^n)]| \leq \mathsf{negl}(n)$.

Proof. In particular, the probability that we sample k, k' such that $k_{>1} = k'_{>1}$ is $\leq \frac{T^2}{2^{n-1}}$ by the birthday bound. Furthermore, the probability that we ever sample k such that $k_{>1} = x$ is $\leq \frac{T}{2^{n-1}}$ by the union bound. Finally, the probability that we ever sample k such that $k = y$ is $\leq \frac{T}{2^n}$ Applying a union bound finishes the argument. $\square$

$\mathbf{G}_6(1^n)$ will be the same as $\mathbf{G}_5(1^n)$, except it guarantees that Mix_n always rejects outputs produced by SG'_n. That is, we replace Mix_n with Mix'_n defined as follows:

1. On input ρ_{ABC}, measure registers A and $B \to (x, y)$.
2. Apply $\widetilde{U}_n^\dagger \rho_C \widetilde{U}_n$ and measure in the standard basis, producing a result z.
3. If $z = (x, 0^n)$ and x was previously produced by SG'_n, output $\widetilde{f}_n(x, y)$.
4. If $z = (y, 0^n)$ and y was previously produced by SG'_n, output $\widetilde{f}_n(x, y)$.
5. Otherwise, output $\perp$.

Claim 9. $|\Pr[1 \leftarrow \mathbf{G}_6(1^n)] - \Pr[1 \leftarrow \mathbf{G}_5(1^n)]| \leq \mathsf{negl}(n)$.

Proof. This follows immediately from Lemma 10. $\square$

Claim 10. $\Pr[1 \leftarrow \mathbf{G}_6(1^n)] \leq \frac{1}{2} + \mathsf{negl}(n)$.

Proof. Let $|\phi_k\rangle = \widetilde{U}_n |k, 0^n\rangle$. The only way for $\mathcal{A}^{SG', Mix', PSPACE}$ to query $\widetilde{f}_n((b, x), y)$ is if $\mathcal{A}$ ever received $|\phi_{b,x}\rangle$ or $|\phi_y\rangle$ from SG'_n. But note that by definition of SG'_n, $\mathcal{A}$ can only ever receive either $|\phi_{0,x}\rangle$ or $|\phi_{1,x}\rangle$ *before* b is sampled, and can receive neither afterwords. Thus, since b is chosen at random, the probability that $\mathcal{A}^{SG', Mix', PSPACE}$ ever queries $\widetilde{f}_n((b, x), y)$ is at most $\frac{1}{2}$.

But if $\mathcal{A}$ doesn't query $\widetilde{f}_n((b, x), y)$, it can predict its value with at most probability $\frac{1}{2^n}$. And so

$$\Pr[1 \leftarrow \mathbf{G}_6(1^n)] \leq \frac{1}{2} + \mathsf{negl}(n).$$

$\square$

Applying the triangle inequality and union bound gives us that

$$\Pr_{U_n,f_n}[\Pr[1 \leftarrow \mathbf{G}_1(1^n)] \leq \mathsf{negl}(n)] \geq 1 - \exp(-2^n) - 2^{-2^{n/4}}.$$

And so taking the union bound over all c gives us that for any $p = \mathsf{poly}(n)$, $\mathcal{A}^{SG,Mix,PSPACE}$ achieves advantage larger than $1/p$ with probability at most $2^{\mathsf{poly}(n)}\left(\exp(-2^n) + 2^{-2^{n/4}}\right) = \mathsf{negl}(n)$. But note that $\sum_{n=1}^{\infty} \mathsf{negl}(n)$ converges, and so by the Borel-Cantelli lemma [ÉB09, Can17], $\mathcal{A}$ achieves negligible advantage for all but finitely many input lengths with probability 1 over $\{U_n\}_{n\in\mathbb{N}}$ and $\{f_n\}_{n\in\mathbb{N}}$. Taking a union bound over adversaries completes the proof. □

5.4 $\mathbf{BQP}^{SG,Mix,PSPACE} = \mathbf{QCMA}^{SG,Mix,PSPACE}$

Theorem 15. *With probability at least* $1 - \mathsf{negl}$ *over* $\{U_n\}_{n\in\mathbb{N}}$ *and* $\{f_n\}_{n\in\mathbb{N}}$*,*

$$\mathbf{BQP}^{SG,Mix,PSPACE} = \mathbf{QCMA}^{SG,Mix,PSPACE}.$$

Proof. Let $\Pi \in \mathbf{QCMA}^{SG,Mix,PSPACE}$, which implies that there exists a **QCMA** verifier $\mathsf{Ver}^{SG,Mix,PSPACE}$ which takes as input $x \in \mathsf{Dom}(\Pi)$ and a classical witness $w \in \{0,1\}^{q(|x|)}$ where $q(|x|) = \mathsf{poly}(|x|)$, with completeness $\frac{99}{100}$ and soundness $\frac{1}{100}$. Without loss of generality, we will assume $q(|x|) \geq |x|$. Let $|x| = n$, and let $T(n)$ be a polynomial upper bound on the running time of Ver on inputs of length n.

We will now describe a QPT algorithm $\mathcal{B}^{SG,Mix,PSPACE}(x)$ such that, with probability $1 - \mathsf{negl}$ over $\{U_n\}_{n\in\mathbb{N}}$ and $\{f_n\}_{n\in\mathbb{N}}$, $\mathcal{B}$ computes Π on all but finitely many inputs $x \in \mathrm{Dom}(\Pi)$. Let $d = \log(3456q(|x|)^2T(|x|)^4 + 2)$ and let $\delta = \frac{1}{200}$.

- For all $\ell \in [d]$, $\mathcal{B}$ performs tomography on each U_ℓ, producing estimates U'_ℓ. Every query made to U_ℓ is simulated by a query to U'_ℓ. The estimates are produced as follows:
  - $\mathcal{B}$ runs SG until, $\forall k \in [2^\ell]$, there exist $t > 2^d \cdot \frac{T}{\delta}$ copies of $|\phi_k\rangle$.
  - For every $k \in [2^\ell]$, run the state tomography algorithm from Theorem 11 on inputs $n = q(n)^2 \cdot T(n), \epsilon = \frac{\delta}{T(n)}$. This produces an estimate $|\phi'_k\rangle$ such that $|\langle\phi_k|\phi'_k\rangle|^2 \geq 1 - \epsilon$ with probability at least $1 - e^{-5q(n)^2T(n)}$.
  - $\mathcal{B}$ defines the map $U'_\ell |k\rangle \rightarrow |\phi'_k\rangle$.
- Further, for all $\ell \in [d]$, $\mathcal{B}$ learns all functions f_ℓ by querying SG, and produces estimates f'_ℓ. Every query made to f_ℓ is simulated by a query to f'_ℓ.
  - $\mathcal{B}$ runs SG until $\forall\ k \in [2^\ell]$, there exists 2^d copies of $|\phi_k\rangle$. With probability $\geq e^{-q(n)^2}$, this terminates in time polynomial in n for all ℓ.
  - For all $x, y \in \{0,1\}^{2\ell}$, $\mathcal{B}$ queries $Mix(x, y, |\phi_x\rangle) \rightarrow f_\ell(x, y)$, and stores the result as $f'_\ell(x, y)$.

Consider a **QCMA** verifier $\widetilde{\mathsf{Ver}}$ which simulates Ver by replacing queries to $\{U_n\}_{n\in\mathbb{N}}$ and $\{f_n\}_{n\in\mathbb{N}}$ as follows:

- For all $\ell \in [d]$, every query to U_ℓ is replaced by a query to U'_ℓ.
- Further, for all $\ell \in [d]$, every query made to f_ℓ is replaced by the output of f'_ℓ on the query.

- For all $\ell \in [d+1, T(n)]$ let $\epsilon = \frac{\delta}{2^{T(n)}2^{T(n)2^{(2n+2)T(n)}}}$, where $\delta = \frac{1}{200}$. Let $m : \mathbb{N} \to \mathbb{N}$ be as in Lemma 5. Let A be the "unitary design sampler" algorithm from Lemma 5 with parameters $\epsilon, n = \ell$, and $t = T$. $\widetilde{\mathsf{Ver}}$ samples $g_\ell : \{0,1\}^\ell \to \{0,1\}^{m(l)}$ from a $2T$-wise independent family of functions. Set $\tilde{U}_{k\ell} = A(g_\ell(k))$. Further, for all $\ell \in [d]$, $\widetilde{\mathsf{Ver}}$ replaces queries to all f_ℓ queries with queries a $2T$ wise independent function.

Now, consider the **QCMA** problem $\widetilde{\Pi}$ corresponding to $\widetilde{\mathsf{Ver}}$ with completeness $\frac{2}{3}$ and soundness $\frac{1}{3}$. Since $\mathbf{QCMA}^A \subseteq \mathbf{PSPACE}^A$ for all classical oracles A, and $\widetilde{\Pi} \in \mathbf{PSPACE}$, $\mathcal{B}$ will then use the **PSPACE** oracle to decide $\widetilde{\Pi}$.

Fix any (x, w). We will show via a sequence of hybrids that with high probability over $\{U_n\}_{n\in\mathbb{N}}$ and $\{f_n\}_{n\in\mathbb{N}}$,

$$\left|\Pr[\widetilde{\mathsf{Ver}}^{PSPACE}(x,w) \to 1] - \Pr[\mathsf{Ver}^{SG,Mix,PSPACE}(x,w) \to 1]\right| \leq \frac{1}{100}.$$

We will show this via a sequence of hybrids. For a verification algorithm V_i, Define $acc(V_i) = \Pr[V_i(w,x) \to 1]$, and let $d = \log(3456|x|T(|x|)^4 + 2)$.

1. Let $V_1 = \mathsf{Ver}^{SG,Mix,PSPACE}$.

2. Let V_2 be the same as V_1 except that, for all $\ell \in [d+1, T(n)]$, queries to U_ℓ are replaced by queries to a freshly sampled Haar random unitary $\tilde{U}_\ell$. Define SG_1 and Mix_1 to be the same as SG and Mix, except, for all $\ell \in [d+1, T(n)]$, queries to U_ℓ are replaced by queries to $\tilde{U}_\ell$.

   Claim 11. *With probability $\geq 1 - \exp\left(-q(n)^2\right)$ over $\{U_n\}_{n\in\mathbb{N}}$,*

   $$\left|acc(V_1) - acc(V_2^{SG_1,Mix_1})\right| \leq 1/12.$$

   Proof. Lemma 4 says that $acc(V_1)$ is $2T$-Lipschitz in the Frobenius norm as a function of $U_{d+1} \oplus \cdots \oplus U_{T(n)}$. But it is clear that $\Pr[V_2 \to 1] = \mathbb{E}_{U_{d+1},\ldots,U_{T(n)}}[\Pr[V_1 \to 1]]$. Invoking the strong concentration of the Haar measure from Theorem 10 with $t = 1/12$ and $L = 2T$ gives us

   $$\Pr_{U_{d+1},\ldots,U_{T(n)}}\left[\left|\Pr_{\tilde{U}_{d+1},\ldots,\tilde{U}_{T(n)}}[V_2 \to 1] - \Pr[V_1 \to 1]\right| \geq t\right] \leq 2\exp\left(-\frac{(2^{2d}-2)t^2}{96T^2}\right) \tag{11}$$

   $$\leq \exp\left(-q(n)^2\right). \tag{12}$$

   □

3. V_3 : This is the same as V_2 except, for all $\ell \in [d]$, all queries to f_ℓ are replaced by queries to f'_ℓ.

   Claim 12. *With probability $\geq 1 - e^{-q(n)^2}$ over the randomness of $\mathcal{B}$, $acc(V_3) = acc(V_2)$.*

   Proof. For all $\ell \in [d]$, $\mathcal{B}$ learns all functions f_ℓ with no error with probability at least $1 - e^{-q(n)^2}$, so $acc(V_3) = acc(V_2)$. □

4. V_4 : This is the same as V_3, except that, for all $\ell \in [d+1, T]$, $\mathcal{B}$ replaces the fixed random function f_ℓ with a bit-fixing oracle, using our concentration theorem for random functions. Let $\mathcal{D}_f$ be the bit-fixing oracle from Theorem 12. In particular, we know that such a $\mathcal{D}_f$ exists with probability $\geq 1 - 2^{-2^{\ell/4}}$. Note furthermore that $\mathcal{D}_f$ fixes at most $\frac{2^{\ell/4}(2^{\ell/4}+4\ell)}{\ell} \leq 2^{3\ell/4}$ points.

Claim 13. *With probability* $\geq 1 - e^{-q(n)^2}$ *over the choice of* $f = \{f'_\ell\}_\ell$*,*

$$|acc(V_4) - acc(V_3)| \leq \mathsf{negl}(n).$$

Proof. From Theorem 12 and observing that f_ℓ is a function from $\{0,1\}^{2\ell} \to \{0,1\}^\ell$, for each ℓ, the difference in acceptance probability is $\leq 2^{-2^{\ell/4}}$. Summing over all $\ell \in [d+1, p(|x|)]$, the difference in acceptance probabilities is at most

$$\sum_{\ell=d+1}^{T(n)} 2^{-2^{\ell/4}} \leq T(n) \cdot 2^{-2^{d/4}} \leq e^{-q(n)^2}.$$

□

5. V_5: This is the same as V_4 except, for all $\ell \in [d+1, T]$, we replace the bit-fixing oracle with a truly random oracle.

   Claim 14.
   $$|acc(V_5) - acc(V_4)| \leq \frac{1}{100}.$$

   Proof. Let us consider the probability that any $k, |\phi_k\rangle$ returned from SG satisfies the property that there is some point $f'(k, \cdot)$ or $f'(\cdot, k)$ fixed by the bit-fixing oracle in V_4. Since there are $\leq 2^{3\ell/4}$ fixed points in each function, the probability that any such point is returned is $\leq 2 \cdot 2^{-\ell/4}$. Corollary 2 gives that with probability $1 - 2 \cdot 2^{-\ell/4} - T \cdot 2^{-2\ell}$, in V_5 the adversary never queries f' on any point fixed by the bit fixing oracle. And so by union bound, the probability the adversary ever queries f' on any point fixed by the bit fixing oracle is $\leq T(n) \cdot (2 \cdot 2^{-d/4} - T \cdot 2^{-2d}) \leq \frac{1}{100}$. □

6. V_6: This is the same as V_5, except for all $\ell \in [d+1, T(n)]$, do the following. Let $\epsilon = \frac{\delta}{2^{T(n)} 2^{T(n)} 2^{(2n+2)T(n)}}$. Let $m : \mathbb{N} \to \mathbb{N}$ be as in Lemma 5. Let A be the "unitary design sampler" algorithm from Lemma 5 with parameters $\epsilon, n = \ell$, and $t = T$. Sample $\ell_x \leftarrow \{0,1\}^{m(\ell)}$. Set $U'_\ell = A(\ell_x)$.

   Claim 15. $|acc(V_6) - acc(V_5)| \leq \delta$*.*

   Proof. Using Lemma 6, the total change in acceptance probability is

   $$\sum_{\ell=d+1}^{T(n)} \sum_{k \in \{0,1\}^\ell} \frac{\delta}{T(n) 2^{T(n)}} \leq \delta\,.$$

   Thus $|acc(V_6) - acc(V_5)| \leq \delta$. □

7. V_7: This is the same as V_6, except, $\forall \ell \in [d]$, U_ℓ is replaced with U'_ℓ.

   Claim 16. *With probability* $\geq 1 - e^{-q(n)^2}$ *over the choice of* $\{U'_\ell\}_{\ell \in [d]}$*,*

   $$|acc(V_7) - acc(V_6)| \leq 4\delta.$$

Proof. From Theorem 11, we have that for each $\ell \in [d]$, for each $k \in \{0,1\}^\ell$, with probability $\geq 1 - e^{-5q(n)T(n)}$, $|\langle \phi_x | \phi'_x \rangle|^2 \geq 1 - \frac{\delta}{T(n)}$. Taking the union bound over all d, k, we get that with probability $\geq 1 - d \cdot 2^d \cdot e^{-5q(n)T(n)} \geq 1 - e^{-q(n)}$, for all $\ell \in [d], k \in \{0,1\}^d$, $|\langle \phi_x \rangle \phi'_x|^2 \geq 1 - \frac{\delta}{T(n)}$.

Consider any $\ell \leq d$ and $x, y \in \{0,1\}^\ell$. Let $\pi_x = |\phi_x\rangle \langle \phi_x|$, and $\pi'_x = |\phi'_x\rangle \langle \phi'_x|$. We know that $|\langle \phi_x \rangle \phi'_x| \geq 1 - \frac{\delta}{T(n)}$. Now, observe that $Mix(x, y, |\phi\rangle)$ can be exactly simulated in the following way:

(a) Apply the measurement $\{\pi_x + \pi_y, I - \pi_x - \pi_y\}$ to $|\phi\rangle$.

(b) If the measurement results in the first outcome, output $f(x, y)$.

(c) Otherwise, output $\perp$.

Let $\mathcal{Q}_{x,y}$ be the quantum channel defined by applying the measurement $\{\pi_x + \pi_y, I - \pi_x - \pi_y\}$. Let $\mathcal{Q}'_{x,y}$ be the quantum channel defined by applying the measurement $\{\pi'_x + \pi'_y, I - \pi'_x - \pi'_y\}$.
Since $|\langle \phi_x | \phi'_x \rangle|^2 \geq 1 - \epsilon$ and $|\langle \phi_y | \phi'_y \rangle|^2 \geq 1 - \epsilon$, we know by Lemma 2 that

$$\mathsf{TD}(\pi_x, \pi'_x) \leq \epsilon$$

and

$$\mathsf{TD}(\pi_y, \pi'_y) \leq \epsilon.$$

By the definition of trace distance, and the triangle inequality we have that for all mixed states ρ,

$$\left|\mathrm{Tr}((\pi_x + \pi_y)\rho) - \mathrm{Tr}((\pi'_x + \pi'_y)\rho)\right| \leq 2\epsilon$$

And so Lemma 3 tells us

$$\|\mathcal{Q}_{x,y} - \mathcal{Q}_{x',y'}\|_\diamond \leq \frac{4\delta}{T(n)}\,.$$

The only difference between V_6 and V_7 is that V_6 makes $T(n)$ queries to $\mathcal{Q}_{x,y}$ while V_7 makes $T(n)$ queries to $\mathcal{Q}'_{x,y}$. It follows, by triangle inequalities and Fact 9, that

$$\left|\Pr[V_6(x, w) = 1] - \Pr[V_7(x, w) = 1]\right| \leq 4\delta\,.$$

□

8. V_8: This is the same as V_7, except the random oracle is replaced by a $2T$ wise independent function f^*.

Claim 17. $acc(V_8) = acc(V_7)$.

Proof. From [Zha12], we know that T queries to a random function are perfectly indistinguishable from queries to a $2T$-wise independent family of functions. Thus, we have

$$acc(V_8) = acc(V_7).$$

□

Now, adding up the differences in acceptance probabilities, we have that, for any (x, w), with probability at least $1 - 4e^{-q(n)^2}$ over the choice of $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$

$$|acc(V_1) - acc(V_8)| \leq \frac{1}{12} + \mathsf{negl}(n) + \frac{1}{100} + 5\delta = \frac{1}{12} + \mathsf{negl}(n) + \frac{1}{100} + \frac{5}{200} \leq \frac{1}{8}.$$

Union bounding over all w, we have that for all w with probability at least $1 - 2^{q(n)} \cdot 4 \cdot e^{-q(n)^2} \geq 1 - e^{-q(n)}$ over the choice of $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$,

$$|acc(V_1) - acc(V_8)| \leq \frac{1}{8}.$$

Observe that V_8 is exactly $\widetilde{Ver}$, so we have shown that

$$\left|\Pr[\widetilde{\mathsf{Ver}}^{PSPACE}(x, w) \to 1] - \Pr[\mathsf{Ver}^{SG,Mix,PSPACE}(x, w) \to 1]\right| \leq \frac{1}{8}.$$

In particular, this means that with probability $\geq 1 - e^{-|x|}$, $\mathcal{B}$ correctly decides $x \in \Pi$. Since $\sum_{i=1}^{\infty} 2^i e^{-i} < \infty$, by the Borel-Cantelli lemma [ÉB09, Can17], $\mathcal{B}$ correctly decides $\Pi(x)$ for all but finitely many x with probability 1 over $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$. Thus, with probability 1 over $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$, there exists some $\mathcal{B}'$ which correctly decides $\Pi(x)$ for all x by hard-coding the points where $\mathcal{B}$ was wrong.

And so, for any fixed $\Pi \in \mathbf{QCMA}^{SG,Mix,PSPACE}$, with probability 1 over $\{U_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}$, $\Pi \in \mathbf{BQP}^{SG,Mix,PSPACE}$.

Taking a union bound over all $\Pi \in \mathbf{QCMA}^{SG,Mix,PSPACE}$, we get that

$$\mathbf{QCMA}^{SG,Mix,PSPACE} \subseteq \mathbf{BQP}^{SG,Mix,PSPACE}.$$

□

5.5 2QKD Exists Relative To $(SG, Mix, PSPACE)$

Theorem 16. *With probability* 1 *over* U, f, *there exists a 2QKD protocol relative to* $SG, Mix, PSPACE$.

Note that every two round QCCC KE protocol is also a 2QKD protocol. But observe that our construction in Theorem 13 is a two round QCCC KE protocol, and so Theorem 16 follows.

6 An Oracle World Where Quantum Lightning Exists But BQP = QMA

We formally define the oracle SG, V, C. Our oracle will be parameterized by a sequence of unitaries $\{U_n\}_{n\in\mathbb{N}}$. Each U_n will be sampled from the Haar distribution over $10n$ qubits.

We define the oracle (SG, V, C) as follows:

1. SG_n: Sample $x \leftarrow \{0,1\}^n$. Output $x, |\phi_x\rangle$. Here $|\phi_x\rangle := U_n |x, 0^{9n}\rangle$ for $x \in \{0,1\}^n$.

2. $V_n(\rho_{AB})$: Measure ρ_A in the standard basis, producing a measurement result a. Measure $U_n^\dagger \rho_B U_n \to z$ (that is, apply $U_n^\dagger$ to the register B and measure in the standard basis, getting a result z). Output 1 if and only if $z = (a, 0^{9n})$.

3. C: a non-trivial oracle independent of $\{U_n\}_{n\in\mathbb{N}}$ serving as a "**QMA** breaker". In particular, this will be the same oracle as the one used in [Kre21].

The following is taken verbatim from [Kre21]:
We construct the language C deterministically and independently of $\{U_n\}_{n\in\mathbb{N}}$. We specify the language in stages: first we define C's behavior on the 1-bit strings, then the 2-bit strings, then the 3-bit strings, and so on. For a string x, we define $C(x) = 1$ if the following all hold:

1. x is a description of a quantum oracle circuit $\mathcal{V}^{\overline{U},C}(|\psi\rangle)$ that takes a quantum state $|\psi\rangle$ as input, and makes queries to a family of quantum oracles $\overline{U} = \{\overline{U}_\ell\}_{\ell\in\mathbb{N}}$ and the classical oracle C. Note that $|\psi\rangle$ and $\overline{U}$ are not part of the description of Ver; they are auxiliary inputs.

2. $\mathcal{V}$ runs in time at most $|x| - 1$, and hence can query C on inputs of length at most $|x| - 1$.

3. The average acceptance probability of Ver (viewed as a QMA verifier) is greater than $1/2$ when averaged over $\overline{U}_\ell \leftarrow \mu_{2^{2\ell}}$. In symbols, we mean precisely:

$$\mathop{\mathbb{E}}_{\overline{U}_\ell\leftarrow\mu_{2^{10\ell}}:\ell\in\mathbb{N}}\left[\max_{|\psi\rangle}\Pr[\mathsf{Ver}^{\overline{U},C}(|\psi\rangle) = 1]\right] > \frac{1}{2}.$$

Theorem 17. *With probability 1 over $\{U_n\}_{n\in\mathbb{N}}$, $\mathbf{BQP}^{SG,V,C} = \mathbf{QMA}^{SG,V,C}$ and quantum lightning exists relative to (SG, V, C).*

6.1 Quantum Lightning Exists Relative To (SG, V, C)

In this section, we show quantum lightning exists relative to the oracle (SG, V, C). We will make use of the following lemma. The proof is deferred to Appendix A.

Lemma 11 (Variant of lemma 4.1 from [AGKL24])**.** *Let $N, M > 0$ be integers with $N \leq M$. Let ρ be the mixed state of dimension NM defined by sampling N states $|\phi_k\rangle \leftarrow \mu^s_M$ for $k \in [N]$. That is,*

$$\rho = \mathop{\mathbb{E}}_{|\phi_k\rangle\leftarrow\mu^s_M,k\in[N]}\left[\bigotimes_{k=1}^{N}|\phi_k\rangle\langle\phi_k|\right].$$

Let ρ' be the same state but where we require that each $|\phi_k\rangle, |\phi_{k'}\rangle$ are orthogonal. Formally, ρ' is the mixed state of dimension NM defined as

$$\rho' = \mathop{\mathbb{E}}_{U\leftarrow\mu_M}\left[\bigotimes_{k=1}^{N}U|k\rangle\langle k|U^\dagger\right].$$

Then $\mathrm{TD}(\rho, \rho') \leq \frac{N(N+1)}{2M}$.

Lemma 12 ([JLS18], Theorem 4)**.** *For $|\phi\rangle$ a quantum state, let $R_{|\phi\rangle}$ be the quantum channel which applies the measurement $\{|\phi\rangle\langle\phi|, I - |\phi\rangle\langle\phi|\}$. For any $\epsilon > 0$, there exists $t = O(1/\epsilon)$ and an efficient algorithm $\mathcal{B}$ such that for all $|\phi\rangle$,*

$$\|R_{|\phi\rangle} - \mathcal{B}(|\phi\rangle^{\otimes t}, \cdot)\|_\diamond \leq \epsilon$$

Remark 4. *This theorem is not identical to Theorem 4 from [JLS18]. In particular, they show a stronger version of this lemma where $R_{|\phi\rangle}$ is the unitary reflecting around $|\phi\rangle$. But it is standard that a reflection oracle can be used to perform this projection, and so our lemma follows.*

Theorem 18. *With probability 1 over $\{U_n\}_{n\in\mathbb{N}}$, quantum lightning exists relative to (SG, V, C).*

Proof. Our construction is as follows.

1. $Mint(1^n)$: Run $SG_n \to k, |\phi_k\rangle$. Output $\sigma = k$, $|\$\rangle = |\phi_k\rangle$.
2. $Ver(1^n, \sigma, |\$\rangle)$: Output $V_n(\sigma, |\$\rangle)$.

Correctness follows by construction. Let $\mathcal{A}$ be an attacker against the quantum lightning security of $(Mint, Ver)$. In particular, we will take $\mathcal{A}^{SG,V,C}(1^n, c)$ to be a uniform algorithm taking in $|c| = \mathsf{poly}(n)$ bits of non-uniform advice and making T queries to its oracles. We will show that for all c, $\Pr[\mathcal{A}$ wins the quantum lightning game$] \leq \mathsf{negl}(n)$.

Let $p(n)$ be any polynomial. We will define a sequence of hybrids, $\mathbf{G}_1, \dots, \mathbf{G}_6$ such that for each i,

$$|\Pr[\mathbf{G}_i \to 1] - \Pr[\mathbf{G}_{i+1} \to 1]| \leq \frac{1}{12p(n)}$$

Furthermore, it will be clear that with high probability over the choice of oracles, $\Pr[1 \leftarrow \mathbf{G}_4] \leq \frac{1}{p(n)}$.

1. $\mathbf{G}_1$ is the quantum lightning game played with $\mathcal{A}(1^n, \cdot)$. Formally, $\mathbf{G}_1$ goes as follows
   (a) Run $\mathcal{A}^{SG,V,C}(1^n, c) \to x, \rho_{AB}$.
   (b) Apply the measurement $\{|\phi_x, \phi_x\rangle\langle\phi_x, \phi_x|, I - |\phi_x, \phi_x\rangle\langle\phi_x, \phi_x|\}$ to ρ_{AB}. Output 1 if measurement results in the first outcome.
2. $\mathbf{G}_2$ replaces all calls to U_n in SG_n and V_n with calls to a freshly sampled $\widetilde{U}_n$. The oracles at all other indices are left alone.
3. $\mathbf{G}_3$ outputs $\perp$ if SG_n ever samples the same k twice.
4. $\mathbf{G}_4$ simulates V_n using Lemma 12. For $\epsilon = \frac{1}{12p(n)\cdot T(n)}$, let $t, \mathcal{B}$ be as in Lemma 12. $\mathbf{G}_4$ replaces all queries to V_n with queries to $\widetilde{V}_n$ defined as follows:
   (a) Measure ρ_A in the standard basis, producing a measurement result a.
   (b) Output $\mathcal{B}(|\phi_a\rangle^{\otimes t}, \rho_B)$.
5. $\mathbf{G}_5$ samples each $|\phi_k\rangle$ independently at random instead of guaranteeing they are all orthogonal. In particular, $\mathbf{G}_5$ samples for each $k \in \{0,1\}^n$, $|\phi'_k\rangle$ randomly from the Haar distribution on $10n$ qubits. It then replaces all queries to SG_n and $\widetilde{V}_n$ with queries to SG'_n and $\widetilde{V}'_n$ defined as follows.
   (a) SG'_n samples $k \leftarrow \{0,1\}^n$ and outputs $(k, |\phi_k\rangle)$.
   (b) $\widetilde{V}'_n(\rho_{AB})$ first measures ρ_A in the computational basis to obtain the result a. It then outputs $\mathcal{B}(|\phi'_a\rangle^{\otimes t}, \rho_B)$
6. $\mathbf{G}_6$ no longer simulates projections by using many copies of $|\phi'_k\rangle$. In particular, it replaces all queries to $\widetilde{V}_n$ with queries to $V'_n(\rho_{AB}$ defined as follows:
   (a) Measure ρ_A in the computational basis to obtain the result a.
   (b) Perform the measurement $\{|\phi'_k\rangle\langle\phi'_k|, I - |\phi'_k\rangle\langle\phi'_k|\}$ on ρ_B. Output 1 if and only if the measurement resulted in the first outcome.

Now we show each hybrids are close with each other.

1. Fix $U_{n'}$ for all $n' \neq n$. We can consider the function $f(U_n) = \Pr[1 \leftarrow \mathbf{G}_1(1^n)]$ as a function of U_n. By Lemma 4, this function is $2(T+1)$-Lipschitz. Furthermore, $\Pr[1 \leftarrow \mathbf{G}_2(1^n)] = \mathbb{E}_{U_n \leftarrow \mu_{2^{10n}}}[f(U_n)]$. And so, by Levy's lemma Theorem 10, with probability $\geq 1 - 2\exp\left(-\frac{(2^{10n}-2)}{2^{2n} \cdot 96(T+1)^2}\right) \geq 1 - 2^{-2^n}$,
$$|\Pr[1 \leftarrow \mathbf{G}_1(1^n)] - \Pr[1 \leftarrow \mathbf{G}_2(1^n)]| \leq \frac{1}{2^n} \leq \frac{1}{12p(n)}.$$

2. We know that $|\Pr[1 \leftarrow \mathbf{G}_2(1^n)] - \Pr[1 \leftarrow \mathbf{G}_3(1^n)]| \leq \Pr[\perp \leftarrow \mathbf{G}_3(1^n)]$. But this is exactly the probability that there is a repeat among $\leq T = \mathsf{poly}(n)$ samples of a random value from $\{0,1\}^n$. By the birthday bound, we get
$$|\Pr[1 \leftarrow \mathbf{G}_2(1^n)] - \Pr[1 \leftarrow \mathbf{G}_3(1^n)]| \leq \frac{T(n)^2}{2^n} \leq \frac{1}{12p(n)}.$$

3. Applying Lemma 12 along with the triangle inequality gives us that
$$|\Pr[1 \leftarrow \mathbf{G}_3(1^n)] - \Pr[1 \leftarrow \mathbf{G}_4(1^n)]| \leq T(n)\frac{1}{12p(n) \cdot T(n)} = \frac{1}{12p(n)}$$

4. Let ρ, ρ' be the states from Lemma 11 with $N = 2^n$ and $M = 2^{10n}$. We observe that the only difference between $\mathbf{G}_4(1^n)$ and $\mathbf{G}_5(1^n)$ is that the adversary in $\mathbf{G}_4(1^n)$ uses $\leq t \cdot T$ copies of ρ to run SG_n and $\widetilde{V}_n$, while the adversary in $\mathbf{G}_5(1^n)$ uses $\leq t \cdot T$ copies of ρ' to run SG'_n and $\widetilde{V}'_n$. But
$$\mathsf{TD}(\rho^{\otimes t \cdot T}, \rho'^{\otimes t \cdot T}) \leq t \cdot T\frac{2^{2n}}{2^{10n}} \leq O(T(n)^2 p(n)) 2^{-5n} \leq \frac{1}{12p(n)}$$
and so
$$|\Pr[1 \leftarrow \mathbf{G}_4(1^n)] - \Pr[1 \leftarrow \mathbf{G}_5(1^n)]| \leq \frac{1}{12p(n)}$$

5. Applying Lemma 12 along with the triangle inequality gives us that
$$|\Pr[1 \leftarrow \mathbf{G}_5(1^n)] - \Pr[1 \leftarrow \mathbf{G}_6(1^n)]| \leq T(n)\frac{1}{12p(n) \cdot T(n)} = \frac{1}{12p(n)}$$

And so by triangle inequality
$$|\Pr[1 \leftarrow \mathbf{G}_1(1^n)] - \Pr[1 \leftarrow \mathbf{G}_6(1^n)]| \leq \frac{1}{2p(n)} \tag{13}$$

We then show that given an adversary $\mathcal{A}$ such that $\Pr[1 \leftarrow \mathbf{G}_6(1^n)] \geq \frac{1}{2^n}$, we can construct an inefficient adversary $\mathcal{A}'$ cloning a single Haar-random state $|\psi\rangle$ with access to a verification oracle $R_{|\psi\rangle}$ which applies the measurement $\{|\psi\rangle\langle\psi|, I - |\psi\rangle\langle\psi|\}$.

$\mathcal{A}'$ behaves as follows.

1. On input $|\psi\rangle$.
2. Sample $x \leftarrow \{0,1\}^n$. Set $|\phi_x\rangle = |\psi\rangle$.
3. For each $k \neq x \in \{0,1\}^n$, sample $|\phi_k\rangle$ uniformly at random.

4. Run $(y, \rho) \leftarrow \mathcal{A}^{SG,V,C}(1^n, c)$.

5. When $\mathcal{A}$ queries SG_n, pick $k \leftarrow \{0,1\}^n$ and respond with $k, |\phi_k\rangle$. If the same k is ever picked twice, fail.

6. When $\mathcal{A}$ queries $V_n(\rho_{AB})$.

   (a) Measure ρ_A in the computational basis to obtain the result k.

   (b) If $k \neq x$, measure ρ_B under $\{|\phi_k\rangle\langle\phi_k|, I - |\phi_k\rangle\langle\phi_k|\}$ and output 1 if measurement results in the first outcome.

   (c) If $k = x$, output $R_{|\phi\rangle}(\rho_B)$.

7. If $y = x$, output ρ.

Note that the view of $\mathcal{A}$ is exactly the same view as in $\mathbf{G}_6$. So as long as $\mathcal{A}'$ correctly guesses the index $\mathcal{A}$ wins at, $\mathcal{A}'$ will also win. That is,

$$\Pr[\mathcal{A}' \text{ wins cloning game}] \geq \frac{1}{2^n} \Pr[1 \leftarrow \mathbf{G}_6(1^n)].$$

[AC12] shows that a counterfeiter needs $\Omega(\sqrt{\epsilon} 2^{10n/2})$ queries to a verification oracle in order clone a $10n$ qubit Haar-random state with success probability ϵ such that $\frac{1}{\epsilon} = o(2^{10n})$. So if we take $\epsilon = \frac{1}{2^n} \Pr[1 \leftarrow \mathbf{G}_4(1^n)] \geq \frac{1}{\mathsf{poly}(n) \cdot 2^n}$, we get that the counterfeiter needs $\Omega(2^{3n})$ queries to succeed. But $\mathcal{A}'$ makes at most $T = \mathsf{poly}(n)$ queries to its verification oracle $R_{|\phi\rangle}(\rho_B)$, so we have a contradiction.

Thus, $\Pr[1 \leftarrow \mathbf{G}_6(1^n)] \leq \frac{1}{2^n}$. And so, for all advice c, with probability $\geq 1 - 2^{-2^n}$ over U,

$$\Pr[\mathcal{A}(1^n, c) \text{ wins the quantum lightning game}] \leq \Pr[1 \leftarrow \mathbf{G}_1(1^n)] \leq \frac{1}{2^n} + \frac{1}{2p(n)} \leq \frac{1}{p(n)}.$$

And so, by a union bound over all advice strings $c \in \{0,1\}^{\mathsf{poly}(n)}$, for all polynomials $p(n)$, $\mathcal{A}$ achieves advantage larger than $1/p$ for any advice string $c \in \{0,1\}^{\mathsf{poly}(n)}$ with probability at most $2^{\mathsf{poly}(n)} \cdot 2^{-2^n} \leq \mathsf{negl}(n)$. But note that since $\sum_{n=1}^{\infty} \mathsf{negl}(n)$ converges, by the Borel-Cantelli lemma [ÉB09, Can17] $\mathcal{A}$ achieves negligible advantage for all but finitely many input lengths $n \in \mathbb{N}$ with probability 1 over $\{U_n\}_{n \in \mathbb{N}}$. In particular, $(Mint, Ver)$ is a quantum lightning scheme. $\square$

6.2 $\mathbf{BQP}^{SG,V,C} = \mathbf{QMA}^{SG,V,C}$

Theorem 19 (Theorem 17 restated)**.** *With probability 1 over $\{U_n\}_{n \in \mathbb{N}}$, $\mathbf{BQP}^{SG,V,C} = \mathbf{QMA}^{SG,V,C}$.*

Proof. Let $\Pi \in \mathbf{QMA}^{SG,V,C}$, which means there exists a polynomial-time verifier $\mathsf{Ver}^{SG,V,C}(x, |\psi\rangle)$ with completeness $\frac{2}{3}$ and soundness $\frac{1}{3}$. Without loss of generality, we can amplify the completeness and soundness probabilities of Ver to $\frac{11}{12}$ and $\frac{1}{12}$, respectively. Let $p(n)$ be a polynomial upper bound on the running time of Ver on inputs x of length n.

We now describe a $\mathbf{BQP}^{SG,V,C}$ algorithm $\mathcal{B}^{SG,V,C}$ such that, with probability 1 over $\{U_n\}_{n \in \mathbb{N}}$, $\mathcal{B}$ computes Π on all but finitely many inputs $x \in \mathrm{Dom}(\Pi)$.

1. Let $d := \log\left(3456|x|p(|x|)^2 + 2\right)$, $\delta = \frac{1}{24}$. For each $n \in [d]$, $\mathcal{B}$ performs process tomography on each U_ℓ in the following way:

   - Run SG until, $\forall k \in [2^\ell]$, there exist $t > 2^d \cdot \frac{T}{\delta}$ copies of $|\phi_k\rangle$.

- For every $k \in [2^\ell]$, run the state tomography algorithm from Theorem 11 on inputs $n = |x|^2 p(|x|)^2, \epsilon = \frac{\delta}{T}$. This produces estimates $|\phi'_k\rangle$ such that $|\langle\phi_k|\phi'_k\rangle|^2 \geq 1 - \epsilon$ with probability at least $1 - e^{-5|x|p(x)^2}$.
- Define the map $U'_\ell \,|k\rangle \to |\phi'_k\rangle$.
- Define SG_1 and V_1 to be the same as SG and V, but for all $\ell \in [d]$, all queries SG_1 and V_1 make to U_ℓ are replaced by U'_ℓ.

We denote the collection of estimates by $U' \coloneqq \{U'_\ell\}_{\ell \in [d]}$.

2. Next, $\mathcal{B}$ constructs a description y of a quantum oracle circuit $\mathcal{V}^{\overline{U},C}(|\psi\rangle\,; x, U')$. $\mathcal{V}$ has x and the unitaries in U' hard-coded into its description, takes an auxiliary input $|\psi\rangle$, and queries oracles $\overline{U}$ and C. On input $|\psi\rangle$, $\mathcal{V}^{\overline{U},C}(|\psi\rangle)$ replicates the behavior of $\mathsf{Ver}^{SG_1,V_1,C}(x, |\psi\rangle)$, except that for each $\ell \in [d]$, queries to U_ℓ are replaced by U'_ℓ, and for each $\ell \in [p(|x|)] \setminus [d]$, queries to U_ℓ are replaced by queries to $\overline{U}_\ell$.

3. Finally, $\mathcal{B}$ queries $C(x)$ and outputs the result.

We now show that for any $x \in \mathrm{Dom}(\Pi)$, with high probability over $\{U_n\}_{n\in\mathbb{N}}$, $\mathcal{B}$ correctly decides Π on x, which is to say that $\Pr\left[\mathcal{B}^{SG,V,C}(x) = \Pi(x)\right] \geq \frac{2}{3}$.

For a fixed x, given sequences of unitaries $U' = \{U_\ell\}_{\ell\in[d]}$ and $\overline{U} = \{\overline{U}_\ell\}_{\ell\in[p(|x|)]\setminus[d]}$, define

$$f(U', \overline{U}) \coloneqq \max_{|\psi\rangle} \Pr\left[\mathcal{V}^{\overline{U},C}(|\psi\rangle\,; x, U') = 1\right].$$

Note that, in this notation, $\mathcal{B}$ outputs 1 if and only if

$$\mathop{\mathbb{E}}_{\overline{U}}\left[f(U', \overline{U})\right] > \frac{1}{2}. \tag{14}$$

By contrast, the **QMA** acceptance probability of Ver itself may be written consistently with this notation as:

$$f(U, U) = \max_{|\psi\rangle} \Pr\left[\mathsf{Ver}^{SG,V,C}(x, |\psi\rangle) = 1\right]. \tag{15}$$

In effect, our goal is to show that Equation (14) gives a good estimator for Equation (15). We will do so in two steps: we first show that replacing U in f's second argument with an average over $\overline{U}$ approximately preserves the **QMA** acceptance probability, and then we argue similarly when replacing U by the estimate U' in f's first argument.

By Lemma 4, f is $p(|x|)$-Lipschitz with respect to the second argument $\overline{U}$, viewed as a direct sum of matrices $\overline{U} \equiv \bigoplus_{n=d+1}^{p(|x|)} \overline{U}_n$.[12] Hence, from Theorem 10 with $N = 3456|x|p(|x|)^2 + 2$, $L = p(|x|)$, and $t = \frac{1}{12}$, we have that:

$$\begin{aligned}
\Pr_{U \leftarrow \mathcal{D}}\left[\left|f(U,U) - \mathop{\mathbb{E}}_{\overline{U}}[f(U,\overline{U})]\right| \geq \frac{1}{12}\right] &\leq 2\exp\left(-\frac{(N-2)t^2}{24L^2}\right) \\
&= 2\exp\left(-\frac{3456|x|p(|x|)^2 \cdot \frac{1}{144}}{24p(|x|)^2}\right) \\
&= 2e^{-|x|}.
\end{aligned} \tag{16}$$

The factor of 2 appears because Theorem 10 applies to one-sided error, but the absolute value forces us to consider two-sided error.

[12] This is because each query to a single $\overline{U}_n$ may be simulated via one query to the entire direct sum.

From Theorem 11 and the union bound we have that with probability at least $1-e^{-|x|}$ (over the randomness of the process tomography algorithm), for all $\ell \in [d]$, $x \in \{0,1\}^\ell$, $|\langle \phi_x|\phi'_x\rangle|^2 \geq 1-\frac{\delta}{T}$. Let $\mathcal{Q}_x, \mathcal{Q}'_x$ be the quantum channels which perform the measurements $\{|\phi_x\rangle\langle\phi_x|, I-|\phi_x\rangle\langle\phi_x|\}$, $\{|\phi'_x\rangle\langle\phi'_x|, I-|\phi'_x\rangle\langle\phi'_x|\}$ respectively. By Lemmas 2 and 3,

$$\|\mathcal{Q}_x - \mathcal{Q}'_x\|_\diamond \leq 2\frac{\delta}{T}$$

Since $T = p(|x|)$ queries are made to U_ℓ,

$$\left| f(U,U) - \mathop{\mathbb{E}}_{\overline{U}}[f(U,\overline{U})] \right| \leq T \cdot \|\mathcal{Q}_x - \mathcal{Q}'_x\|_\diamond$$

and therefore, by Jensen's inequality, with probability $\geq 1-d\cdot e^{-q(|x|)^2}$ (over the randomness of the process tomography)

$$\left| \mathop{\mathbb{E}}_{\overline{U}}[f(U',\overline{U})] - \mathop{\mathbb{E}}_{\overline{U}}[f(U,\overline{U})] \right| \leq T \cdot \|\mathcal{Q}_x - \mathcal{Q}'_x\|_\diamond = 2\delta.$$

Combining with Equation (16), and recalling the acceptance criterion of $\mathcal{B}$ from Equation (14) and setting $\delta = \frac{1}{24}$ we conclude that except with probability at most $2e^{-|x|}$ over U,

$$\Pr_{U \leftarrow \mathbf{Haar}, U' \leftarrow \mathcal{B}}\left[\left| f(U,U) - \mathop{\mathbb{E}}_{\overline{U}}\left[f(U',\overline{U})\right] \right| \geq \frac{1}{6} \right] \leq \frac{1}{3}.$$

This is to say that $\mathcal{B}$ correctly decides $\Pi(x)$, except with probability at most $2e^{-|x|}$ over U. By the Borel–Cantelli lemma [ÉB09, Can17], because $\sum_{i=1}^{\infty} 2^i \cdot 2e^{-i} = \frac{4}{e-2} < \infty$, $\mathcal{B}$ correctly decides $\Pi(x)$ for all but finitely many $x \in \mathrm{Dom}(\Pi)$, with probability 1 over U. Hence, with probability 1 over U, $\mathcal{B}$ can be modified into an algorithm $\mathcal{B}'$ that agrees with Π on every $x \in \mathrm{Dom}(\Pi)$, by simply hard-coding those x on which $\mathcal{B}$ and Π disagree.

Because there are only countably many $\mathbf{QMA}^{SG,V,C}$ machines, we can union bound over all $\Pi \in \mathbf{QMA}^{SG,V,C}$ to conclude that $\mathbf{QMA}^{SG,V,C} \subseteq \mathbf{BQP}^{U,C}$ with probability 1. □

7 OWPuzz Constructions

7.1 OWPuzzs From QCCC KE

Theorem 20. *If there exist QCCC KE, then distributional OWPuzzs exist.*

Proof. Let (A,B) be a QCCC KE protocol. We construct a distributional OWPuzz, Samp, as follows:

1. Run the KE protocol $A \leftrightarrows B$, producing a transcript τ as well as outputs a, b.

2. Output $s = \tau, k = a$.

Define A_τ to be the distribution of A's output a conditioned on the transcript being τ. Observe that A_τ is exactly the distribution over k conditioned on $s = \tau$. Let $a_\tau = \mathsf{argmax}_a \Pr[A_\tau \to a]$. We can see that with probability $\geq \frac{99}{100}$ over τ, $\Pr[A_\tau \to a_\tau] \geq \frac{99}{100}$. Otherwise, A and B would disagree on the key with probability $\geq \frac{1}{10000}$, which violates correctness. In particular, this means that

$$\Pr\left[\ a = a_\tau \ :\ A \leftrightarrows B \to (\tau, a, b)\ \right] \geq \frac{98}{100}. \tag{17}$$

Let us assume towards contradiction that Samp is not a distributional OWPuzz. So there exists a QPT adversary $\mathcal{A}$ such that

$$\Delta((\tau, \mathcal{A}(\tau)), (\tau, a)) \leq \frac{1}{100}. \tag{18}$$

Equations (17) and (18) give us

$$\Pr[a_\tau \leftarrow \mathcal{A}(\tau)] \geq \frac{97}{100} \tag{19}$$

since otherwise we could inefficiently distinguish $(\tau, \mathcal{A}(\tau))$ from (τ, a). Applying the union bound to Equations (18) and (19) then gives us

$$\Pr[\mathcal{A}(\tau) = a : A \leftrightarrows B \rightarrow (\tau, a, b)] \geq \frac{95}{100} \tag{20}$$

and so $\mathcal{A}$ is also an attacker for the KE protocol. Since the KE protocol (A, B) is secure, so is the distributional OWPuzz Samp. □

7.2 OWPuzzs From QCCC Commitments

Theorem 21. *If QCCC commitment schemes exist, then distributional OWPuzzs exist.*

Proof. Let $(\mathsf{Com}, \mathsf{Rec})$ be any QCCC commitment scheme with an opening stage taking t rounds. We will assume that in each round, the receiver sends one message x_i and receives a response d_i.

We will construct a distributional OWPuzz, Samp, as follows:

1. Sample $i \leftarrow [t]$.
2. Sample $m \leftarrow \{0, 1\}$.
3. Run the committing stage on m, producing a transcript z.
4. Run the first i rounds of the opening stage, producing a transcript $x_1, d_1, \ldots, x_i, d_i$.
5. Output $s = (i, z, x_1, d_1, \ldots, x_i), k = d_i$.

Assume that Samp is not a distributional OWPuzzs. So for all constants $c \geq 1$, there exists an adversary $\mathcal{A}$ such that for infinitely many $n \in \mathbb{N}$,

$$\Delta((i, z, x_1, d_1, \ldots, x_i, d_i), (i, z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i))) \leq n^{-c}.$$

But

$$\begin{aligned}
&\Delta((i, z, x_1, d_1, \ldots, x_i, d_i), (i, z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i))) \\
&= \frac{1}{t} \sum_{i \in [t]} \Delta((z, x_1, d_1, \ldots, x_i, d_i), (z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i))) \\
&\geq \frac{1}{t} \Delta((z, x_1, d_1, \ldots, x_i, d_i), (z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i)))
\end{aligned}$$

for each $i \in [t]$. And so for each $i \in [t]$,

$$\Delta((z, x_1, d_1, \ldots, x_i, d_i), (z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i))) \leq t \cdot n^{-c}. \tag{21}$$

We can use $\mathcal{A}$ to build an adversary $\mathcal{B}$ that breaks hiding as follows.

1. Run the committing stage, producing transcript z and receiver state $\widetilde{\rho}^0_{Rec}$.

2. For each $i \in [t]$

   (a) Run $\widetilde{x}_i \leftarrow \mathsf{Rec}(i, z, \widetilde{x}_1, \widetilde{d}_1, \ldots, \widetilde{x}_{i-1}, \widetilde{d}_{i-1}, \widetilde{\rho}^{i-1}_{Rec})$ to produce the receiver's message in the ith round. Note that this will update $\widetilde{\rho_{Rec}}^{i-1}$ to a new state, which we will denote $\widetilde{\rho_{Rec}}^{i}$.

   (b) Run $\widetilde{d}_i \leftarrow \mathcal{A}(i, z, \widetilde{x}_1, \widetilde{d}_1, \ldots, \widetilde{x}_i)$ to simulate the committer's message in the ith round.

3. Produce the receiver's output $m' \leftarrow \mathsf{Rec}(\widetilde{\rho}^t_{Rec})$, and output m'.

Now we show that such m' is equal to the originally committed message m with high probability, and therefore $\mathcal{B}$ breaks hiding. Intuitively, if m' is another message than m, a malicious committer that runs $\mathcal{A}$ can open $m' \neq m$, which breaks the binding. Moreover, we can also show that the probability that $\mathcal{B}$ outputs $\perp$ is bounded.

More precisely, the argument is the following. Note that an adversarial committer could also simulate its messages in this way. That is, define Com' as follows

1. Let z be the transcript of the committing stage, and let $x_1, d_1, \ldots, x_i$ be the current transcript of the opening stage.

2. Run $d_i \leftarrow \mathcal{A}(i, z, x_1, d_1, \ldots, x_i)$.

3. Output d_i.

Observe that the resulting distribution on the output of Rec when interacting with Com' on a message m is identical to the output of our attacker when run on a transcript z generated from $\mathsf{Com}(m) \rightleftarrows \mathsf{Rec}$. Thus, by binding we have that

$$\Pr[\mathcal{B}(z, \rho_{Rec}) \to m' \notin \{m, \perp\} : m \leftarrow \{0,1\}, \mathsf{Com}(m) \rightleftarrows \mathsf{Rec} \to (z, \rho_{Rec})] \leq \mathsf{negl}(n). \quad (22)$$

Consider the mixed state $\mathcal{I}_i = (z, \widetilde{x}_1, \widetilde{d}_1, \ldots, \widetilde{x}_i, \widetilde{d}_i, \widetilde{\rho}^i_{Rec})$ produced by $\mathcal{B}$ after round i. We will also define the mixed state $\mathcal{J}_i = (z, x_1, d_1, \ldots, x_i, d_i, \rho^i_{Rec})$ produced by an honest interaction between Com and Rec after i rounds of the opening stage.

We will show by induction that $\|\mathcal{I}_i - \mathcal{J}_i\|_1 \leq itn^{-c}$. The base case is trivial, since $\mathcal{I}_0 = \mathcal{J}_0$. For the inductive step, assume that $\|\mathcal{I}_i - \mathcal{J}_i\|_1 \leq itn^{-c}$. Define a new distribution $\mathcal{I}'_i$ obtained by applying the ith iteration of $\mathcal{B}$ to a sample from $\mathcal{J}_{i-1}$. In particular, $\mathcal{I}'_i = (z, x_1, d_1, \ldots, d_{i-1}, x'_i, d'_i, \rho^i_{Rec})$ where $(z, x_1, d_1, \ldots, d_{i-1}) \leftarrow \mathcal{J}_{i-1}$, $x'_i = \mathsf{Rec}(\rho^{i-1}_{Rec})$ (producing ρ^i_{Rec}), and $d'_i = \mathcal{A}(i, z, x_1, \ldots, d_{i-1}, x'_i)$. Since trace distance is contractive under CPTP maps, we have

$$\left\|\mathcal{I}_i - \mathcal{I}'_i\right\|_1 \leq \|\mathcal{I}_{i-1} - \mathcal{J}_{i-1}\|_1 \leq (i-1)tn^{-c}. \quad (23)$$

But note that the internal state of the receiver at any point in time is a deterministic function of the transcript. In particular, ρ^i_{Rec} is the state achieved by running the protocol in superposition and then postselecting on the transcript being $(z, x_1, d_1, \ldots, x_i, d_i)$. Furthermore, x'_i and x_i are sampled identically. Thus,

$$\begin{aligned}
&\left\|\mathcal{I}'_i - \mathcal{J}_i\right\|_1 && (24)\\
&= \left\|(z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i), \rho^i_{Rec}) - (z, x_1, d_1, \ldots, x_i, d_i, \rho^i_{Rec})\right\|_1 && (25)\\
&= \|(z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i)) - (z, x_1, d_1, \ldots, x_i, d_i)\|_1 && (26)\\
&= \Delta((z, x_1, d_1, \ldots, x_i, \mathcal{A}(i, z, x_1, d_1, \ldots, x_i)), (z, x_1, d_1, \ldots, x_i, d_i)) \leq tn^{-c}, && (27)
\end{aligned}$$

where the last inequality follows from Equation (21). Combining Equations (23) and (24) gives us

$$\|\mathcal{I}_i - \mathcal{J}_i\|_1 \leq \left\|\mathcal{I}_i - \mathcal{I}'_i\right\|_1 + \left\|\mathcal{I}'_i - \mathcal{J}_i\right\|_1 \leq (i-1)tn^{-c} + tn^{-c} = itn^{-c}. \tag{28}$$

So in particular $\|\mathcal{I}_t - \mathcal{J}_t\|_1 \leq t^2 n^{-c}$. But we know that running Rec on the internal state from $\mathcal{J}_t$ will output something which is not $\perp$ with all but negligible probability by correctness. Thus, as long as $n^{-c} \leq \frac{1}{4t^2}$,

$$\Pr[\perp \leftarrow \mathcal{B}] \leq \frac{1}{4} + \mathsf{negl}(n). \tag{29}$$

So putting together Equations (22) and (29), we get

$$\begin{aligned} \Pr[m \leftarrow \mathcal{B}] &= 1 - \Pr[m \not\leftarrow \mathcal{B}] & (30)\\ &\geq 1 - \Pr[\mathcal{B} \rightarrow m' \notin \{m, \perp\}] - \Pr[\perp \leftarrow \mathcal{B}] & (31)\\ &\geq \frac{3}{4} - \mathsf{negl}(n). & (32) \end{aligned}$$

And so we achieve a contradiction. Thus, Samp is a distributional OWPuzz. □

7.3 2 Round QKD Implies State Puzzles

Definition 12 (State Puzzles [KT24b]). *A state puzzle is defined by a uniform QPT algorithm* $\mathsf{Samp}(1^n)$ *which outputs a classical-quantum state* $(s, |\psi_s\rangle)$ *such that given* s*, it is quantum computationally infeasible to output* ρ *which overlaps notably with* $|\psi_s\rangle$*. Formally, we define* σ_n *to be the mixed state corresponding to the output of* $\mathsf{Samp}(1^n)$*. For a non-uniform QPT* $\mathcal{A}$*, we define* $\sigma_n^{\mathcal{A}}$ *to be the mixed state corresponding to the following process*

$$(s, |\psi_s\rangle) \leftarrow \mathsf{Samp}(1^n)$$
$$\rho \leftarrow \mathcal{A}(1^n, s)$$
$$\textit{Output } (s, \rho)$$

We say that Samp *is a state puzzle if for all QPT* $\mathcal{A}$*,*

$$\mathsf{TD}(\sigma_n, \sigma_n^{\mathcal{A}}) \geq 1 - \mathsf{negl}(n).$$

Remark 5. *Note that this definition at first seems different than the one in [KT24b], which says that the expected overlap between* ρ *and* $|\psi_s\rangle$ *is negligible. Formally,*

$$\mathbb{E}\left[\ \mathrm{Tr}(|\psi_s\rangle\!\langle\psi_s|\,\rho)\ :\ \begin{array}{c}(s, |\psi_s\rangle) \leftarrow \mathsf{Samp}(1^n)\\ \rho \leftarrow A(1^n, s)\end{array}\ \right] \leq \mathsf{negl}(n).$$

However, these definitions are equivalent. In particular, the best distinguisher between σ *and* $\sigma^{\mathcal{A}}$ *is the map which on input* $(s, |\psi_s\rangle)$ *applies the measurement* $\{|\psi_s\rangle\!\langle\psi_s|, I - |\psi_s\rangle\!\langle\psi_s|\}$*. But the advantage of this distinguisher is exactly*

$$1 - \mathbb{E}\left[\ \mathrm{Tr}(|\psi_s\rangle\!\langle\psi_s|\,\rho)\ :\ \begin{array}{c}(s, |\psi_s\rangle) \leftarrow \mathsf{Samp}(1^n)\\ \rho \leftarrow A(1^n, s)\end{array}\ \right].$$

For simplicity of presentation, we will not consider quantum advice. Observing the proof, our result also holds for quantum advice.

Informally, QKD describes a two party protocol where two parties, Alice and Bob, send each other authenticated classical messages and unauthenticated quantum messages. Correctness says that at the end of the protocol, Alice and Bob output the same shared message. Security says that for any adversary Eve who can read all classical messages and tamper with all quantum messages, at the end of the protocol execution

1. (Validity): Either Alice and Bob output the same key, or one of the parties outputs $\perp$.

2. (Security): If either Alice or Bob does not output $\perp$, then Eve cannot learn their final key.

We consider a weaker notion of security, which we call one-sided computational security. In particular, we will only require that if Bob does not output $\perp$, then Eve cannot predict Bob's key. This definition immediately follows from the more standard stronger definition. Since we build OWPuzzs from the weaker definition, it is also possible to build OWPuzzs from the stronger definition.

Definition 13 (Modified from [MW24]). *A two-round quantum key distribution (2QKD) scheme is defined as a tuple of algorithms $(QKDFirst, QKDSecond, QKDDecode)$ with the following syntax.*

1. *$QKDFirst(1^n) \to (msg, \mu, st)$: A QPT algorithm which, on input the security parameter 1^n, outputs a message, consisting of a classical component msg and a quantum state μ, as well as an internal quantum state st. We allow μ to be a mixed state, and in particular μ and st may be entangled.*

2. *$QKDSecond(msg, \mu) \to \{(resp, \eta, k), \perp\}$: A QPT algorithm which, on input the first message (msg, μ), outputs a response consisting of a classical component $resp$ and a quantum mixed state η, as well as a classical key $k \in \{0,1\}^n$. We also allow $QKDSecond$ to output a special symbol $\perp$, denoting rejection.*

3. *$QKDDecode(st, resp, \eta) \to \{k, \perp\}$: A QPT algorithm which, on input the internal state st and the response $(resp, \eta)$, outputs a key $k \in \{0,1\}^n$. We also allow $QKDDecode$ to output a special symbol $\perp$, denoting rejection.*

Definition 14 (2QKD Correctness). *We say a 2QKD protocol $(QKDFirst, QKDSecond, QKDDecode)$ satisfies correctness if there exists a negligible function negl such that for all $n \in \mathbb{N}$,*

$$\Pr[\perp \leftarrow QKDSecond(msg, \mu) : (msg, \mu, st) \leftarrow QKDFirst(1^n)] \leq \mathsf{negl}(n)$$

and

$$\Pr\left[QKDDecode(st, resp, \eta) = k \ : \ \begin{array}{c} (msg, \mu, st) \leftarrow QKDFirst(1^n) \\ (resp, \eta, k) \leftarrow QKDSecond(msg, \mu) \end{array} \right] \geq 1 - \mathsf{negl}(n).$$

Definition 15 (2QKD One-sided Computational Security). *For a family of QPT algorithms $\mathcal{A} = \{\mathcal{A}_n\}_{n\in\mathbb{N}}$, we define the following security game $QKDSec(\mathcal{A}, 1^n)$ as follows.*

1. *Run $(msg, \mu, st) \leftarrow QKDFirst(1^n)$ and send (msg, μ) to $\mathcal{A}_n$.*

2. *$\mathcal{A}_n$ returns a quantum state ρ, which may be entangled with $\mathcal{A}$'s internal state.*

3. *Run $\{(resp, \eta, k), \perp\} \leftarrow QKDSecond(msg, \rho)$.*

4. *Forward the response to $\mathcal{A}_n$.*

5. *$\mathcal{A}_n$ outputs $k' \in \{0,1\}^n$.*

6. *The game outputs 1 if $k' = k \neq \perp$.*

We say a 2QKD protocol $(QKDFirst, QKDSecond, QKDDecode)$ *satisfies one-sided computational security if for all families of QPT algorithms* $\mathcal{A} = \{\mathcal{A}_n\}_n$*, there exists a negligible function* negl *such that*

$$\Pr[1 \leftarrow QKDSec(\mathcal{A}, 1^n)] \leq \mathsf{negl}(n).$$

Theorem 22. *If there exists a 2QKD protocol satisfying correctness and one-sided computational security, then one-way puzzles exist.*

Proof. We will assume without loss of generality that the combined quantum state μ, st is a single pure state $|\phi\rangle_{BC}$ over two registers. We can do this by running $QKDFirst$, deferring all measurements to the end. This will end up with a state over four registers A, B, C, D where A is the output register for msg, B is the output register for μ, C is the output register for st, and D is some ancilla register. We will then set st to be the state over the registers C and D combined. $QKDDecode$ will simply ignore the D register. This leads to a protocol with identical behavior, but where the state μ, st is pure.

In particular, we will build a state puzzle. Samp will be defined as follows.

1. Run $(msg, |\phi\rangle_{BC}) \leftarrow QKDFirst(1^n)$.
2. Output $s = msg$, $|\psi_s\rangle = |\phi\rangle_{BC}$.

We then claim that Samp is a state puzzle. In particular, let p be a polynomial and let $\mathcal{A}$ be a non-uniform QPT adversary breaking state puzzle security of Samp with advantage $\epsilon(n)$. Formally, for $\sigma_n, \sigma_n^{\mathcal{A}}$ from Definition 12

$$\mathsf{TD}(\sigma_n, \sigma_n^{\mathcal{A}}) \leq 1 - \epsilon(n)$$

for infinitely many n.

We will construct an attacker $\mathcal{A}'$ breaking QKD one-sided computational security with advantage dependent on ϵ.

1. On receiving (msg, μ), run $|\phi\rangle_{BC} \leftarrow \mathcal{A}(1^n, msg)$.
2. Output $|\phi\rangle_B$, keeping internal state $|\phi\rangle_C$.
3. On receiving $(resp, \eta)$, run $k' \leftarrow QKDDecode(|\phi\rangle_B, resp, \eta)$.
4. Output k'.

We can construct an inefficient adversary $\mathcal{A}''$ which acts the same as $\mathcal{A}'$, but always perfectly inverts the puzzle. That is, $\mathcal{A}''$ is defined as follows.

1. On receiving (msg, μ), output $|\phi_{msg}\rangle$, the residual state of $QKDFirst$'s quantum output conditioned on the classical message being msg.
2. Output $|\phi_{msg}\rangle_B$, keeping internal state $|\phi_{msg}\rangle_C$.
3. On receiving $(resp, \eta)$, run $k' \leftarrow QKDDecode(|\phi_{msg}\rangle_B, resp, \eta)$.
4. Output k'.

By the definition of $\mathcal{A}$,

$$\mathsf{TD}((msg, \mathcal{A}(1^n, msg)), (msg, |\phi_{msg}\rangle)) \leq 1 - \epsilon(n)$$

and so since the only difference between $QKDSec(\mathcal{A}')$ and $QKDSec(\mathcal{A}'')$ is that $\mathcal{A}(msg) = |\phi\rangle$ is replaced by $|\phi_{msg}\rangle$, we get

$$\left|\Pr[1 \leftarrow QKDSec(\mathcal{A}'', 1^n)] - \Pr[1 \leftarrow QKDSec(\mathcal{A}', 1^n)]\right| \leq 1 - \epsilon(n). \tag{33}$$

But we can also observe that $QKDSec(\mathcal{A}'', 1^n)$ simply runs the honest QKD protocol, with $\mathcal{A}''$ running $QKDDecode$ on the correct state. Thus, by correctness, we have that in $QKDSec(\mathcal{A}'', 1^n)$, the following hold

$$\begin{aligned} \Pr[k = \perp] &\leq \mathsf{negl}(n) \\ \Pr[k' = k] &\geq 1 - \mathsf{negl}(n) \end{aligned} \tag{34}$$

Thus,

$$\Pr[1 \leftarrow QKDSec(\mathcal{A}'', 1^n)] \geq 1 - \mathsf{negl}(n). \tag{35}$$

Equations (33) and (35) together give us

$$\Pr[1 \leftarrow QKDSec(\mathcal{A}', 1^n)] \geq \epsilon(n) - \mathsf{negl}(n)$$

and so by 2QKD security, $\epsilon(n) \leq \mathsf{negl}(n)$. □

Theorem 23 (Corollary of Theorem 13)**.** *There exists an oracle relative to which 2QKD exists, but* $\mathbf{BQP} = \mathbf{QCMA}$*.*

Proof. This oracle is exactly the same as the QCCC KE oracle. In particular, we observe that our QCCC KE protocol consists of exactly two rounds. But any 2 round QCCC KE protocol is trivially also a 2QKD protocol. And so the theorem follows. □

Acknowledgement

Eli Goldin is supported by a NSF Graduate Student Research Fellowship. Tomoyuki Morimae is supported by JST CREST JPMJCR23I3, JST Moonshot JPMJMS20615-1-1, JST FOREST, MEXT QLEAP, the Grant-in Aid for Transformative Research Areas (A) 21H05183, and the Grant-in-Aid for Scientific Research (A) No.22H00522.

References

[AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In Howard J. Karloff and Toniann Pitassi, editors, *44th Annual ACM Symposium on Theory of Computing*, pages 41–60, New York, NY, USA, May 19–22, 2012. ACM Press.

[ACC+22] Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 165–194, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[AGKL24] Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 226–254, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.

[AGL24] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common haar state model: Feasibility results and separations. Cryptology ePrint Archive, Report 2024/1043, 2024.

[BB14] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, dec 2014.

[BCG90] Mihir Bellare, Lenore Cowen, and Shafi Goldwasser. On the structure of secret key exchange protocols (rump session). In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO'89*, volume 435 of *Lecture Notes in Computer Science*, pages 604–605, Santa Barbara, CA, USA, August 20–24, 1990. Springer, New York, USA.

[BCN24] John Bostanci, Boyang Chen, and Barak Nehoran. Oracle separation between quantum commitments and quantum one-wayness. arXiv, 2024, 2410.03358.

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *ITCS 2023: 14th Innovations in Theoretical Computer Science Conference*, volume 251, pages 24:1–24:21, Cambridge, MA, USA, January 10–13, 2023. LIPIcs.

[BMM+24] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. A new world in the depths of microcrypt: Separating OWSGs and quantum money from QEFID. Cryptology ePrint Archive, Paper 2024/1567, 2024.

[BS19] Zvika Brakerski and Omri Shmueli. (Pseudo) random quantum states with binary phase. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 229–250, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland.

[Can17] Francesco Cantelli. Sulla probabilista come limita della frequencza. *Rend. Accad. Lincei*, 26:39, 1917.

[CCS24] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: constructing and separating quantum pseudorandomness. arXiv, 2024, 2404.03295.

[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 227–258, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Cham, Switzerland.

[CGG+23] Bruno Cavalar, Eli Goldin, Matthew Gray, Peter Hall, Yanyi Liu, and Angelos Pelecanos. On the computational hardness of quantum one-wayness. arXiv, 2023, 2312.08363.

[CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 215–248, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[CLM23] Kai-Min Chung, Yao-Ting Lin, and Mohammad Mahmoody. Black-box separations for non-interactive classical commitments in a quantum world. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part I*, volume 14004 of *Lecture Notes in Computer Science*, pages 144–172, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland.

[ÉB09] M. Émile Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, Dec 1909.

[GKKT20] M Guţă, J Kahn, R Kueng, and J A Tropp. Fast state tomography with optimal error bounds. *Journal of Physics A: Mathematical and Theoretical*, 53(20):204001, apr 2020.

[GLLZ21] Siyao Guo, Qian Li, Qipeng Liu, and Jiapeng Zhang. Unifying presampling via concentration bounds. In Kobbi Nissim and Brent Waters, editors, *TCC 2021: 19th Theory of Cryptography Conference, Part I*, volume 13042 of *Lecture Notes in Computer Science*, pages 177–208, Raleigh, NC, USA, November 8–11, 2021. Springer, Cham, Switzerland.

[GYZ17] Sumegha Garg, Henry Yuen, and Mark Zhandry. New security notions and feasibility results for authentication of quantum data. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 342–371, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Cham, Switzerland.

[HKOT23] Jeongwan Haah, Robin Kothari, Ryan O'Donnell, and Ewin Tang. Query-optimal estimation of unitary channels in diamond distance. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 363–390, 2023.

[HR07] Iftach Haitner and Omer Reingold. Statistically-hiding commitment from any one-way function. In David S. Johnson and Uriel Feige, editors, *39th Annual ACM Symposium on Theory of Computing*, pages 1–10, San Diego, CA, USA, June 11–13, 2007. ACM Press.

[IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, Research Triangle Park, NC, USA, October 30 – November 1, 1989. IEEE Computer Society Press.

[INN+22] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. Quantum Search-To-Decision Reductions and the State Synthesis Problem. In Shachar Lovett, editor, *37th Computational Complexity Conference (CCC 2022)*, volume 234 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:19, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, WA, USA, May 15–17, 1989. ACM Press.

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland.

[KMNY24] Fuyuki Kitagawa, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum public-key encryption with tamper-resilient public keys from one-way functions. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 93–125, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[KNY23] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Publicly verifiable deletion from minimal assumptions. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023: 21st Theory of Cryptography Conference, Part IV*, volume 14372 of *Lecture Notes in Computer Science*, pages 228–245, Taipei, Taiwan, November 29 – December 2, 2023. Springer, Cham, Switzerland.

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602, Orlando, FL, USA, June 20–23, 2023. ACM Press.

[Kre21] William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[KT24a] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th Annual ACM Symposium on Theory of Computing*, pages 968–978, Vancouver, BC, Canada, June 24–28, 2024. ACM Press.

[KT24b] Dakshita Khurana and Kabir Tomer. Founding quantum cryptography on quantum advantage, or, towards cryptography from #P-hardness. Cryptology ePrint Archive, Paper 2024/1490, 2024.

[Kuk18] Ryszard Kukulski. Distinguishability of quantum measurements. Presented at Institute of Theoretical and Applied Informatics, Polish Academy of Sciences, 2018.

[LMW24] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th Annual ACM Symposium on Theory of Computing*, pages 979–990, Vancouver, BC, Canada, June 24–28, 2024. ACM Press.

[Mec19] Elizabeth S. Meckes. *The Random Matrix Theory of the Classical Compact Groups.* Cambridge Tracts in Mathematics. Cambridge University Press, 2019.

[MPY23] Tomoyuki Morimae, Alexander Poremba, and Takashi Yamakawa. Revocable quantum digital signatures. Cryptology ePrint Archive, Report 2023/1937, 2023.

[MW24] Giulio Malavolta and Michael Walter. Robust quantum public-key encryption with applications to quantum key distribution. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 126–151, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 269–295, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[MY24] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In Frédéric Magniez and Alex Bredariol Grilo, editors, *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024, September 9-13, 2024, Okinawa, Japan*, volume 310 of *LIPIcs*, pages 4:1–4:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[NC11] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge University Press, USA, 10th edition, 2011.

[Unr07] Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *Advances in Cryptology – CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 205–223, Santa Barbara, CA, USA, August 19–23, 2007. Springer, Berlin, Heidelberg, Germany.

[Unr16] Dominique Unruh. Computationally binding quantum commitments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 497–527, Vienna, Austria, May 8–12, 2016. Springer, Berlin, Heidelberg, Germany.

[Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, January 1983.

[Yam] Shogo Yamada. personal communication.

[Yan22] Jun Yan. General properties of quantum bit commitments (extended abstract). In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 628–657, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland.

[Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd Annual Symposium on Foundations of Computer Science*, pages 679–687, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press.

[Zha19] Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 408–438, Darmstadt, Germany, May 19–23, 2019. Springer, Cham, Switzerland.

A Proof of Lemma 11

Lemma 13 (Lemma 11 restated)**.** *Let $N, N > 0$ be integers with $N \leq M$. Let ρ be the mixed state of dimension NM defined by sampling N states $|\phi_k\rangle \leftarrow \mu_M^s$ for $k \in [N]$. That is,*

$$\rho = \mathop{\mathbb{E}}_{|\phi_k\rangle \leftarrow \mu_M^s, k\in[N]} \left[\bigotimes_{k=1}^{N} |\phi_k\rangle\langle\phi_k| \right].$$

Let ρ' be the same state but where we require that each $|\phi_k\rangle, |\phi_{k'}\rangle$ are orthogonal. Formally, ρ' is the mixed state of dimension NM defined as

$$\rho' = \mathop{\mathbb{E}}_{U \leftarrow \mu_M} \left[\bigotimes_{k=1}^{N} U|k\rangle\langle k|U^\dagger \right].$$

Then $\mathsf{TD}(\rho, \rho') \leq \frac{N(N+1)}{2M}$.

Proof. Define ρ_i to be the state defined by sampling the first k states to all be orthogonal and the last $N - i$ states uniformly. Formally,

$$\rho_i = \mathop{\mathbb{E}}_{U \leftarrow \mu_M, |\phi_k\rangle \leftarrow \mu_M^s} \left(\bigotimes_{k=1}^{i} U|k\rangle\langle k|U^\dagger \right) \otimes \left(\bigotimes_{k=i+1}^{N} |\phi_k\rangle\langle\phi_k| \right).$$

Here $\rho_0 = \rho'$ and $\rho_n = \rho$. We will show that for all i, $\mathsf{TD}(\rho_i, \rho_{i+1}) \leq \frac{i}{M}$. In particular, define

$$S_i = \left\{ \bigotimes_{k=1}^{N} |\psi_k\rangle : \text{for all } j, j' \leq i, \langle \psi_j | \psi_{j'} \rangle = 0 \right\}.$$

Define Π_i to be the projector onto the subspace spanned by S_i. We observe that $\Pi_{i+1}\rho_i$ is proportional to ρ_{i+1}. That is, $I - \Pi_{i+1}$ is the measurement which optimally distinguishes ρ_i and ρ_{i+1}. In particular,

$$\mathsf{TD}(\rho_i, \rho_{i+1}) \leq \mathrm{Tr}((I - \Pi_{i+1})(\rho_i - \rho_{i+1})) = 1 - \mathrm{Tr}(\Pi_{i+1}\rho_i).$$

But $\mathrm{Tr}(\Pi_{i+1}\rho_i)$ is the probability that applying $\{\sum_{k\leq i} |\psi_k\rangle\!\langle\psi_k| , I - \sum_{k<i} |\psi_k\rangle\!\langle\psi_k|\}$ to ρ_i results in the first measurement outcome (here $|\psi_k\rangle = U\,|k\rangle$). So

$$\begin{aligned}
\mathrm{Tr}(\Pi_{i+1}\rho_i) &= \mathbb{E}\left[\ \textstyle\sum_{k\leq i} |\langle\psi_{i+1}| U |k\rangle|^2 \ : \ \begin{matrix} U \leftarrow \mu_m \\ |\psi_{i+1}\rangle \leftarrow \mu_M^s \end{matrix}\ \right] \\
&\leq \sum_{k\leq i} \mathbb{E}\left[\ |\langle\psi_{i+1}| U |k\rangle|^2 \ : \ \begin{matrix} U \leftarrow \mu_m \\ |\psi_{i+1}\rangle \leftarrow \mu_M^s \end{matrix}\ \right] \\
&\leq i \cdot \mathbb{E}\left[|\langle\psi|\phi\rangle|^2 : |\psi\rangle, |\phi\rangle \leftarrow \mu_M^s \right] \\
&= \frac{i}{M}.
\end{aligned}$$

And so, we get that

$$\mathsf{TD}(\rho, \rho') \leq \sum_{i=1}^{N} \frac{i}{M} = \frac{N(N+1)}{2M}.$$

□