The Entangled Quantum Polynomial Hierarchy Collapses

Sabee Grewal*   Justin Yirka*

Abstract

We introduce the entangled quantum polynomial hierarchy QEPH as the class of problems that are efficiently verifiable given alternating quantum proofs that may be entangled with each other. We prove QEPH collapses to its second level. In fact, we show that a polynomial number of alternations collapses to just two. As a consequence, QEPH = QRG(1), the class of problems having one-turn quantum refereed games, which is known to be contained in PSPACE. This is in contrast to the *unentangled* quantum polynomial hierarchy QPH, which contains QMA(2).

We also introduce a generalization of the quantum-classical polynomial hierarchy QCPH where the provers send probability distributions over strings (instead of strings) and denote it by DistributionQCPH. Conceptually, this class is intermediate between QCPH and QPH. We prove DistributionQCPH = QCPH, suggesting that only quantum superposition (not classical probability) increases the computational power of these hierarchies. To prove this equality, we generalize a game-theoretic result of Lipton and Young (1994) which says that the provers can send distributions that are uniform over a polynomial-size support. We also prove the analogous result for the polynomial hierarchy, i.e., DistributionPH = PH. These results also rule out certain approaches for showing QPH collapses.

Finally, we show that PH and QCPH are contained in QPH, resolving an open question of Gharibian et al. (2022).

1 Introduction

The polynomial hierarchy [MS72, Sto76] is a hierarchy of complexity classes that are known to equal P if and only if P = NP. The hierarchy, denoted by PH, is a natural generalization of efficient proof verification and nondeterminism and plays a central role in complexity theory. Given its significance, it is natural to explore quantum generalizations of PH, yet such generalizations remain understudied.

Before discussing quantum polynomial hierarchies, let us first informally define PH. Intuitively, PH is a hierarchy of complexity classes that can solve progressively harder problems, extending beyond both NP and coNP. One can think of PH as a public debate between Alice and Bob, who take turns presenting polynomial-sized proofs (bit strings) to a referee. At the end of the debate, the referee takes the proofs, performs a polynomial-time classical computation, and decides a winner.

More formally, a problem is in the k-th level of the polynomial hierarchy Σ_k^p if there is a deterministic polynomial-time verifier M (the referee) that takes proofs $y_1, \ldots, y_k$ and satisfies the following conditions. On yes-instances, $\exists y_1 \forall y_2 \exists y_3 \ldots$ such that $M(y_1, \ldots, y_k) = 1$, and, on no-instances, $\forall y_1 \exists y_2 \forall y_3 \ldots$ such that $M(y_1, \ldots, y_k) = 0$. PH is comprised of every level Σ_k^p for all natural numbers k, and it is strongly believed that PH is infinite.

Gharibian, Santha, Sikora, Sundaram, and Yirka [GSS+22] studied two quantum generalizations of PH. They generalized the class QCMA to the quantum-classical polynomial hierarchy QCPH, the class of problems for which a quantum verifier can efficiently verify solutions given a constant number of classical proofs from competing provers. Note that this is the same as PH except the verifier can perform a polynomial-time *quantum* computation. In the same work, they generalized the class QMA(2) to the *unentangled* quantum polynomial hierarchy QPH, for which the verifier is still quantum, but the proofs are quantum mixed states and promised to be unentangled from each other. Notably, Gharibian et al. did not introduce a hierarchy in which the proofs can be *entangled*, and they did not establish a relationship between QPH and QCPH (or even QPH and PH), leaving it unclear whether or not QPH was at least as powerful as its classical counterpart.[1] More generally, if QCPH and QPH are indeed more powerful, it prompts the question of why: is it quantum verification, quantum proofs, unentanglement, or some nuanced combination?

In this work, we address all of these questions related to quantum generalizations of the polynomial hierarchy. First, we ask (and answer) what problems admit a PH-style protocol where the provers can send potentially entangled proofs. We show that this new hierarchy—the entangled quantum polynomial hierarchy (QEPH)—behaves drastically differently from what we believe about PH, QCPH, and QPH.

Second, we prove that PH ⊆ QCPH ⊆ QPH, confirming the intuitive relationship between these hierarchies.

Lastly, to understand the power of quantum proofs, we introduce a generalization of QCPH where the provers send probability distributions over classical proofs and denote the class by DistributionQCPH. We prove that DistributionQCPH = QCPH, despite the intuition from game theory that mixed strategies are helpful. Note that the only difference between DistributionQCPH and QPH is that the proofs in QPH involve quantum superposition. Hence, our result establishes that the increased computational power of QPH comes only from the quantum superposition in the proofs.

1.1 Our Results

Our first main result is a characterization of our newly defined hierarchy QEPH (Definition 3.2) via a collapse to its second level. This collapse is in stark contrast to our belief that PH is infinite.

Theorem 1.1 (Combination of Lemma 4.1 and Theorem 4.2). QEPH *collapses to its second level and equals* QRG(1).

This collapse is similar to others known in quantum complexity theory, such as QIP = QIP(3) = QMAM [KW00, MW05], in which the protocols rely on the prover's ability to entangle their messages. We further compare QEPH to other complexity classes involving entangled proofs in Related and Concurrent Work.

We show that QEPH equals QRG(1), the class of problems having one-turn quantum-refereed games.[2] QRG(1) involves a game between two competing players that each privately sends a quantum state to a referee, who then performs a polynomial-time quantum computation to determine a winner. In 2009, Jain and Watrous [JW09] proved QRG(1) ⊆ PSPACE. However, it is conjectured that QRG(1) is strictly less powerful than QRG(2) = PSPACE [GW13, GW23]. Yet despite effort, no

---

[1]While these containments are what one might guess to be true, proving them is nontrivial.

[2]The class QRG(k) and its classical analogue RG(k) have been numbered differently by different authors. We follow recent conventions where the provers' and the referee's messages are counted separately. So, e.g., in QRG(2) the referee sends one message and then the provers each simultaneously send a message.

improved upper bounds on QRG(1) have been proven in over a decade. We suggest a new approach to improving the upper bound on QRG(1) (via the connection to QEPH) in Open Problems.

Our collapse result is stronger than stated above. It is well-known that if one extends PH to a polynomial number of rounds (rather than a constant number), then the resulting class equals PSPACE [AB09, Theorem 4.11]. In contrast, we show that extending QEPH to a polynomial number of rounds does not increase the power of the class.

Theorem 1.2 (Informal version of Corollary 4.3). *Even with a polynomial number of rounds,* QEPH *collapses to its second level.*

One interpretation of our collapse result is that allowing provers to entangle their proofs gives them too much opportunity to cheat. Hence, receiving a single proof from each prover is just as useful as receiving many entangled proofs.

Before this work, it was unclear how the quantum polynomial hierarchies compared to one another, and if QPH even contained PH. In our second result, we establish the following containments between the quantum and classical hierarchies, resolving an open question of Gharibian et al. [GSS⁺22].

Theorem 1.3 (Restatement of Theorem 5.1). PH $\subseteq$ QCPH $\subseteq$ QPH.

We emphasize that even PH $\subseteq$ QPH is not obvious. Placing restrictions on the provers can sometimes increase computational power, as was the case in, e.g., the recent results showing that QMA⁺ = QMA(2)⁺ = NEXP [JW23, BFM23]. Meanwhile, the permissiveness of QEPH, where we allow the provers to entangle their proofs, seems to yield a weaker class than QPH.

In our third result, we show that the power of QCPH does not change if the provers are allowed to send probability distributions (instead of a fixed classical proof).

Theorem 1.4 (Restatement of Corollary 6.5). DistributionQCPH = QCPH.

Our motivation for studying DistributionQCPH is to better understand the power of quantum proofs. In particular, let pureQPH be the same as QPH except the quantum proofs are pure states rather than mixed states.[3] Then the *only* difference between pureQPH and DistributionQCPH is that the former involves proofs that are quantum superpositions over bit strings while the latter involves proofs that are classical distributions over bit strings. Yet DistributionQCPH = QCPH is in the counting hierarchy [GSS⁺22], and pureQPH contains QMA(2) and is contained in EXP^{PP} [AGKR]. Conceptually, our result says that any increase in computational power only comes from the quantum superposition in the proofs.

Theorem 1.4 also goes through for PH.

Theorem 1.5 (Restatement of Theorem 6.1). DistributionPH = PH.

An easy consequence of our result is that DistributionPH collapses if and only if PH collapses.[4] Therefore, any attempts to collapse QCPH, QPH, or pureQPH must not collapse DistributionPH, and so Theorem 1.5 rules out some approaches to collapsing these hierarchies. In particular, one line of attack to showing QMA(2) = NEXP is to show that the $\forall$ quantifier in QΣ_3 does not add any computational power, because QMA(2) $\subseteq$ QΣ_3 $\subseteq$ NEXP [GSS⁺22]. Theorems 1.4 and 1.5 are

---

[3]It is easy to see that QPH $\subseteq$ pureQPH since the provers can send purifications of their mixed proofs.

[4]Distribution$\mathcal{C}$ is not to be confused with the notation Dist$\mathcal{C}$, which has been used in average-case complexity theory, e.g. DistNP and DistPH. Also, similar names like STOCHASTIC SAT or PROBABILISTIC QBF have appeared in the study of randomized quantifiers. These are more similar to the Arthur-Merlin (AM) hierarchy.

evidence that this line of attack won't work, since showing the analogous result for DistributionPH would collapse the polynomial hierarchy.

We give a graphical description of our results and the quantum polynomial hierarchy landscape in Fig. 1.
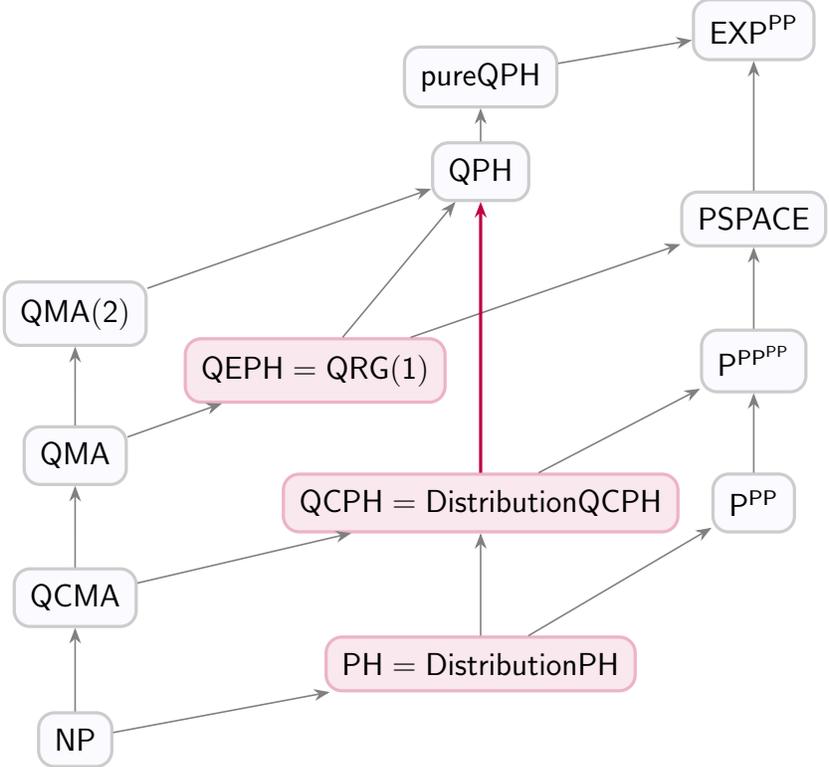

Figure 1: (Color) The quantum polynomial hierarchy landscape in light of our work. The containments and complexity classes shown in grey were previously known, and the containments and complexity classes in red are contributions of this work.

1.2 Main Ideas

Let us begin by explaining QEPH on an intuitive level (see Definition 3.2 for a formal definition). QEPH can be thought of as a constant-round non-interactive game between two competing provers, Alice and Bob, who take turns sending quantum registers, i.e., collections of qubits, to a verifier. Alice and Bob are allowed to entangle their own quantum registers across turns. The verifier then performs a polynomial-time quantum computation, measures a fixed output qubit in the computational basis, and, if the verifier sees 1, they accept (Alice wins), and reject otherwise (Bob wins). QEPH contains the decision problems for which Alice always wins with high probability on yes-instances and Bob always wins with high probability on no-instances. We note that in this game the moves are *public*, which means that Alice knows the state of the quantum registers sent by Bob and vice versa. See Remark 3.4 for further discussion of public vs. private moves in a quantum world.

To highlight the key technique in our proof that QEPH collapses (Lemma 4.1), we explain how to simulate the third level of QEPH, denoted by $QE\Sigma_3$, inside of the second level $QE\Sigma_2$. The proof for higher levels proceeds by induction. As we will explain formally in Section 3, a $QE\Sigma_i$ protocol

can be written as an optimization problem with a value equal to the probability the verifier accepts when both players use optimal strategies. In particular, Alice selects proofs that maximize the probability of the verifier accepting, while Bob selects proofs to minimize that probability. For $\mathsf{QE\Sigma_3}$, given a problem instance in which the verifier's action is encoded by an observable R, the corresponding optimization problem is

$$\max_{\rho_1 \in \mathbf{D}(\mathcal{X}_1)} \min_{\sigma \in \mathbf{D}(\mathcal{Y})} \max_{\rho_2 \in \mathcal{A}} \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma\right)\right),$$

where $\mathbf{D}(\mathcal{H})$ denotes the set of density operators on the Hilbert space $\mathcal{H}$ and $\mathcal{A} \coloneqq \{\rho \in \mathbf{D}(\mathcal{X}_1 \otimes \mathcal{X}_2) \mid \operatorname{tr}_{\mathcal{X}_2}(\rho) = \rho_1\}$. The restriction of the second maximization to the set $\mathcal{A}$ is to enforce that Alice's second move is consistent with her first.

A straightforward analysis shows that when focusing on the inner two operators, a min-max theorem applies, allowing us to swap the ordering of the inner minimization and maximization. Then, because we allow entangled states, we can combine the two sequential maximization operators into one, leaving an optimization problem corresponding to a two-round protocol. Notably, both the three-round and two-round protocols are over the same input and verifier, so the reduction does not increase the problem size or change the error parameters.

It is natural to ask why our technique does not also collapse PH. In short, the above approach fails immediately, since, for one, our collapse theorem relies on the fact that Alice and Bob are choosing quantum proofs from compact and convex sets (see Facts 2.3 and 2.4). In contrast, the set of classical strings is neither compact nor convex.

To show that $\mathsf{QEPH} = \mathsf{QRG}(1)$, we build on a previous characterization of Gharibian et al. [GSS$^+$22] where they showed that the second level of the *unentangled* quantum polynomial, denoted by $\mathsf{Q\Sigma_2}$, equals $\mathsf{QRG}(1)$. We extend their result in Proposition 3.3 to show that $\mathsf{QE\Sigma_2} = \mathsf{Q\Sigma_2} = \mathsf{QRG}(1)$, which yields our characterization that $\mathsf{QEPH} = \mathsf{QE\Sigma_2} = \mathsf{QRG}(1)$. $\mathsf{QE\Sigma_2} = \mathsf{Q\Sigma_2}$ because, after two turns, each prover has only sent a single proof, so there's no distinction yet to be made between the entangled versus entangled hierarchies.

We now turn to the containment $\mathsf{QCPH} \subseteq \mathsf{QPH}$, which are both defined formally in Section 2.3. In QCPH, the verifier receives classical proofs, whereas the proofs in QPH are unentangled quantum mixed states. One naïve approach to simulating QCPH inside of QPH—which does not work—is for the verifier to immediately measure the quantum proofs to get classical strings and then run the QCPH verification protocol. The reason this fails is that the dishonest prover (i.e., the player without a winning strategy) can cheat by sending a quantum state, rather than a classical proof. In more detail, while the honest prover has perfect knowledge of the quantum states sent by the dishonest prover, they do not know which particular classical strings the verifier will observe upon measurement, making it unclear what their response should be. The definition of QCPH guarantees the correct player has an effective response conditioned on any particular proof sent from the other player, but this does not generally guarantee the correct player can succeed against a mixture of potential moves. Unfortunately, the equilibrium point of a zero-sum game which allows for such mixed moves will generally be mixed, rather than pure.

To overcome this, we simulate the i-th level of QCPH in the $2ki$-th level of QPH, for some constant k. We ask the provers to send k copies of each of the proofs they would send in the QCPH protocol, which increases the number of turns by a factor of $2k$. Using the groups of k proofs, we give a simple test to ensure that no player cheats, which works as follows. Measure each of the k proofs in the standard basis. If the outcomes are all equal, then the test passes, and, otherwise, the test fails. We prove that this is enough to force the provers to send computational basis states with high probability.

We remark that this bears some similarity to other protocols involving unentanglement. Harrow and Montanaro [HM13] used unentanglement to force Merlin to send k-partite states, and, recently, Jeronimo and Wu [JW23] use unentanglement to force Merlin to send many copies of (approximately) the same quantum state. Both of these results fundamentally rely on the swap test, which tests for equality between two quantum states [BCWdW01]. In a similar fashion, we use unentanglement to force the provers to send standard basis states, i.e., classical strings. With that, we design a simulation of any QCPH protocol inside of QPH.

Finally, we discuss our proof that DistributionQCPH = QCPH (the same techniques will also show DistributionPH = PH). In DistributionQCPH, the provers send probability distributions over polynomial-length classical proofs. For classical proofs of length m, these distributions can have a support of size exponential in m. The key lemma for this result says that the provers can send *much simpler* distributions without changing the acceptance probability of the verifier too much. In particular, we prove that the distributions sent by the provers can be *uniform* over poly(m) many classical proofs and, even with this simplification, the acceptance probability of the verifier will change by at most a small constant. This simplification lemma (Lemma 6.4) generalizes a result due to Lipton and Young [LY94] and Althöfer [Alt94], where they showed that such a simplification works in the special case of a one-turn two-player game. Our contribution is to generalize their result to the case where the two players alternate sending distributions any constant number of times.

With the simplification lemma, one can prove DistributionQCPH $\subseteq$ QCPH as follows. To send a distribution in QCPH, the provers send every classical string that is in the support of their distribution. By our simplification lemma, there are only a polynomial number of such strings, so all of them can be sent in a polynomially-sized classical proof. Then, since the simplified distributions are uniform, the verifier can randomly sample one of the strings uniformly at random. The other direction DistributionQCPH $\supseteq$ QCPH follows from the same techniques that prove QCPH $\subseteq$ QPH.

1.3 Related and Concurrent Work

Early efforts to define quantum hierarchies include [Yam02, GK12].

We choose to use alternating $\exists$ and $\forall$ quantifiers to define QEPH (as was the case for QCPH and QPH in [GSS+22]). In addition to a quantifier definition, PH can be *equivalently* defined in the oracle model via constant-height towers of the form $\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}^{\cdots}}}$. The oracular definition gives rise to natural definitions of quantum polynomial hierarchies, some of which have been studied recently. Vinkhuijzen [Vin18] and Aaronson, Ingram, and Kretschmer [AIK22] study the "QMA hierarchy", QMAH, which consists of constant-depth towers of the form $\mathsf{QMA}^{\mathsf{QMA}^{\mathsf{QMA}^{\cdots}}}$.[5] [Vin18, Theorem 5] shows that QMAH is contained in the counting hierarchy CH, while the best upperbounds for the quantifier-based hierarchies, QEPH and QPH, are PSPACE and $\mathsf{EXP}^{\mathsf{PP}}$, respectively.

The method of showing equivalence between the quantifier-based and oracle-based definitions of PH does not appear to carry over to QEPH, QPH, or even QCPH. This seems related to the inability to "pull quantumness out of a quantum algorithm" as we can for randomness from randomized algorithms [AIK22] as well as a lack of study of quantum oracle machines. We further discuss questions regarding QMAH vs. QEPH in Open Problems.

There are several quantum complexity classes that involve provers sending possibly entangled proofs to a quantum polynomial-time verifier. We do not attempt to survey them here, but,

---

[5]Vinkhuijzen only allows recursive queries to QMA, whereas Aaronson, Ingram, and Kretschmer allow recursive queries to PromiseQMA.

for convenience, we summarize quantum complexity classes involving entangled proofs (and their classical counterparts) in Table 1.

Our work on DistributionPH builds on previous game-theoretic characterizations in complexity theory (see e.g., [FKS95]). PH-style classes involve a debate with public communication (perfect information), and a non-interacting, passive referee. RG-style classes involve private communication (imperfect information) with provers sending particular strings to the referee (perfect recall). A consequence of imperfect information is that the players must model their competitor's moves as probability distributions (mixed strategies) because they are never sure which move is made. Our class DistributionPH fits into this framework in a nuanced way. Specifically, the distributions sent are public (similar to PH); they represent a mixture of pure moves (similar to RG); but, uniquely, the provers don't know which string will be sampled by the referee (reminiscent of imperfect recall). This is a novel game-theoretic model, and as we discuss further in Section 6, it is naturally motivated by a game of quantum mixed states sent to a non-interacting referee.

Finally, the concurrent and independent work of Agarwal, Gharibian, Koppula, and Rudolph [AGKR] also studies generalizations of the polynomial hierarchy. They prove QCPH ⊆ pureQPH, which is similar to our Theorem 5.1 that QCPH ⊆ QPH. Since QPH ⊆ pureQPH is straightforward (the provers send purifications of their proofs), our Theorem 5.1 implies QCPH ⊆ pureQPH. In this sense, our containment is stronger. However, their containment has the nice (and nontrivial) feature that the k-th level of QCPH is contained in the k-th level of pureQPH, whereas our containment requires blowing up to the ck-th level of QPH for a constant integer c. Besides this, Agarwal et al. contribute several more results including a theorem that if $QC\Sigma_i = QC\Pi_i$ then QCPH collapses (see also [FGN23]); a Karp-Lipton style result that QCMA ⊆ BQP/mpoly implies QCPH collapses; a new upper bound QPH ⊆ pureQPH ⊆ EXP^{PP}, improving on the previous upper bound of EXPH; and a method for one-sided error-reduction of pureQPH.

Table 1: Complexity classes characterizing proof verification that are related to QEPH. "C" means classical and "Q" means quantum. For every class below, multiple provers are always competing, and, for multi-round quantum protocols, the quantum proofs can be entangled across rounds. Public means that the provers have full knowledge of their opponent's previous turns.

Class	# of Rounds	# of Provers	Proofs	Verifier	Interaction from referee?	Public or Private	Equals
NP	1	1	C	C	no	N/A	
QMA	1	1	Q	Q	no	N/A	
IP	poly	1	C	C	yes	N/A	PSPACE [Sha92]
QIP(3)	3	1	Q	Q	yes	N/A	PSPACE [JJUW11]
PH	const	2	C	C	no	pub.	
QEPH	const	2	Q	Q	no	pub.	QRG(1) [This Work]
RG(1)	1	2	C	C	no	priv.	S_2P[Alt94, LY94]
RG(2)	2	2	C	C	yes	priv.	PSPACE [FK97]
RG	poly	2	C	C	yes	priv.	EXP [FK97]
RG(pub)	poly	2	C	C	yes	pub.	PSPACE [FK97]
QRG(1)	1	2	Q	Q	no	priv.	
QRG(2)	2	2	Q	Q	yes	priv.	PSPACE [GW13]
QRG	poly	2	Q	Q	yes	priv.	EXP [GW07]

2 Preliminaries

We introduce notation, definitions, and background that are central to our results. For the most part, we assume familiarity with common concepts and classes in quantum and classical complexity theory as well as quantum computing and quantum information. For a thorough discussion of these topics, see [AB09, Wat18, KSV02, NC10].

We will need the following version of Hoeffding's inequality.

Fact 2.1 (Hoeffding's inequality). *Let $X_1, \ldots, X_n$ be independent random variables subject to $a_i \leq X_i \leq b_i$ for all i. Let $X = \sum_{i=1}^{n} X_i$ and let $\mu = \mathbf{E}[X]$. Then it holds that*

$$\mathbf{Pr}[X - \mu \geq t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right)$$

and

$$\mathbf{Pr}[X - \mu \leq -t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right).$$

2.1 Quantum Information

A *quantum register* refers to a collection of qubits. Associated with each register is a complex Hilbert space, and the state of a quantum register is described by a Hermitian, positive semi-definite matrix with trace one called a *density matrix*. We denote the set of n-qubit density matrices by $\mathbf{D}(n)$, and the sets of linear operators and density matrices on a complex Hilbert space $\mathcal{H}$ by $\mathbf{L}(\mathcal{H})$ and $\mathbf{D}(\mathcal{H})$, respectively.

For two quantum registers (X, Y) with Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$, the combined space is the tensor product space $\mathcal{X} \otimes \mathcal{Y}$. The partial trace $\mathrm{tr}_{\mathcal{Y}} : \mathbf{L}(\mathcal{X} \otimes \mathcal{Y}) \to \mathbf{L}(\mathcal{X})$ is the unique linear map that satisfies $\mathrm{tr}_{\mathcal{Y}}(A \otimes B) = \mathrm{tr}(A)B$ for all $A \in \mathbf{L}(\mathcal{X})$ and $B \in \mathbf{L}(\mathcal{Y})$. If the compound register (X, Y) is in the state $\rho \in \mathbf{D}(\mathcal{X} \otimes \mathcal{Y})$, then the state of register X is $\mathrm{tr}_{\mathcal{Y}}(\rho) \in \mathbf{D}(\mathcal{X})$. That is, operationally speaking, the partial trace is the act of ignoring (or discarding) a quantum register. We note that the partial trace $\mathrm{tr}_{\mathcal{X}}$ can be defined similarly, and, in general, the context in which the partial trace is used should clarify which spaces are being "traced out".

A *quantum measurement* of a quantum register is described by a finite collection of Hermitian, positive semi-definite matrices that sum to identity. Let X be a quantum register with Hilbert space $\mathcal{X}$ whose state is described by ρ. Let $\mathcal{M} = \{E_i \mid i \in \Sigma\}$ be a quantum measurement, where Σ is a finite alphabet. Upon measuring X with $\mathcal{M}$, we observe $i \in \Sigma$ with probability $\mathrm{tr}(E_i \rho)$.

2.2 A Min-Max Theorem

To prove our collapse theorem, we use a weaker version of Sion's min-max theorem.

Theorem 2.2 (A weaker version of Sion's min-max theorem [Sio58]). *Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, let $\mathcal{A} \subseteq \mathcal{X}$ and $\mathcal{B} \subseteq \mathcal{Y}$ be convex and compact subsets, and let $f : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$ be a bilinear function. Then*

$$\max_{a \in \mathcal{A}} \min_{b \in \mathcal{B}} f(a, b) = \min_{b \in \mathcal{B}} \max_{a \in \mathcal{A}} f(a, b).$$

It is a well-known fact that the space of density matrices is compact and convex.

Fact 2.3 ([Wat18, Chapter 1]). *Let $\mathbf{D}(\mathcal{H})$ be the set of density matrices on a complex Hilbert space $\mathcal{H}$. $\mathbf{D}(\mathcal{H})$ is compact and convex.*

It is critical for us that, even if we impose partial trace constraints on the set of density matrices, the set remains compact and convex. We include a proof for completeness.

Fact 2.4. *Let* X, Y *be two quantum registers with Hilbert spaces* $\mathcal{X}$ *and* $\mathcal{Y}$, *respectively, and let* $\mathbf{D}(\mathcal{X} \otimes \mathcal{Y})$ *be the corresponding set of density operators. Let* $\rho' \in \mathbf{D}(\mathcal{X})$ *be some fixed density matrix. Then the set*

$$\mathbf{S} = \{\rho \in \mathbf{D}(\mathcal{X} \otimes \mathcal{Y}) \mid \mathrm{tr}_{\mathcal{Y}}(\rho) = \rho'\}$$

is compact and convex.

Proof. Let $\rho_1, \rho_2 \in \mathbf{S}$, and define $\sigma := \theta\rho_1 + (1 - \theta)\rho_2$ for for $\theta \in [0, 1]$. Then

$$
\begin{aligned}
\mathrm{tr}_{\mathcal{Y}}(\sigma) &= \mathrm{tr}_{\mathcal{Y}}(\theta\rho_1 + (1 - \theta)\rho_2) \\
&= \theta\,\mathrm{tr}_{\mathcal{Y}}(\rho_1) + (1 - \theta)\,\mathrm{tr}_{\mathcal{Y}}(\rho_2) && \text{(By the linearity of the partial trace.)} \\
&= \theta\rho' + (1 - \theta)\rho' && \text{(Because } \rho_1, \rho_2 \in \mathbf{S}.) \\
&= \rho',
\end{aligned}
$$

so $\mathbf{S}$ is convex.

To show that $\mathbf{S}$ is compact, we must show that it is closed and bounded. Without loss of generality, let X be an n-qubit register and Y be an m-qubit register. Then we can identify $\mathbf{S}$ with the vector space $\mathbb{C}^{4^{n+m}}$ and observe that all entries are bounded in magnitude by 1. Therefore, $\mathbf{S}$ is bounded. To see that $\mathbf{S}$ is closed, we need the following definitions. For $x \in \mathbb{C}$, define $f_x : \mathbb{C}^{4^{n+m}} \to \mathbb{C}$ as $f_x(A) = \langle x, Ax \rangle$, which is continuous because the inner product is continuous; define $g : \mathbb{C}^{4^{n+m}} \to \mathbb{C}^{4^{n+m}}$ as $g(A) = A - A^\dagger$, which is a polynomial and therefore continuous; and, finally, define $h : \mathbb{C}^{4^{n+m}} \to \mathbb{C}^{4^n}$ as $h(A) = \mathrm{tr}_{\mathcal{Y}}(A)$, which is a linear map on a finite-dimensional vector space and therefore continuous. Then

$$\mathbf{S} = \bigcap_{x \in \mathbb{C}} f_x^{-1}([0, \infty)) \cap g^{-1}(\{0\}) \cap h^{-1}(\{\rho'\}) \cap \mathrm{tr}^{-1}(\{1\}).$$

The preimage of a continuous function on a closed set is closed, and the intersection of closed sets is closed. Therefore, $\mathbf{S}$ is closed. $\qquad\square$

2.3 Previously Studied Hierarchies

Here, we give formal definitions of the polynomial hierarchy PH, the quantum-classical polynomial hierarchy QCPH, and the unentangled quantum polynomial hierarchy QPH, the latter two of which were both introduced by Gharibian et al. [GSS$^+$22]. These classes will appear again in Section 5 when we prove $\mathsf{QCPH} \subseteq \mathsf{QPH}$ and in Section 6 when we prove $\mathsf{DistributionQCPH} = \mathsf{QCPH}$. We defer definitions of our new classes until later, with QEPH studied in Section 4 and $\mathsf{DistributionQCPH}$ in Section 6.

Definition 2.5 (Σ_i^{p}). A language L is in the i-th level of the polynomial hierarchy Σ_i^{p} if there exists a polynomial-time deterministic Turing Machine M such that for any n-bit input x,

$$
\begin{aligned}
x \in L &\iff \exists y_1 \forall y_2 \exists y_3 \ldots Q_i y_i \text{ such that } M(x, y_1, \ldots, y_i) = 1, \\
x \notin L &\iff \forall y_1 \exists y_2 \forall y_3 \ldots \overline{Q_i} y_i \text{ such that } M(x, y_1, \ldots, y_i) = 0,
\end{aligned}
$$

where Q_i denotes $\exists$ if i is odd and $\forall$ otherwise, $\overline{Q_i}$ denotes the complement of Q_i, and $|y_i| \le p(n)$ for some fixed polynomial p for all i.

Definition 2.6 (The polynomial hierarchy (PH) [Sto76])**.** The Polynomial-time Hierarchy is defined as

$$\mathsf{PH} := \bigcup_{i=0}^{\infty} \Sigma_i^{\mathsf{p}}.$$

Note the union which defines PH is over values of i which are constant, independent of a problem's input size. Observe also that for all i, $\Sigma_i^{\mathsf{p}} \subseteq \Sigma_{i+1}^{\mathsf{p}}$. Additionally, PH is closed under complement, in particular because $\overline{\Sigma}_i^{\mathsf{p}} \subseteq \Sigma_{i+1}^{\mathsf{p}} \subseteq \mathsf{PH}$. The complement of $\overline{\Sigma}_i^{\mathsf{p}}$ is defined to be Π_i^{p}, and for all i we have $\Sigma_i^{\mathsf{p}} \subseteq \Pi_{i+1}^{\mathsf{p}} \subseteq \Pi_{i+2}^{\mathsf{p}}$.

The definition of PH is particularly robust. The class can be defined equivalently by $\Sigma_{i+1}^{\mathsf{p}} = \mathsf{NP}^{\Sigma_i^{\mathsf{p}}}$, giving a constant-height tower of NP oracles. The model of alternating nondeterministic Turing Machines also can be used to define each level of the hierarchy. In another direction, the Sipser–Lautemann theorem shows $\mathsf{BPP} \subseteq \Sigma_2^{\mathsf{p}} \cap \Pi_2^{\mathsf{p}} \subseteq \mathsf{PH}$ [Sip83, Lau83]. So, natural bounded-error or probabilistic definitions of PH collapse to the standard, deterministic definition given above. This is also true for oracle definitions, where we know $\mathsf{MA}^{\mathsf{MA}^{\cdots}} = \mathsf{PH}$.

Even a partial survey of results regarding PH would be impossible to fit here. We finally note that $\Sigma_i^{\mathsf{p}} = \Sigma_{i+1}^{\mathsf{p}}$ or $\Sigma_i^{\mathsf{p}} = \Pi_i^{\mathsf{p}}$ would both "collapse" the hierarchy so that $\mathsf{PH} = \Sigma_i^{\mathsf{p}}$. These two events are analogous to $\mathsf{P} = \mathsf{NP}$ or $\mathsf{NP} = \mathsf{coNP}$. Conversely, if PH collapses to any finite level, it implies analogs of $\mathsf{P} = \mathsf{NP}$ and $\mathsf{NP} = \mathsf{coNP}$ must be true for some degree of nondeterminism, at some level of the hierarchy. So, the strongly-believed conjecture that PH is not equal to any Σ_i^{p} for fixed i is a generalization of those other strongly-believed conjectures.

The uniform circuit model is standard for quantum complexity classes, so we give the definition below.

Definition 2.7 (Polynomial-time uniform family of quantum circuits)**.** A polynomial-time uniform family of quantum circuits is a family $\{V_n\}_{n \in \mathbb{N}}$ such that there exists a polynomial bounded function $t : \mathbb{N} \to \mathbb{N}$ and a deterministic Turing machine M acting as follows. For every n-bit input x, M outputs in time $t(n)$ a description of a quantum circuit V_n, which has a designated output qubit. We say V_n accepts when we observe a 1 upon measuring the designated output qubit in the standard basis.

We generally leave the subscript implicit and just write V. Additionally, we often consider a single problem instance defined by an input x for the full length of an analysis. So instead of writing $V(x, y)$ for input x and proof y, we simply refer to $V(y)$.

As with most quantum complexity classes, we will be working with promise problems. Briefly, a promise problem A is a pair of non-intersecting subsets $(A_{\text{yes}}, A_{\text{no}})$ of $\{0, 1\}^*$. A decision problem, or language, is a promise problem where $A_{\text{yes}} \cup A_{\text{no}} = \{0, 1\}^*$.

We are now ready to define QCPH.

Definition 2.8 ($\mathsf{QC\Sigma_i}$ [GSS$^+$22])**.** A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in i-th level of the quantum-classical polynomial hierarchy $\mathsf{QC\Sigma_i}(\mathsf{c}, \mathsf{s})$ for polynomial-time computable functions $c, s :$ $\mathbb{N} \to [0, 1]$ if there exists a polynomial-time uniform family of quantum circuits $\{V_n\}_{n \in \mathbb{N}}$ such that for every n-bit input x, V_n takes in proofs $y_1, \ldots, y_i \subseteq \{0, 1\}^{m(n)}$ for fixed polynomial m and measures a fixed output qubit to decide to accept or reject, such that

- Completeness: $x \in L_{\text{yes}} \Rightarrow \exists y_1 \forall y_2 \exists y_3 \ldots Q_i y_i$ such that $\mathbf{Pr}\left[V(y_1, \ldots, y_i) \text{ accepts}\right] \geq c$,

- Soundness: $x \in L_{\text{no}} \Rightarrow \forall y_1 \exists y_2 \forall y_3 \ldots \overline{Q_i} y_i$ such that $\mathbf{Pr}\left[V(y_1, \ldots, y_i) \text{ accepts}\right] \leq s$,

where Q_i denotes $\exists$ if i is odd and $\forall$ otherwise, $\overline{Q_i}$ denotes the complement of Q_i, and, for all i, $|y_i| \leq p(n)$ for a fixed polynomially bounded function p. When the completeness and soundness parameters c, s are not specified, define

$$\mathsf{QC\Sigma_i} := \bigcup_{c-s \in \Omega(1/\operatorname{poly}(n))} \mathsf{QC\Sigma_i(c, s)}.$$

Definition 2.9 (The quantum-classical polynomial hierarchy (QCPH) [GSS$^+$22])**.** The quantum-classical polynomial hierarchy is defined as

$$\mathsf{QCPH} := \bigcup_{i=0}^{\infty} \mathsf{QC\Sigma_i}.$$

Observe that $\mathsf{QC\Sigma_0} = \mathsf{BQP}$ and $\mathsf{QC\Sigma_1} = \mathsf{QCMA}$. Gharibian et al. [GSS$^+$22] proved that QCPH is contained in $\mathsf{P^{PP^{PP}}}$, the second level of the counting hierarchy CH.

The definition of $\mathsf{QC\Sigma_i}$ to generically include $\mathsf{QC\Sigma_i(c, s)}$ for all $c - s \geq 1/\operatorname{poly}(n)$ is justified in part by the result of [GSS$^+$22] that for any such c and s, we may reduce the error such that for any polynomially bounded function r, we have $\mathsf{QC\Sigma_i(c, s)} = \mathsf{QC\Sigma_i(1 - 2^{-r}, 2^{-r})}$.

The unentangled quantum polynomial hierarchy QPH is defined similarly. The only difference is that the classical proofs are replaced by unentangled quantum proofs.

Definition 2.10 ($\mathsf{Q\Sigma_i}$ [GSS$^+$22])**.** A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in the i-th level of the unentangled quantum polynomial hierarchy $\mathsf{Q\Sigma_i(c, s)}$ for polynomial-time computable functions $c, s : \mathbb{N} \to [0, 1]$ if there exists a polynomial-time uniform family of quantum circuits $\{V_n\}_{n \in \mathbb{N}}$ such that for every n-bit input x, V_n takes in quantum proofs $\rho_1, \ldots, \rho_i$ and measures a fixed output qubit to decide to accept or reject, such that

- Completeness: $x \in L_{\text{yes}} \Rightarrow \exists \rho_1 \forall \rho_2 \exists \rho_3 \ldots Q_i \rho_i$ such that $\mathbf{Pr}\left[V(\rho_1, \ldots, \rho_i) \text{ accepts}\right] \geq c$,

- Soundness: $x \in L_{\text{no}} \Rightarrow \forall \rho_1 \exists \rho_2 \forall \rho_3 \ldots \overline{Q_i} \rho_i$ such that $\mathbf{Pr}\left[V(\rho_1, \ldots, \rho_i) \text{ accepts}\right] \leq s$,

where Q_i denotes $\exists$ if i is odd and $\forall$ otherwise, $\overline{Q_i}$ denotes the complement of Q_i, and, for all i, ρ_i is a $p(n)$-qubit state for a fixed polynomially bounded function p. When the completeness and soundness parameters c, s are not specified, define

$$\mathsf{Q\Sigma_i} := \bigcup_{c-s \in \Omega(1)} \mathsf{Q\Sigma_i(c, s)}.$$

Definition 2.11 (QPH [GSS$^+$22])**.** The unentangled quantum polynomial hierarchy is defined as

$$\mathsf{QPH} := \bigcup_{i=0}^{\infty} \mathsf{Q\Sigma_i}.$$

Interestingly, $\mathsf{QMA(2)} \subseteq \mathsf{Q\Sigma_3}$, since the verifier can simply ignore the second proof.

Here, we let $\mathsf{Q\Sigma_i} = \mathsf{Q\Sigma_i(c, s)}$ for $c - s \geq \Omega(1)$, rather than $1/\operatorname{poly}(n)$, because we do not currently have an error reduction result for QPH similar to the one known for QCPH (although, [AGKR] recently made progress in this direction).

3 The Entangled Quantum Polynomial Hierarchy

We formally define the entangled quantum polynomial hierarchy. The definition appears more technical than for QCPH and QPH, but this is mostly just an issue of notation.

Definition 3.1 (*i*-th level of the entangled quantum polynomial hierarchy ($\mathsf{QE\Sigma_i}$)). A promise problem $L = (L_{\mathrm{yes}}, L_{\mathrm{no}})$ is in $\mathsf{QE\Sigma_i}(c, s)$ for polynomial-time computable functions $c, s : \mathbb{N} \to [0, 1]$ if there exists a polynomial-time uniform family of quantum circuits $\{V_n\}_{n \in \mathbb{N}}$ such that for every n-bit input x, V_n takes quantum proofs, measures a fixed output qubit to decide to accept or reject, and satisfies

- Completeness: $x \in L_{\mathrm{yes}} \Rightarrow \exists \rho_1 \forall \rho_2 \exists \rho_3 \ldots Q_i \rho_i$ such that $\mathbf{Pr}\left[V(\rho_{i-1}, \rho_i) \text{ accepts}\right] \geq c$,

- Soundness: $x \in L_{\mathrm{no}} \Rightarrow \forall \rho_1 \exists \rho_2 \forall \rho_3 \ldots \overline{Q_i} \rho_i$ such that $\mathbf{Pr}\left[V(\rho_{i-1}, \rho_i) \text{ accepts}\right] \leq s$,

where each ρ_j is chosen from the set

$$
\mathcal{A}_j := \left\{
\begin{array}{ll}
\left\{\rho \in \mathbf{D}\left(\mathcal{X}_1 \otimes \mathcal{X}_3 \otimes \cdots \otimes \mathcal{X}_j\right) \mid \text{if } j > 1, \ \mathrm{tr}_{\mathcal{X}_j}\left(\rho\right) = \rho_{j-2}\right\} & \text{if } j \text{ is odd} \\
\left\{\rho \in \mathbf{D}\left(\mathcal{X}_2 \otimes \mathcal{X}_4 \otimes \cdots \otimes \mathcal{X}_j\right) \mid \text{if } j > 2, \ \mathrm{tr}_{\mathcal{X}_i}\left(\rho\right) = \rho_{j-2}\right\} & \text{if } j \text{ is even}
\end{array}
\right. .
$$

Here, Q_i denotes $\exists$ if i is odd and $\forall$ otherwise, and $\overline{Q_i}$ denotes the complement of Q_i. For all i, the corresponding Hilbert space $\mathcal{X}_i$ is a space of at most $p(n)$ qubits for a fixed polynomial p. When the completeness/soundness parameters are not specified, define

$$
\mathsf{QE\Sigma_i} := \bigcup_{c-s \in \Omega(1/\operatorname{poly}(n))} \mathsf{QE\Sigma_i}(c, s).
$$

Definition 3.2 (The entangled quantum polynomial hierarchy (QEPH)). The entangled quantum polynomial hierarchy is defined as

$$
\mathsf{QEPH} = \bigcup_{i=0}^{\infty} \mathsf{QE\Sigma_i}.
$$

When introducing a complexity class, perhaps the first question one should ask is whether or not the choice of completeness and soundness parameters actually matter. In [GSS$^+$22, Theorem 2.6], it was shown that QCPH is robust to the choice of error parameters, but no such result is known for QPH. In Section 4, we show that the choice of parameters does not matter for any level of QEPH, i.e., for c, s such that $c - s \geq 1/\operatorname{poly}(n)$, $\mathsf{QE\Sigma_i}(c, s) = \mathsf{QE\Sigma_i}(\frac{2}{3}, \frac{1}{3})$ for all $i \in \mathbb{N}$ (see Theorem 4.4).

Let us also make several remarks on our definition. As for PH, the indices i in the definition of QEPH are constants, independent of a problem's input size, and, as one should expect, $\mathsf{BQP} = \mathsf{QE\Sigma_0}$ and $\mathsf{QMA} = \mathsf{QE\Sigma_1}$. One can also define $\mathsf{QE\Pi_i} := \overline{\Sigma}_i^{\mathsf{qe}}$ and $\mathsf{QE\Delta_i} := \mathsf{QE\Sigma_i} \cap \mathsf{QE\Pi_i}$. The players also have no incentive to entangle their moves with their opponent because $\mathsf{QE\Sigma_i}$ can be modeled as a zero-sum game. Therefore, we may assume the even and odd indexed states are unentangled.

Informally, $\mathsf{QE\Sigma_i}$ can be thought of as the following game, where we assume i is even to simplify the exposition. Alice has (possibly entangled) quantum registers $(\mathsf{A}_1, \ldots, \mathsf{A}_{i/2})$, and Bob has (possibly entangled) quantum registers $(\mathsf{B}_1, \ldots, \mathsf{B}_{i/2})$, where each register is a number of qubits that is polynomial in the input size. The game commences as follows. In the first round, Alice reveals the state ρ_1 of A_1, and then Bob reveals the state σ_1 of B_1. In the second round, Alice reveals the state ρ_2 of $(\mathsf{A}_1, \mathsf{A}_2)$, and Bob reveals the state σ_2 of $(\mathsf{B}_1, \mathsf{B}_2)$. To ensure Alice and Bob do not change their "moves" from previous rounds, we demand that $\mathrm{tr}_{\mathsf{A}_2}(\rho_2) = \rho_1$ and $\mathrm{tr}_{\mathsf{B}_2}(\sigma_2) = \sigma_1$.

That is, Alice and Bob cannot modify the state of subsystems that have been revealed in previous rounds. In general, for the i-th round, it must be that $\text{tr}_{\mathsf{A}_i}(\rho_i) = \rho_{i-1}$ and $\text{tr}_{\mathsf{A}_i}(\sigma_i) = \sigma_{i-1}$. The game continues like this until the global states of $(\mathsf{A}_1, \ldots, \mathsf{A}_{i/2})$ and $(\mathsf{B}_1, \ldots, \mathsf{B}_{i/2})$ are known to both players and the referee.

At this point, the referee must accept or reject. The referee's action is determined by a polynomial-time quantum circuit and a single-qubit measurement. This action can be equivalently expressed as a two-outcome quantum measurement $\{R, I - R\}$, where the first observable corresponds to accepting. Then, the probability the referee accepts is equal to $\text{tr}\left(R\left(\rho_{i/2} \otimes \sigma_{i/2}\right)\right)$. We emphasize that we do not intend to actually write the observable R corresponding to some verification circuit V. Rather, the observable R is a convenient way to express the action of the referee.

Alice's goal is to maximize the acceptance probability, and Bob's goal is to minimize the acceptance probability. Therefore, given an instance of an $\mathsf{QE\Sigma_i}$ problem with corresponding observable R, we can express the acceptance probability achieved by both players playing optimal strategies as

$$v = \max_{\rho_1 \in \mathcal{A}_1} \min_{\sigma_1 \in \mathcal{B}_1} \ldots \max_{\rho_{i/2} \in \mathcal{A}_{i/2}} \min_{\sigma_1 \in \mathcal{B}_{i/2}} \text{tr}\left(R\left(\rho_{i/2} \otimes \sigma_{i/2}\right)\right), \tag{1}$$

where $\mathcal{A}_i$ and $\mathcal{B}_i$ are defined as in Definition 3.1, and each alternating max/min operator corresponds to an alternation of quantifiers in Definition 3.1. In this work, we intend to use Eq. (1) as a tool for proving the equality of one game/problem instance to another.

Finally, as an application of Eq. (1), we observe that the second levels of both QEPH and QPH are equal to $\mathsf{QRG}(1)$, which is known to be contained in PSPACE.

Proposition 3.3 (Extension of [GSS$^+$22, Corollary 1.9]).

$$\mathsf{QE\Sigma_2} = \mathsf{QE\Pi_2} = \mathsf{Q\Sigma_2} = \mathsf{Q\Pi_2} = \mathsf{QRG}(1) \subseteq \mathsf{PSPACE}.$$

Proof. In [GSS$^+$22], it was observed that $\mathsf{Q\Sigma_2} = \mathsf{QRG}(1)$. Here, we use the same reasoning to conclude

$$\mathsf{QE\Sigma_2} = \mathsf{QE\Pi_2} = \mathsf{Q\Sigma_2} = \mathsf{Q\Pi_2} = \mathsf{QRG}(1).$$

The equivalence is clear given that the value of a $\mathsf{QRG}(1)$ protocol is described by an expression identical to Eq. (1) when $i = 2$, which corresponds to $\mathsf{QE\Sigma_2}$ (see [JW09] for a formal definition of $\mathsf{QRG}(1)$). Then, note that $\mathsf{QRG}(1)$ is closed under complement, by a min-max theorem, implying $\mathsf{QE\Sigma_2} = \mathsf{QE\Pi_2}$. Second, because entanglement is not a concern until one of the players makes multiple moves, the second levels of the entangled and unentangled hierarchies are equal (similarly, the first levels are equal to each other, as are the zeroeth levels). Finally, the containment of $\mathsf{QRG}(1)$ in PSPACE is due to [JW09, Proposition 4]. $\qquad\square$

The fact that $\mathsf{QE\Sigma_2} = \mathsf{QE\Pi_2} = \mathsf{Q\Sigma_2} = \mathsf{Q\Pi_2}$ is somewhat striking, since such an equality in the classical setting would imply a collapse of PH [AB09, Theorem 5.6].[6]

Remark 3.4 (Public vs. private quantum proofs). While the quantum polynomial hierarchies are well-defined, some may object that the classes are unphysical because the provers have full knowledge of each other's density matrices, even though the verifier only receives a single copy of each proof. The quantum no-cloning theorem also begs the question of how exactly the information is communicated between the provers. This is not an issue for PH because it's trivial to learn classical proofs given a single copy, and, for QRG, this is not an issue because communication is

---

[6]This phenomenon of the second levels being equal is also true for TFPH, the hierarchy generalizing the class TFNP [KKMP21].

private. In every quantum complexity class that we are aware of, the provers are "all-powerful" yet still bound by the laws of quantum mechanics.

Despite being unphysical, we are content with the definition for two reasons. First, it is a well-defined, useful theoretical tool for studying quantum information. Second, in the case of QEPH, we show that it collapses to QEΣ$_2$, where it is known, by a min-max theorem, that public vs. private communication is irrelevant. So, despite starting with an unphysical definition, we show equivalence with a class that adheres entirely to the laws of quantum mechanics.

4   The Entangled Quantum Polynomial Hierarchy Collapses

We prove several results about the entangled quantum polynomial hierarchy. Specifically, we prove that QEPH collapses to its second level, is equal to QRG(1), and that every level of QEPH is robust to the choice of completeness and soundness parameters (i.e., for c, s such that $c - s \geq 1/\operatorname{poly}(n)$, QEΣ$_i(c, s) = $ QEΣ$_i(\frac{2}{3}, \frac{1}{3})$ for all $i \in \mathbb{N}$). We begin by proving that the hierarchy collapses.

Lemma 4.1. *For all constants $i \geq 2$, QEΣ$_2 = $ QEΣ$_i$.*

Proof. Note that for all i, QEΣ$_{i-1}$ is trivially contained in QEΣ$_i$. We will show that for all $i > 2$, QEΣ$_i \subseteq $ QEΣ$_{i-1}$ by an induction argument, beginning with QEΣ$_3 \subseteq $ QEΣ$_2$.

Recall from Eq. (1) in Section 3 that the value of an QEΣ$_3$ protocol is equal to

$$\hat{v} = \max_{\rho_1 \in \mathbf{D}(\mathcal{X}_1)} \quad \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \max_{\rho_2 \in \mathcal{A}} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right),$$

where R is the observable corresponding to the verifier accepting, $\mathcal{X}_1$, $\mathcal{Y}_1$, and $\mathcal{X}_2$ are the Hilbert spaces containing the three proofs, and $\mathcal{A} = \{\rho \in \mathbf{D}(\mathcal{X}_1 \otimes \mathcal{X}_2) \mid \operatorname{tr}_{\mathsf{A}_2}(\rho) = \rho_1\}$, which enforces that Alice's second proof is consistent with her first.

For any choice of $\rho_1 \in \mathbf{D}(\mathcal{X}_1)$, define

$$v(\rho_1) = \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \max_{\rho_2 \in \mathcal{A}} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right),$$

so that $\hat{v} = \max_{\rho_1 \in \mathbf{D}(\mathcal{X}_1)} v(\rho_1)$. Consider that $\mathbf{D}(\mathcal{Y}_1)$ and $\mathcal{A}$ are compact and convex by Facts 2.3 and 2.4. Additionally, the function $\operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right)$ is a composition of bilinear functions and so itself is bilinear in σ_1 and ρ_1. Therefore, by Theorem 2.2, a min-max theorem applies and

$$v(\rho_1) = \max_{\rho_2 \in \mathcal{A}} \quad \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right) = \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \max_{\rho_2 \in \mathcal{A}} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right),$$

changing the optimization problem without changing the value.

Substituting this back into $\hat{v}$, we find

$$
\begin{aligned}
\hat{v} &= \max_{\rho_1 \in \mathbf{D}(\mathcal{X}_1)} \quad v(\rho_1) \\
&= \max_{\rho_1 \in \mathbf{D}(\mathcal{X}_1)} \quad \max_{\rho_2 \in \mathcal{A}} \quad \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right) \\
&= \max_{\rho_2 \in \mathbf{D}(\mathcal{X}_1 \otimes \mathcal{X}_2)} \quad \min_{\sigma_1 \in \mathbf{D}(\mathcal{Y}_1)} \quad \operatorname{tr}\left(R\left(\rho_2 \otimes \sigma_1\right)\right),
\end{aligned}
\tag{2}
$$

where the final equality is clear given the definition of $\mathcal{A}$.

We observe that Eq. (2) matches the characterization of an QEΣ$_2$ protocol given in Eq. (1). Therefore, we have shown the value $\hat{v}$ of an arbitrary QEΣ$_3$ protocol is equivalent to the value of an QEΣ$_2$ protocol. Given an instance of an QEΣ$_3$ problem verified by some polynomial-time uniform

circuit V—corresponding to the observable R above—whether V is satisfiable by an $\mathsf{QE\Sigma_3}$ protocol is equivalent to whether V is satisfiable by an $\mathsf{QE\Sigma_2}$ protocol, i.e. $\mathsf{QE\Sigma_3} \subseteq \mathsf{QE\Sigma_2}$ and indeed they are equal.

By way of induction, assume $\mathsf{QE\Sigma_2} = \mathsf{QE\Sigma_i}$ for some constant $i > 2$. By the same min-max argument as just before, we may show the equivalence of the value of any $\mathsf{QE\Sigma_{i+1}}$ protocol to the value of a $\mathsf{QE\Sigma_i}$ protocol, thus showing the equivalence of the classes. Therefore, the hierarchy QEPH collapses to $\mathsf{QE\Sigma_2}$. $\qquad\square$

The equality between QEPH and $\mathsf{QRG(1)}$ is a straightforward consequence of the collapse lemma.

Theorem 4.2. $\mathsf{QRG(1)} = \mathsf{QEPH} = \mathsf{QE\Sigma_2}$.

Proof. Combining the results $\mathsf{QRG(1)} = \mathsf{QE\Sigma_2}$ from Proposition 3.3 and $\mathsf{QE\Sigma_2} = \mathsf{QEPH}$ from Lemma 4.1 proves the equality. $\qquad\square$

Next, we note that our collapse theorem can be strengthened to $\mathsf{QE\Sigma_i} = \mathsf{QE\Sigma_2}$ for any polynomially bounded i, rather than just constant. Like classical PH, we define QEPH as the union of $\mathsf{QE\Sigma_i}$ for any constant i. This is a natural way of defining PH as it is key to proving that if $\mathsf{P} = \mathsf{NP}$, then PH collapses. However, in contrast to collapse techniques for classical PH, our reduction of $\mathsf{QE\Sigma_i}$ to $\mathsf{QE\Sigma_2}$ does not increase the problem size. In our proof of Lemma 4.1, the $\mathsf{QE\Sigma_2}$ problem in Eq. (2) optimizes over the same quantity as the original $\mathsf{QE\Sigma_i}$ problem. Therefore, our proof applies even to a super-constant number of rounds. The reduction is valid up to a polynomial number of rounds, after which the concatenation of the proof registers would lead to a proof too large for the polynomial-time verifier to accept.

Corollary 4.3. $\mathsf{QE\Sigma_i} = \mathsf{QE\Sigma_2}$ *for any polynomially-bounded i.*

Finally, our results also prove that $\mathsf{QE\Sigma_i}$ is robust to the choice of error parameters.

Theorem 4.4. *For any choice of c, s such that $c - s \geq 1/\operatorname{poly}(n)$, it holds that $\mathsf{QE\Sigma_i}(c,s) = \mathsf{QE\Sigma_i}\left(\frac{2}{3}, \frac{1}{3}\right)$.*

Proof. The reverse containment is trivial, so we focus on proving the forward direction, reducing $\mathsf{QE\Sigma_i}(c,s)$ to $\mathsf{QE\Sigma_i}\left(\frac{2}{3}, \frac{1}{3}\right)$. Again appealing to the fact that our proof of Lemma 4.1 shows that an $\mathsf{QE\Sigma_3}$ problem is equivalent to an $\mathsf{QE\Sigma_2}$ problem with the same game value, we observe that our proof implies $\mathsf{QE\Sigma_2}(c,s) = \mathsf{QE\Sigma_i}(c,s)$. Then, because the equality of $\mathsf{QRG(1)}$ and $\mathsf{QE\Sigma_2}$ (Theorem 4.2) is also based on the optimization definition from Eq. (1), the acceptance probability remains preserved and $\mathsf{QE\Sigma_2}(c,s) = \mathsf{QRG(1)}(c,s)$. We may then appeal to the result of [GW05] that a parallel repetition theorem holds for $\mathsf{QRG(1)}$, so that $\mathsf{QRG(1)}(c,s) = \mathsf{QRG(1)}\left(\frac{2}{3}, \frac{1}{3}\right)$. By the same reasoning as a moment ago, this last class equals $\mathsf{QE\Sigma_2}\left(\frac{2}{3}, \frac{1}{3}\right)$. Contracting this sequence of equalities, we conclude that $\mathsf{QE\Sigma_i}\left(\frac{2}{3}, \frac{1}{3}\right)$ equals our original class $\mathsf{QE\Sigma_i}(c,s)$. $\qquad\square$

5 PH and QCPH Are Contained in QPH

We prove that $\mathsf{QCPH} \subseteq \mathsf{QPH}$. While this result is what one might expect, proving this containment was left as an option question by Gharibian et al. [GSS$^+$22]. It is trivial to see that $\mathsf{PH} \subseteq \mathsf{QCPH}$, and, combining these two containments, we have $\mathsf{PH} \subseteq \mathsf{QCPH} \subseteq \mathsf{QPH}$, establishing that quantifying over unentangled quantum proofs is at least as powerful as quantifying over classical proofs.

The central challenge in proving that $\mathsf{QCPH} \subseteq \mathsf{QPH}$ is that the proofs in QPH are allowed to be quantum states, which, upon measurement, give rise to a distribution over classical strings. A

flawed idea is to simply measure the quantum proofs to get classical proofs, and then run the QCPH verification protocol with no modifications. Suppose, however, that Alice has a winning strategy in the QCPH protocol, so she always has a winning response to any classical proof that Bob sends. When simulating this in QPH, Bob can instead send a *quantum state*—a superposition over many classical proofs—preventing Alice from sending an optimal response. In particular, Alice may not know which response to send, since she doesn't know which classical proof the verifier will observe upon measurement.

We prevent this potential cheating by requiring each player to send multiple copies of each of their proofs. We prove that this is enough to force both players to send classical strings with high probability.

Theorem 5.1. PH $\subseteq$ QCPH $\subseteq$ QPH.

The exact error parameters for Theorem 5.1 are stated in Eq. (3) below. In particular, the reduction is only capable of producing a QPH instance with a constant promise gap. However, the containment does hold for any QCPH instance with at least an inverse-polynomial promise gap, due to known error reduction for QCPH [GSS$^+$22].

Proof. Consider any level QCΣ_i of QCPH. We show that for any integer $k \geq 1$,

$$\mathsf{QC\Sigma_i}(c, s) \subseteq \mathsf{Q\Sigma_{2ki}}\left(c\left(1 - 2^{-k}\right), s + 2^{-k}(1 - s)\right). \tag{3}$$

We simulate any QCΣ_i protocol in QΣ_{2ki} as follows. After the first $2k$ turns, the verifier has k proofs from Alice and k proofs from Bob, and the verifier discards all k proofs from Bob. For the next $2k$ turns, the verifier repeats this process, except they keep Bob's proofs rather than Alice's, which we denote by $\sigma_{1,1}, \ldots, \sigma_{1,k}$. This is repeated i times in total until all $2ki$ turns are over. At the end of the game, the verifier has kept the following ki proofs:

$$\rho_{1,1}, \ldots, \rho_{1,k}, \sigma_{1,1}, \ldots, \sigma_{1,k}, \rho_{2,1}, \ldots, \rho_{2,k}, \ldots .$$

For each chunk of k proofs, the verifier measures each quantum state in the standard basis to get k classical strings. If all k classical strings are equal, we say that the player passed the check, and failed otherwise. If a player fails any check, then the other player is declared the winner. If both players pass all checks, then the verifier keeps one copy of each classical proof from each chunk and runs the QCPH verification procedure to determine the winner.

Let $A = (A_{\text{yes}}, A_{\text{no}})$ be a promise problem in QC$\Sigma_i(c, s)$, and let x be some fixed input. If $x \in A_{\text{yes}}$, then Alice has no incentive to cheat and so we refer to her as the honest prover, while if $x \in A_{\text{no}}$ then we consider Bob the honest prover. We will define a strategy for the honest prover and show that no matter the strategy of the dishonest prover, the honest prover will win high probability. In particular, the honest prover's strategy will be to always send classical proofs, and when replying to a dishonest provers proof $\rho = \sum_j p_j |j\rangle\langle j|$, the honest prover will respond as if only the string $\hat{j}$ with the maximum probability $p_{\hat{j}}$ was sent (we arbitrarily choose to break ties by lexicographic order).

If the dishonest prover fails any check, they lose, so we assume now that the dishonest prover passes every check. Then, since both provers pass every check, the verifier has the i classical proofs $y_1, \ldots, y_i$, where the proofs with odd indices are from Alice and the others are from Bob. In one case, suppose that each of the dishonest prover's moves turns out to be as the honest prover expected. Then the situation is identical to the original QCPH instance, and so the honest prover wins with the probability of the original protocol.

In the second case, at least one chunk of k proofs (sampled independently from k distributions) are equal to each other but not to the proof $\hat{\jmath}$ expected by the honest prover. Any string besides $\hat{\jmath}$ has $p_j \leq 1/2$, so the probability of this case occurring, with all k samples matching, is at most 2^{-k}.

Therefore, in the QPH protocol, if $x \in A_{\text{yes}}$, Alice wins with probability at least $c\left(1 - 2^{-k}\right)$. If $x \in A_{\text{no}}$, then Bob wins with probability at least $(1 - s)\left(1 - 2^{-k}\right)$, so Alice wins with probability at most

$$1 - (1 - s)\left(1 - 2^{-k}\right) = s + 2^{-k}(1 - s).$$

To summarize, the dishonest prover is unable to affect the outcome of the game with more than a small probability. We conclude that $\mathsf{QC\Sigma_i} \subseteq \mathsf{Q\Sigma_{2ki}}$, and therefore $\mathsf{QCPH} \subseteq \mathsf{QPH}$. $\qquad\square$

6 Distribution Hierarchies

We introduce another generalization of the polynomial hierarchy where the provers send probability distributions over bit strings. This gives rise to two new hierarchies: the distributional polynomial hierarchy $\mathsf{DistributionPH}$ and its quantum analogue $\mathsf{DistributionQCPH}$, which is the same as $\mathsf{DistributionPH}$ but with a quantum verifier. We will focus primarily on $\mathsf{DistributionPH}$ since the techniques used to analyze $\mathsf{DistributionPH}$ will work for $\mathsf{DistributionQCPH}$ as well.

$\mathsf{DistributionPH}$ is similar to all of the hierarchies studied in this work. In $\mathsf{DistributionPH}$, the distributions are public (the provers have full knowledge of the distributions that have been sent), but none of the distributions are sampled until every distribution has been sent. One can think of this as a non-interactive game, where the players use public, mixed strategies. Importantly, the distributions are not correlated across rounds.

While $\mathsf{DistributionPH}$ is a classical complexity class, our motivation for studying it is to further understand the quantum polynomial hierarchies. In particular, $\mathsf{DistributionPH}$ involves proofs that are classical mixtures of bit strings. This complements $\mathsf{pureQPH}$, where the proofs are quantum superpositions of bit strings, and QPH, where the proofs are both (classical mixtures of quantum superpositions). Does the computational power of the polynomial hierarchy increase when the proofs only involve classical probability distributions? Or does the increased computational power come only from the quantum superposition allowed in QPH and $\mathsf{pureQPH}$? In this section, we resolve these questions.

Theorem 6.1. $\mathsf{DistributionPH} = \mathsf{PH}$.

That is, if the proofs are distributions over classical proofs, PH does not increase in power. The proof of Theorem 6.1 relies on a technical lemma that says the distributions sent in $\mathsf{DistributionPH}$ can be sparse and uniform. This lemma generalizes a result due to Lipton and Young [LY94] and Althöfer [Alt94].

In the remainder of this section, we will formally define $\mathsf{DistributionPH}$, prove the technical lemma, and prove Theorem 6.1. Finally, we will discuss $\mathsf{DistributionQCPH}$ (the same as $\mathsf{DistributionPH}$ but with a quantum verifier) and the power of classical versus quantum proofs.

We begin by formally defining $\mathsf{DistributionPH}$. Let $\mathcal{D}_m$ denote the set of all probability distributions over $\{0, 1\}^m$. For a computation M which takes length-m strings as input and a distribution $\rho \in \mathcal{D}_m$, let $M(\rho)$ implicitly refer to $M(y)$ for $y \sim \rho$.

Definition 6.2 (*i*-th level of the distribution polynomial hierarchy ($\mathsf{Distribution\Sigma_i}$)). A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in $\mathsf{Distribution\Sigma_i}(c, s)$ for polynomial-time computable functions $c, s : \mathbb{N} \to [0, 1]$ if there exists a classical polynomial-time randomized Turing Machine M such that

- Completeness: $x \in L_{\text{yes}} \Rightarrow \exists \rho_1 \forall \rho_2 \exists \rho_3 \ldots Q_i \rho_i$ such that $\mathbf{Pr}\left[M(\rho_1, \ldots, \rho_i) = 1\right] \geq c$,

- Soundness: $x \in L_{\text{no}} \Rightarrow \forall \rho_1 \exists \rho_2 \forall \rho_3 \ldots \overline{Q_i} \rho_i$ such that $\mathbf{Pr}\left[M(\rho_1, \ldots, \rho_i) = 1\right] \leq s$,

where each ρ_k is a distribution in $\mathcal{D}_m$ for some polynomially-bounded m, and each ρ_k is independent. Q_i is $\exists$ if i is odd and $\forall$ otherwise, and $\overline{Q_i}$ is the complement of Q_i. When the completeness/soundness parameters are not specified, define

$$\mathsf{Distribution\Sigma_i} := \bigcup_{c-s \in \Omega(1)} \mathsf{Distribution\Sigma_i}(c, s).$$

Definition 6.3 (The distribution polynomial hierarchy ($\mathsf{DistributionPH}$)). The distribution polynomial hierarchy is defined as

$$\mathsf{DistributionPH} = \bigcup_{i=0}^{\infty} \mathsf{Distribution\Sigma_i}.$$

We make a few comments on our definition of $\mathsf{DistributionPH}$. If we defined $\mathsf{DistributionPH}$ without the bounded-error condition (i.e., no error probability), then it would be equal to PH. We will also generally leave the input x implicit. Finally, if one prefers, they can equivalently think of the provers sending quantum mixed states that are immediately measured in the computational basis (instead of probability distributions that are immediately sampled). This is why we choose to denote the probability distributions as ρ_i in our definition.

As we discussed in Section 3 for QEPH, one can think of $\mathsf{DistributionPH}$ as a game, where two competing provers take turns sending distributions over bit strings to a verifier. Then the verifier M draws one sample from each distribution and runs a polynomial-time randomized algorithm to determine a winner. Additionally, just like with QEPH, we can express the acceptance probability of the verifier as the following optimization problem:

$$\mathbf{Pr}[M \text{ accepts}] = \max_{\rho_1 \in \mathcal{D}_m} \min_{\rho_2 \in \mathcal{D}_m} \ldots \mathrm{Q}^{\mathrm{i}}_{\rho_i \in \mathcal{D}_m} \mathbf{E}[M(\rho_1, \ldots, \rho_i)],$$

where Q^{i} denotes max if i is odd and min otherwise. The expectation is over the randomness in the distributions $\rho_1, \ldots, \rho_i$. Note that since $M(\rho_1, \ldots, \rho_i)$ is a Bernoulli random variable, $\mathbf{E}[M(\rho_1, \ldots, \rho_i)] = \mathbf{Pr}[M(\rho_1, \ldots, \rho_i) = 1]$.

The distributions sent in $\mathsf{DistributionPH}$ are over $\{0, 1\}^m$ for some polynomially-bounded m, so, in general, the support can be exponentially large in m. We will prove a technical lemma that says the provers can send *uniform* distributions over $\mathrm{poly}(m)$ bit strings without changing the outcome of the game too much.

Lemma 6.4. *For any constant $k \in \mathbb{N}$ and any classical randomized Turing Machine M accepting k length-m inputs, if*

$$\max_{\rho_1 \in \mathcal{D}_m} \min_{\rho_2 \in \mathcal{D}_m} \max_{\rho_3 \in \mathcal{D}_m} \ldots \mathrm{Q}^{\mathrm{k}}_{\rho_k \in \mathcal{D}_m} \mathbf{Pr}\left[M(\rho_1, \ldots, \rho_k) = 1\right] = v,$$

then for any constant $\epsilon > 0$,

$$\max_{\rho_1 \in U_{t_k}} \min_{\rho_2 \in U_{t_{k-1}}} \max_{\rho_3 \in U_{t_{k-2}}} \ldots \mathrm{Q}^{\mathrm{k}}_{\rho_k \in U_{t_1}} \mathbf{Pr}\left[M(\rho_1, \ldots, \rho_k) = 1\right] \in [v - k\epsilon, v + k\epsilon],$$

where $t_i := \lceil m^{2i}/2\epsilon^2 \rceil$, U_t denotes the set of uniform distributions over multi-sets of size at most t of strings in $\{0, 1\}^m$, and Q^{k} denotes max if k is odd and min otherwise. The complement of this result also holds (i.e., when the sequence starts with min instead of max).

18

Proof. We will prove the claim by induction. The base case $k = 2$ is precisely [LY94, Theorem 2] (see also [Alt94]). Our contribution is to generalize their result to larger k.

By way of induction, suppose the claim holds for $k - 1$, and consider an instance with k rounds:

$$v := \max_{\rho_1 \in \mathcal{D}_m} \min_{\rho_2 \in \mathcal{D}_m} \max_{\rho_3 \in \mathcal{D}_m} \dots \underset{\rho_k \in \mathcal{D}_m}{Q^k} \mathbf{Pr}\left[M(\rho_1, \dots, \rho_k) = 1\right].$$

Since the complement of this result (where a min is first instead of a max) follows in the same way, we omit the details.

Fix ρ_1 to a distribution that maximizes the acceptance probability (and think of ρ_1 as hardcoded into the input). Consider the inner $k - 1$ distributions $\rho_2, \dots, \rho_k$. By the inductive hypothesis, we can simplify these distributions to

$$\min_{\rho_2 \in U_{t_{k-1}}} \max_{\rho_3 \in U_{t_{k-2}}} \dots \underset{\rho_k \in U_{t_1}}{Q^k} \mathbf{Pr}\left[M(\rho_1, \dots, \rho_k) = 1\right],$$

while only changing the acceptance probability v by $\pm(k-1)\epsilon$. In particular, we have that

$$v' := \max_{\rho_1 \in \mathcal{D}_m} \min_{\rho_2 \in U_{t_{k-1}}} \max_{\rho_3 \in U_{t_{k-2}}} \dots \underset{\rho_k \in U_{t_1}}{Q^k} \mathbf{Pr}\left[M(\rho_1, \dots, \rho_k) = 1\right] \in [v - (k-1)\epsilon, v + (k-1)\epsilon].$$

We want to show that we can simplify the first distribution ρ_1 in a similar fashion. Specifically, we want to show

$$v'' := \max_{\rho_1 \in U_{t_k}} \min_{\rho_2 \in U_{t_{k-1}}} \max_{\rho_3 \in U_{t_{k-2}}} \dots \underset{\rho_k \in U_{t_1}}{Q^k} \mathbf{Pr}\left[M(\rho_1, \dots, \rho_k) = 1\right] \in [v - k\epsilon, v + k\epsilon].$$

Observe that choosing ρ_1 from U_{t_k} instead of $\mathcal{D}_m$ can only hurt the maximizing player. That is, the probability that M accepts can only *decrease*, so $v'' \leq v' + \epsilon \leq v + k\epsilon$ is trivial. All that remains is to show that $v'' \geq v - k\epsilon$. To prove this, it suffices to show that $v'' \geq v' - \epsilon$.

Let $\rho_1^* \in \mathcal{D}_m$ be a distribution that maximizes the acceptance probability of M. Form a multiset S by drawing t_k independent samples from ρ_1^*. Consider a string $y \in S$. This gives rise to a random variable on the interval $[0, 1]$:

$$\mathop{\mathbf{E}}_{\rho_2, \dots, \rho_k}[M(y, \rho_2, \dots, \rho_k)],$$

where we are taking the expectation over optimal choices of $\rho_2, \dots, \rho_k$. In expectation over ρ_1^*, we have

$$\mathop{\mathbf{E}}_{y \sim \rho_1^*}\left[\mathop{\mathbf{E}}_{\rho_2, \dots, \rho_k}[M(y, \rho_2, \dots, \rho_k)]\right] = v'.$$

Therefore, by Hoeffding's inequality (Fact 2.1),

$$\mathbf{Pr}\left[\frac{1}{|S|}\sum_{y \in S} \mathop{\mathbf{E}}_{\rho_2, \dots, \rho_k}[M(y, \rho_2, \dots, \rho_k)] \leq v' - \epsilon\right] \leq \exp\left(-2t_k\epsilon^2\right).$$

To complete the proof, we must count the number of sequences of distributions the minimizing player can send. The minimizing player sends at most $k/2$ of the distributions $\rho_2, \dots, \rho_k$, each of which is a uniform distribution over at most t_{k-1}-sized subsets of $\{0, 1\}^m$. Therefore, in total, there are at most

$$\left(\sum_{i=1}^{t_{k-1}} \binom{2^m}{i}\right)^{k/2} \leq \left(\sum_{i=1}^{t_{k-1}} 2^{im}\right)^{k/2} \leq \left(t_{k-1}2^{mt_{k-1}}\right)^{k/2} = t_{k-1}^{k/2}2^{kmt_{k-1}/2}$$

possible sequences. We want to choose t_k so that

$$\exp\left(-2t_k\epsilon^2\right) < \frac{1}{t_{k-1}^{k/2}2^{kmt_{k-1}/2}}, \tag{4}$$

which would imply that strictly less than 1 of the minimizing player's sequences of distributions can decrease v' by more than ϵ. Or, more directly, it would imply that there are no sequences the minimizing player can send to decrease v' by more than ϵ. We will show that choosing $t_k = m^{2k}/2\epsilon^2$ suffices. By plugging in the definitions of t_k and t_{k-1}, Eq. (4) becomes

$$\exp\left(-m^{2k}\right) < \frac{\epsilon^k}{m^{k(k-1)}}2^{\frac{k}{2}-\frac{km^{2k-1}}{4\epsilon^2}} \iff \exp\left(-m^{2k}\right)\frac{m^{k(k-1)}}{\epsilon^k}2^{\frac{km^{2k-1}}{4\epsilon^2}-\frac{k}{2}} < 1. \tag{5}$$

We show that the inequality in Eq. (5) holds, which proves that our setting of t_k is correct.

$$\begin{aligned}
\exp\left(-m^{2k}\right)\frac{m^{k(k-1)}}{\epsilon^k}2^{\frac{km^{2k-1}}{4\epsilon^2}-\frac{k}{2}} &< \exp\left(-m^{2k}\right)m^{k^2}2^{\frac{km^{2k-1}}{4\epsilon^2}-\frac{k}{2}} \\
&< m^{k^2}2^{\frac{km^{2k-1}}{4\epsilon^2}-\frac{k}{2}-m^{2k}} \\
&= m^{k^2}2^{m^{2k-1}\left(\frac{k}{4\epsilon^2}-\frac{k}{2m^{2k-1}}-m\right)} \\
&< m^{k^2}2^{-m^{2k-1}} \\
&< 1.
\end{aligned}$$

The first inequality holds because $m^k > \epsilon^{-k}$ for constant $\epsilon > 0$. The second-to-last inequality holds because $\left(\frac{k}{4\epsilon^2}-\frac{k}{2m^{2k-1}}-m\right) < -1$ for constant $\epsilon > 0$.

We conclude that $v'' \geq v' - \epsilon \geq v - k\epsilon$, which completes the proof. □

We can now prove that DistributionPH = PH.

Proof of Theorem 6.1. PH $\subseteq$ DistributionPH follows from the proof that PH $\subseteq$ QPH. This only achieves containment in DistributionPH with constant promise gap, and it puts the k-th level of PH in some higher level of DistributionPH (see Theorem 5.1 for more detail).

To show DistributionPH $\subseteq$ PH, we use Lemma 6.4. Set $\epsilon < \frac{1}{12k}$. For DistributionΣ_k, Lemma 6.4 implies that

$$\max_{\rho_1\in U_{t_k}}\min_{\rho_2\in U_{t_{k-1}}}\max_{\rho_3\in U_{t_{k-3}}}\ldots \underset{\rho_k\in U_{t_1}}{Q^k}\mathbf{Pr}\left[M(\rho_1,\ldots,\rho_k)\right] \in [v-k\epsilon, v+k\epsilon] \subseteq \left[v-\frac{1}{12},v+\frac{1}{12}\right].$$

Given the DistributionΣ_k promise gap of $\frac{2}{3},\frac{1}{3}$, this modified game has a promise gap of $\frac{7}{12},\frac{5}{12}$.

We simulate this in PH as follows. To send the distribution ρ_i, the prover sends every string in the support of ρ_i, which is only poly(n) many bits by Lemma 6.4. The verifier can then take the list of strings and sample one uniformly at random. This completes the proof since PH can simulate randomness [Sip83, Lau83]. □

One can also define DistributionQCPH in the same way, and it follows from Theorem 6.1 that this class is equal to QCPH.

Corollary 6.5. DistributionQCPH = QCPH.

The *only* difference between DistributionQCPH and pureQPH is that the former involves proofs that are classical distributions over bit strings and the latter involves proofs that are quantum superpositions over bit strings. DistributionQCPH = QCPH is in the counting hierarchy [GSS$^+$22], while the best known upper bound for pureQPH is EXP^{PP} [AGKR] and it contains QMA(2) and QPH. The conceptual takeaway is that it is only the quantum superposition in the proofs that gives the quantum hierarchies more computational power.

We also remark that if one allows the distributions in DistributionPH and DistributionQCPH to be correlated, then the techniques in Lemma 4.1 can be used to collapse the resulting hierarchies to the second level. The correlated version of DistributionPH collapses to S$_2$P. The correlated version of DistributionQCPH collapses to a quantum-classical version of QRG(1), which, to our knowledge, has never been studied.

7   Open Problems

It is well-known that PH can equivalently be defined via oracle Turing machines. This suggests oracular definitions of quantum polynomial hierarchies, such as QMAH discussed in Section 1.3. One could similarly define QCMAH as QCMA$^{QCMA^{QCMA^{\cdots}}}$ and QMA(2)H as QMA(2)$^{QMA(2)^{QMA(2)^{\cdots}}}$. We ask how these oracular hierarchies compare to the quantifier-based ones.

Question 7.1. *Does* QEPH = QMAH*?* QPH = QMA(2)H*?* QCPH = QCMAH*?*

It is unclear if these hierarchies are equal, as in the classical world, or if one version would be stronger than the other. One immediate obstacle is the fact that QEPH and QPH are quantifying over quantum states, so perhaps it is easier to begin with QCPH, which still quantifies over classical bits. Alas, it is still unclear if an oracle machine definition of QCPH would be equal to a quantifier definition, since, in the oracular case, queries can be made in superposition.

Answering Question 7.1 could yield progress towards characterizing QRG(1). Jain and Watrous showed that QRG(1) ⊆ PSPACE in 2009 [JW09], and, since then, no improved upper bounds have been proven despite effort [GW23]. Our work shows that QRG(1) = QEPH, and, if one can show QEPH ⊆ QMAH, then that would imply QRG(1) ⊆ CH, since [Vin18] showed that QMAH ⊆ CH.

More broadly, proving better upper or lower bounds on the quantum polynomial hierarchies and finding more connections to other parts of complexity theory are important directions for future work. For example, does any level of QPH contain PSPACE? Can one improve the containment QPH ⊆ EXP^{PP}? Or, how can these hierarchies be used to better understand the relationships between QCMA, QMA, and QMA(2)?

Acknowledgements

References

[AB09]   Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, 2009. `doi:10.1017/CBO9780511804090`. [pp. 3, 8, 13]

[AGKR]   Avantika Agarwal, Sevag Gharibian, Venkata Koppula, and Dorian Rudolph. Quantum Polynomial Hierarchies: Karp-Lipton, error reduction, and lower bounds. To appear. [pp. 3, 7, 11, 21]

[AIK22]   Scott Aaronson, DeVon Ingram, and William Kretschmer. The Acrobatics of BQP. In *37th Computational Complexity Conference (CCC 2022)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 20:1–20:17, 2022. `doi:10.4230/LIPIcs.CCC.2022.20`. [p. 6]

[Alt94]   Ingo Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355, 1994. Special Issue Honoring Ingram Olkin. `doi:10.1016/0024-3795(94)90357-3`. [pp. 6, 7, 17, 19]

[BCWdW01]   Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. `doi:10.1103/PhysRevLett.87.167902`. [p. 6]

[BFM23]   Roozbeh Bassirian, Bill Fefferman, and Kunal Marwaha. Quantum Merlin-Arthur and proofs without relative phase, 2023. `arXiv:2306.13247`. [p. 3]

[FGN23]   Chirag Falor, Shu Ge, and Anand Natarajan. A Collapsible Polynomial Hierarchy for Promise Problems, 2023. `arXiv:2311.12228`. [p. 7]

[FK97]   Uriel Feige and Joe Kilian. Making Games Short (Extended Abstract). In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 506–516, New York, NY, USA, 1997. Association for Computing Machinery. `doi:10.1145/258533.258644`. [p. 7]

[FKS95]   J. Feigenbaum, D. Koller, and P. Shor. A Game-Theoretic Classification of Interactive Complexity Classes. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 227–237, 1995. `doi:10.1109/SCT.1995.514861`. [p. 7]

[GK12]   Sevag Gharibian and Julia Kempe. Hardness of approximation for quantum problems. In *International Colloquium on Automata, Languages, and Programming*, pages 387–398. Springer, 2012. `doi:10.1007/978-3-642-31594-7_33`. [p. 6]

[GSS+22]   Sevag Gharibian, Miklos Santha, Jamie Sikora, Aarthi Sundaram, and Justin Yirka. Quantum generalizations of the Polynomial Hierarchy with applications to QMA(2). *Computational Complexity*, 31(2):13, 2022. `doi:10.1007/s00037-022-00231-8`. [pp. 2, 3, 5, 6, 9, 10, 11, 12, 13, 15, 16, 21]

[GW05]   Gus Gutoski and John Watrous. Quantum interactive proofs with competing provers. In *STACS 2005: 22nd Annual Symposium on Theoretical Aspects of Computer Science*, pages 605–616. Springer, 2005. `doi:10.1007/978-3-540-31856-9_50`. [p. 15]

[GW07]   Gus Gutoski and John Watrous. Toward a General Theory of Quantum Games. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, pages 565—-574, New York, NY, USA, 2007. Association for Computing Machinery. `doi:10.1145/1250790.1250873`. [p. 7]

[GW13]   Gus Gutoski and Xiaodi Wu. Parallel Approximation of Min-Max Problems. *Computational Complexity*, 22:385–428, 2013. `doi:10.1007/s00037-013-0065-9`. [pp. 2, 7]

[GW23]   Soumik Ghosh and John Watrous. Complexity limitations on one-turn quantum refereed games. *Theory of Computing Systems*, 67(2):383–412, 2023. `doi:10.1007/s00224-022-10105-9`. [pp. 2, 21]

[HM13]   Aram W. Harrow and Ashley Montanaro. Testing Product States, Quantum Merlin-Arthur Games and Tensor Optimization. *J. ACM*, 60(1), 2013. `doi:10.1145/2432622.2432625`. [p. 6]

[JJUW11] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Journal of the ACM (JACM)*, 58(6):1–27, 2011. `doi:10.1145/2049697.2049704`. [p. 7]

[JW09]   Rahul Jain and John Watrous. Parallel Approximation of Non-interactive Zero-sum Quantum Games. In *24th Annual IEEE Conference on Computational Complexity*, pages 243–253, 2009. `doi:10.1109/CCC.2009.26`. [pp. 2, 13, 21]

[JW23]   Fernando Granha Jeronimo and Pei Wu. The power of unentangled quantum proofs with non-negative amplitudes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 1629—-1642, New York, NY, USA, 2023. Association for Computing Machinery. `doi:10.1145/3564246.3585248`. [pp. 3, 6]

[KKMP21] Robert Kleinberg, Oliver Korten, Daniel Mitropolsky, and Christos Papadimitriou. Total Functions in the Polynomial Hierarchy. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 44:1–44:18, 2021. `doi:10.4230/LIPIcs.ITCS.2021.44`. [p. 13]

[KSV02]   Alexei Y. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Soc., 2002. `doi:10.1090/gsm/047`. [p. 8]

[KW00]   Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pages 608–617, 2000. `doi:10.1145/335305.335387`. [p. 2]

[Lau83]   Clemens Lautemann. BPP and the polynomial hierarchy. *Information Processing Letters*, 17(4):215–217, 1983. `doi:10.1016/0020-0190(83)90044-3`. [pp. 10, 20]

[LY94]   Richard J. Lipton and Neal E. Young. Simple Strategies for Large Zero-Sum Games with Applications to Complexity Theory. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 734–740, 1994. `doi:10.1145/195058.195447`. [pp. 6, 7, 17, 19]

[MS72]   A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *13th Annual Symposium on Switching and Automata Theory (SWAT 1972)*, pages 125–129, 1972. `doi:10.1109/SWAT.1972.29`. [p. 1]

[MW05]   Chris Marriott and John Watrous. Quantum Arthur–Merlin Games. *Computational Complexity*, 14(2):122–152, 2005. `doi:10.1007/s00037-005-0194-x`. [p. 2]

[NC10]   Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge University Press, 2010. `doi: 10.1017/CBO9780511976667`. [p. 8]

[Sha92]   Adi Shamir. IP = PSPACE. *J. ACM*, 39(4):869–877, 1992. `doi:10.1145/146585. 146609`. [p. 7]

[Sio58]   Maurice Sion. On General Minimax Theorems. *Pacific Journal of Mathematics*, 1958. `doi:10.2140/pjm.1958.8.171`. [p. 8]

[Sip83]   Michael Sipser. A Complexity Theoretic Approach to Randomness. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, pages 330–335. Association for Computing Machinery, 1983. `doi:10.1145/800061.808762`. [pp. 10, 20]

[Sto76]   Larry J. Stockmeyer. The Polynomial-Time Hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976. `doi:10.1016/0304-3975(76)90061-X`. [pp. 1, 10]

[Vin18]   Lieuwe Vinkhuijzen. A Quantum Polynomial Hierarchy and a Simple Proof of Vyalyi's Theorem. Master's thesis, Leiden University, 2018. URL: `https://theses.liacs. nl/1505`. [pp. 6, 21]

[Wat18]   John Watrous. *The Theory of Quantum Information.* Cambridge University Press, 2018. `doi:10.1017/9781316848142`. [p. 8]

[Yam02]   Tomoyuki Yamakami. Quantum NP and a Quantum Hierarchy. In *Foundations of Information Technology in the Era of Networking and Mobile Computing*, volume 96 of *IFIP — The International Federation for Information Processing*, pages 323–336, Boston, MA, 2002. Springer. `doi:10.1007/978-0-387-35608-2_27`. [p. 6]