The Complexity of Entangled Games

by

Thomas Vidick

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Umesh V. Vazirani, Chair
Professor Satish Rao
Associate Professor Ashvin Vishwanath

Fall 2011

The Complexity of Entangled Games

Abstract

The Complexity of Entangled Games

by

Thomas Vidick

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Umesh V. Vazirani, Chair

Entanglement is at the heart of quantum mechanics. The nonlocal correlations that can be obtained from space-time separated measurements on an entangled state are a central feature which provably distinguish it from local theories. This dissertation studies entanglement through a computational viewpoint. We develop new insights into the complex nature of entanglement by studying its role in multiplayer games, in which cooperating, but non-communicating, players interact with a referee in an attempt to win a pre-specified game. On the one hand, the nonlocal correlations that entanglement allows may enable players using it to develop new colluding strategies, defeating previously secure protocols. On the other, the richness of this new resource may also be exploited in order to design new protocols, providing solutions to problems previously deemed impossible. We explore both aspects of this dual nature of entanglement, putting limits on its strength while at the same time showing how it can be put to profit to solve new computational problems.

A major unresolved question on the computational complexity of multiplayer entangled games is the power of MIP*, the class of languages having entangled multi-prover interactive proofs: how does it relate to its purely classical analogue MIP, which was completely characterized through the fundamental equation MIP = NEXP? Since the players may use entanglement to increase their odds at colluding against the verifier, MIP* could potentially be a much *weaker* class than MIP. Indeed, for a long time it has been an open question whether two entangled provers are more useful than a single prover.

In this thesis we resolve this question by showing that the class of languages having multiprover interactive proofs with entangled provers is at least as large as its classical counterpart: NEXP $\subseteq$ MIP*. At the heart of this result is an analysis of the multilinearity test of Babai, Fortnow, and Lund in the presence of entanglement. The fact that this test remains sound gives a systematic way for a verifier to impose strong limits on the ability of entangled provers to collude against the verifier.

Gap amplification is a fundamental primitive in the study of classical multiplayer games. While sequential repetition of a game always decreases the prover's maximum success proba-

bility at an exponential rate, the fact that parallel repetition also achieves a gap amplification is a highly non-trivial fact. We show that gap amplification can be performed in parallel even in the presence of entanglement between the provers. We adapt a technique which was originally introduced by Feige and Kilian and results in a polynomial rate of amplification.

The phenomenon of monogamy of entanglement states, in first approximation, that if two parties are maximally entangled then they cannot simultaneously be entangled with a third party. We use this phenomenon in two distinct results. In the first, we show that the bits generated in our randomness-expansion protocol are certifiably random even from the point of view of a quantum adversary who may share prior entanglement with the provers. In addition, we prove the security against quantum adversaries of a randomness-efficient extractor construction originally due to Trevisan. This lets us transform the high-entropy bits that are generated in our protocol into ones that are almost indistinguishable from uniform by any adversary.

More generally, we show how the monogamy of entanglement can be exploited to design multi-prover interactive proof systems that are partially entanglement-resistant. Quantitative bounds on the monogamy of entanglement have generally been elusive, and the analysis of our protocol demonstrates such a bound in a new context.

The nonlocal correlations that can be created by entangled players provide a statistical means of differentiating them from classical, unentangled players. This is the main idea behind Bell inequalities, the violation of which demonstrates the nonlocality of quantum mechanics. We show how this phenomenon may be exploited to design a protocol in which the bits produced by successful players necessarily contain a large quantity of fresh randomness. The presence of randomness is guaranteed irrespective of the provers' actual strategy, as long as the sole constraint of no signaling is respected. Hence a statistical certification for the presence of randomness, a feat easily seen to be impossible to achieve classically.

In order to manipulate the random bits produced in our protocol, and make them useful in cryptography, we give the first proof of security of a poly-logarithmic seed extractor secure against quantum adversaries. To achieve this we adapt the reconstruction paradigm originally introduced by Trevisan to the quantum setting.

We study other ways in which entanglement may be used in interactive proof systems by also allowing a quantum interaction between the referee and the players. We show that, using entanglement, the class of QMIP* proof systems can be parallelized to only three rounds of interaction, and made public-coin, a property that does not hold in the absence of entanglement between the players.

À ma grand-mère,

Qui aurait été si fière.

Contents

Acknowledgments

First and foremost I would like to thank my advisor Umesh Vazirani for his support, encouragements, and endless ability to ensure that I always left his office infused with a level of motivation and enthusiasm incomparable to what they were before entering it. I am most grateful for the incredible environment of creativity and intellectual curiosity that Umesh inspires to those around him. The last few months I spent at Berkeley were especially formative and I am indebted to Umesh for his guidance and support during that period.

Julia Kempe first introduced me to quantum computing, and is the one to blame for starting all this. She has provided me with infaillible support and advice all along, and I am deeply indebted to her generosity, warmth, and scientific talent.

I would like to extend my warmest thanks to all my co-authors, many of which also served as hosts over extended periods of time throughout these past four years. Martin Roetteler led me through a highly enjoyable internship at NEC labs in Princeton during my first summer of grad school. Julia Kempe and Oded Regev hosted me both at Tel-Aviv University in the summer of 2009, and at LRI, then LIAFA, in Paris. Harry Buhrman and Ronald de Wolf welcomed me at CWI in Amsterdam, also in the summer of 2009, and Stephanie Wehner hosted me for a wonderful month of July 2010 at CQT in Singapore.

I thank them, as well as my other co-authors, for their patience and generosity in guiding my first steps in research areas ranging from communication complexity to quantum cryptography and oracle classes. I am also grateful to my more junior collaborators, among which Jop Briet, Anindya De, Joshua Brody and Tsuyoshi Ito, for sharing their ideas and their patience in listening to mine.

Soda Hall has proved a great place to not work, and for this I thank my office-mates, Greg, James and Yaron, as well as the regular Soda Hall occupants without whom these four years wouldn't have been the same: Grant, Madhur, Lorenzo, Anindya, Falk, Yi-Kai, Anupam, Urmila, Siu Man, Siu On, and all the others. Whatever the time of day or night they always made Soda Hall the most welcoming place in Berkeley.

I would like to extend a specially warm thank you to Zeph Landau, who served as an always welcoming and stabilizing presence in the quantum computing group. Though I did my best to hide this from him, Zeph taught me many things, not least of which the un-matched deliciousness of the avocado and tomato baguette sandwich.

My stay in the Bay Area was greatly enlightened by the friends I made here. Most of all I thank Joseph and Pui-Wa for their continual presence, and Joseph for (trying) to teach me to be free.

I would not have reached this point without the constant support of my family. I especially thank my parents for their unwavering belief in my aptitude to carry this through. I

owe them more than they know. If anyone reads this thesis it will be my dad, and I hope that it finally helps him remember its title. I am also grateful to my brother and sister for helping me keep my feet on the ground by turning a blind eye the minute I would start talking about my research.

Chapter 1

Introduction

Entanglement is arguably the most counterintuitive aspect of quantum mechanics — it plays a crucial role in the exponential speed-ups of quantum computers. Einstein, whose skepticism about quantum mechanics rested in part on thought experiments involving entanglement, derided it as "spukhafte Fernwirkung", or "spooky action at a distance". In his landmark 1965 paper [15], Bell proved that entanglement had a testable consequence: simultaneous measurements on a pair of space-like separated entangled particles could lead to outcomes correlated in a way that no classical local hidden variable theory could explain. However, these correlations are still limited by the *no-signaling principle*, which states that no information can be communicated from one particle to the other. Enormous efforts have since been devoted to a systematic investigation of the statistical aspects of nonlocal correlations. Understanding and quantifying the precise nature of these correlations is one of the basic goals of quantum information theory, and has led to a full characterization of the "nonlocal polytope" in two dimensions, the case of three dimensions still being open (see e.g. [47] for a survey).

This dissertation takes a different, more computational approach towards understanding the nature of entanglement. The strength of this approach is in taking entanglement out of the static context commonly used in physics, and putting it into a dynamic setting. The result is a new understanding of the limits (beyond no-signaling) of the power of entanglement, as well as the discovery of new tasks that are made possible using it, but are impossible in a classical world. Moreover, the benefit is reciprocal: the introduction of a powerful new element, entanglement, in the theory of computation promises to bring about a new understanding of some of the basic techniques of complexity theory, such as parallel repetition or multilinearity testing.

The first theme of this dissertation, introduced in Section 1.2, consists in using computational complexity to study entanglement: what are the computational consequences of the nonlocal correlations that it generates? A second theme, introduced in Section 1.3, explores the following question: what are computational tasks that are made possible by the presence of entanglement? Before describing these themes, in the next section we introduce the main computational model that our work is based on, multiplayer games.

1.1 Entanglement as a nonlocal resource

The following experiment, originally introduced by Clauser, Horne, Shimony and Holt [25], gives the simplest demonstration of the strength of entanglement as a nonlocal resource. Consider two distant parties, each holding one half of an entangled pair of particles. Each of the parties receives a single bit $x, y \in \{0,1\}$, chosen uniformly at random, as input. They are allowed to perform arbitrary measurements on their particle, but are not allowed to communicate.[1] Their goal is to produce outputs a, b that satisfy the *CHSH condition* $a \oplus b = x \wedge y$.

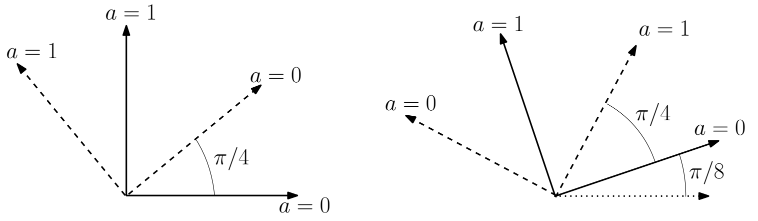

Figure 1.1: The bases used in the CHSH game. Plain lines correspond to the basis used on input $x = 0$ (left) and $y = 0$ (right). Dotted lines correspond to the basis used on input $x = 1$ (left) and $y = 1$ (right). Pairs of vectors corresponding to valid outputs always make an angle of $\pi/8$.

It is not hard to see that, if the particles are only classically correlated (i.e. they play the role of shared randomness) then the best strategy will lead to a success probability of 3/4 — in fact, systematically outputting 0 is the best one can do. However, there are quantum measurements on a 2-qubit entangled state that allow one to obtain a strictly higher success probability of $\cos^2 \pi/8 \approx 0.85$. The entangled state is a *Bell pair*

$$|\Psi\rangle \;=\; \frac{1}{\sqrt{2}}\big(\,|0\rangle|0\rangle + |1\rangle|1\rangle\,\big).$$

This notation describes the joint state of a pair of two-dimensional systems that are in an equal superposition of two identical states: the $|0\rangle|0\rangle$ state and the $|1\rangle|1\rangle$ state. Each party will measure its own half of $|\Psi\rangle$ using one of two possible choices of basis, depending on the input bit. These bases are such that, out of the 4 pairs of bases, those corresponding to input pairs (x, y) such that $x \wedge y = 0$ make an angle $\pi/8$ with each other, while the pair corresponding to the inputs $(1, 1)$ make an angle $\pi/2 - \pi/8$ (see Figure 1.1 for an illustration).

[1] In the original version of the experiment the no-communication assumption was enforced through space-time separation of the two parties.

The resulting measurements have the property that, for *every* possible pair of inputs, the pair of outputs obtained by making the corresponding measurements on both halves of $|\Psi\rangle$ will be correct with probability $\cos^2 \pi/8$. While using these measurements the marginal distribution of outcomes obtained by either party will be uniform, their *joint* distribution depends on the local choice of basis. The possibility to obtain such correlations is the source of the strength of entanglement.

Multiplayer games. Multiplayer games generalize the setting of the CHSH experiment by framing it as an *interactive game* between a referee, who runs the game, and two or more players. In such a game the players co-operate in an attempt to win the game arbitrated by the referee. While they have unlimited computational power, they are crucially not allowed to communicate with each other once their interaction with the referee has started. This no-communication, or *no-signaling*, assumption is at the heart of the richness of multiplayer games.

Given a game, the key quantity associated to it is its *value*: the maximum winning probability that any players can obtain in the game.[2] This quantity lets us frame games as a model of computation: the input is a description of the game itself, and the output is its value. One may think of the players as computing this output for the referee: if it is close to 1 then the players have a high chance of winning the game, and if it close to 0 it is unlikely that they will succeed.

The introduction of entanglement in multiplayer games is all but natural — in a quantum mechanical universe, the sharing of entanglement between resource-unbounded players *cannot be physically avoided*. We will call such games *entangled games*, and their value the *entangled value*, thereby referring to the additional resource that the players may share. The language of games lets us study entanglement in a new context, going well beyond the simple non-interactive setting of the EPR paradox and Bell inequalities. For instance we may ask, does entanglement strengthen or weaken the types of computations that can be performed using multiplayer games? We have seen that entanglement may *increase* the value of a game: what are the computational consequences of this fact? If the increase in value was systematically bounded, these consequences would be minimal. The following example shows that this is not the case.

The Magic Square game. As a further example demonstrating the strength of the non-local correlations of entanglement, consider the Mermin-Peres Magic Square game [83, 91]. In this game there are two players, the row player and the column player. The row player is trying to convince the referee that the cells of an imaginary 3×3 square can be labeled with bits in $\{0, 1\}$ so that the bits in each row have *even* parity, while the column player is trying to convince him that the square can be labeled in such a way that the bits in each column have *odd* parity. In order to catch them, the referee asks the row player (resp. the

[2] This probability is taken over all random choices made in the game: the referee's and the players'.

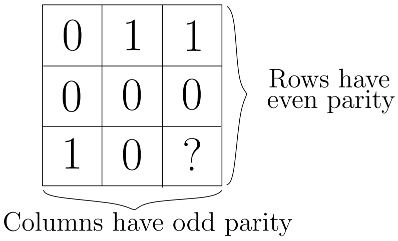

0	1	1
0	0	0
1	0	?

Figure 1.2: The Magic Square game. Cells in any given row should have even parity, while in any given column their parity should be odd. Any labeling of the remaining cell will violate either a row or a column constraint.

column player) for the values that he would assign to the three cells of a randomly chosen row (resp. randomly chosen column). He then verifies that the parity of each player's answers is correct, *and* that the two players are consistent in the value that they assign to the unique cell in which the chosen row and column intersect. A moment's thought will convince the reader that the players *cannot* win this game with certainty — indeed, the square's overall parity must be either odd or even, so that one of the players has to be wrong (cf. Figure 1.2 for an illustration). In stark contrast, Aravind [8] demonstrated the existence of a simple entangled strategy succeeding with *certainty*, dashing all hopes of entanglement providing only a bounded advantage over classical players in general.

This striking example can be pushed even further. As we will see in the next section, the fact that entangled players can *collude perfectly* in the Magic Square game has dramatic consequences on the computational complexity of multiplayer games. Indeed, it demonstrates that the *soundness* property of certain proof systems (e.g. some proof systems used in connection with the PCP theorem to show hardness of approximation of constraint satisfaction problems) can completely fail in the presence of entanglement.

1.2 The computational complexity of entangled games

The introduction of multiplayer games as a model of computation in the late 80s had a profound impact on classical complexity theory. It was the natural result of a revolutionary line of work expanding on the definition of the class NP by adding a layer of *randomization* and *interaction*. In that context multiplayer games are often referred to as *interactive proof systems*: a polynomial-time referee (or verifier) interacts with the players (or provers) in order to verify the validity of a certain statement. One may think of the provers as holding a detailed *proof* of the statement. That proof may be very long and complex, and the verifier can only ask specific questions about it. The provers are computationally unbounded, but not allowed to communicate. They will always attempt to convince the verifier to accept, irrespective of the truth of the statement he is attempting to verify. It is therefore crucial that such proof systems have a good *soundness* property: if the statement is false, then the interaction should be such that no answers from the provers could convince the verifier otherwise (except with small probability).

The corresponding class MIP of languages having multiprover interactive proofs [17] was fully characterized in the celebrated result MIP = NEXP [12], where NEXP stands for non-deterministic exponential time. This characterization demonstrates the impressive computational power of multiple provers. It should be contrasted with the power of a single prover, as expressed in the result that IP = PSPACE [80, 106], the set of languages recognizable in polynomial space. At the same time, the discovery that MIP = NEXP started a long, extremely fruitful line of work exploring the properties of multiplayer games, eventually leading to a proof of the PCP theorem [9, 10] and to the subsequent exploration of its deep connections with hardness of approximation [45].

The introduction of entanglement leads to the natural extension of multiprover *entangled* interactive proofs, in which the provers may share an arbitrary entangled state. In spite of the no-signaling principle, which shows that the players cannot use entanglement to exchange information, its introduction can have a profound effect on the properties of certain proof systems. As an example, consider the following simple protocol used to verify that a given 3XOR formula has a large fraction of its clauses satisfiable. A 3XOR formula is given by a list of clauses $x_i \oplus y_i \oplus z_i = a_i$, where x_i, y_i, z_i are variables and $a_i \in \{0, 1\}$. The verifier picks two clauses (i, j) at random, under the constraint that they share at least one variable, say $x_i = x_j$. He sends the three variables in the first clause to a first prover, and the variables from the second clause to a second prover. Each prover should answer him with an assignment to the three variables it was sent. The verifier checks that the assignments he receives satisfy the clauses, and are *consistent*: both provers should assign the same value to the shared variable x_i.

It is possible to relate the value of this proof system to the maximum number of clauses that can be simultaneously satisfied in the formula: if there is an assignment satisfying a large fraction of clauses then the provers have a successful strategy, and conversely any successful strategy implies the existence of a good assignment. The key point is that the consistency check made by the verifier prevents the provers from using a cheating strategy that would assign different values to the same variable, depending on the clause they are being asked: such a strategy will fail because the provers do not know which variable they share in common.

The Magic Square game shows that such a relationship *no longer holds* in the presence of entanglement between the provers: there are examples of 3XOR formulas that are far from satisfiable, but such that entangled players have a perfect winning strategy. This example shows that the *soundness* property of certain interactive proof systems may be broken by the introduction of entanglement. It raises a fundamental question:

What is the computational complexity of entangled games?

This question can be stated more precisely by introducing the “entangled” analogue of the complexity class MIP, MIP* [26]. The question then becomes: what is the relationship between MIP and MIP*? To show the inclusion MIP $\subseteq$ MIP*, one has to show that the soundness property of an interactive proof system is preserved: if there is no *classical* strategy achieving a high success probability, then there is no *entangled* strategy achieving a much higher success probability. But the example discussed above shows that such a relationship does not hold! Hence the impossibility of a *direct* reduction between the two classes.

Cleve, Høyer, Toner and Watrous [26] pushed this observation further by showing that entanglement could indeed lead to the collapse of a whole complexity class. More specifically, they study a class, $\oplus$MIP, of languages having a certain restricted type of two-prover interactive proofs. While it follows from work of Håstad [52] that this class equals NEXP (and is thus as powerful as the whole of MIP), Cleve & al. show that the corresponding class with entanglement, $\oplus$MIP*, *collapses* to EXP. This result shows that, in the setting of $\oplus$MIP proof systems, whatever the verifier's attempts to prevent entangled provers from colluding, they will have a strategy that fools him.

In spite of this negative result, the question of the complexity of general entangled interactive proofs remained open. Indeed, it could be that by allowing the verifier to interact with the provers in a less constrained way, one may devise more complex proof systems that are immune to the kind of behavior that caused the collapse of $\oplus$MIP*. Despite intense efforts on this question, for a long time little was known. The best lower bound on MIP* consisted in the trivial observation that multiple entangled provers are at least as powerful as a single prover, hence PSPACE $\subseteq$ MIP*.

1.2.1 Contributions on the complexity of entangled games

We prove three results putting strong limits on the ability of entangled players to use their entanglement in order to collude against the referee in a multiplayer game.

NEXP $\subseteq$ MIP*. We prove the inclusion NEXP = MIP $\subseteq$ MIP*, answering a longstanding open question [72] and establishing the fact that multi-prover interactive proof systems with entanglement are at least as expressive as their classical counterparts.

We prove our result by adapting Babai, Fortnow and Lund's [12] original proof that NEXP $\subseteq$ MIP to the entangled setting. A key component in this proof is a *multilinearity test*, by which one ensures that the provers are answering the verifier's questions according to an arbitrary multilinear function $f : \mathbb{F}^n \to \mathbb{F}$, where $\mathbb{F}$ is a finite field. The test is very simple: the verifier picks a triple of axis-aligned points $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathbb{F}^n$, and he checks that the provers provide him with answers $a, b, c \in \mathbb{F}$ that are correspondingly aligned. Babai, Fortnow and Lund showed that if three *deterministic* provers had a high probability of success in this test, then it must be the case that each prover computes his answer using a *function* $a = f(\boldsymbol{x})$ that is linear in each of its n coordinates.

The main difficulty in extending this test to the quantum setting is that *there is no underlying function*: the strength of an entangled-prover strategy is in the *correlations* that are generates by the provers' respective measurements on their shared entangled state, and there is no meaning to either provers' strategy when taken in isolation. As such, the situation is similar to, but more complex than, a scenario in which the provers would be using shared randomness but one would *not* allow for the usual "convexity" argument stating that this randomness could be "fixed". We explain this difficulty, and the way in which we overcome it, in more detail in Section 1.2.2 below. We also give a more detailed introduction to the problem of linearity testing with entangled provers in Chapter 2.

Our main result is that the multilinearity test is sound even in the presence of entangled players. This demonstrates that even entangled players cannot escape the strong linear structure imposed in this test, making it impossible for them to gain more than a negligible advantage from their shared entanglement. This result is presented in Chapter 5.

Our result leaves open the intriguing possibility that entanglement may lead to a *larger*, more expressive class of proof systems: is $\text{MIP}^* \subseteq \text{MIP}$? Since the presence of entanglement seemingly only increases the power of the *provers*, weakening the soundness guarantees of existing protocols, it may seem like this inclusion should *de facto* hold. But there is an intriguing possibility: the presence of entanglement between the provers may also increase the power of the *verifier* by allowing him to devise new protocols, enabling the verification of more complex classes of languages. This could result in MIP^* being a *larger* class than MIP. In the second part of this thesis (cf. Section 1.3.2) we will show that entanglement can indeed be used to perform certain tasks that are impossible in its absence.

The monogamy of entanglement. Monogamy is a genuinely quantum phenomenon. At an intuitive level, monogamy dictates that the correlations obtained from entanglement can only be shared successfully between two parties, not more. Indeed, if two players are *maximally entangled*, then neither of them can be simultaneously entangled with a third player. Unfortunately this appealing property is very difficult to quantify precisely, and it can be expressed in many distinct ways through the use of different entanglement measures [118, 71].

Multiplayer games provide a concrete way to understand the monogamy of entanglement: in a three-player game, constraining two of the players to be strongly correlated should limit their ability to collude with the third. We show how this idea can be put to profit by transforming any game in a way that, even if in the original game entangled players were able to use their entanglement to collude perfectly, in the modified game this is no longer possible. Our transformation consists in introducing a third player, sending him the same question as to one of the original players, and checking that he provides the same answer.

Using this transformation, we are able to give the first hardness of approximation result for three-player one-round entangled games: it is NP-hard to approximate the value of such a game to within a factor that is inverse polynomial in the size of the game, as measured by the number of possible questions. This result is incomparable to the previous one: while the hardness factor is weaker (inverse polynomial compared to constant), it applies to a more restricted class of games, in which there are only three provers and a single round of interaction. Moreover, the transformation that we introduce here is generic, and can be used to enhance the "entanglement-resistance" capabilities of any two-player game.

Parallel repetition of entangled games. Gap amplification is a fundamental primitive in complexity theory. In the context of interactive proofs, one often arrives at a situation where one has designed a specific game such that one can show that *either* there exists a strategy for the players with success probability 1, *or* no players can succeed with probability larger than, say, 0.99 — without knowing which is the case. It is then required to *amplify* this distinction in order to make it more robust, while altering the properties of the game the least possible.

In the setting of two-player classical games, an important result of Raz [96] shows that gap amplification can be performed *in parallel*, without increasing the number of rounds of interaction: the referee simultaneously sends independent pairs of questions to the players, receives all their answers together, and checks that all pairs of answers are valid for the corresponding questions. This result shows a limitation of classical non-communicating provers: they cannot take advantage of the fact that, in a parallel repeated game, they are allowed to see all their questions before sending back their answers.

In Chapter 7 we show that gap amplification can also be performed with entangled players, albeit only at a polynomial rate, by adapting the "miss/match" technique introduced by Feige and Kilian [42]. While the polynomial rate we obtain is sub-optimal, this technique has the benefit of proving *more* than amplification. Indeed, one can show that any players with a reasonable (at least inverse polynomial) success probability in the repeated game must be using a strategy taking a very specific *sequential* form, which may be useful to analyze repeated games in more detail.

1.2.2 Proof strategy

The three results discussed above all rely on the analysis of certain multiplayer games, or multi-prover interactive protocols, designed to achieve a specific goal. The main hurdle in the analysis of such a game is to prove its soundness: given a game, show that no players can have a significantly higher success probability than could "honest" players, playing according to an "ideal", well-behaved strategy. In the classical setting, a strategy for the players is specified by a function f from the question set to the answer set. One usually reasons by contrapositive, showing that high success in the game imposes strict constraints on f, up to the point where one manages to show that f must be "close" to the ideal strategy. In order to make this analysis possible one has at one's disposal all the modern tools of computer science, including error-correcting codes, Fourier analysis, and many others.

The analysis of entangled strategies, however, poses an immediate challenge: *there is no underlying function.* This difficulty already arises in the presence of shared randomness, used by the players to coordinate in selecting one of many possible functions before producing their answer. The solution, in that case, is well-known: there is always an "optimal" shared random string which can be fixed, reducing the analysis to the deterministic case. Entanglement, however, cannot be "fixed". This is a consequence of the non-commuting nature of the players' measurements, making it impossible in general to extract even a joint *global* distribution describing the players' choice of answers for every possible pair of questions.[3]

One is therefore constrained to work *directly* with the players' strategies, as represented by their measurements and potentially very high-dimensional entangled state. In the setting of classical players using shared randomness, this would correspond to carrying out the whole analysis *in superposition*, without fixing the randomness — a challenging task by itself.

As in the classical case, our goal is, starting from the assumption that the provers have a high success probability in the game, to deduce constraints on the structure of the strategy they could be using, relating it to a honest, "ideal" strategy whose success we can easily bound — for instance because it can be modeled using shared randomness, thus reducing the analysis to the classical case.

In order to carry out this approach we develop techniques geared at analyzing and constraining entangled-player strategies. One such technique is a quantum counterpart to the celebrated linearity test of Blum, Luby, and Rubinfeld [23], and we present it in an accessible way in Chapter 2. An extension of this test to multilinearity testing is at the heart of our proof that NEXP $\subseteq$ MIP*, described in Chapter 5.

Another technique exploits the monogamy of entanglement to obtain *almost-commuting* conditions on the players' measurements. This condition lets us apply a weak *rounding* procedure from entangled strategies to classical strategies. We relate the success of this rounding technique to a conjecture about "almost-commuting" versus "nearly-commuting" measurements, which is discussed in Chapter 4.

Our last technique, used in the proof of our result on parallel repetition, consists in exploiting *consistency* constraints on the players' answers to deduce that their measurements must obey certain *orthogonality* conditions. These conditions are used to derive a *direct product test* for entangled strategies.

1.3 Using entanglement in multiplayer games

The second theme of this thesis explores the new possibilities that are afforded by entanglement in multiplayer games in order to present tasks that can only be accomplished in the *presence* of entanglement. These complement the results in the first part of the thesis, in which entanglement was seen as a negative resource used by the players to collude against the referee. We give two main applications: the first to cryptography and the second to interactive proof systems with quantum messages.

[3] This very fact is what makes the existence of games such as the Magic Square possible: in this game entangled players have a perfect winning strategy even though there is no deterministic assignment of answers to all questions that satisfies the constraints imposed by the referee.

1.3.1 Generating certified randomness

A source of independent random bits is a basic resource in many computational tasks, such as cryptography, game theoretic protocols, algorithms and physical simulations. However, constructing a physical source of randomness is an unexpectedly tricky task. What makes this task particularly challenging is the following: how can one even test whether one has succeeded? In other words, suppose we are given a box that claims to output perfectly random bits; is there a test to verify that claim? On the face of it, this task is impossible: a perfect random number generator must output every n-bit sequence with equal probability $1/2^n$, and there seems to be no basis on which to reject any particular output in favor of any other.

Entanglement provides a surprising way out of this conundrum. Recall the CHSH game outlined earlier. This game has the property that classical players can achieve a success probability at most $p_{\texttt{CHSH}} \leq \frac{3}{4}$, while for any number $\frac{3}{4} < p_{\texttt{CHSH}} \leq \cos^2 \pi/8 \approx 0.85$ there is a quantum strategy achieving exactly that success probability. Hence we may define the *quantum regime* for the CHSH game as this range of probabilities: for any value in that range there is a simple quantum-mechanical strategy, obeying the no-signaling condition, which achieves that success probability.

These well-known facts have a striking consequence, first made explicit in Colbeck's Ph.D. thesis [28] (see also [29] for an expanded version): any players producing correlations that fall in the quantum regime *must be randomized*! Indeed, deterministic players are inherently classical, so that their success probability must fall in the classical regime $p_{\texttt{CHSH}} \leq 3/4$. By checking that the players produce answers that are more strongly correlated than could any deterministically chosen answers, the referee is in effect implementing a *statistical test* for randomness.

This idea was quantitatively analyzed in work by Pironio & al. [92]. They showed, using a protocol based on the CHSH game, that one could achieve a quadratic expansion of randomness: while the protocol requires the referee to use $\sqrt{n}$ uniformly distributed bits in order to select his questions, n bits of randomness are generated.

The work in [92] left open two important questions. First, what is the best expansion factor achievable? Is quadratic optimal? Among the many applications for random bits, some of the most prominent pertain to the area of cryptography. For instance, the most widely-studied key distribution protocol, BB84 [18], requires a large number of uniformly distributed bits in order to make an initial choice of basis. Hence a second question: Are the random bits produced secure for cryptographic uses? That is, could any information about them potentially be leaked to an adversary, who may share prior entanglement with the players?

Our contribution. We give an answer to both these questions. In Chapter 9 we introduce a protocol that only requires the referee to use $O(\log n)$ random bits, and still results in the generation of n bits of certified randomness: an exponential expansion. In our protocol, an experimenter (the referee) has a sequential interaction with a pair of unknown devices (the players). The experimenter repeatedly selects inputs to the devices, and collects outputs. He then verifies that the chosen inputs, together with the outputs obtained, verify a certain constraint (in this case, the CHSH condition). If so, he accepts the outputs, and if not he rejects them. This protocol is such that, if the devices share entanglement, there is a simple strategy that will lead them to produce outputs that are accepted by the referee with near-certainty (and are highly random). Moreover, any pair of devices satisfying the no-signaling condition — whether or not their inner workings can be described by quantum mechanics — will either be rejected by the experimenter with high probability, or produce bits that contain large amounts of entropy. Hence the certified presence of randomness does not depend on any assumption on the physical nature of the devices — it is guaranteed by a simple statistical test, together with the no-signaling condition.

We also show that the bits produced in our protocol appear random even to the point of view of a *quantum* adversary, who may herself be entangled with the two players used in the protocol. This condition is crucial for the use of the random bits in cryptography, as well as for composability of the protocol. Our proof of this additional security guarantee exploits some key features of a specific construction of *quantum-proof extractor*. Specifically, suppose that the conditional min-entropy of the devices' outputs B, conditioned on the adversary's system E, is much smaller than the number of random bits we claim the devices produce: $H_\infty(B|E) \ll n$. The key observation is then that, if we were to apply an extractor to B in an attempt to extract *more* bits than its conditional min-entropy, then certainly the extractor's output would not be secure: Eve would be able to distinguish it from a uniformly random string. By exploiting key features of the security proof of a *specific* extractor, based on a construction paradigm due to Trevisan [119], we are able to use this argument to derive strong conditions on the adversary, eventually leading to a contradiction with the no-signaling condition. We explain our results one extractors next.

Quantum-proof extractors. Extractors are pseudorandom constructions that transform a high-entropy string of bits (the *source*) into one that is close to uniform (but shorter). In order to achieve this, they typically require an additional input, the *seed*, that is uniformly distributed. An *adversary* to the extractor is given access to the seed, and to the *output* of the extractor. The goal of the adversary is to distinguish this output from a uniformly distributed bit string. If the adversary succeeds then the extractor is not accomplishing its task. *Quantum* adversaries may be further (weakly) correlated with the source of the extractor, and use this quantum side information in order to help them distinguish the output of the extractor from uniform.

Showing that an extractor is secure against quantum adversaries is a challenging task, and the first such proof of security is due to Renner [100]. Other constructions were proven secure on a case-by-case basis [117, 74, 76], but for a long time no extractor construction with poly-logarithmic seed was shown secure against quantum adversaries. The first such result came in work by Ta-Shma [112], who provided an analysis of a variant of Trevisan's extractor [119] in the quantum bounded-storage model. That variant had a poly-logarithmic seed, but an output length with a poor dependence on the adversary's memory size.

In Chapter 8 we show that Trevisan's extractor is secure in the presence of a quantum adversary, in the most general model of security. Moreover, we show that the parameters of the extractor are essentially the same as in the classical setting. Our proof technique adapts the so-called "reconstruction paradigm" to the quantum setting. Adapting this technique poses the unique challenge of overcoming a fundamental property of quantum measurements, which is that they perturb the state on which they are performed. The resulting "quantum reconstruction paradigm", which derives its key ingredient from work of Koenig and Terhal [76], plays an important role in the proof of security of our randomness-expansion protocol.

1.3.2 Interactive proofs with quantum messages.

Entanglement plays a dual role in multilayer games. As we have seen, it can be used by the players in order to collude against the referee. But it may also potentially be useful to the *referee*: by exploiting the presence of entanglement between the players the referee may require them to perform more complex tasks, that even "honest" but un-entangled players would not be able to achieve. This intriguing possibility, together with the absence of a quantitative bound on the amount of entanglement that may be required of the players to play even near-optimally in a given game, helps explain why there is currently no upper bound known on the class MIP*. (See [38] for an upper bound on the related class of languages having quantum *commuting-prover* interactive proofs.)

In order to explore this question we study multiplayer games in which the referee is allowed to exchange *quantum* messages with the players. This natural modification may open the way to stronger forms of interaction: for instance, the referee may himself send entangled questions to the players, potentially making it harder for them to collude — or at least making it *necessary* for them to share entanglement in order to succeed in the game.

The corresponding complexity class, QMIP*, was introduced by Kobayashi and Matsumoto [72]. Kobayashi and Matsumoto show that, in the absence of entanglement between the players, QMIP = NEXP, while if a *polynomial* number of qubits of entanglement are allowed then the inclusion QMIP* $\subseteq$ NEXP still holds. Their first result demonstrates that quantum messages are no more useful than classical messages in the context of multiprover interactive proofs *without* entanglement.[4]

Our contribution. We show that the presence of entanglement in multiprover interactive proofs with quantum messages can be used beneficially by the verifier. Indeed, we show

[4] An analogous result was very recently shown for the class of languages having *single-prover* quantum interactive proofs: QIP = IP = PSPACE [64]. While the inclusion IP $\subseteq$ QIP is not hard to see, it is the proof of the reverse containment QIP $\subseteq$ PSPACE = IP that required major work and the development of radically new techniques [63, 62].

that using entanglement QMIP* systems can be parallelized to three rounds of interaction, and made *public-coin*: the verifier's sole message to the provers is the broadcast of a single random bit.

This second property does not hold in the absence of entanglement: in the classical case, public-coin multi-prover interactive proofs are only as powerful as single-prover interactive proofs — since every prover receives the same question from the verifier, every prover knows how other provers will behave and the joint strategy of the provers can therefore simulate any strategy of a single prover. Hence, these systems cannot be as powerful as general classical multi-prover interactive proofs unless NEXP = PSPACE.

In contrast, our result shows that in the quantum case, public-coin QMIP systems *are* as powerful as general QMIP systems. The non-triviality of public-coin QMIP systems may be explained as follows: even if every quantum prover knows how other quantum provers will behave, still each quantum prover can only apply local transformations over a part of some state that may be entangled among the provers, which is not enough to simulate every possible strategy a single quantum prover could follow.

1.4 Bibliographical remarks

This thesis is based on seven different papers. The result NEXP $\subseteq$ MIP* is joint work with T. Ito [60]. The parallel repetition of entangled games was studied in joint work with J. Kempe [67]. The results exploiting the monogamy of entanglement in three-player games appear in joint work with J. Kempe, H. Kobayashi, K. Matsumoto and B. Toner [68]. The randomness-expansion protocol is joint work with U. V. Vazirani [121], while the security of Trevisan's extractor against quantum adversaries was first shown in the bounded storage model in joint work with A. De [33], and extended to the most general setting in work with A. De, C. Portmann and R. Renner [34]. The results on the structure of QMIP* are taken from joint work with J. Kempe, H. Kobayashi and K. Matsumoto [69].

Chapter 2

Working with entangled provers: the example of linearity testing

The results in this thesis bring together techniques coming from two distinct areas, quantum computing and classical multiplayer games. In this chapter we give a gentle introduction to some of the key concepts from both in a simple but fundamental context: the analysis of the celebrated *linearity test* of Blum, Luby and Rubinfeld [23] in the presence of entangled players. While being simple enough to afford an elementary description, this analysis captures many of the key ideas that will be used in later chapters.[1] The results in this chapter are joint work with T. Ito [60].

We first recall the definition of the linearity test, and give a brief proof of its soundness for the case of classical players, in Section 2.1. In Section 2.2 we give a "beginner's introduction" to the quantum formalism used to describe entangled players.[2] In Section 2.3 we state the "entangled-prover linearity test", taking the opportunity to explain some of the challenges that arise in the analysis of entangled games. Finally, in Section 2.3.2 we give a proof of the soundness of that test, emphasizing the important tools and techniques that it makes use of.

2.1 The linearity test

Blum, Luby and Rubinfeld's linearity test is a game played with three provers. The verifier's questions are elements of $\mathbb{F}_2^n$, for some integer n, and he expects answers in $\mathbb{F}_2$. The test is designed to verify that each prover answers the verifier's question according to a *linear* function $f : \mathbb{F}_2^n \to \mathbb{F}_2$, i.e. one that can be written as $f(x) = u \cdot x$ for some $u \in \mathbb{F}_2^n$. The test is as follows:

[1] We will come back to these ideas in more detail in Chapter 4, in which we'll give a more complete introduction to the tools and techniques used in this thesis.

[2] A more complete description of the useful notions from quantum information theory, as well as definitions related to games and interactive proofs, will be given in Chapter 3.

Linearity test. Perform either of the following with probability 1/2 each:

1. *(Consistency.)* Select a random $x \in \mathbb{F}_2^n$, and send x to each of the provers. Accept if and only if all three provide the same answer.
2. *(Linearity.)* Select two random points $x, y \in \mathbb{F}_2^n$, and set $z := x + y$. Send x to the first prover, y to the second, and z to the third. Expect three answers $a, b, c \in \mathbb{F}_2$, and accept if and only if $a + b = c$.

This test has *perfect completeness*: if the provers answer according to the same linear function, then they succeed in the test with certainty. Note also that the marginal distribution on each prover's questions being the same in both parts of the test, there is no way for the provers to determine locally which part they are being tested on by the verifier: they must use the same strategy in both cases. BLR show the following.

Theorem 1 (BLR). *Suppose that three* deterministic *provers succeed in the linearity test with probability* $1 - \varepsilon$*, and let* $f_1, f_2, f_3 : \mathbb{F}_2^n \to \mathbb{F}_2$ *be the functions describing their respective strategies. Then there is an* $u \in \mathbb{F}_2^n$ *such that, for each* $i \in \{1, 2, 3\}$*,* $f_i(x) = u \cdot x$ *for all but a fraction at most* $8\,\varepsilon$ *of* $x \in \mathbb{F}_2^n$*.*

Theorem 1 starts from the assumption that the three provers are deterministic. In the case of classical, randomized provers this always holds without loss of generality by "fixing the randomness": among all shared random strings that the provers may use, there is always one that gives them at least as good a success probability as on the average, and we may as well assume that they are using the corresponding deterministic strategies. In the case of entangled players this step will not be possible.

Analysis of the linearity test.

We refresh our reader's memory by giving a classic Fourier-analytic proof of Theorem 1. The proof for the case of entangled players will follow the same outline, and readers new to linearity testing may find it useful to familiarize themselves with the classical proof first.

By assumption the provers succeed with probability $1 - \varepsilon$ in the linearity test, so they must succeed in the "consistency" and "linearity" parts of the test with probability at least $1 - 2\varepsilon$ each. From the "consistency" part we can infer that there is at most a 2ε fraction of $x \in \mathbb{F}_2^n$ such that $f_2(x) \neq f_1(x)$ or $f_3(x) \neq f_1(x)$; call them "bad" x. In the "linearity" part of the test, the probability that either of the two questions y or z is bad is at most 4ε. Hence we may as well assume that all three provers answer according to the same function $f := f_1$, in which case their success in the linearity part of the test should be at least $1 - 2\varepsilon - 4\varepsilon = 1 - 6\,\varepsilon$.

Instead of working directly with f, it will be convenient to introduce the function $g : \mathbb{F}_2^n \to \{-1, 1\}$ defined as $g(x) = (-1)^{f(x)}$ for every $x \in \mathbb{F}_2^n$. For any $u \in \mathbb{F}_2^n$, define the

Fourier coefficient of g at u as $\widehat{g}(u) = \mathrm{E}_x(-1)^{u\cdot x}g(x)$. Parseval's identity states that

$$\sum_u \big(\widehat{g}(u)\big)^2 = \sum_u \mathrm{E}_{x,y}\left[(-1)^{u\cdot(x+y)}g(x)g(y)\right] = \mathrm{E}_x\left[g(x)^2\right] = 1.$$

It is not hard to see that for any $\varepsilon' > 0$ the provers having success probability at least $1-\varepsilon'$ in the "linearity" part of the test is equivalent to the following equation on g:

$$\frac{1}{2} + \frac{1}{2}\,\mathrm{E}_{x,y}\left[\,g(x)g(y)g(x+y)\,\right] \;\geq\; 1 - 2\,\varepsilon'. \tag{2.1}$$

The key claim in the proof of Theorem 1 is the following.

Claim 2. *Suppose deterministic provers applying the same function f succeed in the linearity test with probability at least $1-\varepsilon'$, and let $g = (-1)^f$. Then*

$$\sum_u \big(\widehat{g}(u)\big)^3 \;\geq\; 1 - 2\,\varepsilon'. \tag{2.2}$$

Proof. Expand

$$\begin{aligned}\sum_u \big(\widehat{g}(u)\big)^3 &= \sum_u \mathrm{E}_{x,y,z}\left[\,(-1)^{u\cdot(x+y+z)}g(x)g(y)g(z)\,\right]\\ &= \mathrm{E}_{x,y}\left[\,g(x)g(y)g(x+y)\,\right],\end{aligned}$$

since $\sum_u (-1)^{u\cdot(x+y+z)} = 0$ whenever $z \neq x+y$ in $\mathbb{F}_2^n$. The claim then follows directly from (2.1). □

As a consequence of the bound proven in Claim 2, one can see that g must have a large Fourier coefficient:

$$1 - 12\,\varepsilon \;\leq\; \sum_u \big(\widehat{g}(u)\big)^3 \;\leq\; \Big(\max_u\,\big|\widehat{g}(u)\big|\Big)\Big(\sum_u \big(\widehat{g}(u)\big)^2\Big) \;=\; \max_u\,\big|\widehat{g}(u)\big|$$

by Parseval's identity. Let u_0 be such that $\big|\widehat{g}(u_0)\big| \geq 1 - 12\varepsilon$. Then by definition

$$\big|\mathrm{E}_x\left[\,(-1)^{u_0\cdot x}g(x)\,\right]\big| \;=\; \big|\widehat{g}(u_0)\big| \;\geq\; 1 - 12\,\varepsilon.$$

Recalling the definition of g, this bound immediately implies that the functions f and $x \mapsto (u_0 \cdot x)$ can differ on at most a fraction 6ε of coordinates, proving Theorem 1.

Recapitulating, the proof of Theorem 1 has three main steps. First we argued that, since the provers had a high success probability in the "consistency" part of the test, we could assume that they were using the same function f to compute their answers. Then we proved Claim 2, which puts a lower bound on the sum of the third powers of the Fourier coefficients of any function that passes the "linearity" part of the test with high probability.

Finally, from that bound we deduced that there must exist a *linear* function ℓ (the function $x \mapsto u_0 \cdot x$) that differs from the prover's function f on at most a small fraction of questions, implying that if we *replaced* the provers by ones answering according to ℓ rather than f then the verifier would see little difference — even if this replacement is done as part of a larger protocol in which the linearity test is only a subroutine (provided that, in the larger protocol, the marginal distribution of the questions to each of the "linear" provers is uniform in $\mathbb{F}_2^n$).

2.2 Entangled strategies

In this section we introduce some notation and concepts from quantum information theory that are needed to describe *what an entangled strategy is.* Even though the linearity test uses three provers, for clarity we focus on the setting of two provers; everything that we say here has a natural extension to the case of more provers.

The reader may already be familiar with *pure* quantum states, which are described by unit vectors $|\Psi\rangle \in \mathbb{C}^d$,[3] for some dimension d which is usually a power of 2 (if the system is represented by qubits, then the number of qubits is $\log_2 d$). The reader may also have prior experience with orthogonal measurements: a measurement is described by the choice of an orthonormal basis $\{|e_i\rangle,\ i = 1, \ldots, d\}$ for the space $\mathbb{C}^d$. Upon measuring in that basis the i-th outcome is observed with probability $|\langle e_i|\Psi\rangle|^2$, while the state of the system is projected to its *post-measurement state* $|e_i\rangle$.

In the remainder of this section we introduce generalizations of these two fundamental concepts that will be necessary to describe entangled-prover strategies.

Density matrices. While pure states provide a convenient way to describe isolated systems, such as the joint state of *all* provers in a multiplayer game, we will sometimes need to work with more general, *non-isolated* systems, such as the first prover's subsystem alone. In full generality, a quantum system is described by a *probabilistic mixture* of pure states $(p_i, |\Psi_i\rangle)$. This mixture is represented by a corresponding *density matrix* $\rho = \sum_i p_i |\Psi_i\rangle\langle\Psi_i|$.[4] Since the p_i are a distribution, and the $|\Psi_i\rangle$ are normalized, ρ is a positive matrix with trace 1.

How does one compute the density matrix representing the first prover's subsystem, given that both provers are jointly in the pure state $|\Psi\rangle$? Suppose the second prover was to measure his system using an arbitrary fixed orthogonal measurement, described by an orthonormal

[3] Recall that Dirac's very convenient "ket" notation indicates how we think of the vector Ψ: a ket $|\Psi\rangle$ is a column vector, while a bra $\langle\Psi|$ is a line vector: $\langle\Psi| = |\Psi\rangle^\dagger$.

[4] The careful reader may have noticed that different distributions can give rise to the same density matrix. But quantum mechanics states that the formalism of density matrices *does* give a full description of a given subsystem's state: different distributions giving rise to the same density matrix are *indistinguishable*.

basis $\{|e_i\rangle,\, i = 1,\dots,d\}$. Given such a basis, one may always express $|\Psi\rangle$ as

$$|\Psi\rangle \,=\, \sum_i \sqrt{p_i}|\Psi_i\rangle|e_i\rangle,$$

where the p_i are a distribution, and the $|\Psi_i\rangle$ normalized (but not necessarily orthogonal). Hence once the second prover measures, he obtains outcome i with probability p_i, and the *whole* system gets projected in the state $|\Psi_i\rangle|e_i\rangle$, implying that the first prover is then in state $|\Psi_i\rangle$. But of course, by the no-signaling principle, the state of the first prover should be *independent* of whether the second prover makes the measurement or not — hence the first prover can be accurately described as being in state $|\Psi_i\rangle$ with probability p_i, i.e. the state of his subsystem is represented by the density matrix

$$\rho_1 \,=\, \sum_i p_i|\Psi_i\rangle\langle\Psi_i|.$$

The skeptical reader should check that this matrix is *independent* of our choice of basis for the measurement on the second prover's subsystem, as it should be.

The operation of finding a description of a subsystem given a description of the whole is called "tracing out", and it will be important for us. It can be performed in the same way as described above even starting from a density matrix representation of the whole system: if both provers are in a joint state σ, then the first prover's reduced state, obtained by tracing out the second prover, is given by

$$\rho_1 \,=\, \mathrm{Tr}_2(\sigma) \,=\, \sum_i (\mathrm{Id}\otimes\langle e_i|)\,\sigma\,(\mathrm{Id}\otimes|e_i\rangle). \tag{2.3}$$

Generalized measurements. Just as we generalized our notion of quantum state from pure states to density matrices, we'll need to generalize measurements to allow "high-rank" measurements, which have fewer possible outcomes than the system's dimension. While we used to think of a measurement as projecting a state on a *basis*, a *projective measurement* corresponds to projecting a state on a *subspace*. Hence a projective measurement is given by a set of *orthogonal projectors* $A_1,\dots,A_k$ such that $A_1+\ldots+A_k = \mathrm{Id}$. The "measurement rule" is that when $\{A_i\}$ is performed on the quantum state ρ, outcome i will be observed with probability $\mathrm{Tr}(A_i\rho)$, the "overlap" of the density ρ on A_i. Using a decomposition $\rho = \sum_j p_j|\Psi_k\rangle\langle\Psi_j|$, this probability also reads $\mathrm{Tr}(A_i\rho) = \sum_j p_j\langle\Psi_j|A_i|\Psi_j\rangle$, which is the average, according to the distribution $\{p_j\}$, of the overlap of the state $|\Psi_j\rangle$ on the subspace on which A_i projects.

Entangled strategies. We are ready to put our newly-learned notions from quantum information theory to practice in describing an entangled-prover strategy in a multiplayer game. Consider a two-prover game, in which the first (resp. second) prover is sent a question

x (resp. y), and has to provide an answer a (resp. b). For simplicity, assume that the provers' answers are bits, as will be the case in the linearity test. Let σ be the density matrix describing the joint state of the two provers. Upon receiving his question x, the first prover measures his subsystem using a two-outcome measurement, described by a pair of orthogonal projectors A_x^0 and A_x^1 such that $A_x^0 + A_x^1 = \mathrm{Id}$; this measurement can also be succinctly described through the corresponding *observable* $A_x = A_x^0 - A_x^1$.[5] As a result, the prover obtains an outcome $a \in \{0,1\}$, and he sends it back to the verifier as his answer.[6] Similarly, upon receiving y the second prover makes the measurement $\{B_y^0, B_y^1\}$ on his share of σ, and sends his outcome back to the verifier.

In order to complete our picture we need to give the rule describing the *joint* probability $p(a,b|x,y)$ that the two provers answer the questions (x,y) with (a,b). The way to compute this is to imagine the two provers as making a single, joint measurement, described by four projectors obtained by taking the tensor product of both prover's measurement operators: $A_x^0 \otimes B_y^0$, $A_x^0 \otimes B_x^1$, etc. This leads to the following definition:

$$p(a,b|x,y) \,:=\, \mathrm{Tr}\big(\big(A_x^a \otimes B_y^b\big)\,\sigma\big) \;=\; \sum_{(r,s),(r',s')} \big(A_x^a\big)_{r,r'}\,\big(B_y^b\big)_{s,s'}\,\sigma_{(r,s),(r',s')}\,.$$

To make sure one understands this equation, one can think of how one would describe classical provers using shared randomness in this formalism. In that case, σ would be a diagonal matrix representing the prior probability $q(r)$ of each shared random string $r \in [d]$: σ's rows can be indexed by pairs $(r,s) \in [d]^2$, and it would have a coefficient $q(r)$ in the diagonal entry corresponding to the (r,r) row; all other coefficients (diagonal and otherwise) would be 0.[7] Let f_r (resp. g_r) be the function that the first (resp. second) prover would use if the shared random string was r. Then for any question x and answer a the prover's measurement matrix A_x^a (resp. B_y^b) would also be diagonal, and contain a 1 in each diagonal entry (r,r) such that $f_r(x) = a$ (resp. $g_r(y) = b$). One can now check that with this setup

$$\mathrm{Tr}\big(\big(A_x^a \otimes B_y^b\big)\,\sigma\big) \;=\; \sum_{r\in[d]} q(r)\,1_{f_r(x)=a}\,1_{g_r(y)=b}\,,$$

as should be the case. Of course, general entangled strategies will differ from the one constructed above in a key aspect: the prover's measurements corresponding to different questions will not in general commute, hence they will not be diagonal in the same basis.

[5] Going from the pair (A_x^0, A_x^1) to A_x is the quantum analogue of going from a $\{0,1\}$-valued function f to the $\{-1,1\}$-valued function $g = (-1)^f$.

[6] The measurement formalism that we have described also encompasses seemingly more complex strategies, in which the prover would make a measurement, then do some classical processing on the outcomes, maybe another measurement, etc. — all these operations are taken into account by the projectors A_x^0 and A_x^1, which describe his final answer.

[7] As a side remark, note that choosing $\sigma = d^{-2}\mathrm{Id}$ would not correspond to shared randomness — indeed, that state has a tensor product form $(d^{-1}\mathrm{Id}_d) \otimes (d^{-1}\mathrm{Id}_d)$, so that it corresponds to a setting where there would be no correlations between the provers at all.

2.3 Linearity testing of entangled provers

The first difficulty in extending the linearity test to the case of entangled provers is to determine what its precise statement should be. Indeed, in the presence of entanglement (as in the presence of shared randomness between the provers), there is no hope of extracting a *single* linear function from the prover's strategy, as was done in Theorem 1. What does it even mean for entangled provers to be *linear*, if their strategy cannot be tied to a single function? The following informal theorem suggests an answer:

Theorem 3 (Entangled-prover linearity test, informal)**.** *Suppose that three entangled provers, using a strategy described by measurements $\{A_x^a\}$, $\{B_y^b\}$, $\{C_z^c\}$ and a shared entangled state σ, succeed in the linearity test with probability $1-\varepsilon$. For $u \in \mathbb{F}_2^n$, let*

$$\widehat{A}_u = \mathrm{E}_x(-1)^{x\cdot u}\big(A_x^0 - A_x^1\big)$$

be the matrix Fourier coefficient associated to the first prover's strategy. For every u, define

$$M^u \;:=\; \big(\widehat{A}_u\big)^2. \tag{2.4}$$

Then $\{M^u\}$ is a proper quantum measurement. Moreover, the original prover's strategy is almost indistinguishable *from that of three "oblivious" provers whom would behave as follows:*

1. *Measure their share of the entangled state using $\{M^u\}$, each obtaining an outcome u_i, for $i \in \{1,2,3\}$,*

2. *Upon receiving the prover's question x, answer with $u_i \cdot x$.*

The key point in Theorem 3 is Eq. (2.4), which, rather than defining a *single* linear function to which the provers would be close, constructs a *global* measurement, *independent* of the prover's questions, and then claims that this measurement faithfully reproduces the original provers' strategy. It does this by considering new, "oblivious" provers, who *first* measure according to the constructed measurement, obtaining the label u of a linear function, and *then* apply that function to their question. Hence the measurement might as well have been made before the start the protocol: the oblivious provers are effectively reduced to being classical, using shared randomness to determine the linear function they will use.[8]

The definition of M^u has a simple interpretation in the special case where the provers are classical, but may use shared randomness. It corresponds to the following definition of an "oblivious" strategy: the prover simply looks at his random string r, pointing to a function f_r according to which the "original" prover would answer his questions. Instead,

[8] This observation shows that success in the linearity test imposes a very strong constraint on the structure of even entangled provers, thus reduced to using their entanglement in only a very limited way. A generalization of this strong constraint will be at the heart of our proof that NEXP $\subseteq$ MIP*, given in Chapter 5.

the oblivious prover samples a *function* $\ell : x \mapsto u \cdot x$, where u is chosen according to the distribution suggested by f_r's Fourier spectrum.[9] When his question x arrives, he answers with $\ell(x) = u \cdot x$.

Why is this a good strategy? Recall that in the soundness proof for the classical, deterministic case we had proved Eq. (2.12), which stated that the Fourier coefficients of $g = (-1)^f$ were sharply concentrated. This justified our "rounding" of the provers' strategy to the linear function corresponding to the largest Fourier coefficient. In the case of a randomized strategy it will still be the case that most functions f_r have a large Fourier coefficient. Even though *which* coefficient might depend on the random string r, the "linear" strategy defined above intrinsically accounts for that possibility, and it is not hard to see that it will indeed faithfully reproduce the original prover's actions.

Returning to Theorem 3, what is maybe more surprising is that essentially the *same* definition of a new strategy will also work in the case where the original provers use an arbitrary entangled strategy! Before showing that this is indeed the case, we should make precise what we mean by two entangled-prover strategies being "almost indistinguishable". In the classical case this was taken to mean "the new provers' strategy differs from the original one in a fraction at most $O(\varepsilon)$ of questions", and we give a definition formalizing a similar intuition in the case of entangled provers in the next section. We then give a more precise statement of Theorem 3, as well as its proof, in Section 2.3.2.

2.3.1 Measuring the distance between provers

In the classical analysis of the linearity test, from three functions (f, g, h) having high success in the test one constructs a linear function ℓ such that f, g, h differ from ℓ in a small fraction of points $x \in \mathbb{F}_2^n$. This ensures that the linearity test can be performed as part of a bigger protocol, possibly involving a larger number of provers: its goal is to constrain a subset of the provers to answer according to a linear function. Provided in the larger protocol the verifier's question to each prover is distributed as in the linearity test (uniformly in $\mathbb{F}_2^n$), replacing them with their linear approximation will not affect the provers' success probability in the overall protocol much, while potentially making the analysis much simpler.

We would like to achieve the same result in the case of entangled provers. In order to make a meaningful statement, we need an appropriate measure of what it means for two distinct strategies of a single prover to be "indistinguishable". Any measure that we use should be *strong enough* that a small distance implies that the type of "prover replacement" described above does not affect the provers' success probability in the overall protocol too much, while still being *weak enough* that one is able to prove bounds on that distance simply from the fact that the provers have a high success in the linearity test.[10]

[9] Letting $g_r = (-1)^{f_r}$, this is the distribution induced by the $|\widehat{g_r}(u)|^2$ (Parseval's identity shows that this is indeed a distribution).

[10] This implies for instance that the operator norm on the provers' measurements would *not* be appropriate,

Consider two distinct measurements, $\{A_x^a\}$ and $\{\tilde{A}_x^a\}$, that the first prover could apply on his share of the entangled state, which we'll take to be a pure state $|\Psi\rangle$ for simplicity. Fix a question $x \in \mathbb{F}_2^n$. Quantum mechanics dictates that, once the first prover has applied the measurement $\{A_x^0, A_x^1\}$ on his share of $|\Psi\rangle$ and obtained an outcome $a \in \{0,1\}$, the entangled state gets projected to $|\Psi'\rangle = \big(A_x^a \otimes \mathrm{Id} \otimes \mathrm{Id}\big)|\Psi\rangle$, where the identity terms are meant to indicate that the other two provers have not performed any action yet.[11] This means that, for this fixed x, the global states resulting from the first prover measuring using either $\{A_x^0, A_x^1\}$ or $\{\tilde{A}_x^0, \tilde{A}_x^1\}$, and conditioning on either answer being obtained, will be close if the following quantity is small:

$$\begin{aligned}\sum_a \big\|(A_x^a \otimes \mathrm{Id} \otimes \mathrm{Id})|\Psi\rangle - (\tilde{A}_x^a \otimes \mathrm{Id} \otimes \mathrm{Id})|\Psi\rangle\big\|^2 &= \sum_a \langle\Psi|\big(\big(A_x^a - \tilde{A}_x^a\big)^2 \otimes \mathrm{Id} \otimes \mathrm{Id}\big)|\Psi\rangle \\ &= \sum_a \mathrm{Tr}\big(\big(A_x^a - \tilde{A}_x^a\big)^2 \rho\big), \qquad (2.5)\end{aligned}$$

where ρ is the reduced density of the state $|\Psi\rangle$ on the first prover's register (cf. Eq. (2.3) for a definition). The magic of this last equation is that the systems corresponding to the second and third provers have disappeared; this was made possible by the fact that they act on subsystems separated from the first prover's. Nevertheless, we have shown that if the quantity in (2.5) small, then the provers' shared state is almost the same *after* the first prover has measured his subsystem using *either* A or $\tilde{A}$, and obtained an answer a. Hence whatever happens in the remainder of the protocol, the probabilities that arise from other provers' measurements will be essentially the same irrespective of which of A or $\tilde{A}$ the first prover applied.

Incorporating the choice of the question x, we are ready to define our distance measure on strategies: we'll say that the strategies described by $\{A_x^a\}$ and $\{\tilde{A}_x^a\}$, together with the entangled state σ with reduced density ρ on the first prover's subsystem, are δ-close if

$$d_\rho(A,B) := \Big(\mathrm{E}_x \, \mathrm{Tr}\big(\big(A_x^a - \tilde{A}_x^a\big)^2 \rho\big)\Big)^{1/2} \leq \delta.$$

We will explore properties of d_ρ in more detail in Chapter 4. In particular, it is not too hard to see that d_ρ is indeed a distance measure (it is non-negative and satisfies the triangle inequality), and that if $\{A_x^a\}$ and $\{\tilde{A}_x^a\}$ are measurements then it is bounded between 0 and $\sqrt{2}$.

We argued that the distance measure d_ρ was strong enough to ensure that switching from one of two strategies close in that distance to the other would have a small effect on

as success in the test does not put constraints directly on the provers' measurements themselves, but only on their probability of obtaining certain outcomes *when applied on the entangled state* σ.

[11] Observe that the squared norm of the state after the prover's measurement is $\|(A_x^a \otimes \mathrm{Id} \otimes \mathrm{Id})|\Psi\rangle\|^2 = \mathrm{Tr}\big(((A_x^a)^2 \otimes \mathrm{Id} \otimes \mathrm{Id})(|\Psi\rangle\langle\Psi|)\big) = \langle\Psi|(A_x^a \otimes \mathrm{Id} \otimes \mathrm{Id}|\Psi\rangle$, since A_x^a is a projector: it is the probability of obtaining outcome a when measuring Ψ with $\{A_x^a\}$. As such, the post-measurement state is not normalized.

the overall performance of the provers in any entangled game. The analysis of the linearity test in the next section will show that it is also weak enough, in the sense described above: one can place bounds on d_ρ from the seemingly weak assumption that provers have a high success probability in a certain well-chosen test that the verifier plays with them.

2.3.2 The quantum analysis

Before stating our theorem, we observe that, given that the linearity test is symmetric under permutation of the three provers, one may assume without loss of generality that the following hold of the provers' strategies:

1. All three provers are applying the same set of measurements $\{A_x^a\}$,
2. The provers' shared state σ is invariant with respect to any permutation of its three subsystems.

This observation follows from a symmetrization argument; it is re-stated in more detail and proved as Lemma 13 in Chapter 4. In practice, it means that the probabilities $p(a, b, c|x, y, z)$ do not depend on which prover applied the measurement corresponding to each question:

$$\mathrm{Tr}\big(\big(A_x^a \otimes A_y^b \otimes A_z^c\big)\,\sigma\big) \;=\; \mathrm{Tr}\big(\big(A_y^b \otimes A_z^c \otimes A_x^a\big)\,\sigma\big) \;=\; \ldots \;=\; \mathrm{Tr}\big(\big(A_z^c \otimes A_y^b \otimes A_x^a\big),$$

a convenient property we will make frequent use of.

The following theorem, a precise reformulation of Theorem 3, is the main result of this chapter.

Theorem 4 (Linearity test with entangled provers)**.** *Suppose three entangled provers succeed in the* linearity test *with probability at least* $1-\varepsilon$ *using a symmetric strategy*[12] *with measurements* $\big\{A_x^a\big\}$ *and entangled state* σ*. Then there exists a measurement* $\big\{B^u\big\}$*, independent of* x *and indexed by outcomes* $u \in \mathbb{F}_2^n$*, such that if we let* $B_x^a := \sum_{u:\, u\cdot x = a} B^u$ *then*

$$\big(d_\rho(A,B)\big)^2 \;=\; \mathrm{E}_x \sum_a \mathrm{Tr}\Big(\Big(A_x^a - \sum_{u:\, u\cdot x=a} B^u\Big)^2 \rho\Big) \;\le\; 8\sqrt{\varepsilon}. \tag{2.6}$$

As discussed in Section 2.3.1, through Eq. (2.6) the theorem asserts that "oblivious" provers, who would first measure according to $\{B^u\}$, obtain an outcome u, and then answer their question x with $u \cdot x$, are almost *indistinguishable* from the original provers, in the

[12] Even though this symmetry implies that all three provers are using the same set of measurements $\{A_x^a\}$, and their entangled state is invariant with respect to any permutation of the provers' subsystems, this assumption alone is not sufficient to guarantee that they will succeed with certainty in the "consistency" part of the linearity test. Indeed, the assumption of symmetry does not for instance preclude randomized strategies in which the random strings would be triples of bits (a, b, c) and the i-th prover would answer with the bit contained in the i-th position.

strong sense that one may *replace* the original provers by the new ones while only affecting their success probability in *any* overall protocol,[13] of which the linearity test might only be a subroutine, by $O\big(\sqrt{\varepsilon}\big)$.

We now turn to the proof of Theorem 4. Following the classical proof given in Section 2.1, we will use Fourier analysis directly on the prover's observables $A_x := A_x^0 - A_x^1$: for every $u \in \mathbb{F}_2^n$ one may define

$$\widehat{A}_u \,:=\, \mathrm{E}_x\,[\,(-1)^{u\cdot x} A_x\,]\,.$$

In general, $\widehat{A}_u$ is Hermitian, with eigenvalues in $[-1,1]$. Indeed, a variant of Parseval's identity also holds in this setting:

$$\sum_u \big(\widehat{A}_u\big)^2 \,=\, \sum_u \mathrm{E}_{x,y}\,\big[\,(-1)^{u\cdot(x+y)} A_x A_y\,\big] \,=\, \mathrm{E}_x\,\big[\,A_x^2\,\big] \,=\, \mathrm{Id}, \tag{2.7}$$

where the last equality uses that the A_x are observables.

As in the classical case (cf. (2.1)), the following two equations re-formulate the fact that the provers must succeed in each of the "consistency" and the "linearity" parts of the tests with probability at least $1-2\varepsilon$.[14] (Recall that we use $p(a,b,c|x,y,z)$ to denote the probability that the provers answer (a,b,c) to questions (x,y,z).)

$$\mathrm{E}_x \sum_{a,b} p(a,a,b|x,x,x) \,=\, \mathrm{E}_x\,\mathrm{Tr}\big(\big(A_x\otimes A_x\otimes \mathrm{Id}\big)\,\sigma\big) \,\geq\, 1-4\varepsilon, \tag{2.8}$$

$$\mathrm{E}_{x,y} \sum_{a,b} p(a,b,a+b|x,y,z) \,=\, \mathrm{E}_{x,y}\,\mathrm{Tr}\big(\big(A_x\otimes A_y\otimes A_{x+y}\big)\,\sigma\big) \,\geq\, 1-4\varepsilon. \tag{2.9}$$

The proof of both equations is exactly similar to the classical case, and we postpone it until Section 2.3.3. Still in complete analogy with the classical setting, one can translate Eqs. (2.8) and (2.9) into conditions on the Fourier coefficients $\widehat{A}_u$ that we associated with the observables A_x:

$$\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u\otimes \widehat{A}_u\otimes \mathrm{Id}\big)\,\sigma\big) \geq 1-4\varepsilon, \tag{2.10}$$

$$\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u\otimes \widehat{A}_u\otimes \widehat{A}_u\big)\,\sigma\big) \geq 1-4\varepsilon. \tag{2.11}$$

The proof of both equations follows from the definition of $\widehat{A}_u$ and Eqs. (2.8) and (2.9). In the classical setting, (2.11) would already be a proof of Claim 2: since the "entangled state"

[13] As we already mentioned, this is only possible provided the marginal distribution on the "linear" provers' questions in the overall protocol is as it is in the linearity test, uniform over $\mathbb{F}_2^n$.

[14] Note that, for the first equation, we write the probability of the provers succeeding as if the verifier only checked that two out of the three answers were consistent; a weaker but sufficient requirement for our purposes.

in that case is one-dimensional, the tensor product becomes a product, and the sum of the third powers of the Fourier coefficients appears.

In the presence of entanglement, however, it seems that we are stuck: each prover acts on his own subsystem only; how could one bring different subsystems together? The following claim shows that this is, in fact, possible as a consequence of the "consistency" part of the linearity test: a measurement $\{A_x^a\}$ performed on the first subsystem can be "replaced" by the same measurement performed on the second subsystem, without affecting the provers' shared state, after that measurement has been performed, by much.[15] Note that this property is distinct from that of the strategy's symmetry: while symmetry dictates that the *distribution* of outcomes should be the same irrespective of which measurement is performed on which subsystem, here we are showing that the whole post-measurement state is almost the same in both cases. Claim 5 provides an important tool to manipulate entangled strategies, and we will subsequently see how it lets us deduce an analogue of Claim 2 from (2.11).

Claim 5. *Suppose the provers succeed in the consistency test with probability* $1-\varepsilon$. *Then*

$$\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u \otimes \mathrm{Id} \otimes \mathrm{Id} - \mathrm{Id} \otimes \widehat{A}_u \otimes \mathrm{Id}\big)^2 \sigma\big) \;\leq\; 8\varepsilon,$$

and the same holds under arbitrary permutation of the registers.

Proof. It suffices to expand the expression on the left-hand side as

$$\begin{aligned}
&\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u \otimes \mathrm{Id} \otimes \mathrm{Id} - \mathrm{Id} \otimes \widehat{A}_u \otimes \mathrm{Id}\big)^2 \sigma\big)\\
&\qquad = \sum_u \Big(\mathrm{Tr}\big(\big(\widehat{A}_u^2 \otimes \mathrm{Id} \otimes \mathrm{Id}\big)\,\sigma\big) + \mathrm{Tr}\big(\big(\mathrm{Id} \otimes \widehat{A}_u^2 \otimes \mathrm{Id}\big)\,\sigma\big) - 2\mathrm{Tr}\big(\big(\widehat{A}_u \otimes \widehat{A}_u \otimes \mathrm{Id}\big)\,\sigma\big)\Big)\\
&\qquad \leq 2 - 2(1-4\varepsilon),
\end{aligned}$$

where we used Parseval's identity (2.7) to compute the first two terms, and (2.10) to lower-bound the last term. □

Writing

$$\begin{aligned}
\widehat{A}_u^3 \otimes \mathrm{Id} \otimes \mathrm{Id} - \widehat{A}_u \otimes \widehat{A}_u \otimes \widehat{A}_u &= \big(\widehat{A}_u^2 \otimes \mathrm{Id} \otimes \mathrm{Id}\big)\big(\widehat{A}_u \otimes \mathrm{Id} \otimes \mathrm{Id} - \mathrm{Id} \otimes \widehat{A}_u \otimes \mathrm{Id}\big)\\
&\quad + \big(\widehat{A}_u \otimes \widehat{A}_u \otimes \mathrm{Id}\big) \cdot \big(\widehat{A}_u \otimes \mathrm{Id} \otimes \mathrm{Id} - \mathrm{Id} \otimes \mathrm{Id} \otimes \widehat{A}_u\big),
\end{aligned}$$

Claim 5 together with Eq. (2.11) and the Cauchy-Schwarz inequality let us obtain the following quantum analogue of Claim 2.

[15] The claim proves this statement for the Fourier operator $\widehat{A}_u$, but a similar bound can be proven directly for the observable A_x itself.

Claim 6. *The following holds*

$$\sum_u \mathrm{Tr}\big(\widehat{A}_u^3\,\rho\big) \geq 1 - 8\sqrt{\varepsilon}. \tag{2.12}$$

While in the classical setting Eq. (2.2), together with Parseval's identity, immediately implied the existence of a single large Fourier coefficient for g, in the presence of entanglement Eq. (2.12) does not imply such a strong statement. Indeed, even in the case of provers using shared randomness, (2.12) only states that *for most random strings* there should be a corresponding large Fourier coefficient — but which coefficient it is may well depend on the random string itself.

Instead, as described at the beginning of Section 2.3 we define a measurement $\{M^u\}$ as

$$M^u \,:=\, \big(\widehat{A}_u\big)^2.$$

Each M^u is non-negative, and Parseval's identity shows that $\sum_u M^u = \mathrm{Id}$: the M^u form a proper quantum measurement. To show that they satisfy the requirement of the theorem, let $C_x = A_x - \sum_u (-1)^{u\cdot x} M^u$, and observe that the Fourier coefficient of C_x at u is

$$\widehat{C}_u \,=\, \widehat{A}_u - \mathrm{E}_x \sum_u (-1)^{u\cdot x} M^u \,=\, \widehat{A}_u - \big(\widehat{A}_u\big)^2,$$

so that by Parseval's identity

$$\mathrm{E}_x\, C_x^2 \,=\, \sum_u \big(\widehat{C}_u\big)^2 = \sum_u \widehat{A}_u^2\big(\mathrm{Id} - \widehat{A}_u\big)^2 \,\leq\, 2\sum_u \widehat{A}_u^2\big(\mathrm{Id} - \widehat{A}_u\big) \,=\, 2\sum_u \big(\mathrm{Id} - \widehat{A}_u^3\big),$$

again as a consequence of Parseval's identity. Hence

$$\begin{aligned}\mathrm{E}_x\, \mathrm{Tr}\Big(\Big(A_x - \sum_u (-1)^{u\cdot x} M^u\Big)^2 \rho\Big) &\leq 2 - 2\sum_u \mathrm{Tr}\big(\widehat{A}_u^3\big) \\ &\leq 16\sqrt{\varepsilon}\end{aligned}$$

by (2.12). This proves the theorem since

$$A_x - \sum_u (-1)^{u\cdot x} M^u = \frac{1}{2}\big(\mathrm{Id} + A_x^0\big) - \frac{1}{2}\Big(\mathrm{Id} + \sum_{u\cdot x=0} M^u\Big) = \frac{1}{2}\Big(A_x^0 - \sum_{u\cdot x=0} M^u\Big),$$

and a similar equation holds after replacing $'0'$ by $'1'$.

2.3.3 Omitted proofs

We give details of the proofs that were omitted from our analysis of the linearity test in the presence of entangled strategies.

Proof of Eqs. (2.8) and (2.9). The player's success probability in the consistency test is

$$\mathrm{E}_{x,y}\sum_a \frac{1}{3}\Big(\mathrm{Tr}\big(\big(A_x^a\otimes A_x^a\otimes \mathrm{Id}\big)\,\sigma\big)+\mathrm{Tr}\big(\big(A_x^a\otimes \mathrm{Id}\otimes A_x^a\big)\,\sigma\big)+\mathrm{Tr}\big(\big(\mathrm{Id}\otimes A_x^a\otimes A_x^a\big)\,\sigma\big)\Big)\geq 1-2\varepsilon \tag{2.13}$$

By symmetry, all three terms inside the summation are the same. By definition of $A_x = A_x^0 - A_x^1$, we have

$$\mathrm{Tr}\big(\big(A_x\otimes A_x\otimes \mathrm{Id}\big)\,\sigma\big)=\sum_{a=a'}\mathrm{Tr}\big(\big(A_x^a\otimes A_x^{a'}\otimes \mathrm{Id}\big)\,\sigma\big)-\sum_{a\neq a'}\mathrm{Tr}\big(\big(A_x^a\otimes A_x^{a'}\otimes \mathrm{Id}\big)\,\sigma\big).$$

Using that $A_x^0 + A_x^1 = \mathrm{Id}$, the sum (instead of the difference) of the two terms on the right-hand side is 1. Combining this observation with (2.13) proves (2.8), and (2.9) is proved in a similar way. □

Proof of Claim 6.

$$\begin{aligned}
&\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u^3\otimes \mathrm{Id}\otimes \mathrm{Id}-\widehat{A}_u\otimes\widehat{A}_u\otimes\widehat{A}_u\big)\,\sigma\big)\\
&=\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u^3\otimes \mathrm{Id}\otimes \mathrm{Id}-\widehat{A}_u^2\otimes\widehat{A}_u\otimes \mathrm{Id}\big)\,\sigma\big)+\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u^2\otimes\widehat{A}_u\otimes \mathrm{Id}-\widehat{A}_u\otimes\widehat{A}_u\otimes\widehat{A}_u\big)\,\sigma\big)\\
&=\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u^2\otimes \mathrm{Id}\otimes \mathrm{Id}\big)\cdot\big(\widehat{A}_u\otimes \mathrm{Id}\otimes \mathrm{Id}-\mathrm{Id}\otimes\widehat{A}_u\otimes \mathrm{Id}\big)\,\sigma\big)\\
&\qquad+\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u\otimes\widehat{A}_u\otimes \mathrm{Id}\big)\cdot\big(\widehat{A}_u\otimes \mathrm{Id}\otimes \mathrm{Id}-\mathrm{Id}\otimes \mathrm{Id}\otimes\widehat{A}_u\big)\,\sigma\big)\\
&\leq\Big(\sum_u \mathrm{Tr}\big(\big(\widehat{A}_u\otimes \mathrm{Id}\otimes \mathrm{Id}-\mathrm{Id}\otimes\widehat{A}_u\otimes \mathrm{Id}\big)^2\sigma\big)\Big)^{1/2}\Big(\sum_u \mathrm{Tr}\big(\widehat{A}_u^4\,\rho\big)+\mathrm{Tr}\big(\big(\widehat{A}_u^2\otimes\widehat{A}_u^2\otimes \mathrm{Id}\big)\,\sigma\big)\Big)^{1/2}\\
&\leq\sqrt{8\varepsilon}\cdot\sqrt{2},
\end{aligned}$$

where the first inequality is the Cauchy-Schwarz inequality, and the last uses Parseval's identity $\sum_u \hat{A}_u^2 = \mathrm{Id}$. Eq. (2.11) lets us conclude the proof. □

Chapter 3

Preliminaries

This chapter describes the notation that is used throughout this dissertation, and collects some useful definitions relating to quantum computing, entangled games and interactive proofs. We start with some notation in Section 3.1, before giving a brief overview of the relevant notions from quantum computing in Section 3.2 and from quantum information theory in Section 3.3. We then define multiplayer games in Section 3.4.1, and introduce the complexity classes built on them in Section 3.4.2.

3.1 Notation

Sets and indices. We write $[N]$ for the set of integers $\{1, \ldots, N\}$. If $x \in \{0,1\}^n$ is a string of length n, $i \in [n]$ an integer, and $S \subseteq [n]$ a set of integers, we write x_i for the i^{th} bit of x, and x_S for the string formed by the bits of x at the positions given by the elements of S. We also use the shorthands $x_{<i}$ for $x_{[1..i-1]}$, $x_{\geq i}$ for $x_{[i..n]}$, etc. Given two n-bit strings x, y we let $d_H(x,y) = \frac{1}{n}\sum_{i=1}^{n} \delta_{x_i,y_i}$ denote their relative Hamming distance.

Algebra. $\mathbb{R}$ and $\mathbb{C}$ denote the fields of real and complex numbers respectively, and $\|\cdot\|_2$ the Euclidean norm. A calligraphic letter $\mathcal{H}$ will usually denote a finite-dimensional complex Hilbert state. $\mathcal{M}_d(\mathbb{C})$ is the set of all $d \times d$ matrices with coefficients in $\mathbb{C}$. We will use $\mathrm{Tr}(A) := \sum_i A_{ii}$ to denote the trace, and $A^\dagger := \left(\overline{A_{ji}}\right)_{ij}$ to denote the conjugate-transpose. $\mathcal{M}_d(\mathbb{C})$ forms a Hilbert space when equipped with the inner-product $(A, B) \mapsto \mathrm{Tr}(AB^\dagger)$, and the resulting norm is the Frobenius norm $\|A\|_F = \sqrt{\mathrm{Tr}(AA^\dagger)}$. We also let $\|A\|_1 = \mathrm{Tr}\sqrt{AA^\dagger}$ be the Schatten 1-norm, also called the trace norm $\|A\|_{\mathrm{tr}} := \|A\|_1$ in the context of quantum computing, and $\|A\|_\infty$ be the operator norm. We denote by $\mathcal{P}(\mathcal{H})$ the set of positive semi-definite operators on $\mathcal{H}$. We define the set of normalized quantum states $\mathcal{S}(\mathcal{H}) := \{\rho \in \mathcal{P}(\mathcal{H}) : \mathrm{Tr}\,\rho = 1\}$ and the set of sub-normalized quantum states $\mathcal{S}_{\leq}(\mathcal{H}) := \{\rho \in \mathcal{P}(\mathcal{H}) : \mathrm{Tr}\,\rho \leq 1\}$.

Random variables. We use capital letters X, Y, Z to denote random variables, and $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ for the sets that they range in. If a classical random variable X takes the value $x \in \mathcal{X}$ with probability p_x, it can be represented by the state $\rho_X = \sum_{x\in X} p_x |x\rangle\langle x|x$, where $\{|x\rangle\}_{x\in\mathcal{X}}$ is an orthonormal basis of a Hilbert space $\mathcal{H}_X$. If the classical system X is part of a composite system XB, any state of that composite system can be written as $\rho_{XB} = \sum_{x\in\mathcal{X}} p_x |x\rangle\langle x|x \otimes \rho_B^x$. Such a state is called a cq-state (cq stands for classical-quantum).

Other notation. $\mathbb{F}$ will be used to denote a finite field. We use poly as a short-hand for any fixed polynomial. We use the script M to denote a register containing a quantum state.

3.2 Quantum computing

We refer the reader to Section 2.2 in the previous chapter for a gentle introduction to quantum states and measurements in the context of entangled games. Here we collect some of the most important definitions uses throughout this dissertation, referring the reader to a book such as Nielsen and Chuang's [85] for additional details.

States and measurements. A d-dimensional quantum state is a vector $|\Psi\rangle \in \mathbb{C}^d$. A d-dimensional density matrix is a positive matrix $\rho \in \mathcal{M}_d(\mathbb{C})$ with trace 1. A k-outcome positive operator-valued measurement (POVM in short) is given by a set of k (possibly rectangular) matrices A_i such that $\sum_i A_i^\dagger A_i = \mathrm{Id}$. The $A_i^\dagger A_i$ are called the *POVM elements*. A *projective* measurement is a POVM in which each element is a projector, i.e. $(A_i^\dagger A_i)^2 = A_i^\dagger A_i$. When the POVM $\{A_i\}$ is performed on a state ρ, the outcome i is observed with probability $\mathrm{Tr}(A_i^\dagger A_i \rho)$, and the state is projected onto the *post-measurement state*

$$\rho_i = \frac{A_i \rho A_i^\dagger}{\mathrm{Tr}\big(A_i^\dagger A_i \rho\big)}.$$

Distance measures. The usual distance measure on states is the trace norm, defined for any matrix A as

$$\|A\|_1 := \mathrm{Tr}\sqrt{A^\dagger A}.$$

Another useful distance is the Fidelity, defined for a pair of density matrices ρ and σ as

$$F(\rho,\sigma) := \Big(\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\Big)^2.$$

We will use the following property of the fidelity.

Lemma 7 ([109, 84]). *For any density operators ρ, σ, ξ over a Hilbert space $\mathcal{H}$,*

$$F(\rho,\sigma)^2 + F(\sigma,\xi)^2 \leq 1 + F(\rho,\xi).$$

3.3 Quantum information theory

In this section we introduce a few basic information-theoretic notions that will be used mostly in Chapters 8 and 9. Given a random variable $X \in \{0,1\}^n$, its min-entropy is

$$H_{\min}(X) = -\log \max_x \Pr(X = x).$$

For two distributions p, q on a domain D, their statistical distance is

$$\|p - q\|_1 := (1/2) \sum_{x \in D} \big|p(x) - q(x)\big|_1.$$

This notion of distance can be extended to random variables with the same range in the natural way. In the context of randomness extraction, we need a "robust" version of the min-entropy, the *smooth conditional min-entropy*, to measure how much randomness a source contains and can be extracted. In the classical setting it is given by

$$H^{\varepsilon}_{\min}(X) = \sup_{Y,\, \|Y - X\|_1 \le \varepsilon} H_{\min}(Y),$$

where $\varepsilon > 0$ is a small parameter. The following simple claim will be useful.

Claim 8. *Let $\alpha, \varepsilon > 0$ and X a random variable such that $H^{\varepsilon}_{\min}(X) \le \alpha$. Then there exists a set B such that $\Pr(X \in B) \ge \varepsilon$ and for every $x \in B$, it holds that $\Pr(X = x) \ge 2^{-\alpha}$.*

Proof. Let B be the set of x such that $\Pr(X = x) \ge 2^{-\alpha}$, and suppose $\Pr(X \in B) < \varepsilon$. Define Y so that $\Pr(Y = x) = \Pr(X = x)$ for every $x \notin B$, $\Pr(Y = x) = 0$ for every $x \in B$. In order to normalize Y, introduce new values z such that $\Pr(X = z) = 0$, and extend Y by defining $\Pr(Y = z) = 2^{-\alpha-1}$ until it is properly normalized. Then $\|Y - X\|_1 < \varepsilon$ and $H_{\min}(Y) > \alpha$, contradicting the assumption on the smooth min-entropy of X. □

A quantum analogue of the smooth conditional min-entropy was first introduced by Renner [100]. It represents the optimal measure for randomness extraction in the presence of quantum adversaries, in the sense that it is always possible to extract that amount of almost-uniform (from the point of view the adversary) randomness from a source (with which the adversary may be correlated), but never more. We first define the (non-smooth) quantum conditional min-entropy.

Definition 9 (conditional min-entropy [100])**.** *Let $\rho_{AB} \in \mathcal{S}_{\le}(\mathcal{H}_{AB})$. The* min-entropy *of A conditioned on B is defined as*

$$H_{\min}(A|B)_{\rho} := \max\{\lambda \in \mathbb{R} : \exists \sigma_B \in \mathcal{S}(\mathcal{H}_B) \text{ s.t. } 2^{-\lambda} \mathbb{1}_A \otimes \sigma_B \ge \rho_{AB}\}.$$

We will often drop the subscript ρ when there is no doubt about what underlying state is meant.

This definition has a simple operational interpretation when the first system is classical, which is the case we will consider. König et al. [75] showed that for a state $\rho_{XB} = \sum_{x\in\mathcal{X}} p_x|x\rangle\langle x|x \otimes \rho_B^x$ classical on X,

$$H_{\min}(X|B)_\rho = -\log p_{\text{guess}}(X|B)_\rho, \tag{3.1}$$

where $p_{\text{guess}}(X|B)$ is the maximum probability of guessing X given B, namely

$$p_{\text{guess}}(X|B)_\rho := \max_{\{E_B^x\}_{x\in\mathcal{X}}} \left(\sum_{x\in\mathcal{X}} p_x \operatorname{Tr}(E_B^x \rho_B^x) \right),$$

where the maximum is taken over all POVMs $\{E_B^x\}_{x\in\mathcal{X}}$ on B. If the system B is empty, then the min-entropy of X reduces to the standard definition, $H_{\min}(X) = -\log \max_{x\in\mathcal{X}} p_x$ (sometimes written $H_\infty(X)$). In this case the connection to the guessing probability is particularly obvious: when no side information is available, the best guess we can make is simply the value $x \in \mathcal{X}$ with highest probability.

In terms of randomness extraction the conditional min-entropy is not quite optimal, in the sense that it is sometimes possible to extract more randomness. However, the *smooth* min-entropy is optimal. In analogy with the classical setting, this information measure consists in maximizing the min-entropy over all sub-normalized states ε-close to the actual state ρ_{XB} of the system considered. Thus by introducing an extra error ε, we have a state with potentially much more entropy.

Definition 10 (smooth min-entropy [100, 116])**.** *Let $\varepsilon \geq 0$ and $\rho_{AB} \in \mathcal{S}_\leq(\mathcal{H}_{AB})$, then the* ε-smooth min-entropy *of A conditioned on B is defined as*

$$H_{\min}^{\varepsilon}(A|B)_\rho := \max_{\tilde{\rho}_{AB}\in\mathcal{B}^\varepsilon(\rho_{AB})} H_{\min}(A|B)_{\tilde{\rho}},$$

where $\mathcal{B}^\varepsilon(\rho_{AB}) \subseteq \mathcal{S}_\leq(\mathcal{H}_{AB})$ is a ball of sub-normalized states of radius ε around ρ_{AB}.[1]

3.4 Games and complexity classes

Multiplayer games and interactive proofs give two different languages to study the same object, and each is adapted to a slightly different context. A game is usually thought of

[1] Theoretically any distance measure could be used to define an ε-ball. We use the *purified distance*, $P(\rho,\sigma) := \sqrt{1 - F(\rho,\sigma)^2}$, where $F(\cdot,\cdot)$ is the fidelity, since this measure has some advantages over other metrics such as the trace distance. The only property of the purified distance we will need is that it is larger than the trace distance, i.e., $P(\rho,\sigma) \geq \frac{1}{2}\|\rho - \sigma\|_{\text{tr}}$. We refer to [116] for a formal definition of the purified distance (and fidelity) on sub-normalized states and a discussion of its advantages.

as being specified *explicitly*, through a description of the verifier's set of possible questions to the players, the probability distribution with which he chooses them, and the predicate with which he decides to accept or reject the provers' answers. Moreover, the language of games is often used to describe the simplest setting of a single round of interaction between the verifier and two provers (also sometimes called players in this context), even though in principle they can also involve more rounds of interaction or more provers.

In contrast, when using the language of interactive proofs one usually thinks of the protocol as being given *implicitly*, through the specification of a polynomial-time randomized Turing machine describing the verifier's behavior in the protocol, given access to an input $x \in \{0,1\}^*$. As such, interactive proofs may involve a number of possible questions to the provers that is exponential in the input size, but such that the verifier can sample from the corresponding distribution in polynomial time. Interactive proofs can involve a polynomial number of rounds of interaction between the verifier and a polynomial number of provers.

We proceed with the formal definitions.

3.4.1 Games

Let Q and A be finite sets and let k be a positive integer. We distinguish three types of games.

Classical game: A classical game is given by a distribution $\pi\colon Q^k \to [0,1]$ and a function $V\colon A^k \times Q^k \to \{0,1\}$.[2] The verifier samples questions $(q_1,\dots,q_k)$ according to π, and sends q_i to prover i from whom he then receives an answer a_i. He accepts those answers if and only if $V(a_1,\dots,a_k \mid q_1,\dots,q_k) = 1$. The *value* of the game is

$$\omega(G) = \max \Big[\sum_{\substack{(q_1,\dots,q_k)\in Q^k \\ (a_1,\dots,a_k)\in A^k}} \pi(q_1,\dots,q_k) \Pr(a_1,\dots,a_k \mid q_1,\dots,q_k) \\ \times V(a_1,\dots,a_k \mid q_1,\dots,q_k) \Big],$$

where the maximum is taken over all the provers' strategies W_i for $i \in \{1,\dots,k\}$, i.e., functions $W_i\colon Q \times R \to A$ for some domain R ("shared randomness"), and

$$\Pr(a_1,\dots,a_k \mid q_1,\dots,q_k) = \Pr_{r\in R}\Big(W_1(q_1,r) = a_1,\dots,W_N(q_k,r) = a_k\Big).$$

In fact we can assume the strategies to be *deterministic*: there is always some $r \in R$ that maximizes the winning probability and we can fix it in advance.

Classical entangled game: A classical entangled game is similar to a classical game, except that the provers are now allowed to share an arbitrary state $|\Psi\rangle$ of arbitrary

[2] We write $V(\cdot,\cdot)$ as $V(\cdot \mid \cdot)$ to clarify the role of the inputs.

dimension. This increases the set of possible strategies to quantum operations performed on the prover's share of the entangled state. Note that no restrictions on $|\Psi\rangle$ (such as $|\Psi\rangle$ consisting of EPR pairs, or $|\Psi\rangle$ having bounded dimension) are currently known to hold without loss of generality[3]. By standard purification techniques (see, e.g., [26]) one can assume that for each question q each prover performs a projective measurement $\mathcal{W}_q = \{W_q^a\}_{a \in A}$ with outcomes in A. We will use a superscript "$*$" to indicate entangled-prover games. The value $\omega^*(G)$ of such a game is given by[4]

$$\omega^*(G) = \sup \Big[\sum_{\substack{(q_1,\ldots,q_k) \in Q^k \\ (a_1,\ldots,a_k) \in A^k}} \pi(q_1, \ldots, q_k) \Pr(a_1, \ldots, a_k \mid q_1, \ldots, q_k) \\ \times V(a_1, \ldots, a_k \mid q_1, \ldots, q_k) \Big],$$

where the supremum is taken over all a priori shared states $|\Psi\rangle$ and all projective measurements $(\mathcal{W}_i)_q = \{(W_i)_q^a\}_{a \in A}$ for $i \in \{1, \ldots, k\}$ and $q \in Q$, and the probability now is

$$\Pr(a_1, \ldots, a_k \mid q_1, \ldots, q_k) = \langle \Psi | (W_1)_{q_1}^{a_1} \otimes \cdots \otimes (W_k)_{q_k}^{a_k} | \Psi \rangle.$$

Quantum entangled game: A quantum entangled game is a game in which both the verifier and the provers are quantum, and they exchange quantum messages. We usually denote such a game by G_{q}. The verifier holds k message registers of size $\mathrm{poly}(\log |Q|)$ each, in addition to a private register of size $\mathrm{poly}(\log |Q|)$, all initialized to the state $|0 \cdots 0\rangle$. He applies a unitary V_1 to all the registers, then sends the message registers to the corresponding provers. By purification we can assume that the j-th prover performs a unitary transformation U_j on his message register and his part of the entangled state $|\Psi\rangle$ and then sends the message register back to the verifier. The verifier performs a quantum operation V_2 on the message registers and his private space, followed by a measurement $\{\Pi_{\mathrm{acc}}, \Pi_{\mathrm{rej}}\}$ of his first qubit. The value of a quantum entangled game, ω_{q}^*, is given by

$$\omega_{\mathrm{q}}^*(G_{\mathrm{q}}) = \sup_{|\Psi\rangle, U_1, \ldots, U_k} \mathrm{Tr}(\Pi_{\mathrm{acc}} V_2 U V_1 |\Psi\rangle\langle\Psi|\Psi \otimes |0 \cdots 0\rangle\langle 0 \cdots 0 | 0 \cdots 0 V_1^\dagger U^\dagger V_2^\dagger),$$

where $U = U_1 \otimes \cdots \otimes U_k$.

[3] In fact, there are games known in which the maximum success probability of the provers goes to 1 with the dimension of their entangled state [77]. Note however that these games involve quantum messages, and are thus quantum entangled games in our terminology.

[4] We use a supremum because the optimal strategies might not be finite in the case of entangled provers.

Input size. We measure the size of a game through the cardinality of the question set Q. All the other components of the game's distribution (the distribution π, the answer size, the verifier's circuits V_1 and V_2 in the quantum case) will always be of size polynomial in $|Q|$.[5]

Special classes of games. In our work on parallel repetition (cf. Chapter 7) we will distinguish the following classes of games.

Definition 11. *A two-player game $G = (V, \pi)$ is called a*

- Projection game *if for every $q', q \in Q$ and $a' \in A$, there is a unique $a \in A$ such that $V(a', a|q', q) = 1$.*
- Free game *if $\pi = \pi_A \times \pi_B$ is a product distribution.*
- Symmetric game *if π is symmetric, and for any q', q, a', a we have $V(a', a|q', q) = V(a, a'|q, q')$.*

Symmetric games

More generally, we will say that a k-prover entangled game is *symmetric*, or *permutation-invariant*, if, for any tuple of strategies $(P_1, \dots, P_k)$ of the provers, the verifier accepts $(P_1, \dots, P_k)$ with exactly the same probability as he accepts any permutation $(P_{\sigma(1)}, \dots, P_{\sigma(k)})$.

In this section we show that any entangled game can be turned into an equivalent symmetric game, in the sense that the maximum success probability of any provers in either game is the same. Moreover, if a game is symmetric then the provers always have an optimal symmetric strategy as defined below. Symmetry is a useful simplifying assumption in two respects: first it lets one assume that the set of POVMs used by both provers is the same. Second, and most important, it implies that the prover's shared entangled state is also permutation-invariant. This property will be essential in much of our analysis of entangled games (cf. e.g. the very definition of the ρ-norms in Section 4.1 in Chapter 4). In its simplest form we will often make use of it to argue that the specific register on which each prover makes his or her measurement is unimportant.

Definition 12. *Let $(P_1, \dots, P_k, |\Psi\rangle)$ be a k-prover strategy.*[6] *We say that this strategy is symmetric, or permutation-invariant, if $P_1 = \dots = P_k$ and $|\Psi\rangle$ is invariant with respect to any permutation of the subsystems corresponding to each prover.*

[5] In fact all games we consider also have circuits of size poly(log $|Q|$) to prepare the questions and check the answers.

[6] We think of P_i as an arbitrary representation of the set of all of prover i's POVMs.

The following lemma shows that one can always assume without loss of generality that a game is symmetric, and that in any symmetric game there is an optimal symmetric strategy for the provers.

Lemma 13. *For every k-prover game G there is a k-prover game G' of the same value and k times as many questions that is permutation-invariant. Moreover, G and G' have the same value, and given any strategy $P_1, \dots, P_k$ with entangled state $|\Psi\rangle$ that wins with probability p in G, there exists a strategy $P'_1, \dots, P'_k$ with entangled state $|\Psi'\rangle$ and success probability p in G' such that $P'_1 = \dots = P'_k$ and $|\Psi'\rangle$ is permutation-invariant. In addition, if $|\Psi\rangle$ was a maximally entangled state then $|\Psi'\rangle$ is also.*

Proof. The verifier V' in game G' picks a uniformly random permutation $\pi \in \mathfrak{S}_k$, where $\mathfrak{S}_k$ is the set of permutations of $\{1, \dots, k\}$, at the start of the protocol. He then behaves exactly as in G, except that for every question q he would have sent to the i-th prover, where $i \in \{1, \cdots, k\}$, he sends the question (q, i) to prover $\pi(i)$ instead. Any answer he receives from prover $\pi(i)$ he treats as an answer received from prover i in the original protocol.

By appropriately padding with extra qubits, assume that all k registers of $|\Psi\rangle$ have the same dimension. Define strategies $P'_1, \dots, P'_k$ as follows: the provers share the entangled state $|\Psi'\rangle = \sum_{\sigma \in \mathfrak{S}_k} |\sigma(1)\rangle \otimes \dots \otimes |\sigma(k)\rangle \otimes |\Psi^\sigma\rangle$, where the register containing $|\sigma(i)\rangle$ is given to prover i and $|\Psi^\sigma\rangle$ is obtained from $|\Psi\rangle$ by permuting its registers according to σ. For $1 \leq i \leq k$ prover i measures the register containing $|\sigma(i)\rangle$ and behaves as in the strategy $P_{\sigma(i)}$. By symmetry of π and V this new strategy has the same success probability p, and $|\Psi'\rangle$ has the required symmetry properties.

This also shows that the value of G' is at least that of G. Conversely, if $P_1, \dots, P_k$ is a strategy in G', one constructs a strategy with at least the same value for G by choosing the best out of $(P_{\sigma(1)}, \dots, P_{\sigma(k)})$ over all permutations $\sigma \in \mathfrak{S}_k$. □

Lemma 13 has the following trivial but useful consequence.

Claim 14. *Let $(P_1, \dots, P_k, |\Psi\rangle)$ be a symmetric strategy, and for every $i \in \{1, \dots, k\}$, $\{A_i^a\}_a$ a POVM for the i-th prover in that strategy. Then for every permutation σ on $\{1, \dots, k\}$, and every $(a_1, \dots, a_k)$,*

$$\langle \Psi | A_1^{a_1} \otimes \dots \otimes A_k^{a_k} | \Psi \rangle \;=\; \langle \Psi | A_{\sigma(1)}^{a_{\sigma(1)}} \otimes \dots \otimes A_{\sigma(k)}^{a_{\sigma(k)}} | \Psi \rangle,$$

i.e. the register on which each of the POVMs is performed does not matter up to an arbitrary permutation.

The CHSH game

We conclude this section with a brief description of the most famous entangled game, the CHSH game, originally introduced by Clause, Horne, Shimony and Holt [25] to demonstrate the non-locality of quantum mechanics. This game will be important for our results in

Chapter 9. There are two players, Alice and Bob. Each is given a bit $x, y \in \{0,1\}$ distributed uniformly at random. Their goal is to produce bits a, b respectively such that $a \oplus b = x \wedge y$. It is not hard to see that classical parties (possibly using shared randomness) have a maximum success probability of 3/4 in this game. In contrast, quantum mechanics predicts that the following strategy, which we will sometimes refer to as the "honest" strategy, achieves a success probability of $\cos^2(\pi/8) \approx 0.85$. Alice and Bob share an EPR pair $|\Psi\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$. Upon receiving her input, Alice measures either in the computational ($x = 0$) or the Hadamard ($x = 1$) basis. Bob measures in the computational basis rotated by either $\pi/8$ ($y = 0$) or $3\pi/8$ ($y = 1$). One can then verify that, for every pair of inputs (x, y), this strategy produces a pair of correct outputs with probability exactly $\cos^2(\pi/8)$.

3.4.2 Interactive proofs with multiple provers

In this section we define the three main complexity classes that this dissertation is concerned with: Multi-Prover Interactive Proof Systems (MIP systems), Multi-Prover Interactive Proof Systems *with Entanglement* (MIP* systems), and *Quantum* Multi-Prover Interactive Proof Systems *with Entanglement* (MIP* systems). They are based respectively on multiplayer games, multiplayer entangled games and quantum multiplayer games.

As before, k denotes the number of provers and m denotes the number of turns (a turn consists in a single step of interaction in the protocol, in which either a message is sent from the verifier to each of the provers, or a message is sent from each of the provers to the verifier; a round is made of two turns). All of these are from the set of polynomially bounded functions in the input size $|x|$, denoted by poly. Further, c and s denote polynomial-time computable functions of the input size into $[0,1]$ corresponding to completeness and soundness. For notational convenience in what follows we will omit the arguments of these functions.

Multi-Prover Interactive Proof Systems (MIP systems): MIP systems were first introduced in [17]. A k-prover interactive proof system consists of a verifier V and k provers $P_1, \ldots, P_k$. The verifier is a probabilistic polynomial-time Turing machine, while the provers are computationally unbounded. Each of them has a read-only input tape, a private work tape and a random tape. In addition, the provers share an infinite read-only random tape of 0's and 1's (their shared randomness). Each prover has a write-only communication tape in which he writes messages to the verifier. The verifier has k write-only communication tapes, on which he writes messages to each of the provers.

The protocol proceeds in m turns. A turn is either a turn for the verifier or a turn for the provers. A turn for the verifier consists in the verifier reading the message tapes from each prover, and writing new messages to the provers on the corresponding write-only tape. A turn of the provers is similar, except now the provers individually read the messages they received from the verifier, and write their answers on the corresponding tape. The last turn is always a turn of the provers. The verifier then gets to read the last messages he received from the provers, and he produces a special output bit, designating acceptance or rejection.

Definition 15. *A language L is in* $\mathrm{MIP}(k, m, c, s)$ *iff there exists an m-turn polynomial-time verifier V for k-prover interactive proof systems such that, for every input x:*

(Completeness) *if $x \in L$, there exist m-turn provers $P_1, \ldots, P_k$ such that the interaction protocol of V with $(P_1, \ldots, P_k)$ results in the verifier accepting with probability at least c,*

(Soundness) *if $x \notin L$, for any m-turn provers $P'_1, \ldots, P'_k$, the probability that the interaction protocol of V with $(P_1, \ldots, P_k)$ results in the verifier accepting is at most s.*

If the parameters k or m are not specified they will be taken to be fixed polynomials in the input size. If the parameters c (resp. s) are not specified they will be taken as $c = 2/3$ (resp. $s = 1/3$). Hence $\mathrm{MIP} = \mathrm{MIP}(\mathrm{poly}, \mathrm{poly}, 2/3, 1/3)$.

Multi-Prover Interactive Proof Systems with Entanglement (MIP^* systems): MIP^* systems were first introduced in [26], and they are defined analogously to MIP systems. The only difference is that now the provers are allowed to be *quantum*, while the verifier (and communication) remains bounded in classical probabilistic polynomial-time. This implies the provers may share an arbitrary entangled state $|\Psi\rangle$ in-between themselves before the protocol starts. A turn of the provers in the protocol is defined as a turn in which the provers individually read the messages they received from the verifier, perform a measurement on their share of the entangled state, and send back the outcome to the verifier.[7] Formally,

Definition 16. *A language L is in* $\mathrm{MIP}^*(k, m, c, s)$ *iff there exists an m-turn polynomial-time verifier V for k-prover interactive proof systems such that, for every input x:*

(Completeness) *if $x \in L$, there exist m-turn provers $P_1, \ldots, P_k$ and a state $|\Psi\rangle$ such that the interaction protocol of V with $(P_1, \ldots, P_k)$ results in the verifier accepting with probability at least c,*

(Soundness) *if $x \notin L$, for any m-turn provers $P'_1, \ldots, P'_k$ and entangled state $|\Psi\rangle$, the probability that the interaction protocol of V with $(P_1, \ldots, P_k)$ results in the verifier accepting is at most s.*

[7] Any classical post-processing by the prover can be incorporated as part of the description of his measurement.

Quantum Multi-Prover Interactive Proof Systems with Entanglement (QMIP* **systems):** As in earlier work [125, 70, 72], we define QMIP* systems in terms of quantum circuits. It is assumed that our circuits consist of unitary gates, which is sufficient since non-unitary and unitary quantum circuits are equivalent in computational power [4]. To avoid unnecessary complication, however, the descriptions of protocols often include non-unitary operations (measurements). Even in such cases, it is always possible to construct unitary quantum circuits that essentially achieve the same outcome. A notable exception is in the definition of the public-coin quantum verifier, where we want to define the public coin-flip to be a classical operation. This requires a non-unitary operation for the verifier, the (classical) public coin-flip.

A quantum k-prover interactive proof system consists of a *verifier* V with private quantum register V and k *provers* $P_1, \ldots, P_k$ with private quantum registers $\mathsf{P}_1, \ldots, \mathsf{P}_k$, as well as quantum message registers $\mathsf{M}_1, \ldots, \mathsf{M}_k$, which without loss of generality are assumed to have the same number of qubits, denoted by q_{M}. One of the private qubits of the verifier is designated as the *output qubit*. At the beginning of the protocol, all the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$ are initialized to $|0 \cdots 0\rangle$, and the qubits in $(\mathsf{P}_1, \ldots, \mathsf{P}_k)$ are in some *a priori shared state* $|\Phi\rangle$ prepared by the provers in advance (and hence possibly entangled), which without loss of generality can be assumed to be pure. No direct communication between the provers is allowed after that. The protocol consists of alternating turns of the provers and of the verifier, starting with the verifier, if m is even, and with the provers otherwise. At a turn of the verifier, V applies some polynomial-size circuit to the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$, and then sends each register M_i to prover P_i. At a turn of the provers each prover P_i applies some transformation to the registers $(\mathsf{P}_i, \mathsf{M}_i)$ for $1 \leq i \leq k$ and sends M_i back to the verifier. The last turn is always a turn for the provers. After the last turn the verifier applies a polynomial-size circuit to the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$, and then measures the output qubit in the standard basis, accepting if the outcome is $|1\rangle$ and rejecting otherwise.

Formally, an *m-turn polynomial-time quantum verifier V* for k-prover QMIP* systems is a polynomial-time computable mapping from input strings x to a set of polynomial-time uniformly generated circuits $\{V^1, \ldots, V^{\lceil (m+1)/2 \rceil}\}$, and a partition of the space on which they act into registers $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$, which consist of polynomially many qubits. Similarly an *m-turn quantum prover P* is a mapping from x to a set of circuits $\{P^1, \ldots, P^{\lceil (m+1)/2 \rceil}\}$ each acting on registers (P, M). No restrictions are placed on the complexity of this mapping or the size of P. We will denote the i-th prover, his registers and transformations with a subscript i. We will always assume that each prover P_i is *compatible* with the verifier, i.e., that the corresponding register M_i is the same for the verifier and the prover for $1 \leq i \leq k$.

The *protocol* $(V, P_1, \ldots, P_k, |\Phi\rangle)$ is the alternating application of the circuits of the provers and the verifier to the initial state $|0 \cdots 0\rangle \otimes |\Phi\rangle$ in registers $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{P}_1, \ldots, \mathsf{P}_k)$. For odd m, circuits $P_1^1 \otimes \cdots \otimes P_k^1$, V^1, $P_1^2 \otimes \cdots \otimes P_k^2$, V^2 and so on are applied in sequence terminating with $V^{(m+1)/2}$. If m is even, the sequence begins with V^1 followed by $P_1^1 \otimes \cdots \otimes P_k^1$ and so on up to $V^{(m+2)/2}$. We say that $(V, P_1, \ldots, P_k, |\Phi\rangle)$ accepts x if the designated output qubit in V is measured in $|1\rangle$ at the end of the protocol and call the probability with which this happens $p_{\mathrm{acc}}(x, V, P_1, \ldots, P_k, |\Phi\rangle)$.

Definition 17. *A language L is in* QMIP$^*(k, m, c, s)$ *iff there exists an m-turn polynomial-time quantum verifier V for quantum k-prover interactive proof systems such that, for every input x:*

(Completeness) *if $x \in L$, there exist m-turn quantum provers $P_1, \ldots, P_k$ and an a priori shared state $|\Phi\rangle$ such that $p_{\mathrm{acc}}(x, V, P_1, \ldots, P_k, |\Phi\rangle) \geq c$,*

(Soundness) *if $x \notin L$, for any m-turn quantum provers $P'_1, \ldots, P'_k$ and any a priori shared state $|\Phi'\rangle$, $p_{\mathrm{acc}}(x, V, P'_1, \ldots, P'_k, |\Phi'\rangle) \leq s$.*

Finally, we introduce the notions of *public-coin* quantum verifiers and *public-coin* QMIP* systems, that will be used in Chapter 10. These are natural generalizations of the corresponding notions in the single-prover case introduced by [81]. Intuitively, a quantum verifier for quantum multi-prover interactive proof systems is public-coin if, at each of his turns, after receiving the message registers from the provers, he first flips a fair classical coin at most a polynomial number of times, and then simply broadcasts the result of these coin-flips to all the provers. No other messages are sent from the verifier to the provers. At the end of the protocol, the verifier applies some quantum operation to the messages received so far, and decides acceptance or rejection.

Formally, an m-turn polynomial-time quantum verifier for k-prover interactive proof systems is *public-coin* if each of the circuits $V^1, V^2, \ldots, V^{\lceil (m-1)/2 \rceil}$ implements the following procedure: V receives the message registers M_i from the provers, stores them in his private space, and then flips a classical fair coin at most q_{M} times to generate a public string r_j, records r_j in his private space, and broadcasts r_j to all the provers. The circuit $V^{\lceil (m+1)/2 \rceil}$ is some unitary transformation controlled by all the recorded random strings r_j for $1 \leq j \leq \lceil (m-1)/2 \rceil$. A QMIP* system is *public-coin* if the associated verifier is public-coin, and we define QMIP$^*_{\mathrm{pub}}(k, m, c, s)$ to be the class of languages in QMIP$^*(k, m, c, s)$ with a public-coin verifier.

Chapter 4

Techniques

In this chapter we give a technical introduction to some of the important tools that are used to analyze entangled strategies throughout the remainder of this dissertation. Our goal is to find ways to manipulate entangled provers by designing tests that the verifier can perform and which enforce specific structural constraints on the provers' strategies.

The first tool we introduce, in Section 4.1, is an appropriate distance measure on entangled strategies. It is designed to measure how much arbitrary provers differ from "ideal" provers in a multiplayer game.[1] Such a distance should be strong enough to imply that close strategies induce a similar behavior in the game (their success probabilities are close), but weak enough that one can obtain bounds on it from the sole assumption that the provers have a high probability of succeeding in a certain game.

In Section 4.2 we introduce an important test in the analysis of entangled strategies, the *consistency test*, and we study the structure of strategies having a high probability of passing that test. Finally, in Section 4.3 we give two general-purpose results building on the previous sections. The first one is the *orthogonalization lemma*, which states that consistent strategies are close to orthogonal (projective) strategies. This lemma will be used in the proof of our parallel repetition result in Chapter 7. The second is the *almost-commuting vs. nearly-commuting* conjecture, which provides an attempt at limiting the provers' use of entanglement in a multiplayer game. It will be explored further in Chapter 6, and we explain how it would imply a hardness result for the value of entangled games.

Notation. To simplify the presentation, in this chapter we will for the most part focus on the case of two provers, Alice and Bob, applying only one POVM each: $\{A^a\}$ for Alice and $\{B^b\}$ for Bob. We will assume their entangled state $|\Psi\rangle_{AB}$ is symmetric with respect to permutation of the A and B subsystems (this assumption is justified by Lemma 13 from

[1] Recall that the soundness analysis of a game usually proceeds by showing that arbitrary successful provers must be "close" to ideal, honest provers, whose maximum success probability can be bounded.

Chapter 3), and denote its reduced density on either subsystem by

$$\rho := \mathrm{Tr}_B\big(|\Psi\rangle\langle\Psi|\big) = \mathrm{Tr}_A\big(|\Psi\rangle\langle\Psi|\big).$$

Moreover, we will choose a basis in which to represent matrices that is such that ρ is real (such as its diagonalization basis).

The maximally entangled state. Many of the results in this section can be greatly simplified, and sometimes even trivialized, by assuming that the prover's entangled state $|\Psi\rangle$ is the maximally entangled state $|\Psi\rangle = d^{-1/2}\sum_{i=1}^{d}|i\rangle|i\rangle$ of arbitrary dimension d, with reduced density $\rho = d^{-1}\mathrm{Id}$. Whenever possible we will first introduce the results in that setting, and then explain what needs to be done for the case of a general entangled state.

4.1 Distance measures

In this section we introduce two key distance measures on *single-prover* strategies, the *univariate* and *bivariate* ρ-norms. In the case where the prover's entangled state is maximally entangled these two norms collapse to a single one, which we introduce in Section 4.1.1. The univariate and bivariate ρ-norms are then introduced in Section 4.1.2. We relate these distance measures to more standard measures from quantum information theory, such as the trace norm, in Section 4.1.3.

4.1.1 Prelude: the case of the maximally entangled state

Let $\{A_1^a\}$ and $\{A_2^a\}$ be two distinct measurements on Alice's share of the maximally entangled state $|\Psi\rangle$. We think of A_1 and A_2 as two possible strategies that she could apply: what is an appropriate way to measure the distance between these two strategies in the context of a multiplayer game?

We can always think of Alice as making her measurement first, obtaining an outcome a that she sends to the verifier as her answer; then Bob will make his own measurement. Once Alice has applied the measurement A_i, for $i \in \{1,2\}$, conditioned on the outcome a she obtained Bob's share of the state is projected onto the *post-measurement* state

$$\rho_{B,i}^a := \mathrm{Tr}_A\Big(\big(\sqrt{A_i^a}\otimes\mathrm{Id}_B\big)\,|\Psi\rangle\langle\Psi|\,\big(\sqrt{A_i^a}\otimes\mathrm{Id}_B\big)\Big) = \sqrt{A_i^a}^{T}\rho_B\sqrt{A_i^a}^{T} = \frac{1}{d}(A_i^a)^T,$$

where for the last two equalities we made use of the fact that $|\Psi\rangle$ was the maximally entangled state, so that for any X we have $(X\otimes\mathrm{Id})|\Psi\rangle = (\mathrm{Id}\otimes X^T)|\Psi\rangle$, and $\rho_B = \mathrm{Tr}_A|\Psi\rangle\langle\Psi| = d^{-1}\mathrm{Id}$. A sufficient condition for the provers' behavior in the game to be roughly similar irrespective of whether Alice measured using A_1 or A_2 is that the corresponding post-measurement states be *close*, *on average* over the outcome a.[2] The appropriate measure of distance between two quantum states is the trace norm, so that a good candidate measure of the distance between the strategies A_1 and A_2, in case the underlying state is maximally entangled, would be

$$d_1(A_1, A_2) \;:=\; \sum_a \big\|\rho^a_{B,1} - \rho^a_{B,2}\big\|_1 \;=\; \sum_a \frac{1}{d}\mathrm{Tr}\big|A^a_1 - A^a_2\big|.$$

The difficulty in working with the distance d_1 is the absolute value in the last term, which is hard to compute in general. Instead of d_1, we will use its much more malleable Euclidean equivalent d_2, defined as

$$d_2(A_1, A_2) \;:=\; \Big(\sum_a \frac{1}{d}\mathrm{Tr}\big(\big(\sqrt{A^a_1} - \sqrt{A^a_2}\big)^2\big)\Big)^{1/2}.$$

Note that d_2 corresponds to measuring the amount by which the provers' entangled state was moved in the *Euclidean distance*, as a result of Alice applying either A_1 or A_2:

$$d_2(A_1, A_2) \;=\; \Big(\sum_a \big\|(\sqrt{A^a_1}\otimes \mathrm{Id})|\Psi\rangle - (\sqrt{A^a_2}\otimes \mathrm{Id})|\Psi\rangle\big\|_2^2\Big)^{1/2}.$$

In fact, in the specific setting of the maximally entangled state d_2 is simply a dimension-normalized variant of the Frobenius norm: $\big(d_2(A_1, A_2)\big)^2 = \sum_a d^{-1}\big\|\sqrt{A^a_1} - \sqrt{A^a_2}\big\|_F^2$.

The measure d_2 will be more convenient than d_1 for two reasons. The first is that it derives from an inner product: if we define

$$\langle A, B\rangle \;:=\; \sum_a \frac{1}{d}\mathrm{Tr}\big(A^a(B^a)^\dagger\big)$$

then $\langle\cdot,\cdot\rangle$ is an inner product on $\big(\mathcal{M}_d(\mathbb{C})\big)^k$, where k is the number of distinct outcomes a, and $d_2(A_1, A_2)^2 = \langle\sqrt{A_1} - \sqrt{A_2}, \sqrt{A_1} - \sqrt{A_2}\rangle$. As a result, the Cauchy-Schwarz inequality will be an important tool to manipulate distances measured according to d_2.

The second reason is that the provers' success in a game often naturally translates into constraints on the distance d_2. We will explore this connection in much more detail in Section 4.2, but to give an idea of its flavor we show the following.

Lemma 18. *Suppose that, when Alice and Bob respectively apply their measurements $\{A^a\}$ and $\{B^a\}$ on their share of the maximally entangled state $|\Psi\rangle$, they obtain the same outcome with probability $1-\varepsilon$. Then*

$$\big(d_2(A, \overline{B})\big)^2 \;=\; \sum_a \frac{1}{d}\big\|A^a - \overline{B^a}\big\|_F^2 \;\le\; 2\varepsilon,$$

showing that consistent measurements are close in the d_2 distance.

[2] If an outcome has a very small probability of occurring then it does not matter whether the post-measurement states differ much; the verifier will still not notice.

Proof. Rephrasing the fact that Alice and Bob's measurements give consistent answers,

$$\begin{aligned} 1-\varepsilon &\leq \sum_a \langle\Psi|A^a \otimes B^a|\Psi\rangle \\ &= \sum_a \frac{1}{d}\mathrm{Tr}\big(A^a(B^a)^T\big) \\ &\leq \sum_a \frac{1}{d}\mathrm{Tr}\big(\sqrt{A^a}\sqrt{(B^a)^T}\big) \\ &= \langle\sqrt{A},\sqrt{B}^T\rangle, \end{aligned} \tag{4.1}$$

where the first equality uses that $|\Psi\rangle$ is maximally entangled, and the second $0 \leq A^a \leq \sqrt{A^a} \leq \mathrm{Id}$, and similarly for $(B^a)^T$, since A^a is a POVM element. Given that $\langle\sqrt{A},\sqrt{A}\rangle = 1$ as a result of $\sum_a A^a = \mathrm{Id}$, (4.1) proves the lemma (since the B^a are Hermitian). □

We end this section by proving a lemma showing that the distances d_1 and d_2 are closely related, implying that not only is d_2 a more convenient distance to work with than d_1, but it also provides a good bound on the distance between two entangled strategies in a multiplayer game.

Lemma 19. *Let $\{A_1^a\}$ and $\{A_2^a\}$ be two POVMs. Then the following relations hold:*

$$\frac{1}{2}\,d_2(A_1,A_2)^2 \;\leq\; d_1(A_1,A_2) \;\leq\; 2\,d_2(A_1,A_2).$$

Proof. For any a it holds that

$$\big(\sqrt{A_1^a}-\sqrt{A_2^a}\big)^2 \;\leq\; \big|\sqrt{A_1^a}+\sqrt{A_2^a}\big|\,\big|\sqrt{A_1^a}-\sqrt{A_2^a}\big| \;\leq\; 2\big|\sqrt{A_1^a}-\sqrt{A_2^a}\big|,$$

since both $\sqrt{A_1^a}, \sqrt{A_2^a} \leq \mathrm{Id}$; the first inequality follows. For the second, write

$$\begin{aligned} d_1(A_1,A_2) &= \sum_a \frac{1}{d}\big\|A_1^a - A_2^a\big\|_1 \\ &\leq \sum_a \frac{1}{d}\Big(\big\|\sqrt{A_1^a}\big(\sqrt{A_1^a}-\sqrt{A_2^a}\big)\big\|_1 + \big\|\big(\sqrt{A_1^a}-\sqrt{A_2^a}\big)\sqrt{A_2^a}\big\|_1\Big) \\ &\leq \sum_a \frac{1}{d}\big\|A_1^a - A_2^a\big\|_F\,\big(\big\|\sqrt{A_1^a}\big\|_F + \big\|\sqrt{A_2^a}\big\|_F\big) \\ &\leq \Big(\sum_a \frac{1}{d}\big\|A_1^a - A_2^a\big\|_F^2\Big)^{1/2}\Big(\sum_a \frac{1}{d}\big(\big\|\sqrt{A_1^a}\big\|_F + \big\|\sqrt{A_2^a}\big\|_F\big)^2\Big)^{1/2} \\ &\leq 2\,d_2(A_1,A_2), \end{aligned}$$

where the first inequality is the triangle inequality, and the second and third each follow from the Cauchy-Schwarz inequality. □

4.1.2 The ρ-norms

We turn to the case where the provers' entangled state $|\Psi\rangle$ is a general *symmetric* bipartite state, with reduced density ρ on either prover's subsystem. We will introduce two distinct distance measures, each deriving from a norm: the *univariate* and *bivariate* ρ-norms. In case $|\Psi\rangle$ is the maximally entangled state both distances reduce to the distance d_2 introduced in the previous section, but in general they are distinct and will have different uses. To motivate their introduction, first recall Ando's identity:

Fact 20 (Ando's identity)**.** *For any matrices A and B and symmetric $|\Psi\rangle = \sum_i \lambda_i |u_i\rangle|v_i\rangle$, let $K = \sum_i \sqrt{\lambda_i}|u_i\rangle\langle v_i|$. Then*

$$\langle\Psi|A\otimes B|\Psi\rangle = \mathrm{Tr}\big(AKB^TK^\dagger\big), \tag{4.2}$$

and the reduced density

$$\rho \;=\; \mathrm{Tr}_A|\Psi\rangle\langle\Psi| \;=\; \mathrm{Tr}_B|\Psi\rangle\langle\Psi| \;=\; KK^\dagger \;=\; K^\dagger K.$$

Proof. By symmetry it holds that $\rho = \sum_i \lambda_i|u_i\rangle\langle u_i| = \sum_j \lambda_j|v_j\rangle\langle v_j|$, and

$$\begin{aligned}\langle\Psi|A\otimes B|\Psi\rangle &= \sum_{i,j}\sqrt{\lambda_i}\sqrt{\lambda_j}\langle u_i|A|v_j\rangle\langle u_i|B|v_j\rangle\\ &= \mathrm{Tr}\big(K^\dagger AKB^T\big).\end{aligned}$$

□

Suppose given two measurements $\{A_1^a\}$ and $\{A_2^a\}$ for Alice, and a measurement $\{B_2^b\}$ for Bob. For a fixed pair (a,b), the probability that Alice and Bob obtain that outcome is $\langle\Psi|A_1^a\otimes B^b|\Psi\rangle$ if Alice measures using A_1, and it is $\langle\Psi|A_2^a\otimes B^b|\Psi\rangle$ if she measures using A_2. By Ando's identity, the difference between these two probabilities is

$$\big|\langle\Psi|A_1^a\otimes B^b|\Psi\rangle - \langle\Psi|A_2^a\otimes B^b|\Psi\rangle\big| \;=\; \big|\mathrm{Tr}\big((A_1^a - A_2^a)K(B^b)^TK^\dagger\big)\big|. \tag{4.3}$$

In an attempt to bound this quantity by one that depends only on Alice's two measurements, one may use the Cauchy-Schwarz inequality in two different ways. The first sees the quantity on the right-hand side as the matrix inner product between $(A_1^a - A_2^a)K$ and $K(B^b)^T$, in which case it can be upper-bounded as

$$\big|\mathrm{Tr}\big((A_1^a - A_2^a)K(B^b)^TK^\dagger\big)\big| \;\le\; \mathrm{Tr}\big((A_1^a - A_2^a)^2\rho\big)^{1/2}\mathrm{Tr}\big(\big((B^b)^T\big)^2\rho\big)^{1/2}. \tag{4.4}$$

The second interprets the right-hand side of (4.3) as the matrix inner-product of $K_1(A_1^a - A_2^a)K_1^\dagger$ and $K_2(B^b)^TK_2^\dagger$, where K_1, K_2 are such that $K_1^\dagger K_2 = K$, in which case it can be bounded as

$$\big|\mathrm{Tr}\big((A_1^a - A_2^a)K(B^b)^TK^\dagger\big)\big| \;\le\; \mathrm{Tr}\big((A_1^a - A_2^a)\rho^{1/2}(A_1^a - A_2^a)\rho^{1/2}\big)^{1/2}\mathrm{Tr}\big((B^b)^T\rho^{1/2}(B^b)^T\rho^{1/2}\big)^{1/2}, \tag{4.5}$$

where $\rho^{1/2} = (KK^\dagger)^{1/2} = K_1^\dagger K_1 = K_2^\dagger K_2$. Both ways of bounding (4.3) have their uses, and they give rise to the two distance measures on entangled strategies that we now introduce.[3]

The univariate ρ-norm

The univariate ρ-norm is the one that arises from the bound (4.4). As we will see, it is the strongest of the two norms that we introduce, and is defined as follows.

Definition 21. *Let A be any matrix, and $\rho \geq 0$. The* univariate ρ-norm *of A is*

$$\big\|A\big\|_{\mathbf{u},\rho} := \sqrt{\mathrm{Tr}\big(AA^\dagger\rho\big)}.$$

Remark. *The choice of* $\mathrm{Tr}\big(AA^\dagger\rho\big)$ *instead of* $\mathrm{Tr}\big(A^\dagger A\rho\big)$ *in the definition of the norm is arbitrary, but the two choices are* not *equivalent, as is seen by taking e.g. $A = |0\rangle\langle 1|$ and $\rho = |0\rangle\langle 0|$.*

Based on this norm one can define a corresponding distance on POVMs: if $\{A_1^a\}$ and $\{A_2^a\}$ are two POVMs for Alice,[4] we let

$$d_{\mathbf{u},\rho}(A_1, A_2) := \Big(\sum_a \big\|\sqrt{A_1^a} - \sqrt{A_2^a}\big\|_{\mathbf{u},\rho}^2\Big)^{1/2}.$$

We note that this is the same distance as was already introduced in Section 2.3.1 of Chapter 2.

An important tool in using the univariate ρ-norm is that, while the map $(A, B) \mapsto \mathrm{Tr}(AB^\dagger\rho)$ is not quite an inner-product in general (it is not symmetric), it still obeys a Cauchy-Schwarz inequality.

Claim 22 (Cauchy-Schwarz inequality for the univariate ρ-norm)**.** *Let A, B be any two matrices, and $\rho \geq 0$. Then*

$$\big|\mathrm{Tr}\big(AB^\dagger\rho\big)\big| \leq \big\|A\big\|_{\mathbf{u},\rho}\big\|B\big\|_{\mathbf{u},\rho}.$$

Proof. Apply the Cauchy-Schwarz inequality for the inner-product $(A', B') \mapsto \mathrm{Tr}\big(A'(B')^\dagger\big)$, with $A' = \rho^{1/2}A$ and $B' = \rho^{1/2}B$. □

[3] The attentive reader might have already observed that (4.5) provides a tighter bound than (4.4), since for any matrices A, B it holds that $\mathrm{Tr}(ABAB) \leq \mathrm{Tr}(A^2B^2)$.

[4] We'll assume that A_1 and A_2 have the same set of outcomes, as otherwise they are incomparable.

The bivariate ρ-norm

The bivariate ρ-norm arises from the bound (4.5), and we define it as follows.

Definition 23. *Let A be any matrix, and $\rho \geq 0$. The* bivariate ρ-norm *of A is*

$$\big\|A\big\|_{\mathbf{b},\rho} := \sqrt{\operatorname{Tr}\big(A\rho^{1/2}A^\dagger\rho^{1/2}\big)}.$$

Remark. *If $\rho = d^{-1}\mathrm{Id}$ is the totally mixed state, then both the univariate and the bivariate ρ-norms are the same: $\big\|A\big\|_{\mathbf{u},\rho} = \big\|A\big\|_{\mathbf{b},\rho} = d^{-1/2}\|A\|_F$. In general this equality no longer holds, and we will explore the relationship between the two norms in more detail in Section 4.1.3.*

Claim 24. *$A \mapsto \big\|A\big\|_{\mathbf{b},\rho}$ is a (semi-)norm,[5] and it derives from the (semi)-inner-product*

$$(A,B) \in \mathcal{M}_d(\mathbb{C}) \;\mapsto\; \langle A,B\rangle_\rho := \operatorname{Tr}\big(A\rho^{1/2}B^\dagger\rho^{1/2}\big) = \langle\Psi|A\otimes\overline{B}|\Psi\rangle.$$

Proof. $(A,B) \mapsto \langle A,B\rangle_\rho$ is sesquilinear, symmetric by cyclicity of the trace, and non-negative: it is a semi-inner product. Hence $A \mapsto \sqrt{\langle A,A\rangle_\rho}$ is a semi-norm. □

As a consequence, the bivariate ρ-norm satisfies a Cauchy-Schwarz inequality analogue to the one proved in Claim 22. We end this section with another useful inequality showing that the bivariate ρ-norm is an upper-bound on a measurement's *self-consistency*.

Claim 25. *For any A and symmetric state $|\Psi\rangle$,*

$$\langle\Psi|A\otimes A|\Psi\rangle \;\leq\; \big\|A\big\|^2_{\mathbf{b},\rho}.$$

Proof. Using Ando's identity (4.2),

$$\langle\Psi|A\otimes A|\Psi\rangle \;=\; \operatorname{Tr}\big(AKA^TK^\dagger\big),$$

and the claim follows from the Cauchy-Schwarz inequality as in (4.5). □

4.1.3 Relationships between norms

In this section we explore the relationships between the two norms introduced in the previous section and other distance measures on quantum operations or states. We first show that, in the case of a Hermitian A, the univariate ρ-norm is an upper-bound on the bivariate ρ-norm.

Claim 26. *For any Hermitian A and $\rho \geq 0$,*

$$\big\|A\big\|_{\mathbf{b},\rho} \;\leq\; \big\|A\big\|_{\mathbf{u},\rho}.$$

[5] It is a norm if ρ is positive-definite.

Proof. The claim follows from the Cauchy-Schwarz inequality:

$$\big\|A\big\|_{\mathbf{b},\rho}^2 \,=\, \mathrm{Tr}\big((A\rho^{1/2})(\rho^{1/2}A)^\dagger\big) \,\leq\, \mathrm{Tr}\big(A\rho A^\dagger\big)^{1/2}\mathrm{Tr}\big(\rho^{1/2}AA^\dagger\rho^{1/2}\big)^{1/2} \,=\, \big\|A\big\|_{\mathbf{u},\rho}^2,$$

where for the last equality we used that A was Hermitian. □

Remark. *In case A is not Hermitian, but still satisfies $AA^\dagger \leq \mathrm{Id}$ (and we have $\mathrm{Tr}(\rho) \leq 1$), the proof of the previous claim only shows the weaker relationship $\big\|A\big\|_{\mathbf{b},\rho}^2 \leq \big\|A\big\|_{\mathbf{u},\rho}$. This bound is tight in general: if $A = |0\rangle\langle 1|$ and $\rho = (1-\varepsilon)|1\rangle\langle 1| + \varepsilon|0\rangle\langle 0|$ then $\big\|A\big\|_{\mathbf{u},\rho}^2 = \varepsilon$, while $\big\|A\big\|_{\mathbf{b},\rho}^2 = \sqrt{\varepsilon(1-\varepsilon)}$. Choosing $A = |1\rangle\langle 0|$ instead shows that one cannot hope for any nontrivial inequality in the other direction.*

The next lemma shows that the univariate ρ-norm is at least as strong as the trace norm. It provides a partial analogue to Lemma 19, which was proven in case ρ was totally mixed.

Lemma 27. *Let ρ be a density matrix and $\big\{A^a\big\}$, $\big\{B^a\big\}$ two POVMs. Then*

$$\sum_a \Big\|\sqrt{A^a}\rho\sqrt{A^a} - \sqrt{B^a}\rho\sqrt{B^a}\Big\|_1 \,\leq\, 2\sqrt{\sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|_{\mathbf{u},\rho}^2} \,=\, 2\, d_{\mathbf{u},\rho}(A,B).$$

Proof. The inequality follows from a calculation similar to Ogawa and Nagaoka's proof [90] of Winter's "gentle measurement lemma" (Lemma 9 in [129]). By the triangle inequality, for any a

$$\big\|\sqrt{A^a}\rho\sqrt{A^a} - \sqrt{B^a}\rho\sqrt{B^a}\big\|_1 \leq \big\|\sqrt{A^a}\rho\big(\sqrt{A^a} - \sqrt{B^a}\big)\big\|_1 + \big\|\big(\sqrt{A^a} - \sqrt{B^a}\big)\rho\sqrt{B^a}\big\|_1.$$

Applying the Cauchy-Schwarz inequality (cf. Theorem 116 in Appendix A),

$$\begin{aligned}\big\|\sqrt{A^a}\rho\big(\sqrt{A^a} - \sqrt{B^a}\big)\big\|_1 &\leq \big\|\sqrt{A^a}\rho^{1/2}\big\|_F\big\|\rho^{1/2}\big(\sqrt{A^a} - \sqrt{B^a}\big)\big\|_F\\ &= \big\|\sqrt{A^a}\big\|_{\mathbf{u},\rho}\big\|\sqrt{A^a} - \sqrt{B^a}\big\|_{\mathbf{u},\rho}\end{aligned}$$

The claim follows by another application of Cauchy-Schwarz, together with the fact that $\sum_a \big\|\sqrt{A^a}\big\|_{\mathbf{u},\rho}^2 = \sum_a \mathrm{Tr}(A^a\rho) = 1$. □

Unfortunately, no bound similar to the one in Lemma 27 holds for the bivariate ρ-norm. The following claim shows that a relationship does hold between the two norms when one considers measurements that are *consistent*, a property that will be discussed in more detail in Section 4.2.

Claim 28. *Let $\{A^a\}$ and $\{B^a\}$ be two POVMs, $|\Psi\rangle$ a symmetric entangled state with reduced density ρ, and let*

$$\varepsilon \,:=\, 1 - \sum_a \big\|A^a\big\|_{\mathbf{b},\rho}^2.$$

Then

$$\sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|_{\mathbf{u},\rho}^2 \,\leq\, \sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|_{\mathbf{b},\rho}^2 + 6\sqrt{\varepsilon}.$$

Remark. *Claim 25 shows that the assumption placed on A in Claim 28 is* weaker *than the more natural one that would be placed by defining ε directly through A's self-consistency as a measurement by defining $\varepsilon := 1 - \sum_a \langle\Psi|A^a \otimes A^a|\Psi\rangle$.*

Proof. Expand

$$\begin{aligned}\sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|^2_{\mathbf{u},\rho} &= \sum_a \operatorname{Tr}\big((\sqrt{A^a} - \sqrt{B^a})^2\rho\big)\\ &= \sum_a \Big(\operatorname{Tr}\big((A^a + B^a)\rho\big) - 2\operatorname{Tr}\big(\sqrt{A^a}\sqrt{B^a}\rho\big)\Big). \qquad (4.6)\end{aligned}$$

The first term inside the summation adds up to 2. In order to lower-bound the second, we first bound

$$\begin{aligned}\sum_a \operatorname{Tr}\big((\sqrt{A^a}\rho^{1/2} - \rho^{1/2}\sqrt{A^a})^2\big) &= 2 - 2\sum_a \big\|\sqrt{A^a}\big\|^2_{\mathbf{b},\rho}\\ &\le 2\varepsilon, \qquad (4.7)\end{aligned}$$

where we used $\big\|\sqrt{A^a}\big\|_{\mathbf{b},\rho} \ge \big\|A^a\big\|_{\mathbf{b},\rho}$ (since $\sqrt{A^a} \ge A^a$), together with the definition of ε. Hence

$$\begin{aligned}&\sum_a \big|\operatorname{Tr}\big(\sqrt{A^a}\sqrt{B^a}\rho - \sqrt{A^a}\rho^{1/2}\sqrt{B^a}\rho^{1/2}\big)\big|\\ &\qquad\le \sum_a \operatorname{Tr}\big(B^a\rho\big)^{1/2}\operatorname{Tr}\big((\sqrt{A^a}\rho^{1/2} - \rho^{1/2}\sqrt{A^a})^2\big)^{1/2}\\ &\qquad\le \sqrt{2\varepsilon}, \qquad (4.8)\end{aligned}$$

where both inequalities are by Cauchy-Schwarz, and the second also uses (4.7). From (4.6) we get

$$\begin{aligned}\sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|^2_{\mathbf{u},\rho} &\le 2 - 2\sum_a \operatorname{Tr}\big(\sqrt{A^a}\rho^{1/2}\sqrt{B^a}\rho^{1/2}\big) + 2\sqrt{2\varepsilon}\\ &\le \sum_a \Big(\operatorname{Tr}\big(\sqrt{A^a}\rho^{1/2}\sqrt{A^a}\rho^{1/2}\big) + \operatorname{Tr}\big(\sqrt{B^a}\rho^{1/2}\sqrt{B^a}\rho^{1/2}\big)\\ &\qquad - 2\operatorname{Tr}\big(\sqrt{A^a}\rho^{1/2}\sqrt{B^a}\rho^{1/2}\big)\Big) + 2\varepsilon + 2\sqrt{2\varepsilon}\\ &\le \sum_a \big\|\sqrt{A^a} - \sqrt{B^a}\big\|^2_{\mathbf{b},\rho} + 6\sqrt{\varepsilon},\end{aligned}$$

where the first inequality is by (4.8), the second uses the definition of ε, and the last is by definition of the bivariate ρ-norm. □

4.2 Consistency

In a classical symmetric multiplayer game it is often natural to enforce that the provers are *consistent*, by checking, with some probability, that they provide the same answer when simultaneously presented with the same question. This was done, for instance, in the linearity test that we presented in Chapter 2. Indeed, consistency is at the heart of the connection between multiplayer games and probabilistically checkable proofs, as it implies the existence of a single underlying "proof" collecting all of the provers' answers to the verifier's possible questions.

In the entangled setting, however, consistency is a more stringent requirement, and does not always hold naturally. To see the subtlety, observe that even if all provers apply the *same*, *orthogonal*, *projective* measurement, they need not obtain identical answers — this will only be true of all such measurements if their entangled state is the maximally entangled state.

Nevertheless, in most cases one still expects that honest provers should provide consistent answers (indeed, the "ideal", honest prover is often a classical one, answering his questions deterministically), so that one is willing to incorporate the following test as part of the game being played:

> **Consistency test.** Pick a question at random[6] and send it to both players. Accept if and only if they provide the same answer.

In this section we explore the consequences that can be drawn of entangled players passing the consistency test, and show that this seemingly weak requirement induces strong constraints on the type of strategies they may use.

Lemma 29. *Let $\{A^a\}$ and $\{B^a\}$ be two POVMs indexed by the same answer set, and $|\Psi\rangle$ a symmetric bipartite state. Suppose further that A and B pass the consistency test with success probability $1-\varepsilon$:*

$$\sum_a \langle\Psi|A^a\otimes B^a|\Psi\rangle \,\geq\, 1-\varepsilon. \tag{4.9}$$

Then the following hold:

1. *The POVM $\{A^a\}$ itself is self-consistent, in the sense that $\sum_a \big\|A^a\big\|_{\mathbf{b},\rho}^2 \geq 1-2\varepsilon$, and the same holds of $\{B^a\}$,*

2. *If, moreover, $\{B^a\}$ is self-consistent in the stronger sense that $\sum_a \langle\Psi|B^a\otimes B^a|\Psi\rangle \geq 1-\delta$, then the POVMs with elements A^a and B^a are close in the univariate ρ-norm:*

$$\sum_a \big\|\sqrt{A^a}-\sqrt{B^a}\big\|_{\mathbf{u},\rho}^2 \,\leq\, 8\sqrt{\varepsilon},$$

[6] For the test to be effective, the distribution according to which the question is chosen should correspond to the marginal distribution arising from the overall game of which we plan to use the consistency test as a part of; if the game is symmetric then the marginals on different players will be identical.

3. *The density ρ is close to being invariant under application of $\{A^a\}$, in the sense that*

$$\Big\| \sum_a \sqrt{A^a}\rho\sqrt{A^a} - \rho \Big\|_1 \leq 2\,\varepsilon + \sqrt{2\,\delta},$$

and the same holds for $\{B^a\}$.

Remark. *This lemma generalizes Lemma 18, that was proved for the case of the maximally entangled state: in that case condition 3. is trivial, and condition 1. follows if we were only to impose that $\{A^a\}$ is an orthogonal measurement. Indeed, one can think of the consistency test as a way to force the provers to behave* as if *their shared state was maximally entangled, an exceedingly convenient assumption in many a protocol's analysis. We will draw more consequences of the same vein in Lemma 30 below.*

Proof. For the first item, we use Ando's identity (4.2) and the Cauchy-Schwarz inequality to write, for any a,

$$\begin{aligned}\langle\Psi|A^a\otimes B^a|\Psi\rangle &= \mathrm{Tr}\big(A^aK\overline{B^a}K^\dagger\big)\\ &\leq \mathrm{Tr}\big(A^a\rho^{1/2}A^a\rho^{1/2}\big)^{1/2}\mathrm{Tr}\big(B^a\rho^{1/2}B^a\rho^{1/2}\big)^{1/2}\\ &\leq \frac{1}{2}\Big(Tr\big(A^a\rho^{1/2}A^a\rho^{1/2}\big) + \mathrm{Tr}\big(B^a\rho^{1/2}B^a\rho^{1/2}\big)\Big),\end{aligned}$$

where we used that B^a was Hermitian, and ρ real. Hence from (4.9) and the fact that $\sum_a \mathrm{Tr}\big(A^a\rho^{1/2}A^a\rho^{1/2}\big) \leq 1$, and the same holds for B, we obtain

$$\sum_a \mathrm{Tr}\big(A^a\rho^{1/2}A^a\rho^{1/2}\big) \geq 1-2\varepsilon \qquad \text{and} \qquad \sum_a \mathrm{Tr}\big(B^a\rho^{1/2}B^a\rho^{1/2}\big) \geq 1-2\varepsilon, \tag{4.10}$$

proving the first item. For the second, first bound

$$\begin{aligned}\sum_a \mathrm{Tr}\big(\big(\sqrt{B^a}K - K\sqrt{B^a}^T\big)\big(\sqrt{B^a}K - K\sqrt{B^a}^T\big)^\dagger\big) &\leq 2 - 2\sum_a \mathrm{Tr}\big(\sqrt{B^a}K\sqrt{B^a}^T K^\dagger\big)\\ &\leq 2\delta \qquad (4.11)\end{aligned}$$

by Ando's identity and our assumption on B's self-consistency. Hence

$$\begin{aligned}\sum_a \big\|\sqrt{A^a}-\sqrt{B^a}\big\|_{\mathbf{u},\rho}^2 &= \sum_a \Big(\big\|\sqrt{A^a}\big\|_{\mathbf{u},\rho}^2 + \big\|\sqrt{B^a}\big\|_{\mathbf{u},\rho}^2 - 2\,\mathrm{Tr}\big(\sqrt{A^a}\sqrt{B^a}\rho\Big)\\ &= 2 - 2\sum_a \mathrm{Tr}\big(\sqrt{A^a}K\sqrt{B^a}^T K^\dagger\big) + \sqrt{2\delta}\\ &\leq 2\varepsilon + \sqrt{2\delta},\end{aligned}$$

where the second equality uses (4.11) together with the Cauchy-Schwarz inequality, and the inequality is by our assumption (4.9). This proves item 2.

Regarding the last item, by monotonicity of the trace norm we have

$$\begin{aligned}
&\Big\|\sum_a \sqrt{A^a}\rho\sqrt{A^a}-\rho\Big\|_1\\
&\quad\leq \Big\|\sum_{a,a'} \sqrt{A^a}\otimes\sqrt{B^{a'}}|\Psi\rangle\langle\Psi|\sqrt{A^a}\otimes\sqrt{B^{a'}}-\sum_{a'}\mathrm{Id}\otimes\sqrt{B^{a'}}|\Psi\rangle\langle\Psi|\mathrm{Id}\otimes\sqrt{B^{a'}}\Big\|_1\\
&\quad\leq \Big\|\sum_{a} \sqrt{A^a}\otimes\sqrt{B^{a}}|\Psi\rangle\langle\Psi|\sqrt{A^a}\otimes\sqrt{B^{a}}-\sum_{a}\mathrm{Id}\otimes\sqrt{B^{a}}|\Psi\rangle\langle\Psi|\mathrm{Id}\otimes\sqrt{B^{a}}\Big\|_1\\
&\qquad+\sum_{a\neq a'}\langle\Psi|A^a\otimes B^{a'}|\Psi\rangle\\
&\quad\leq 2\Big(\sum_a \big\|\sqrt{A^a}\otimes\sqrt{B^a}-\mathrm{Id}\otimes\sqrt{B^a}\big\|_{\mathbf{u},\rho}^2\Big)^{1/2}+\varepsilon,
\end{aligned}$$

where the second inequality uses the triangle inequality, and the last is by Lemma 27. It remains to bound

$$\begin{aligned}
\sum_a \big\|\sqrt{A^a}\otimes\sqrt{B^a}-\mathrm{Id}\otimes\sqrt{B^a}\big\|_{\mathbf{u},\rho}^2 &= \sum_a\Big(\langle\Psi|A^a\otimes B^a|\Psi\rangle+\langle\Psi|\mathrm{Id}\otimes B^a|\Psi\rangle\\
&\qquad\qquad -2\langle\Psi|\sqrt{A^a}\otimes B^a|\Psi\rangle\Big)\\
&\leq 2-2\sum_a\langle\Psi|A^a\otimes B^a|\Psi\rangle\\
&\leq 2\varepsilon,
\end{aligned}$$

where for the first inequality we used that $0\leq A^a\leq\mathrm{Id}$ and hence $\sqrt{A^a}\geq A^a$. This proves the last item in the lemma. □

From Lemma 29 one may derive even further conditions: the POVM $\{A^a\}$ is *almost-orthogonal*, and $\{A^a\}$ and $\{B^a\}$ *almost-commute*. As we will see in the next section, the orthogonality condition can be used to transform $\{A^a\}$ into a *projective* POVM. A key observation is that this together with item 3. in the lemma shows that we have almost reduced the provers to using their state as shared randomness, at least *locally* — a question at a time. In Section 4.3.2 we will discuss the extent to which this locality can be extended by using a global variant of the "almost-commute" condition.

Lemma 30. *Under the same assumptions as in Lemma 29, it further holds that*

1. *The POVM elements $\{A^a\}$ are almost-orthogonal:*

$$\sum_{a\neq a'}\mathrm{Tr}\big(\sqrt{A^a}A^{a'}\sqrt{A^a}\rho\big)\leq 15\sqrt{\varepsilon},$$

2. *If, in addition,* $\{B^a\}$ *is strongly consistent:* $\sum_a \langle\Psi|B^a\otimes B^a|\Psi\rangle \geq 1-\varepsilon$*, then the POVMs* $\{A^a\}$ *and* $\{B^a\}$ *almost-commute:*

$$\sum_a \big\|\sqrt{A^a}\sqrt{B^a} - \sqrt{B^a}\sqrt{A^a}\big\|_{\mathbf{u},\rho}^2 = O\big(\sqrt{\varepsilon}\big).$$

Proof. For the first item,

$$\begin{aligned}\sum_{a\neq a'} \mathrm{Tr}\big(\sqrt{A^a}A^{a'}\sqrt{A^a}\rho\big) &= \sum_{a\neq a',a''} \langle\Psi|\sqrt{A^a}A^{a'}\sqrt{A^a}\otimes B^{a''}|\Psi\rangle\\ &\leq \sum_{a\neq a'}\langle\Psi|\sqrt{A^a}A^{a'}\sqrt{A^a}\otimes B^{a}|\Psi\rangle + \varepsilon\\ &\leq \sum_{a\neq a',a''} \langle\Psi|\sqrt{A^{a''}}A^{a'}\sqrt{A^{a''}}\otimes B^{a}|\Psi\rangle + 2\varepsilon\\ &\leq \sum_{a\neq a'}\langle\Psi|A^{a'}\otimes B^a|\Psi\rangle + 12\sqrt{\varepsilon} + 2\varepsilon\\ &\leq 15\sqrt{\varepsilon},\end{aligned}$$

where the first two inequalities hold by the consistency assumption on $\{A^a\}$ and $\{B^a\}$, the third uses item 3. from Lemma 29, and the last uses the consistency assumption again.

Regarding the second item, we simply sketch its proof, since we will not use it formally. Expanding $\big\|\sqrt{A^a}\sqrt{B^a} - \sqrt{B^a}\sqrt{A^a}\big\|_{\mathbf{u},\rho}^2$ by using the definition of the univariate ρ-norm, one sees that it will suffice to show that $\sum_a \mathrm{Tr}\big(\sqrt{A^a}\sqrt{B^a}\sqrt{A^a}\sqrt{B^a}\rho\big)$ is close enough to 1. Using the Cauchy-Schwarz inequality, one can show that this term is $O(\sqrt{\varepsilon})$ from $\sum_a \mathrm{Tr}\big(A^aB^a\rho\big)$. This last term can in turn be lower-bounded by using B^a's self-consistency (cf item 1. from Lemma 29) and assumption (4.9). □

4.3 Applications

4.3.1 The orthogonalization lemma

Item 1. from Lemma 30 in the previous section shows that simply by requiring that a strategy be *self-consistent*, one can obtain strong conditions showing that the measurements applied as part of that strategy should satisfy *almost-orthogonality* relations. In this section we show that one can take the step from *almost* to *exactly* orthogonal, proving that an almost-orthogonal POVM is close, in the univariate ρ-norm, to one that is exactly orthogonal. A variant of this result will be used in Chapter 7, in which we prove our result on the parallel repetition of entangled games. The following is a statement of the lemma in the simplest setting.

Lemma 31 (Orthogonalization lemma). *Let $\varepsilon > 0$, ρ a density matrix, and $\{A^a\}$ a POVM, such that*

$$\sum_{a\neq a'} \mathrm{Tr}\big(\big(\sqrt{A^a}A^{a'}\sqrt{A^a}\big)\,\rho\big) \,\leq\, \varepsilon. \tag{4.12}$$

Then there exists an orthogonal *measurement $\{B^a\}$ such that*

$$\big(d_{\mathbf{u},\rho}(A,B)\big)^2 \,=\, \sum_a \big\|\sqrt{A^a} - B^a\big\|_{\mathbf{u},\rho}^2 \,=\, O\big(\varepsilon^{1/4}\big).$$

The strength of the lemma is in the bound it proves on the distance between the POVM $\{A^a\}$ and the orthogonal measurement $\{B^a\}$ being *independent* of the number of outcomes of these POVMs. Indeed, it is not too hard to obtain a bound depending on the number of outcomes by adopting an iterative orthogonalization procedure, such as a variant of the Gram-Schmidt process. To avoid this dependence one needs to find an appropriate *global* procedure, moving each of the POVM elements by only a small amount on average.

The idea in the present case is to use the singular value decomposition, or SVD, of the matrix having as its columns the eigenvectors of A^a whose corresponding eigenvalue is large enough. This method is based on a variant of Schöneman's solution to the "orthogonal Procrustes problem".[7] Given any d-dimensional square matrices A and B, this is the problem of finding the orthogonal matrix Ω which minimizes

$$\Omega := \text{ argmin } \frac{1}{d}\,\|A - B\Omega\|_F^2.$$

Schöneman [102] showed that the optimal Ω is $\Omega = UV^\dagger$, where $U\Sigma V^\dagger$ is the singular value decomposition of B^TA.[8] Indeed, given unit vectors $|u_1\rangle,\dots,|v_k\rangle$, one can let A be the matrix with columns the $|u_i\rangle$, and B the identity. In this case, the orthogonal Procruste's problem consists in finding the best rigid rotation which maps the canonical basis of space to the vectors $|v_i\rangle$, where the error is measured in the least squares sense — the columns of the corresponding orthogonal matrix will then form an orthonormal family close to the $|u_i\rangle$.

We carry out this solution precisely in Claim 32 below, which contains all the intuition necessary to solve the original problem on POVMs. Unfortunately, the solution to the latter is made more involved technically by the matrices not being of rank 1, and the slightly unorthodox (and, in particular, not rotationally invariant) norm used to measure the error. The proof of Lemma 31 itself is given in Appendix A.

Claim 32. *Let $|u_1\rangle,\dots,|u_k\rangle \in \mathbb{C}^k$ be unit vectors such that $\frac{1}{k}\sum_{i\neq j}\langle u_i,u_j\rangle^2 \leq \varepsilon$. Then there exist orthogonal unit vectors $|v_1\rangle,\dots,|v_k\rangle \in \mathbb{C}^k$ such that $\frac{1}{k}\sum_i \big\|\,|u_i\rangle - |v_i\rangle\,\big\|^2 \leq \varepsilon$.*

[7] According to Wikipedia, Procrustes, or "the stretcher", a figure from Greek mythology, was a rogue smith and bandit from Attica who physically attacked people, stretching them, or cutting off their legs so as to make them fit an iron bed's size.

[8] We are grateful to the user "ohai" of MathOverflow.net for pointing out the connection between this problem and that of the robust orthonormalization of almost-orthogonal vectors.

Proof. Let X be the $k \times k$ matrix whose columns are made of the vectors $|u_i\rangle$, expressed in the canonical basis. The SVD of X is $X = U\Sigma V^\dagger$, where U, V are unitary and Σ is diagonal with the singular values s_i of M on the the diagonal. Then

$$\frac{1}{k}\sum_{i=1}^{k}(1-s_i^2)^2 \;=\; \|\Sigma^\dagger\Sigma - \mathrm{Id}\|_F^2 \;=\; \|X^\dagger X - \mathrm{Id}\|_F^2 \;=\; \frac{1}{k}\sum_{i\neq j}\big|\langle u_i, u_j\rangle\big|^2 \;\leq\; \varepsilon \tag{4.13}$$

where for the first equality we used the unitary invariance of the Frobenius norm, and the second is by definition of X and uses the fact that the $|u_i\rangle$ have unit norm. Let $Y = UV^\dagger$. Y is a unitary matrix so its column vectors $|v_i\rangle$ form an orthonormal family. We have

$$\frac{1}{k}\sum_{i=1}^{k}\big\|\,|u_i\rangle - |v_i\rangle\,\big\|_2^2 \;=\; \|X-Y\|_F^2 \;=\; \|\mathrm{Id}-\Sigma\|_F^2 \;=\; \frac{1}{k}\sum_{i=1}^{k}(1-s_i)^2$$

which can be bounded by (4.13) since $(1-s_i)^2 \leq (1-s_i)^2(1+s_i)^2 = (1-s_i^2)^2$. □

4.3.2 The "almost-commuting vs. nearly-commuting" conjecture

In the previous section we showed that, if a POVM's elements were almost-orthogonal, then the POVM could be transformed into one that was exactly orthogonal, while only moving each POVM element by a small amount on average. In particular, the bound we obtained was independent of the number of outcomes of the POVM. In this section we introduce a similar-looking problem: given a set of *almost-commuting* projectors, does there exist a corresponding set of *exactly commuting* projectors that are close?

To motivate this question, we first show how an *almost-commuting* constraint on the prover's strategies arises naturally from a simple transformation that one can apply to any two-player entangled game. We then state the conjecture, and explain why a satisfactory answer would imply a hardness result for the complexity of approximating the value of two-player one-round entangled games, an important open problem.

Obtaining commutation relations in a two-player game

A technique to obtain almost-commuting constraints on prover strategies was first introduced by Kobayashi & al. [68] in the context of entangled three-prover games, as well as entangled two-player games with quantum messages (see Chapter 6 for more details on the technique for the case of three provers). The specific variant that we present here, adapted to the case of two-prover entangled games, is due to Ito & al. [59].

Consider a two-prover game G in which there are q questions per prover. Suppose the game is symmetric and, for simplicity, assume that the distribution on pairs of questions in the game is uniform. We introduce a modified game G', in which the verifier performs the following:

1. Sample two questions $(i,j) \in [q]^2$ uniformly at random.
2. Send the first question i to Alice, and send the (unordered) pair $\{i,j\}$ to Bob.
3. Upon receiving Alice and Bob's answers, check that they are consistent, *and* that Bob's answers satisfy the corresponding constraint from G.

This transformation of G into G' is in general known as the *oracularization* technique. In our context, the important point is that it introduced a variant of the *consistency test* in the original game in a way that will let us establish almost-commuting relations on the provers' measurements, provided they have a high probability of succeeding in the game G'.

Let $\{A_i^a\}$, $\{B_{ij}^{ab}\}$ and $|\Psi\rangle$ be a projective strategy in G' that has success probability at least $1-\varepsilon$, for some $\varepsilon > 0$. By making the game symmetric (permuting the role of Alice and Bob at random, telling the provers whose role they are supposed to play), we may assume that $|\Psi\rangle$ is a symmetric bipartite state, with reduced density ρ on either prover. The fact that this strategy succeeds with probability $1-\varepsilon$ implies that the following *consistency* condition must hold:

$$\mathrm{E}_{i,j} \sum_{a,b} \langle \Psi | A_i^a \otimes B_{ij}^{ab} | \Psi \rangle \;\geq\; 1-\varepsilon. \tag{4.14}$$

Additional constraints hold due to the checking of the original game's predicate, but we ignore those here and instead focus on the implications of (4.14). For simplicity we will write $B_{ij}^a := \sum_b B_{ij}^{ab}$, and symmetrically for B_{ij}^b (we will always use a to represent the answer to question i, and b for the answer to question j, so that there is no ambiguity). The following lemma is adapted from [59].

Lemma 33. *Suppose that (4.14) holds. Then*

$$\mathrm{E}_{i,j} \sum_{a,b} \Big\| \, [\, A_i^a, A_j^b \,] \, \Big\|_{\mathbf{u},\rho}^2 \;\leq\; 16\,\varepsilon.$$

Proof. First observe that (4.14) implies that

$$\begin{aligned}\mathrm{E}_{i,j} \sum_b \langle \Psi | \big(A_j^b \otimes \mathrm{Id} - \mathrm{Id} \otimes B_{ij}^b \big)^2 | \Psi \rangle &= 2 - 2\,\mathrm{E}_{i,j} \sum_b \langle \Psi | A_j^b \otimes B_{ij}^b | \Psi \rangle \\ &\leq 2\,\varepsilon, \end{aligned} \tag{4.15}$$

while

$$\begin{aligned}\mathrm{E}_{i,j} \sum_{a,b} \langle \Psi | \big(A_i^a \otimes B_{ij}^b - \mathrm{Id} \otimes B_{ij}^{ab} \big)^2 | \Psi \rangle &= 2 - 2\,\mathrm{E}_{i,j} \sum_{a,b} \langle \Psi | A_i^a \otimes B_{ij}^{ab} | \Psi \rangle \\ &\leq 2\,\varepsilon. \end{aligned} \tag{4.16}$$

This lets us bound

$$\begin{aligned}
&\mathrm{E}_{i,j} \sum_{a,b} \langle\Psi|\big(A_j^b A_i^a \otimes \mathrm{Id} - A_i^a \otimes B_{ij}^b\big)\big(A_j^b A_i^a \otimes \mathrm{Id} - A_i^a \otimes B_{ij}^b\big)^\dagger|\Psi\rangle \\
&\quad = \mathrm{E}_{i,j} \sum_{a,b} \langle\Psi|\big(A_j^b \otimes \mathrm{Id} - \mathrm{Id} \otimes B_{ij}^b\big)\big(A_i^a \otimes \mathrm{Id}\big)\big(A_j^b \otimes \mathrm{Id} - \mathrm{Id} \otimes B_{ij}^b\big)^\dagger \mathrm{Id}|\Psi\rangle \\
&\quad = \mathrm{E}_{i,j} \sum_{b} \langle\Psi|\big(A_j^b \otimes \mathrm{Id} - \mathrm{Id} \otimes B_{ij}^b\big)\big(A_j^b \otimes \mathrm{Id} - \mathrm{Id} \otimes B_{ij}^b\big)^\dagger|\Psi\rangle \\
&\quad \leq 2\,\varepsilon, \qquad (4.17)
\end{aligned}$$

where the last inequality is by (4.15). Using a similar calculation one can show that

$$\mathrm{E}_{i,j} \sum_{a,b} \langle\Psi|\big(A_i^a A_j^b \otimes \mathrm{Id} - A_j^b \otimes B_{ij}^a\big)\big(A_i^a A_j^b \otimes \mathrm{Id} - A_j^b \otimes B_{ij}^a\big)^\dagger|\Psi\rangle \;\leq\; 2\,\varepsilon. \qquad (4.18)$$

Finally, combining (4.17) and (4.18) with (4.16) and the triangle inequality, one obtains

$$\mathrm{E}_{i,j} \sum_{a,b} \langle\Psi|\big(A_i^a A_j^b \otimes \mathrm{Id} - A_j^b A_i^a \otimes \mathrm{Id}\big)^2|\Psi\rangle \;\leq\; 16\,\varepsilon,$$

proving the lemma. □

Making almost-commuting projectors nearly-commute

Consider the following conjecture.

Conjecture 34 ((ε,δ) nearly-commuting conjecture)**.** *Let $0 \leq \varepsilon < 1$, ρ a density matrix, and $P_1,\dots,P_k$ be d-dimensional projectors such that*

$$\frac{1}{k^2}\sum_{i,j=1}^{k} \Big\| \,[\,P_i,\, P_j\,]\, \Big\|_{\mathbf{u},\rho}^2 \;\leq\; \varepsilon.$$

Then there exists a $\delta > 0$ and projectors $Q_1,\dots,Q_k$ such that $[Q_i,Q_j] = 0$ for every $i \neq j$, and

$$\frac{1}{k}\sum_{i=1}^{k} \big\|P_i - Q_i\big\|_{\mathbf{u},\rho}^2 \;\leq\; \delta.$$

The assumption of the conjecture is that the projectors are *almost-commuting*, i.e. the norm of their commutators is small. The conclusion is that they must be *nearly-commuting*: they are close to exactly commuting projectors.

Of course for any $\varepsilon > 0$ there will always be a $\delta > 0$ such that the conjecture holds. Moreover, it is not too hard to see that it can be shown to hold for a δ that is dimension-independent. But can δ be made independent of the number k of projectors? If not, what is the best dependence of δ on k that one could hope for? This natural question is, as far as we know, completely open, even in the case where ρ is the maximally entangled state and the norm $\big\| \cdot \big\|_{\mathbf{u},\rho}$ is correspondingly replaced by the normalized Frobenius norm $d^{-1/2}\| \cdot \|_F$ in the equations above.

The history of the conjecture, and its implications, are discussed further in Chapter 6, and we refer the reader there for additional details. Here we simply point out that, if the conjecture was resolved in a satisfactory way, with a dependence $\delta = O(\text{poly} \log k \cdot \text{poly}(\varepsilon))$ on the number of projectors that is at most poly-logarithmic, then Lemma 33, which extracts an almost-commuting condition on the prover's POVMs, would also imply that the prover's measurements must nearly-commute — i.e. their strategy should be close to a classical strategy.

Chapter 5

Hardness of entangled games

In this chapter we show that multi-prover interactive proofs with entangled provers as at least as powerful as their classical counterparts, resolving a long-standing open question [72].

Theorem 35. *The following inclusion holds:*

$$\mathrm{NEXP} \subseteq \mathrm{MIP}^*(4, \mathrm{poly}, 1, 1 - 1/\mathrm{poly}).$$

After sequential repetition of the protocol, this implies that $\mathrm{NEXP} \subseteq \mathrm{MIP}^*(4, \mathrm{poly}, 1, 2^{-p})$ *for all* $p \in \mathrm{poly}$.

Prior to this result the best lower bound known was $\mathrm{PSPACE} \subseteq \mathrm{MIP}^*$, which follows by having the verifier interact with a single prover (in which case entanglement plays no role) and using IP = PSPACE [106]. Together with the celebrated result NEXP = MIP [12], Theorem 35 shows that entanglement does not weaken the power of general entangled-prover proof systems.

The proof of Theorem 35 is based on Babai, Fortnow and Lund's original proof [12] that $\mathrm{NEXP} \subseteq \mathrm{MIP}$. Key to its extension to the entangled-prover setting is an analysis of the *multilinearity test* that is at the heart of their proof. Our main contribution is to give an analysis of this test in the presence of entanglement between the provers. The challenges of such an analysis were already exposed in a simpler setting in Chapter 2. In that chapter we analyzed a much simpler variant of the multilinearity test, the *linearity test* of Blum, Luby and Rubinfeld [23].

We start by giving a brief overview of our proof strategy in the following section. After giving some necessary preliminaries, in Section 5.3 we describe the protocol used to prove Theorem 35. In Section 5.4 we show how the proof of the theorem reduces to the analysis of an appropriate multilinearity test with entangled provers. This analysis, which constitutes the core of the proof of Theorem 35, is given in Section 5.5.

5.1 Overview

The multi-prover protocol for NEXP introduced in [12] consists of two main components. In the first part of the protocol, the verifier performs a polynomial-round low-degree sum-check protocol with a single prover. In the second part of the protocol, he has a one-round interaction with each of the remaining three provers. This interaction is randomly chosen between two possibilities. The first consists in asking each of the three provers a single uniformly distributed question, and checking the answers they provide against the results of the sum-check protocol. The second consists in performing a one-round *multilinearity test* with the three provers. This test is designed to enforce that the function according to which each of the three provers chose their answers is the same, and has the property of being multilinear from $\mathbb{F}^n$ to $\mathbb{F}$, where $\mathbb{F}$ is a field in which range the questions and answers in the protocol.

The most challenging part in adapting this protocol to the case of entangled provers consists in the analysis of the multilinearity test, and we now outline the main steps of that analysis. The test is very simple. Let $\mathbb{F}$ be a field of characteristic 2, and n an integer. With probability 1/2 the verifier picks three axis-aligned points $\boldsymbol{x}, \boldsymbol{x} + \alpha\, \boldsymbol{e}_i, \boldsymbol{x} + \beta\, \boldsymbol{e}_i \in \mathbb{F}^n$, and checks that the provers' answers $a, b, c \in \mathbb{F}$ are correspondingly aligned. With the remaining probability he sends all three provers the same question, checking they reply with the same answer.

The analysis of the multilinearity test, in case the provers are *deterministic*, shows that if they have success probability at least $1 - \varepsilon$ then there must *exist* a multilinear function g such that g agrees with each prover's own answers on a large fraction of points $\boldsymbol{x} \in \mathbb{F}^n$ (provided ε was small enough). In the presence of entanglement, however, one cannot hope for such a statement: the provers may well be using their entanglement as shared randomness, e.g. to coordinate in selecting one of many possible multilinear functions to compute their answers. This observation points to a fundamental difficulty: how does one "extract" a multilinear function from arbitrary entangled-prover strategies? What does it even mean for such strategies to be "close to multilinear"?

Our work addresses these questions. A key insight, in retrospect a necessity, is to *directly* manipulate the provers' strategies themselves, without *explicitly* trying to relate them to a classical strategy. More precisely, for $\boldsymbol{x} \in \mathbb{F}_p^n$ let $\left\{A_{\boldsymbol{x}}^a\right\}_{a\in\mathbb{F}_p}$ be the POVM applied by the provers[1] to determine their answer a, upon receiving question $\boldsymbol{x}$. We will show how one can remove the dependence of $\{A_{\boldsymbol{x}}^a\}$ on $\boldsymbol{x}$, one coordinate at a time, by constructing a sequence of POVMs $\left\{B_{x_{k+1},\dots,x_n}^g\right\}_g$ with outcomes g in the set of multilinear functions $\mathbb{F}_p^k \to \mathbb{F}_p$. These POVMs will have the following key property: the respective strategies corresponding to (i) measuring according to $\{A_{\boldsymbol{x}}^a\}$, and answering a or (ii) measuring according to $\{B_{x_{k+1},\dots,x_n}^g\}$ and answering $g(a_1, \dots, x_k)$ are *consistent*, in the sense that if two distinct provers use either strategy then with high probability they obtain the same answer. Setting $k = n$ will give the

[1] A standard symmetry argument lets us assume that all provers use the same measurement.

final result, by further proving that consistency implies closeness in an appropriate distance measure. We will have shown that the original provers' strategy is indistinguishable from one in which the provers perform a measurement *independent* of their question, obtaining a multilinear function g from which they compute their answer by evaluating $g(\boldsymbol{x})$. This effectively reduces the provers to using their entanglement as shared randomness, performing their measurement even before the protocol starts.

A word on how the B measurements are defined. As we already pointed out, it is essential that they are constructed as a *global* function of the original POVMs $\{A^a_{\boldsymbol{x}}\}$. We define them inductively, and only explain the one-dimensional case here. Our definition is very intuitive: $\{B^\ell\}$ is the POVM which corresponds to measuring using $\{A^a_{x_1}\}$ successively using two randomly chosen values of x_1, and returning the unique linear function which interpolates between the two outcomes obtained. Once one has settled on this definition, it is in fact not too hard to show that success of $\{A^a_{x_1}\}$ in the multilinearity test implies that the strategies corresponding to the $\{A^a_{x_1}\}_a$ and $\{B^\ell\}_\ell$ POVMs are consistent.

An additional major hurdle arises as a result of the inductive argument sketched above (and this difficulty is already present in Babai, Fortnow and Lund's classical analysis of the test): the quality of the approximation between the A and B strategies blows up *exponentially* with k. In order to control this error, one has to perform an additional step of *self-correction* on the B strategy. Making this step work requires substantially more work in the case of entangled strategies than it does in the classical setting, and is one of our main technical contributions.

5.2 Preliminaries

Notation. For a field $\mathbb{F}$, a linear function $g\colon \mathbb{F} \to \mathbb{F}$ is a function such that there exists $a, b \in \mathbb{F}$, $g(x) = ax + b$. A multilinear function $g\colon \mathbb{F}^k \to \mathbb{F}$ is a function that is linear in each of its coordinates. $\mathrm{ML}(\mathbb{F}^k, \mathbb{F})$ will denote the set of all multilinear functions from $\mathbb{F}^k$ to $\mathbb{F}$. We will denote tuples using bold symbols such as $\boldsymbol{x}$ and $\boldsymbol{b}$. Given a tuple $\boldsymbol{x} = (x_1, \dots, x_n)$ and $k \in [n]$, let $\boldsymbol{x}_{\leq k} = (x_1, \dots, x_k)$, $\boldsymbol{x}_{>k} = (x_{k+1}, \dots, x_n)$ and $\boldsymbol{x}_{\neg k} = (x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_n)$.

Given a positive matrix ρ and an arbitrary matrix A, we let $\mathrm{Tr}_\rho(A) := \mathrm{Tr}(A\rho)$. In case ρ is a matrix on two Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$, and A is a matrix on $\mathcal{H}_1$, we will sometimes abuse notation and also write $\mathrm{Tr}_\rho(A)$ for $\mathrm{Tr}_\rho(A \otimes \mathrm{Id}_2)$. We also let

$$\|A\|_\rho^2 \,:=\, \mathrm{Tr}\big(AA^\dagger\rho\big),$$

and observe that $A \mapsto \|A\|_\rho$ is a semi-norm (it is definite if ρ is invertible). It was introduced as the "univariate ρ-norm" $\big\|A\big\|_{\mathbf{u},\rho}$ in Chapter 4. In that chapter we proved the following Cauchy-Schwarz inequality: for any A, B,

$$\mathrm{Tr}_\rho\big(AB^\dagger\big) \,\leq\, \|A\|_\rho \, \|B\|_\rho.$$

We will work with incomplete POVMs $\{P_i\}$, which are simply a collection of positive matrices such that $\sum_i P_i \leq \mathrm{Id}$ (a complete POVM, in contrast, satisfies $\sum_i P_i = \mathrm{Id}$).

5.2.1 NEXP-complete problems

We will use the following NEXP-complete problem, as stated in Proposition 4.2 of Ref. [12]:

Problem 1: Oracle-3-satisfiability.
Instance. Integers $r, s \in \mathbb{N}$ in unary and a Boolean formula $B(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, a_1, a_2, a_3)$ in variables $\boldsymbol{z} \in \{0,1\}^r$, $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3 \in \{0,1\}^s$ and $a_1, a_2, a_3 \in \{0,1\}$.

Question. Does there exist a mapping $A\colon \{0,1\}^s \to \{0,1\}$ such that $B(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, A(\boldsymbol{b}_1), A(\boldsymbol{b}_2), A(\boldsymbol{b}_3))$ = 1 simultaneously for all $\boldsymbol{z} \in \{0,1\}^r$ and $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3 \in \{0,1\}^s$?

Using the standard technique of arithmetization (e.g. Proposition 3.1 and Lemma 7.1 of Ref. [12]), one can show that the following problem is also NEXP-complete.

Problem 2: Oracle-3-satisfiability, arithmetized version.
Instance. Integers $r, s \in \mathbb{N}$ in unary and an arithmetic expression[2] for a polynomial $f(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, a_1, a_2, a_3)$, where $\boldsymbol{z}$ represents r variables and each of $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3$ represents s variables.

Yes-promise. There exists a mapping $A\colon \{0,1\}^s \to \{0,1\}$ such that for all $\boldsymbol{z} \in \{0,1\}^r$ and all $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3 \in \{0,1\}^s$, it holds that

$$f(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, A(\boldsymbol{b}_1), A(\boldsymbol{b}_2), A(\boldsymbol{b}_3)) = 0 \tag{5.1}$$

in $\mathbb{Z}$ (and therefore in every field).

No-promise. For every pair $(\mathbb{F}, A)$ of a field $\mathbb{F}$ and a mapping $A\colon \{0,1\}^s \to \mathbb{F}$, there exist $\boldsymbol{z} \in \{0,1\}^r$ and $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3 \in \{0,1\}^s$ such that Eq. (5.1) is not satisfied in $\mathbb{F}$.

We note that the degree of the polynomial f represented by the arithmetic expression can be at most the size of the arithmetic expression, and is therefore bounded by the input size.

5.2.2 Summation test

Let $\mathbb{F}$ be a finite field of characteristic two. If $|\mathbb{F}| = 2^k$, an encoding scheme of elements in $\mathbb{F}$ is specified by k and a primitive polynomial $f(t)$ over $\mathbb{F}_2$ of degree k. It is well-known that given 1^k, $f(t)$, and the complete factorization of $2^k - 1$ along with the certificate that each factor in the factorization is indeed a prime (such as the Pratt certificate), it is possible to check that k and $f(t)$ form a valid encoding scheme of the field $\mathbb{F}$ in polynomial time.

Consider the following promise problem, which has both an explicit and an implicit input.

[2] An *arithmetic expression* is a rooted tree whose internal nodes represent either addition or multiplication and whose leaves represent either variables or an integer constant. The size of an arithmetic expression is the number of nodes plus the sum of the number of bits required to represent the integer for each constant node.

Problem 3: Summation Test Problem.
Explicit input. Integers $m, d \in \mathbb{N}$ in unary, and an encoding scheme of a finite field $\mathbb{F}$ of characteristic two.

Implicit input. A mapping $h\colon \mathbb{F}^m \to \mathbb{F}$.

Promise. The given encoding scheme is valid, and the mapping $h\colon \mathbb{F}^m \to \mathbb{F}$ is a polynomial function of degree at most d in each variable.

Question. Is

$$\sum_{\boldsymbol{x}\in\{0,1\}^m} h(\boldsymbol{x}) = 0 \quad (\text{in } \mathbb{F})? \tag{5.2}$$

In a (single-prover) interactive proof system for a problem with an implicit input, the implicit input is given to the verifier as an oracle.[3] The following variant of the summation test of Lund, Fortnow, Karloff and Nisan [80] is a special case of Lemma 3.5 in Ref. [12].

Lemma 36 (Summation test [12])**.** *Suppose that $|\mathbb{F}| \geq 2dm$. Then there exists a single-prover interactive proof system for the Summation Test Problem with perfect completeness and soundness error at most $dm/|\mathbb{F}|$. Moreover, in this interactive proof system, the verifier behaves as follows. First he chooses $\boldsymbol{q} \in \mathbb{F}^m$ uniformly at random. Then he interacts with the prover. At the same time, he reads the value $h(\boldsymbol{q})$ from the implicit input. Finally he accepts or rejects depending on $\boldsymbol{q}$, $h(\boldsymbol{q})$, and the interaction with the prover.*[4]

To apply the summation test to Problem 2, we have to consider exponentially many constraints instead of one.

Problem 4: AND Test Problem.
Explicit input. Integers $k, d \in \mathbb{N}$ in unary, and an encoding scheme of a finite field $\mathbb{F}$ of characteristic two.

Implicit input. A mapping $h\colon \mathbb{F}^k \to \mathbb{F}$.

Promise. The given encoding scheme is valid, and the mapping $h\colon \mathbb{F}^k \to \mathbb{F}$ is a polynomial function of degree at most d in each variable.

Question. Is $h(\boldsymbol{i}) = 0$ (in $\mathbb{F}$) for all $\boldsymbol{i} \in \{0,1\}^k$?

The idea for the following corollary is already explained in Section 7.1 of Ref. [12]. We will give a proof in Appendix A.2.1 for the sake of completeness.

[3] In Ref. [12], the authors refer to the interactive proof system for the Summation Test Problem as an "interactive oracle-protocol," viewing the mapping h as an exponentially long certificate string which is given to the verifier as an oracle. However, for our purposes it will be more convenient to treat h as part of the input.

[4] In particular, this implies that the verifier reads only one value $h(\boldsymbol{q})$ from the implicit input and the position $\boldsymbol{q} \in \mathbb{F}^m$ to read is chosen uniformly in $\mathbb{F}^m$. Together with the soundness guarantee, this in turn implies that if the implicit input is δ-close to a polynomial function $\tilde{h}$ of degree at most d in each variable and $\tilde{h}$ fails to satisfy the equation (5.2), then the verifier accepts with probability at most $\delta + dm/|\mathbb{F}|$ no matter what the prover does.

Corollary 37. *There exists a polynomial* $p\colon \mathbb{N}\times\mathbb{N}\to\mathbb{N}$ *for which the following holds. There exists a single-prover interactive proof system for the AND Test Problem with perfect completeness and soundness error at most* $5/8 + p(k,d)/|\mathbb{F}|$. *Moreover, in this interactive proof system, the verifier behaves as follows. First he chooses* $\boldsymbol{i}\in\mathbb{F}^k$ *uniformly and independently at random. Then he interacts with the prover. At the same time, he reads the value* $h(\boldsymbol{i})$ *from the implicit input. Finally he accepts or rejects depending on* $\boldsymbol{i}$, $h(\boldsymbol{i})$, *and the interaction with the prover.*

5.3 Protocol

In order to prove Theorem 35 we construct a four-prover poly-round proof system for Problem 2 which has perfect completeness with classical provers and soundness error at most $1-1/\text{poly}$ with entangled provers. Our protocol follows that of [12] very closely. We replace the three calls made in their protocol to the Oracle by calls to three distinct provers.

Label the provers as P, X_1, X_2, X_3. The protocol will be symmetric under any permutation of the three provers X_1, X_2, X_3. Let (r,s,f) be an instance of Problem 2, as described in Section 5.2.1. Let d_f be the maximum degree of f in any one variable. Let $m = r+3s$ and $d = 2d_f$. Let N be the smallest power of two such that $N > 8p(d,m)$, where p is the polynomial appearing in the statement of Corollary 37. Let $\mathbb{F}$ be the finite field of size N.

The verifier first receives an encoding scheme for $\mathbb{F}$ and its certificate from P, and rejects if it is not valid. In the rest of the protocol, all arithmetic operations in $\mathbb{F}$ are performed using this encoding scheme. If it is valid, the verifier proceeds to one of the following three tests chosen uniformly at random:

- *Consistency test.* He chooses $\boldsymbol{x}\in\mathbb{F}^s$ uniformly at random and sends the same question $\boldsymbol{x}$ to the provers X_1, X_2, X_3. He expects each prover to answer with an element of $\mathbb{F}$, and accepts if and only if all the answers are equal.

- *Linearity test.* He chooses $i\in\{1,\dots,s\}$, $\boldsymbol{x}\in\mathbb{F}^s$ and $y_i, z_i\in\mathbb{F}$ uniformly at random, and sets $y_j = z_j = x_j$ for every $j\in\{1,\dots,s\}\setminus\{i\}$. He sends $\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}$ to X_1, X_2, X_3, in random order. He receives integers a, b, c, and accepts if and only if

$$\frac{b-a}{y_i-x_i} = \frac{c-b}{z_i-y_i} = \frac{c-a}{z_i-x_i}.$$

- *Summation test.* The verifier simulates the interactive proof system in Corollary 37 with the explicit input (m,d) and P. When the verifier in Corollary 37 tries to read the value $h(\boldsymbol{z},\boldsymbol{b}_1,\boldsymbol{b}_2,\boldsymbol{b}_3)$ in the implicit input, where $\boldsymbol{z}\in\mathbb{F}^r$ and $\boldsymbol{b}_1,\boldsymbol{b}_2,\boldsymbol{b}_3\in\mathbb{F}^s$, our verifier simulates this by sending $\boldsymbol{b}_1,\boldsymbol{b}_2,\boldsymbol{b}_3$ to X_1, X_2, X_3. Upon obtaining answers a_1, a_2, a_3 to his queries, he evaluates $f(\boldsymbol{z},\boldsymbol{b}_1,\boldsymbol{b}_2,\boldsymbol{b}_3,a_1,a_2,a_3)$ and uses the result as the value of $h(\boldsymbol{z},\boldsymbol{b}_1,\boldsymbol{b}_2,\boldsymbol{b}_3)$.

Note that in each of the three tests, each of the X provers is asked a question $\boldsymbol{x} \in \mathbb{F}^s$ distributed uniformly at random.

5.3.1 Completeness

To prove completeness, let (r, s, f) be a yes-instance of Problem 2. Then there exists a mapping $A\colon \{0,1\}^s \to \{0,1\}$ such that Eq. (5.1) is satisfied for all $\boldsymbol{z} \in \{0,1\}^r$ and all $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3 \in \{0,1\}^s$ simultaneously. Let g be the unique extension of A to a multilinear function $g\colon \mathbb{F}^s \to \mathbb{F}$. Each of X_1, X_2, X_3 answers $g(\boldsymbol{b})$ on question $\boldsymbol{b} \in \mathbb{F}^s$, while P behaves as it should in the AND test. Then it is clear that this deterministic strategy is accepted with certainty in the consistency test and the linearity test. In the summation test, note that the value of $h(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3)$ which the verifier uses is given by

$$h(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3) = f(\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3, g(\boldsymbol{b}_1), g(\boldsymbol{b}_2), g(\boldsymbol{b}_3)),$$

which is a polynomial in $\boldsymbol{z}, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3$ of degree at most $2d_f = d$ in each variable. Therefore, the promise in the AND test is satisfied and prover P has a strategy which makes the verifier accept with certainty.

5.3.2 Soundness

The soundness analysis is divided in two parts. First we analyze the consistency and linearity tests, and show that success in those tests implies the following. (We refer the reader to Section 5.2 for some relevant notation and definitions.)

Theorem 38. *There exist universal constants $0 < c_0, c_1 < 1$, $C_1 > 1$ such that the following holds. Let $\mathbb{F}$ be a finite field. Suppose $|\Psi\rangle$ and $\{A^a_{\boldsymbol{x}}\}_a$ form a symmetric strategy which passes both the consistency and the linearity tests with probability $1-\varepsilon$. Assume furthermore that $|\mathbb{F}|^{-1} \le \varepsilon$ and $\varepsilon^{c_0/2} \le s^{-1}$. Then there exists a POVM $\{V^g\}$, indexed by multilinear $g\colon \mathbb{F}^s \to \mathbb{F}$, such that*

$$\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big((A^a_{\boldsymbol{x}} - V^a_{\boldsymbol{x}})^2\big) \;\le\; C_1\, \varepsilon^{c_1}, \tag{5.3}$$

where we defined $V^a_{\boldsymbol{x}} := \sum_{g\colon g(\boldsymbol{x})=a} V^g$.

Theorem 38 is proved in Section 5.4. Assuming the theorem, we prove that our proof system has soundness error at most

$$1 - 3\, s^{-c_0/2},$$

provided s is larger than an absolute constant depending on c_1 and C_1. Let (r, s, f) be a no-instance. Suppose that the provers have a symmetric[5] entangled strategy S whose acceptance

[5] This is without loss of generality by Lemma 13.

probability is $1-\varepsilon/3$. Let $|\Psi\rangle \in \mathcal{P} \otimes \mathcal{X}_1 \otimes \mathcal{X}_2 \otimes \mathcal{X}_3$ be the state used in the strategy S and $(A_{\boldsymbol{x}}^a)_{a\in\mathbb{F}}$ be the projective measurements used by each X prover in the strategy S.

The verifier can be viewed as performing the multilinearity game with the players X_1, X_2, X_3 with probability 2/3 and performing something else (the summation test) with probability 1/3. Therefore, the marginal strategy $S|_{X_1,X_2,X_3}$ has winning probability at least $1-\varepsilon/2 \geq 1-\varepsilon$ in the multilinearity test. By Theorem 38, there exists a POVM $\left\{V^g\right\}_{g\in\mathrm{ML}(\mathbb{F}^s,\mathbb{F})}$ such that inequality (5.3) holds, where ρ is the reduced state of $|\Psi\rangle\langle\Psi|$ on $\mathcal{X}_1$. Let

$$V_{\boldsymbol{x}}^a = \sum_{\substack{g\in\mathrm{ML}(\mathbb{F}^s,\mathbb{F})\\ g(\boldsymbol{x})=a}} V^g.$$

For $0 \leq i \leq 3$, let S_i be the entangled strategy obtained from S by replacing the POVM for the first i provers $X_1, \dots, X_i$ by $V_{\boldsymbol{x}}^a$. Note $S_0 = S$. In the strategy S_i, the provers $X_1, \dots, X_i$ can be implemented so that they measure the prior entanglement without looking at their question. In particular, in the strategy S_3, every prover except for P measures the prior entanglement without looking at the question, and therefore S_3 can be implemented by shared randomness.

For $0 \leq i \leq 3$, let p_i be the probability that the strategy S_i is accepted in the four-prover protocol. By definition, $p_0 = 1-\varepsilon/3$. We prove the following.

Claim 39. *For $i = 1, \dots, 3$, it holds that $|p_{i-1} - p_i| \leq \sqrt{C_1}\varepsilon^{c_1/2}$.*

Proof. The only difference between the strategies S_{i-1} and S_i is the measurements used by the prover X_i. We call the message from the verifier to X_i as register $\mathcal{A}$, and call everything other than $\mathcal{A}$ and the private space $\mathcal{X}_i$ for prover X_i as register $\mathcal{B}$. Register $\mathcal{A}$ is classical, but we treat it as a quantum register which always contains a state in the computational basis. Let σ be the global state before the prover X_i performs his measurement, and σ_A (resp. σ_V) be the global state after the prover X_i performs the measurement $A_{\boldsymbol{x}}$ (resp. V) on his share of the state, and then discards the post-measurement state. Since the marginal distribution on the question to X_i is uniform, the state σ has the following form:

$$\sigma = \mathrm{E}_{\boldsymbol{x}\in\mathbb{F}^s} |\boldsymbol{x}\rangle\langle\boldsymbol{x}|_{\mathcal{A}} \otimes \sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}},$$

where $\mathrm{Tr}_{\mathcal{B}}\sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}} = \sigma^{\mathcal{X}_i} = \rho$ is independent of $\boldsymbol{x}$. We want to bound $(1/2)\|\sigma_W - \sigma_M\|_1$, where

$$\sigma_W = \mathrm{Tr}_{\mathcal{X}_i}\left[\mathrm{E}_{\boldsymbol{x}\in\mathbb{F}^s} |\boldsymbol{x}\rangle\langle\boldsymbol{x}|_{\mathcal{A}} \otimes \sum_{a\in\mathbb{F}} |a\rangle\langle a|_{\mathcal{C}} \otimes (A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}})\sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}}(A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}})\right],$$

$$\sigma_M = \mathrm{Tr}_{\mathcal{X}_i}\left[\mathrm{E}_{\boldsymbol{x}\in\mathbb{F}^s} |\boldsymbol{x}\rangle\langle\boldsymbol{x}|_{\mathcal{A}} \otimes \sum_{a\in\mathbb{F}} |a\rangle\langle a|_{\mathcal{C}} \otimes (\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}})\sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}}(\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}})\right],$$

and $\mathcal{C}$ denotes the register used for prover i's answers. For $\boldsymbol{x} \in \mathbb{F}^s$, define isometries $U_{\boldsymbol{x}}, V_{\boldsymbol{x}} \colon \mathcal{X}_i \otimes \mathcal{B} \to \mathcal{X}_i \otimes \mathcal{B} \otimes \mathcal{C}$ by

$$\begin{aligned} U_{\boldsymbol{x}} &= \sum_{a \in \mathbb{F}} A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}} \otimes |a\rangle_{\mathcal{C}}, \\ V_{\boldsymbol{x}} &= \sum_{a \in \mathbb{F}} \sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}} \otimes |a\rangle_{\mathcal{C}}. \end{aligned}$$

Then,

$$\begin{aligned} &\|\rho_W - \rho_M\|_1 \\ &\leq \|[\| \Bigg] \mathrm{E}_{\boldsymbol{x} \in \mathbb{F}^s} |\boldsymbol{x}\rangle\langle \boldsymbol{x}|_{\mathcal{A}} \otimes \sum_{a \in \mathbb{F}} |a\rangle\langle a|_{\mathcal{C}} \otimes \Big((A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}}) \sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}} (A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}}) - (\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}}) \sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}} (\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}}) \Big)_1 \\ &\leq \mathrm{E}_{\boldsymbol{x} \in \mathbb{F}^s} \|[\| \Bigg] \sum_{a \in \mathbb{F}} |a\rangle\langle a|_{\mathcal{C}} \otimes \Big((A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}}) \sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}} (A_{\boldsymbol{x}}^a \otimes I_{\mathcal{B}}) - (\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}}) \sigma_{\boldsymbol{x}}^{\mathcal{X}_i\mathcal{B}} (\sqrt{V_{\boldsymbol{x}}^a} \otimes I_{\mathcal{B}}) \Big)_1 \\ &\leq 2\mathrm{E}_{\boldsymbol{x} \in \mathbb{F}^s} \sqrt{\sum_{a \in \mathbb{F}} \mathrm{Tr}\big((A_{\boldsymbol{x}}^a - \sqrt{V_{\boldsymbol{x}}^a})^2 \rho \big)} \\ &\leq 2 \sqrt{\mathrm{E}_{\boldsymbol{x} \in \mathbb{F}^s} \sum_{a \in \mathbb{F}} \mathrm{Tr}\big((A_{\boldsymbol{x}}^a - \sqrt{V_{\boldsymbol{x}}^a})^2 \rho \big)} \\ &\leq 2\sqrt{C_1 \varepsilon^{c_1}}, \end{aligned}$$

where the third inequality is by Lemma 117, the fourth is by convexity and the last by (5.3). Therefore, we have that $|p_{i-1} - p_i| \leq \sqrt{C_1}\varepsilon^{c_1/2}$ as claimed. □

By the triangle inequality, Claim 39 implies that $|p_0 - p_3| \leq 3\sqrt{C_1}\varepsilon^{c_1/2}$, and therefore

$$p_3 \geq p_0 - 3\sqrt{C_1}\varepsilon^{c_1/2} = 1 - \varepsilon - 3\sqrt{C_1}\varepsilon^{c_1/2} \geq 1 - 4\sqrt{C_1}\varepsilon^{c_1/2},$$

where the last inequality uses $c_1 \leq 1$.

If not all of the provers choose the same multilinear function, then they pass in the consistency test with probability at most $s/|\mathbb{F}| \leq 1/6$ by Schwartz's lemma [103]. In the strategy S_3, they pass in the consistency test with probability at least $1 - 21\sqrt{C_1}\varepsilon^{c_1/2}$. Therefore, they choose the same multilinear function with probability at least $1 - 21\sqrt{C_1}\varepsilon^{c_1/2}/(1 - 1/6) \geq 1 - 26\sqrt{C_1}\varepsilon^{c_1/2}$. This implies that if an oracle chooses a multilinear function in the same way as the prover X_1 and use it for all three queries, the distribution on their answers will differ by at most $52\sqrt{C_1}\varepsilon^{c_1/2}$. Therefore, this oracle (which always implements a multilinear function) together with the prover P is accepted in the interactive proof system of Corollary 37 with probability at least $1 - 21\sqrt{C_1}\varepsilon^{c_1/2} - 52\sqrt{C_1}\varepsilon^{c_1/2} = 1 - 73\sqrt{C_1}\varepsilon^{c_1/2}$.

Because of our choice of the finite field $\mathbb{F}$, the acceptance probability in the interactive proof system of Corollary 37 is less than 3/4. Comparing this with the lower bound in the previous paragraph, we obtain

$$1 - 73\sqrt{C_1}\varepsilon^{c_1/2} < \frac{3}{4},$$

which implies

$$\varepsilon > \frac{1}{(292^2 \cdot C_1)^{1/c_1}},$$

contradicting our assumption that $\varepsilon \leq s^{-c_0/2}$ as soon as s is large enough. Since we obtained this lower bound on ε from the assumption that there exists an entangled strategy with acceptance probability $1 - \varepsilon/3$, we have proved the claimed soundness guarantee against entangled provers.

5.4 Analysis of the multilinearity game

In this section we analyze the combination of the consistency test and the linearity test described in Section 5.3 as a stand-alone game played by three players. For convenience, we restate the test here. The game is parametrized by two integers n and p, where p is a power of 2 and $\mathbb{F}$ a finite field of order $|\mathbb{F}| = p$,[6] and is performed with three players X_1, X_2, X_3 treated symmetrically. The referee performs the following two tests with probability 1/2 each:

- *Consistency test.* The referee chooses $\boldsymbol{x} \in \mathbb{F}^n$ uniformly at random and sends the same question $\boldsymbol{x}$ to the players X_1, X_2, X_3. He expects each player to answer with an element of $\mathbb{F}$, and accepts if and only if all the answers are equal.
- *Linearity test.* The referee chooses $i \in \{1, \dots, s\}$, $\boldsymbol{x} \in \mathbb{F}^n$ and $y_i, z_i \in \mathbb{F}$ uniformly at random, and sets $y_j = z_j = x_j$ for every $j \in \{1, \dots, n\} \setminus \{i\}$. He sends $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ to a random permutation of the players, receives integers a, b, c, and accepts if and only if
$$\frac{b-a}{y_i - x_i} = \frac{c-b}{z_i - y_i} = \frac{c-a}{z_i - x_i}.$$

The following definition will be useful in our analysis.

Definition 40. *Let $\big\{T^g_{\boldsymbol{x}_{\geq k}}\big\}_{g \in \mathrm{ML}(\mathbb{F}^{k-1},\mathbb{F})}$ and $\big\{V^h_{\boldsymbol{x}_{\geq \ell}}\big\}_{h \in \mathrm{ML}(\mathbb{F}^{\ell-1},\mathbb{F})}$ be two families of (possibly incomplete) POVMs. Let $\delta > 0$. We say that T and V are* δ-consistent *if the following holds:*

$$\mu(T,V) := \mathrm{E}_{\boldsymbol{x}} \sum_{g,h:\, g(\boldsymbol{x}_{<k}) \neq h(\boldsymbol{x}_{<\ell})} \mathrm{Tr}_\rho\big(T^g_{\boldsymbol{x}_{\geq k}} \otimes V^h_{\boldsymbol{x}_{\geq \ell}}\big) \leq \delta.$$

[6] The dimension parameter n was called s, and p was called N, in the previous section.

We also let $\mu(T) := \mu(T,T)$: $\mu(T)$ measures the probability that two distinct provers simultaneously measuring their share of the entangled state using $\{T^g_{\boldsymbol{x}_{\geq k}}\}$ obtain functions g, g' that differ at (a random) $\boldsymbol{x}_{<k}$. Finally, by a slight abuse of notation we write

$$\mathrm{Tr}_\rho(T) := \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(T^g_{\boldsymbol{x}_{\geq k}}\big).$$

The main result of this section is the following.

Theorem 41. *There exists universal constants $0 < c_0 < 1$, $C_0 > 1$ such that the following holds. Suppose $|\Psi\rangle$ and $\{A^a_{\boldsymbol{x}}\}_a$ form a permutation-invariant strategy which passes both the consistency and the linearity tests with probability $1 - \varepsilon$. Assume furthermore that $p = |\mathbb{F}| \geq \varepsilon^{-1}$ and $\varepsilon^{c_0/2} \leq n^{-1}$. Then there exists an incomplete POVM $\{V^h\}_{h \in \mathrm{ML}(\mathbb{F}^n, \mathbb{F})}$, indexed by multilinear $h : \mathbb{F}^n \to \mathbb{F}$, such that*

1. *V is $O(\varepsilon^{c_0})$-consistent with A,*
2. *$\mathrm{Tr}_\rho(V) \geq 1 - C_0\,\varepsilon^{c_0}$.*

The first condition in the theorem intuitively guarantees that, if a prover measures according to V, obtaining a multilinear function h as outcome, then the value of that function at $\boldsymbol{x}$ will correspond to the answer a that another prover applying the original POVM $\{A^a_{\boldsymbol{x}}\}$ would have obtained. The second condition states that the POVM V is “not too incomplete”.

We will show how Theorem 41 implies Theorem 38 in Section 5.4.2, while Theorem 41 will be proved in Section 5.5. In the following section we show a weaker version of the multilinearity test, the “linearity test”, which implies Theorem 41 for $n = 1$. While that claim avoids a lot of the difficulties of the overall proof, it demonstrates some of its basic ideas and techniques, and will be used as an important building block in the proof of the general case.

5.4.1 Preliminary analysis

As an immediate consequence of $\{A^a_{\boldsymbol{x}}\}$ succeeding in the multilinearity test, we get that the following relations hold:

$$\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big(A^a_{\boldsymbol{x}} \otimes A^a_{\boldsymbol{x}}\big) \geq 1 - \varepsilon, \quad (5.4)$$

$$\forall i \in [n], \quad \mathop{\mathrm{E}}_{x_i, x_i', x_i'' \boldsymbol{x}_{\neg i}} \sum_{\frac{a'-a}{x_i'-x_i} = \frac{a''-a'}{x_i''-x_i'} = \frac{a''-a}{x_i''-x_i}} \mathrm{Tr}_\rho\big(A^a_{x_i, \boldsymbol{x}_{\neg i}} \otimes A^{a'}_{x_i', \boldsymbol{x}_{\neg i}} \otimes A^{a''}_{x_i'', \boldsymbol{x}_{\neg i}}\big) \geq 1 - n\varepsilon$$
$$\geq 1 - \sqrt{\varepsilon}, \quad (5.5)$$

where the last inequality follows from our assumption that $n \leq \varepsilon^{-c_0/2} \leq \varepsilon^{-1/2}$. We note that here as elsewhere we abuse notation and use ρ to denote the prover’s density matrix on all three provers' systems, any two of the provers', or any one of them. Which will always be clear from the context (recall that ρ is permutation-invariant).

The following claim proves the "linearity" part of the multilinearity test. It generalizes Theorem 4 from Chapter 2 to field sizes $p > 2$.

Claim 42. *Suppose that $\{A_{\boldsymbol{x}}^a\}$ is a projective strategy passing the consistency test with probability $1-\varepsilon$, and the linearity test in the i-th direction with success $1-\sqrt{\varepsilon}$. Then there exists a POVM $\big\{B_{\boldsymbol{x}_{\neg i}}^{u,v}\big\}_{(u,v)\in\mathbb{F}^2}$ such that*

$$\mathrm{E}_{\boldsymbol{x}}\sum_a \Big\|A_{\boldsymbol{x}}^a - \sum_{u,v:ux_i+v=a} B_{\boldsymbol{x}_{\neg i}}^{u,v}\Big\|_\rho^2 = O\big(\sqrt{\varepsilon}\big).$$

Proof. Define

$$B_{\boldsymbol{x}_{\neg i}}^{u,v} := \mathrm{E}_{x_i\neq x_i'} A_{x_i,\boldsymbol{x}_{\neg i}}^{ux_i+v} A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} A_{x_i,\boldsymbol{x}_{\neg i}}^{ux_i+v}.$$

Then $\big\{B_{\boldsymbol{x}_{\neg i}}^{u,v}\big\}_{u,v}$ is a well-defined POVM: for fixed $x_i \neq x_i'$, as (u,v) ranges over $\mathbb{F}^2$ both ux_i+v and $ux_i'+v$ independently range over $\mathbb{F}$. Using the definition of the norm $\|\cdot\|_\rho$, we can expand

$$\begin{aligned}\mathrm{E}_{\boldsymbol{x}}\sum_a \Big\|A_{\boldsymbol{x}}^a - \sum_{u,v:ux_i+v=a} B_{\boldsymbol{x}_{\neg i}}^{u,v}\Big\|_\rho^2 &= \mathrm{E}_{\boldsymbol{x}}\sum_a \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a\big) + \mathrm{E}_{\boldsymbol{x}} \sum_{\substack{(u,v),\,(u',v')\\ ux_i+v=u'x_i+v'}} \mathrm{Tr}_\rho\big(B_{\boldsymbol{x}_{\neg i}}^{u,v} B_{\boldsymbol{x}_{\neg i}}^{u',v'}\big)\\ &\quad - 2\,\mathrm{E}_{\boldsymbol{x}} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a B_{\boldsymbol{x}_{\neg i}}^{u,v}\big).\end{aligned} \tag{5.6}$$

By Lemma 120, using the consistency of $\{A_{\boldsymbol{x}}^a\}$ (and letting $B_{\boldsymbol{x}}^a := \sum_{(u,v):\,ux_i+v=a} B_{\boldsymbol{x}_{\neg i}}^{u,v}$ play the role of B_i in that lemma), in order to lower-bound the last term it will suffice to lower-bound $\mathrm{E}_{\boldsymbol{x}} \sum_{(u,v):ux_i+v=a} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes B_{\boldsymbol{x}_{\neg i}}^{u,v}\big)$. Using the definition of $B_{\boldsymbol{x}_{\neg i}}^{u,v}$, we have

$$\begin{aligned}&\mathrm{E}_{\boldsymbol{x}} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes B_{\boldsymbol{x}_{\neg i}}^{u,v}\big)\\ &\qquad = \mathrm{E}_{\boldsymbol{x},x_i'\neq x_i''} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} A_{x_i'',\boldsymbol{x}_{\neg i}}^{ux_i''+v} A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v}\big)\\ &\qquad = \mathrm{E}_{\boldsymbol{x},x_i'\neq x_i''} \sum_{a,u,v:ux_i+v=a} \sum_{a'} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} A_{x_i'',\boldsymbol{x}_{\neg i}}^{ux_i''+v} A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{a'}\big)\\ &\qquad \leq \mathrm{E}_{\boldsymbol{x},x_i'\neq x_i''} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} A_{x_i'',\boldsymbol{x}_{\neg i}}^{ux_i''+v} A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v} \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v}\big) + \varepsilon\\ &\qquad \leq \mathrm{E}_{\boldsymbol{x},x_i'\neq x_i''} \sum_{a,u,v:ux_i+v=a} \sum_{a'} \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{a'} A_{x_i'',\boldsymbol{x}_{\neg i}}^{ux_i''+v} A_{x_i',\boldsymbol{x}_{\neg i}}^{a'} \otimes A_{x_i',\boldsymbol{x}_{\neg i}}^{ux_i'+v}\big) + 2\,\varepsilon\end{aligned} \tag{5.7}$$

where the first equality simply uses that the $A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}}$ sum to identity over a', and the two inequalities both use (5.4) on the last two registers (together with $A^a_{\boldsymbol{x}} \leq \mathrm{Id}$). Moreover, by Claim 119, together with (5.4), we know that

$$\mathrm{E}_{\boldsymbol{x}} \Big\| \sum_a \big(A^a_{\boldsymbol{x}} \otimes \mathrm{Id} \otimes \mathrm{Id}\big)\, \rho\, \big(A^a_{\boldsymbol{x}} \otimes \mathrm{Id} \otimes \mathrm{Id}\big) - \rho \Big\|_1 = O\big(\sqrt{\varepsilon}\big),$$

where we used that the $A^a_{\boldsymbol{x}}$ are projectors. Hence

$$\begin{aligned}
&\mathrm{E}_{\boldsymbol{x},x'_i \neq x''_i} \sum_{a,u,v:ux_i+v=a} \sum_{a'} \mathrm{Tr}_\rho\big(A^a_{\boldsymbol{x}} \otimes \big(A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}} A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} - A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}}\big) \otimes A^{ux'_i+v}_{x'_i,\boldsymbol{x}_{\neg i}}\big) \\
&\quad = \mathrm{E}_{\boldsymbol{x},x'_i \neq x''_i} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}\Big(\big(A^a_{\boldsymbol{x}} \otimes A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}} \otimes A^{ux'_i+v}_{x'_i,\boldsymbol{x}_{\neg i}}\big) \cdot \\
&\qquad\qquad \Big(\sum_{a'} \big(\mathrm{Id} \otimes A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} \otimes \mathrm{Id}\big)\, \rho\, \big(\mathrm{Id} \otimes A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} \otimes \mathrm{Id}\big)\Big) - \rho\Big)\Big) \\
&\quad \leq \mathrm{E}_{\boldsymbol{x},x'_i} \big\| \sum_{a'} A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} \rho A^{a'}_{x'_i,\boldsymbol{x}_{\neg i}} - \rho \big\|_1 = O\big(\sqrt{\varepsilon}\big),
\end{aligned}$$

where for the inequality we used that for every $\boldsymbol{x}$, $\sum_{a,u,v:ux_i+v=a} A^a_{\boldsymbol{x}} \otimes A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}} \otimes A^{ux'_i+v}_{x'_i,\boldsymbol{x}_{\neg i}} \leq \mathrm{Id}$, and monotonicity of the trace distance. Combining this last bound with (5.7), we obtain

$$\begin{aligned}
\mathrm{E}_{\boldsymbol{x}} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A^a_{\boldsymbol{x}} \otimes B^{u,v}_{\boldsymbol{x}_{\neg i}}\big) &= \mathrm{E}_{\boldsymbol{x},x'_i \neq x''_i} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A^a_{\boldsymbol{x}} \otimes A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}} \otimes A^{ux'_i+v}_{x'_i,\boldsymbol{x}_{\neg i}}\big) + O\big(\sqrt{\varepsilon}\big) \\
&= \mathrm{E}_{\boldsymbol{x},x'_i \neq x''_i} \sum_{u,v} \mathrm{Tr}_\rho\big(A^{ux_i+v}_{\boldsymbol{x}} \otimes A^{ux''_i+v}_{x''_i,\boldsymbol{x}_{\neg i}} \otimes A^{ux'_i+v}_{x'_i,\boldsymbol{x}_{\neg i}}\big) + O\big(\sqrt{\varepsilon}\big).
\end{aligned}$$

Up to the missing $x'_i = x''_i$ terms, the last summation above is exactly the probability that the $A^a_{\boldsymbol{x}}$ pass the linearity test along the first coordinate, hence is at least $1 - \sqrt{\varepsilon}$ by (5.5). Terms in the expectation for which $x'_i = x''_i$ contribute at most $1/p \leq \varepsilon$, hence we have shown:

$$\mathrm{E}_{\boldsymbol{x}} \sum_{a,u,v:ux_i+v=a} \mathrm{Tr}_\rho\big(A^a_{\boldsymbol{x}} \otimes B^{u,v}_{\boldsymbol{x}_{\neg i}}\big) \geq 1 - O\big(\sqrt{\varepsilon}\big).$$

Since the first two terms in (5.6) are at most 1 each, this implies the bound in the claim. □

5.4.2 Proof of Theorem 38

In this section we show how Theorem 38, which is the result we need in order to analyze the overall protocol from Section 5.3, follows from Theorem 41. Theorem 41 is proved in Section 5.5.

Proof of Theorem 38. Given that A passes the consistency test, we have $\mu(A) \geq 1 - \varepsilon$ (cf. Eq. (5.4)) and hence

$$\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big((A_{\boldsymbol{x}}^a \otimes \mathrm{Id} - \mathrm{Id} \otimes A_{\boldsymbol{x}}^a)^2\big) \;=\; 2 - 2\sum_a \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes A_{\boldsymbol{x}}^a\big) \;\leq 2\,\varepsilon. \tag{5.8}$$

Let $\big\{V^h\big\}$ be the incomplete POVM guaranteed by Theorem 41. Expanding

$$\begin{aligned}\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big((A_{\boldsymbol{x}}^a - V_{\boldsymbol{x}}^a)^2\big) &= \mathrm{E}_{\boldsymbol{x}} \sum_a \Big(\mathrm{Tr}_\rho\big((A_{\boldsymbol{x}}^a)^2\big) + \mathrm{Tr}_\rho\big((V_{\boldsymbol{x}}^a)^2\big) - 2\,\mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a V_{\boldsymbol{x}}^a\big)\Big)\\ &\leq 2 - 2\,\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a V_{\boldsymbol{x}}^a\big),\end{aligned}$$

it will suffice to show that this last expectation is close to 1. By Lemma 120 and self-consistency of A as per (5.8), it suffices to lower-bound

$$\mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\big(A_{\boldsymbol{x}}^a \otimes V_{\boldsymbol{x}}^a\big) = \mathrm{Tr}_\rho(V) - \mu(V, A) \;\geq\; 1 - C_0\varepsilon^{c_0} - O\big(\varepsilon^{c_0}\big),$$

assuming the bounds claimed in Theorem 41. This proves Theorem 38. □

5.5 Proof of Theorem 41

In this section we prove our main result on the analysis of the multilinearity test in the presence of entanglement between the provers, Theorem 41. Its proof proceeds by induction, and is based on two lemmas. The first, an analogue of the "pasting lemma" from [12], lets us remove the dependence of the provers' strategy on the verifier's question one coordinate at a time.

Lemma 43 (Pasting lemma). *Let $\{A_{\boldsymbol{x}}^a\}_a$, where $\boldsymbol{x} \in \mathbb{F}^n$ and $a \in \mathbb{F}$, be a strategy that is accepted with probability at least $1 - \varepsilon$ in both the linearity test and the consistency test. Let $k \in [n]$ and $\big\{T_{\boldsymbol{x}_{\geq k}}^h\big\}_h$, where $h : \mathbb{F}^{k-1} \to \mathbb{F}$ is multilinear, be a (incomplete) POVM such that T satisfies properties 1.–3. in the conclusion of Lemma 44, for some $n^{-4} \geq \delta \geq \sqrt{\varepsilon}$, and assume $p \geq n^3$. Then there exists a (incomplete) POVM $\big\{V_{\boldsymbol{x}_{>k}}^g\big\}_g$, where $g : \mathbb{F}^k \to \mathbb{F}$, such that*

1. *V is consistent with A: $\mu(V, A) = O\big(\delta^{1/4}\big)$,*
2. *V has large trace: $\mathrm{E}_{\boldsymbol{x}_{>k}} \sum_g \mathrm{Tr}_\rho\big(V_{\boldsymbol{x}_{>k}}^g\big) \geq \mathrm{E}_{\boldsymbol{x}_{\geq k}} \sum_h \mathrm{Tr}_\rho\big(T_{\boldsymbol{x}_{\geq k}}^h\big) - O\big(\delta^{1/4}\big)$.*

We note that in this lemma we keep very little information on how V relates to T, but rather connect V directly to A through the consistency property. Indeed, it can be thought of as a weaker variant of Claim 42, but one that we will be able to carry out throughout the induction.

The second lemma is the quantum analogue of the "self-improvement lemma" of [12]. It shows that, if a strategy $\{A_{\boldsymbol{x}}^a\}$ is weakly consistent with a multilinear POVM $\{R_{\boldsymbol{x}_{\geq k}}^g\}$, *and* it passes the consistency and linearity tests with high probability, then it must be highly consistent with an "improved" multilinear POVM $\{T_{\boldsymbol{x}_{\geq k}}^g\}$. (The extra conditions in the conclusion of the lemma are not ultimately needed, but are helpful to carry out the induction procedure.)

Lemma 44 (Self-improvement lemma)**.** *There exists a universal constant* $0 < c < 1$ *such that the following holds. For every* $\boldsymbol{x} \in \mathbb{F}^n$*, let* $\{A_{\boldsymbol{x}}^a\}_a$ *be a POVM indexed by* $a \in \mathbb{F}$*. Let* $1/n \geq \delta \geq \sqrt{\varepsilon} > 0$ *be such that the following hold:*

1. $\{A_{\boldsymbol{x}}^a\}_{\boldsymbol{x}}$ *passes both tests with probability* $1 - \varepsilon$*,*
2. *There exists a (incomplete) POVM* $\{R_{\boldsymbol{x}_{\geq k}}^g\}_g$*, indexed by multilinear* $g : \mathbb{F}^{k-1} \to \mathbb{F}$*, such that* $\mu(R, A) \leq \delta$*.*

Then there exists a (incomplete) POVM $\{T_{\boldsymbol{x}_{\geq k}}^g\}_g$*, where* $g : \mathbb{F}^{k-1} \to \mathbb{F}$*, such that* T *has a factorization* $T_{\boldsymbol{x}_{\geq k}}^g = \hat{T}_{\boldsymbol{x}_{\geq k}}^g \big(\hat{T}_{\boldsymbol{x}_{\geq k}}^g\big)^\dagger$*, and for every* $\boldsymbol{x}$ *there is a* $S_{\boldsymbol{x}}^g = \hat{S}_{\boldsymbol{x}}^g \big(\hat{S}_{\boldsymbol{x}}^g\big)^\dagger$ *such that the following hold:*

1. *For every* $\boldsymbol{x}$ *and* a*,* $\sum_{g:g(\boldsymbol{x}_{<k})=a} S_{\boldsymbol{x}}^g \leq A_{\boldsymbol{x}}^a$*,*
2. $\mathrm{E}_{\boldsymbol{x}} \sum_g \big\|\hat{T}_{\boldsymbol{x}_{\geq k}}^g - \hat{S}_{\boldsymbol{x}}^g\big\|_\rho^2 = O(\varepsilon^c)$*,*
3. $\mu(T, A) = O(\varepsilon^c)$*,*
4. $\mathrm{Tr}_\rho(T) \geq \mathrm{Tr}_\rho(R) - O\big(\delta^{1/4}\big)$*.*

Based on these two lemmas, we can give a proof by induction of Theorem 41.

Proof of Theorem 41. We show the following by induction on $0 \leq k \leq n$:

> IH(k): There exists universal constants $0 < c < 1$, $C, C' > 1$, and a POVM $\{V_{\boldsymbol{x}_{>k}}^h\}$, indexed by multilinear $h : \mathbb{F}^k \to \mathbb{F}$, such that
>
> 1. V is $C'\varepsilon^{c/4}$-consistent with A,
> 2. $\mathrm{Tr}_\rho(V) \geq 1 - C(k+1)\,\varepsilon^{c/8}$.

IH(0) is the base case, and is trivial: just set $V = A$; properties 1 and 2 both hold as a consequence of A passing the tests. IH(n) implies the theorem.

Suppose that IH($k-1$) is true, prove IH(k). Let $\big\{R_{\boldsymbol{x}_{\geq k}}^h\big\}$ be the POVM guaranteed by the induction hypothesis, where $h : \mathbb{F}^{k-1} \to \mathbb{F}$ is multilinear. First apply Lemma 44 to A

and R. Let C'' be an upper bound on all constants implied by the $O(\cdot)$ in the conclusion of the lemma. We obtain a POVM $\left\{T^g_{\boldsymbol{x}\geq k}\right\}_g$ such that $\mathrm{Tr}_\rho(T) \geq \mathrm{Tr}_\rho(R) - C''(C'\varepsilon^{c/4})^{1/4}$, and T satisfies the other three conclusions of Lemma 44, with an error $\delta = C''\varepsilon^c$.

Apply Lemma 43 to T. Let C''' be an upper bound on all constants implied by the $O(\cdot)$ notation in the conclusion of the lemma. We obtain a POVM $\left\{V^g_{\boldsymbol{x}>k}\right\}$, where now $g : \mathbb{F}^k \to \mathbb{F}$ is multilinear, such that V is $C'''(C''\varepsilon^c)^{1/4}$-consistent with A and $\mathrm{Tr}_\rho(R) \geq \mathrm{Tr}_\rho(V) - C'''(C''\varepsilon^c)^{1/4}$. This proves IH($k$) provided C, C' are chosen such that $C' \geq C'''(C'')^{1/4}$ and $C \geq C'''(C'')^{1/4} + C''(C')^{1/4}$. □

5.6 The self-improvement lemma

In this section we prove Lemma 44, which shows that the consistency with A of a given (incomplete) POVM $\left\{R^g_{\boldsymbol{x}\geq k}\right\}_g$ can be improved, while not increasing its "incompleteness", as measured by $\mathrm{Tr}_\rho(R)$, too much. Before proceeding with the details, we give some intuition and a high-level overview of how we will proceed.

Think about the following simplified situation in two dimensions. Consider also the case where $p = 2$, so that the prover's answers are simply bits. For every $\boldsymbol{x} \in \mathbb{F}^2$ we are given a binary projective measurement $(A^0_{\boldsymbol{x}}, A^1_{\boldsymbol{x}})$: picture two orthogonal "planes" of dimension $d/2$ each. Our goal is to find a global "refinement" of these planes: a single projective measurement $\{T^g\}$, with outcomes in the set of bilinear functions $g : \mathbb{F}^2 \to \mathbb{F}$, such that at every $\boldsymbol{x}$ the approximation $A^a_{\boldsymbol{x}} \approx_\varepsilon \sum_{g:\, g(\boldsymbol{x})=a} T^g$ holds.[7] In order to achieve this, we start from two assumptions:

1. There exists another measurement $\{R^g\}$ which achieves an approximation of weaker quality, up to some $\delta \gg \varepsilon$, than the one we are looking for,

2. The $A^a_{\boldsymbol{x}}$ are very close to *linear*: for every line $(x_1, \cdot)$ (resp. $(\cdot, x_2)$) there is measurement $B^\ell_{x_1}$ with outcomes in the set of linear functions $\mathbb{F} \to \mathbb{F}$ such that $A^a_{(x_1,x_2)} \approx_\varepsilon \sum_{\ell:\, \ell(x_2)=a} B^\ell_{x_1}$.

The goal is to use the high quality of the approximation along lines to improve the quality of the overall "bilinear" approximation. Let's trust that an ideal measurement $\{T^g\}$, achieving an approximation of order ε, exists. How can $\{R^g\}$ differ from this ideal measurement? We may think of $\{R^g\}$ as a perturbation of $\{T^g\}$. There are two ways $\{T^g\}$ can be perturbed: the first is by applying an arbitrary (but not too large) rotation (or, in general, unitary transformation) on the whole space. The second is by "mis-labeling" some of the POVM elements: e.g. for some g, a subspace of the space on which R^g projects could have been labeled as a subspace of $T^{g'}$ for some $g' \neq g$.

[7] At this point we are being purposely vague as to how the approximation is measured. It will be done in a relatively weak sense, through the consistency $\mu(T, A)$.

These are the two main types of errors to keep in mind. We devise a procedure which "corrects" the first type of error, but not the second: indeed, if some parts of the measurement $\{R^g\}$ are mis-labeled, and hence produce outcomes which agree very poorly with the original $\{A^a_{\boldsymbol{x}}\}$, there is no generic way to recover from such errors. This type of error is unique to the quantum setting, and is the main reason why the measurements we construct "shrink" at every step of the induction: any mislabeled portions of space will have to be thrown out. Since we cannot recover from these errors, it is crucial that they do not add up to too much throughout the whole induction procedure.

To correct the first type of error, we introduce the following procedure:

1. For every $\boldsymbol{x}$, find the POVM $\{S^g_{\boldsymbol{x}}\}_g$ which is closest to $\{R_g\}$ (for some appropriate measure of distance between POVMs) while being *perfectly* consistent with $\{A^a_{\boldsymbol{x}}\}$: that is, $\sum_{g:g(\boldsymbol{x})=a} S^g_{\boldsymbol{x}} = A^a_{\boldsymbol{x}}$. This is possible only because $\boldsymbol{x}$ is fixed. We obtain $\{S^g_{\boldsymbol{x}}\}$ as the optimum solution of a convex program (5.9).

2. Show that $\{S^g_{\boldsymbol{x}}\}$ in fact does not depend too much on $\boldsymbol{x}$, so that defining $T^g := \mathrm{E}_{\boldsymbol{x}} S^g_{\boldsymbol{x}}$ leads to the consistent measurement we are looking for.

The second step is key: why would the $\{S^g_{\boldsymbol{x}}\}$ be (almost) independent of $\boldsymbol{x}$? Here the linearity property comes into play. Using the perfect consistency of S and A, together with the linearity of A, we are able to conclude that the $\{S^g_{\boldsymbol{x}}\}$ should not vary too much *along any line*. That is, $S^g_{(x_1,x_2)} \approx_\varepsilon S^g_{(x_1,x_2')}$ for any x_1 and x_2, x_2' (and similarly in the other direction). This step uses properties of the specific optimization problem that we introduced in order to define $\{S^g_{\boldsymbol{x}}\}_g$. The invariance along lines, together with the (reasonably) good expansion properties of the hypercube, lets us conclude that the $\{S^g_{\boldsymbol{x}}\}$ are in fact globally invariant, leading to the required POVM $\{T^g\}$.[8]

We now proceed with the details. In the following section we introduce the optimization procedure that is used to define the operators $\left\{S^g_{\boldsymbol{x}}\right\}_g$. In Section 5.6.2 we show that the $\{S^g_{\boldsymbol{x}}\}$ are close to being independent of $\boldsymbol{x}$, leading to a POVM $\{T^g\}$ having the properties claimed in Lemma 44.

[8] There will be a loss of a factor n in the quality of the approximation, but this is ok: they key point is that the quality of the approximation of the T^g depends on ε (and n) only, not on δ.

5.6.1 A convex optimization problem

For every $\boldsymbol{x}$ and multilinear g, let $\{\hat{S}^g_{\boldsymbol{x}}\}_{\boldsymbol{x}}$ be a solution to the following convex optimization problem:

$$\begin{aligned}\delta_1 \,:=\, &\min \,\mathrm{E}_{\boldsymbol{x}} \sum_g \Big\| \hat{S}^g_{\boldsymbol{x}} - \sqrt{R^g_{\boldsymbol{x}_{\geq k}}} \Big\|_\rho^2 \\ &\forall \boldsymbol{x}, g,\, \hat{S}^g_{\boldsymbol{x}} = A^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} \hat{S}^g_{\boldsymbol{x}}; \qquad \forall \boldsymbol{x}, a, \sum_{g: g(\boldsymbol{x}_{\leq k})=a} \hat{S}^g_{\boldsymbol{x}} (\hat{S}^g_{\boldsymbol{x}})^\dagger \leq A^a_{\boldsymbol{x}},\end{aligned} \tag{5.9}$$

where $\sqrt{R^g_{\boldsymbol{x}_{\geq k}}}$ is the positive square root of the (incomplete) POVM elements $\big\{R^g_{\boldsymbol{x}_{\geq k}}\big\}_g$ promised in the assumptions of Lemma 44. Let $S^g_{\boldsymbol{x}} = \hat{S}^g_{\boldsymbol{x}}\big(\hat{S}^g_{\boldsymbol{x}}\big)^\dagger$.[9] Our first claim shows that the optimum of (5.9) is bounded as a function of the consistency of R and A.

Claim 45. *Suppose that the $\big\{R^g_{\boldsymbol{x}_{\geq k}}\big\}_g$ are POVMs satisfying the assumptions of Lemma 44. Then the optimum of (5.9) is at most $O\big(\sqrt{\mu(A,R)} + \sqrt{\varepsilon}\big)$.*

Proof. We construct a feasible solution achieving the objective value claimed. Let $\hat{S}^g_{\boldsymbol{x}} := A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}$. Then by definition $\{\hat{S}^g_{\boldsymbol{x}}\}$ is a feasible solution to (5.9). To upper-bound its value, we first evaluate

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \Big(\mathrm{Tr}_\rho\big(\hat{S}^g_{\boldsymbol{x}} \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}\big) - \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}}\big)\Big) = \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) R^g_{\boldsymbol{x}_{\geq k}}\big)$$

We can bound

$$\begin{aligned}&\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) R^g_{\boldsymbol{x}_{\geq k}} \otimes \big(\mathrm{Id} - A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big)\big) \\ &\quad\leq \Big(\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}} \otimes \big(\mathrm{Id} - A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big)\big)\Big)^{1/2} \Big(\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) R^g_{\boldsymbol{x}_{\geq k}} \big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big)\big)\Big)^{1/2} \\ &\quad\leq \sqrt{\mu(A,R)},\end{aligned} \tag{5.10}$$

where the first inequality is Cauchy-Schwarz and the second follows by definition of $\mu(A,R)$. Similarly,

$$\begin{aligned}&\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) R^g_{\boldsymbol{x}_{\geq k}} \otimes A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big) \\ &\quad\leq \Big(\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}} \otimes A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big)\Big)^{1/2} \Big(\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) R^g_{\boldsymbol{x}_{\geq k}} \big(A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} - \mathrm{Id}\big) \otimes A^{g(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big)\Big)^{1/2} \\ &\quad\leq \sqrt{\mu(A,A)} \,=\, O\big(\sqrt{\varepsilon}\big),\end{aligned} \tag{5.11}$$

[9] We will usually use a hat, as in $\hat{S}$, to denote matrices which we think of as factorizations of POVM elements, and hence are not necessarily positive. In general, the relation between $\hat{X}$ and X will always be that $X = \hat{X}\hat{X}^\dagger$.

by consistency of A. Together, Eq. (5.10) and (5.11) show that

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\hat{S}^g_{\boldsymbol{x}} \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}\big) \geq \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}}\big) - O\big(\sqrt{\mu(R,A)} + \sqrt{\varepsilon}\big).$$

Expanding out $\big\|\hat{S}^g_{\boldsymbol{x}} - \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}\big\|_\rho^2$, to conclude it suffices to show that also

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\hat{S}^g_{\boldsymbol{x}}(\hat{S}^g_{\boldsymbol{x}})^\dagger\big) \leq \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}}\big) + O\big(\sqrt{\varepsilon}\big),$$

which follows from arguments similar to the one above. □

5.6.2 Constructing a POVM independent of $\boldsymbol{x}_{\leq k}$

As a first step in showing that any optimal solution to (5.9) must be close to being independent of $\boldsymbol{x}$, we show that such an optimal solution must be close to another feasible solution which is furthermore close to being invariant along the direction of any line. Precisely, we have the following.

Claim 46. *Assume $p^{-1} \leq \varepsilon$. For every $i < k$ there exists a feasible solution $\big\{\hat{Z}^g_{\boldsymbol{x}}\big\}$ to (5.9), with objective value at most $\delta_1 + O\big(\varepsilon\big)^{1/4}$, such that*

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \big\|\hat{Z}^g_{\boldsymbol{x}} - \mathrm{E}_{x_i'} \hat{Z}^g_{\boldsymbol{x}_{\neg i}, x_i'}\big\|_\rho^2 = O\big(\sqrt{\varepsilon}\big).$$

Proof. Let $\{\hat{S}^g_{\boldsymbol{x}}\}$ be an optimal solution to (5.9), and for any $i < k$ let

$$\hat{Y}^g_{\boldsymbol{x}_{\neg i}} := B^{g|\ell_i(\boldsymbol{x})}_{\boldsymbol{x}_{\neg i}} \mathrm{E}_{x_i} \hat{S}^g_{\boldsymbol{x}},$$

where $\ell_i(\boldsymbol{x})$ is the line going through $\boldsymbol{x}$ and parallel to the i-th axis, and $\{B^\ell_{\boldsymbol{x}_{\neg i}}\}_\ell$ is the "lines" POVM introduced in Claim 42. We first claim that the $\hat{Y}^g_{\boldsymbol{x}_{\neg i}}$, while not strictly feasible, achieve an objective value in (5.9) of at most $\delta_1 + O(\varepsilon)$.

Towards proving this, we first show that $B^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} \hat{S}^g_{\boldsymbol{x}}$ is close to $\hat{S}^g_{\boldsymbol{x}}$. This follows from the relation $A^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} \hat{S}^g_{\boldsymbol{x}} = \hat{S}^g_{\boldsymbol{x}}$, and the closeness of A and B. Indeed, we have

$$\begin{aligned}
\mathrm{E}_{\boldsymbol{x}} \sum_g \big\|B^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} \hat{S}^g_{\boldsymbol{x}} - \hat{S}^g_{\boldsymbol{x}}\big\|_\rho^2 &= \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\Big(\big(B^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} - A^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}}\big) S^g_{\boldsymbol{x}} \big(B^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} - A^{g(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}}\big)\Big) \\
&\leq \mathrm{E}_{\boldsymbol{x}} \sum_a \big\|B^a_{\boldsymbol{x}} - A^a_{\boldsymbol{x}}\big\|_\rho^2 \\
&= O\big(\sqrt{\varepsilon}\big) \qquad (5.12)
\end{aligned}$$

by Claim 42. By convexity, the following (infeasible) operators

$$\tilde{Y}^g_{\boldsymbol{x}_{\neg i}} := \mathrm{E}_{x_i} B_{\boldsymbol{x}}^{g(\boldsymbol{x}_{\leq k})} \hat{S}^g_{\boldsymbol{x}}$$

also achieve a value $\delta_1 + O(\sqrt{\varepsilon})$ in (5.9).

Next we show that the $\tilde{Y}^g_{\boldsymbol{x}_{\neg i}}$ are close to the

$$\hat{Y}^g_{\boldsymbol{x}_{\neg i}} := B_{\boldsymbol{x}_{\neg i}}^{g_{|\ell_i(\boldsymbol{x})}} \mathrm{E}_{x_i} \hat{S}^g_{\boldsymbol{x}},$$

which are invariant along lines in the i-th direction. Indeed, from the definition

$$\tilde{Y}^g_{\boldsymbol{x}_{\neg i}} = B_{\boldsymbol{x}_{\neg i}}^{g_{|\ell_i(\boldsymbol{x})}} \mathrm{E}_{x_i} \hat{S}^g_{\boldsymbol{x}} + \mathrm{E}_{x_i} \sum_{\substack{ux_i+v=g(\boldsymbol{x}_{\leq k}) \\ (u,v)\neq g_{|\ell_i(\boldsymbol{x})}}} B^{u,v}_{\boldsymbol{x}_{\neg i}} \hat{S}^g_{\boldsymbol{x}}.$$

The norm of the second term can be expanded as follows:

$$\begin{aligned}
\mathrm{E}_{\boldsymbol{x}_{\neg i}} \sum_g \Big\| \mathrm{E}_{x_i} \sum_{\substack{ux_i+v=g(\boldsymbol{x}_{\leq k}) \\ (u,v)\neq g_{|\ell_i(\boldsymbol{x})}}} B^{u,v}_{\boldsymbol{x}_{\neg i}} \hat{S}^g_{\boldsymbol{x}} \Big\|_\rho^2 \\
&= \mathrm{E}_{\boldsymbol{x}_{\neg i}} \sum_g \mathrm{E}_{x_i,y_i} \sum_{\substack{ux_i+v=g(\boldsymbol{x}_{\leq k}) \\ (u,v)\neq g_{|\ell_i(\boldsymbol{x})}}} \sum_{\substack{u'y_i+v'=g(\boldsymbol{x}_{\leq k}) \\ (u',v')\neq g_{|\ell_i(\boldsymbol{x})}}} \mathrm{Tr}_\rho\big(B^{u,v}_{\boldsymbol{x}_{\neg i}} \hat{S}^g_{\boldsymbol{x}_{\neg i},x_i} (\hat{S}^g_{\boldsymbol{x}_{\neg i},y_i})^\dagger B^{u',v'}_{\boldsymbol{x}_{\neg i}} \big)
\end{aligned}$$

Eq. (A.2) from Lemma 120 shows that terms such that $(u,v) \neq (u',v')$ contribute at most $O\big(\sqrt{\mu(B,B)}\big) = O\big(\varepsilon^{1/4}\big)$. But the only possibility for $(u,v) = (u',v')$ is that also $x_i = y_i$, since two distinct lines intersect in at most one point. Hence we have that

$$\begin{aligned}
\mathrm{E}_{\boldsymbol{x}_{\neg i}} \sum_g \Big\| \mathrm{E}_{x_i} \sum_{\substack{ux_i+v=g(\boldsymbol{x}_{\leq k}) \\ (u,v)\neq g_{|\ell_i(\boldsymbol{x})}}} B^{u,v}_{\boldsymbol{x}_{\neg i}} \hat{S}^g_{\boldsymbol{x}} \Big\|_\rho^2 &= \mathrm{E}_{\boldsymbol{x}_{\neg i}} \sum_g \frac{1}{p} \mathrm{E}_{x_i} \sum_{\substack{ux_i+v=g(\boldsymbol{x}_{\leq k}) \\ (u,v)\neq g_{|\ell_i(\boldsymbol{x})}}} \mathrm{Tr}_\rho\big(B^{u,v}_{\boldsymbol{x}_{\neg i}} S^g_{\boldsymbol{x}} B^{u,v}_{\boldsymbol{x}_{\neg i}} \big) + O\big(\varepsilon^{1/4}\big) \\
&\leq \frac{1}{p} + O\big(\varepsilon^{1/4}\big).
\end{aligned}$$

This in particular implies that the $\hat{Y}^g_{\boldsymbol{x}_{\neg i}}$, while still not necessarily feasible, achieve an objective value in (5.9) of $\delta_1 + O\big(\varepsilon^{1/4}\big)$.

Finally, define $\hat{Z}^g_{\boldsymbol{x}} := A_{\boldsymbol{x}}^{g(\boldsymbol{x}_{\leq k})} B_{\boldsymbol{x}_{\neg i}}^{g_{|\ell_i(\boldsymbol{x})}} \mathrm{E}_{x_i} \hat{S}^g_{\boldsymbol{x}}$. Then the $\big\{\hat{Z}^g_{\boldsymbol{x}}\big\}$ are feasible in (5.9), and the fact that

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \big\| \hat{Z}^g_{\boldsymbol{x}} - \hat{Y}^g_{\boldsymbol{x}_{\neg i}} \big\|_\rho^2 = O\big(\sqrt{\varepsilon}\big) \tag{5.13}$$

follows from arguments similar to those used to prove Eq. (5.12). Hence the $\{\hat{Z}_{\boldsymbol{x}}^g\}$ are a feasible solution to (5.9) with objective value $\delta_1 + O(\varepsilon^{1/4})$. Finally, by convexity (5.13) implies that

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \big\| \mathrm{E}_{x_i} \hat{Z}_{\boldsymbol{x}}^g - \hat{Y}_{\boldsymbol{x}_{\neg i}}^g \big\|_\rho^2 = O\big(\sqrt{\varepsilon}\big),$$

which together with the triangle inequality and (5.13) shows that the $\hat{Z}$ are close to their expectation on any axis-parallel line in the i-th direction, proving the claim. □

Using convexity of $X \to \|X - A\|_\rho^2$ for fixed A, the following follows from Claims 45 and 46.

Claim 47. *Let $\{\hat{S}_{\boldsymbol{x}}^g\}$ be an optimal solution to (5.9). Then*

$$\mathrm{E}_{\boldsymbol{x},i<k} \sum_g \|\hat{S}_{\boldsymbol{x}}^g - \mathrm{E}_{x_i'} \hat{S}_{\boldsymbol{x}_{\neg i} x_i'}^g\|_\rho^2 = O\big(\varepsilon^{1/4}\big).$$

Proof. We show that the two solutions thus constructed to (5.9), $\{\hat{S}_{\boldsymbol{x}}^g\}$ and $\{\hat{Z}_{\boldsymbol{x}}^g\}$ from Claim 46, must be close:[10]

$$\mathrm{E}_{\boldsymbol{x},i} \sum_g \big\| \hat{Z}_{\boldsymbol{x}}^g - \hat{S}_{\boldsymbol{x}}^g \big\|_\rho^2 = O\big(\varepsilon^{1/4}\big). \tag{5.14}$$

The claim will follow by combining this bound with the fact, proved in Claim 46, that the $\hat{Z}_{\boldsymbol{x}}^g$ themselves are close to their expectation along any axis-parallel line in the i-th direction.

Eq. (5.14) essentially follows from the fact that both $\{\hat{S}_{\boldsymbol{x}}^g\}$ and $\{\hat{Z}_{\boldsymbol{x}}^g\}$ are almost-optimal solutions to (5.9), which is the minimization of a "Euclidean-like" distance to a convex set — hence they must be close. Precisely, since the feasible set of (5.9) is convex, for any $0 \le t \le 1$ the elements $(1-t)\hat{S}_{\boldsymbol{x}}^g + t\hat{Z}_{\boldsymbol{x}}^g$ also constitute a feasible solution. By optimality of $\{\hat{S}_{\boldsymbol{x}}^g\}$, its objective value must be at least δ_1: for every $0 \le t \le 1$,

$$\begin{aligned}
\mathrm{E}_{\boldsymbol{x}} \sum_g \Big\| \hat{S}_{\boldsymbol{x}}^g - \sqrt{R_{\boldsymbol{x}_{\ge k}}^g} \Big\|_\rho^2 &\le \mathrm{E}_{\boldsymbol{x}} \sum_g \Big\| (1-t)\hat{S}_{\boldsymbol{x}}^g + t\hat{Z}_{\boldsymbol{x}}^g - \sqrt{R_{\boldsymbol{x}_{\ge k}}^g} \Big\|_\rho^2 \\
&= t^2\, \mathrm{E}_{\boldsymbol{x}} \sum_g t^2 \Big\| \hat{Z}_{\boldsymbol{x}}^g - \hat{S}_{\boldsymbol{x}}^g \Big\|_\rho^2 + \mathrm{E}_{\boldsymbol{x}} \sum_g \Big\| \hat{S}_{\boldsymbol{x}}^g - \sqrt{R_{\boldsymbol{x}_{\ge k}}^g} \Big\|_\rho^2 \\
&\qquad + 2\,t\, \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\Big(\big(\hat{Z}_{\boldsymbol{x}}^g - \hat{S}_{\boldsymbol{x}}^g\big)\big(\hat{S}_{\boldsymbol{x}}^g - \sqrt{R_{\boldsymbol{x}_{\ge k}}^g}\big)^\dagger \Big)
\end{aligned}$$

[10] Note that $\hat{Z}_{\boldsymbol{x}}^g$ implicitly depends on i, and the following equation is measuring the distance on average over the $k-1$ different constructions of $\hat{Z}_{\boldsymbol{x}}^g$ obtained for all $1 \le i < k$.

Using the known objective values, re-arranging and making $t \to 0$, we obtain that

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\Big(\hat{Z}^g_{\boldsymbol{x}} - \hat{S}^g_{\boldsymbol{x}}\big)\big(\sqrt{R^g_{\boldsymbol{x}_{\geq k}}} - \hat{S}^g_{\boldsymbol{x}}\big)^\dagger\Big) = O\big(\varepsilon^{1/4}\big).$$

Hence

$$\begin{aligned}\mathrm{E}_{\boldsymbol{x}} \sum_g \big\|\hat{S}^g_{\boldsymbol{x}} - \hat{Z}^g_{\boldsymbol{x}}\big\|_\rho^2 &= \mathrm{E}_{\boldsymbol{x}} \sum_g \Big(\Big\|\hat{Z}^g_{\boldsymbol{x}} - \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}\Big\|_\rho^2 - \Big\|\hat{S}^g_{\boldsymbol{x}} - \sqrt{R^g_{\boldsymbol{x}_{\geq k}}}\Big\|_\rho^2 \\ &\qquad + 2\,\mathrm{Tr}_\rho\Big(\big(\hat{Z}^g_{\boldsymbol{x}} - \hat{S}^g_{\boldsymbol{x}}\big)\big(R^g_{\boldsymbol{x}_{\geq k}} - S^g_{\boldsymbol{x}}\big)^\dagger\Big)\Big) \\ &= O\big(\varepsilon^{1/4}\big),\end{aligned}$$

proving (5.14). □

Claim 47 shows that the $\{\hat{S}^g_{\boldsymbol{x}}\}_g$ do not vary much along any axis-parallel line in the i-th direction. Using the expansion properties of the hypercube, we can deduce that the $\{\hat{S}^g_{\boldsymbol{x}}\}_g$ are close (in the $\|\cdot\|_\rho^2$ norm) to a single operator, independent of the first $(k-1)$ coordinates.

Claim 48. *For every $\boldsymbol{x}_{\geq k}$ and g, let $\hat{T}^g_{\boldsymbol{x}_{\geq k}} := \mathrm{E}_{\boldsymbol{x}_{<k}} \hat{S}^g_{\boldsymbol{x}}$. Then*

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \big\|\hat{S}^g_{\boldsymbol{x}} - \hat{T}^g_{\boldsymbol{x}_{\geq k}}\big\|_\rho^2 \;=\; O\big(n\varepsilon^{1/4}\big)$$

Proof. This is a direct consequence of Claim 121. □

The only points remaining to be proven in Lemma 44 are items 3 and 4. Both follow from the two previous items in a similar way. Item 4 follows as

$$\begin{aligned}\mathrm{Tr}_\rho(T) &= \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\hat{T}^g_{\boldsymbol{x}_{\geq k}}\big(\hat{T}^g_{\boldsymbol{x}_{\geq k}}\big)^\dagger\big) \\ &= \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(\hat{S}^g_{\boldsymbol{x}}\big(\hat{S}^g_{\boldsymbol{x}}\big)^\dagger\big) + O\big(\varepsilon^{1/8}\big) \\ &= \mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(R^g_{\boldsymbol{x}_{\geq k}}\big) + O\big(\delta^{1/4}\big) + O\big(\varepsilon^{1/8}\big),\end{aligned}$$

where the second equality follows from Cauchy-Schwarz and Claim 48, and the third follows from Cauchy-Schwarz and the bound proved in Claim 45.

5.7 The pasting lemma

In this section we prove Lemma 43. Let T be the POVM whose existence is promised in the lemma's assumptions. Our goal is to define a new POVM V, depending on one less coordinate of $\boldsymbol{x}$ than T, but such that V is still reasonably consistent with A and is not "too incomplete": its trace should not be much smaller than that of T. The idea is to define V as (roughly) corresponding to the sequential application of T twice. This will produce two $(k-1)$-multilinear functions h and h', from which a k-multilinear function g can be recovered by interpolation. This is essentially the same method as was used to define the "line" operators B from the "point" operators A in Claim 42. Here the added difficulty is that we are working with incomplete POVMs, and special care must be taken to ensure the overall trace does not decrease too much.

For every $\boldsymbol{x}_{\geq k}$ and h choose a factorization of each POVM element $T^h_{\boldsymbol{x}_{\geq k}} = \hat{T}^h_{\boldsymbol{x}_{\geq k}}\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}}\big)^\dagger$ that is such that $\hat{T}^h_{x_k\boldsymbol{x}_{>k}}\big(\hat{T}^{h'}_{x'_k\boldsymbol{x}_{>k}}\big)^\dagger = 0$ whenever $h \neq h'$ or $x_k \neq x'_k$. While this factorization may not be the one that appears in Lemma 44, all properties of this lemma are independent of the specific factorization chosen (provided it is reflected appropriately in the factorization for $S^h_{\boldsymbol{x}}$).

Define $\hat{T}_{\boldsymbol{x}_{\geq k}} = \sum_h \hat{T}^h_{\boldsymbol{x}_{\geq k}}$ and let $\hat{T}_{\boldsymbol{x}_{>k}} = \mathrm{E}_{x_k}\hat{T}_{\boldsymbol{x}_{\geq k}}$. Let $\hat{T}_{\boldsymbol{x}_{>k}} = U\Sigma V^\dagger$, where Σ is diagonal with non-negative coefficients, be the SVD (to lighten the notation we leave the dependence of U, Σ and V on $\boldsymbol{x}_{>k}$ implicit). Let $\tilde{\Sigma}$ be Σ with all eigenvalues less than ε^d rounded down to 0, where $d > 0$ is a small constant to be specified later. Let $\tilde{T}_{\boldsymbol{x}_{>k}} = U\tilde{\Sigma}V^\dagger$, and let $\tilde{T}^{-1}_{\boldsymbol{x}_{>k}} = V\tilde{\Sigma}^{-1}U^\dagger$ be the pseudo-inverse.

For every $\boldsymbol{x}_{>k}$ and multilinear $g : \mathbb{F}^k \to \mathbb{F}$ define

$$V^g_{\boldsymbol{x}_{>k}} := \mathrm{E}_{y_k}\mathrm{E}_{x_k\neq y_k, x'_k\neq y_k}\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x'_k}_{x'_k,\boldsymbol{x}_{>k}}\big)^\dagger. \tag{5.15}$$

We first verify that the $\big\{V^g_{\boldsymbol{x}_{>k}}\big\}_g$ form an (incomplete) POVM.[11]

Claim 49. *For every $\boldsymbol{x}_{>k}$, it holds that $V^g_{\boldsymbol{x}_{>k}} \geq 0$ for every g and $\sum_g V^g_{\boldsymbol{x}_{>k}} \leq \big(1 + O(\varepsilon^{-2d}/p)\big)\,\mathrm{Id}$.*

Proof. The non-negativity part is clear, as

$$V^g_{\boldsymbol{x}_{>k}} = \mathrm{E}_{y_k}\Big(\mathrm{E}_{x_k\neq y_k}\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\Big)\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\Big(\mathrm{E}_{x_k\neq y_k}\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\Big)^\dagger.$$

Using $AB^\dagger + BA^\dagger \leq AA^\dagger + BB^\dagger$, for any (x_k, y_k, x'_k) we have

$$\begin{aligned}&\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x'_k}_{x'_k,\boldsymbol{x}_{>k}}\big)^\dagger + \hat{T}^{g|x'_k}_{x'_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\big)^\dagger\\ &\quad\leq \hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\big)^\dagger + \hat{T}^{g|x'_k}_{x'_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x'_k}_{x'_k,\boldsymbol{x}_{>k}}\big)^\dagger.\end{aligned}$$

[11] The claim shows that the $V^g_{\boldsymbol{x}_{>k}}$ sum to "barely more" than Id. To make them into a POVM it will suffice to eventually scale them by the appropriate amount. This will not affect the properties that need to be proved in Lemma 43 by much.

Hence

$$\begin{aligned}\sum_g V^g_{\boldsymbol{x}_{>k}} &\leq \sum_g \mathrm{E}_{y_k}\mathrm{E}_{x_k\neq y_k}\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}\big(\hat{T}^{g|x_k}_{x_k,\boldsymbol{x}_{>k}}\big)^\dagger\\ &= \big(1+1/p\big)\mathrm{E}_{x_k}\hat{T}_{x_k,\boldsymbol{x}_{>k}}\big(\hat{T}_{x_k,\boldsymbol{x}_{>k}}\big)^\dagger\\ &\leq \big(1+1/p\big)\,T_{\boldsymbol{x}_{>k}},\end{aligned}$$

where in the last inequality the extra $(1/p)$Id accounts for the terms for which $x_k = y_k$ that are missing, and we used $\tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T_{\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}} \leq \mathrm{Id}$ by definition. □

Now that we defined an (incomplete) POVM $\big\{V^g_{\boldsymbol{x}_{>k}}\big\}_g$, we show in the next two sections that it satisfies items 1. and 2. in the conclusion of Lemma 43.

5.7.1 Consistency of V and A

Our next claim shows that V is consistent with A.

Claim 50. *Let $\big\{V^g_{\boldsymbol{x}_{>k}}\big\}$ be the (incomplete) POVM defined in (5.15). Then*

$$\mu(V,A) = \mathrm{E}_{\boldsymbol{x}}\sum_{g,a\neq g(\boldsymbol{x}_{\leq k})}\mathrm{Tr}_\rho\big(V^g_{\boldsymbol{x}_{>k}}\otimes A^a_{\boldsymbol{x}}\big) \leq O\big(\varepsilon^{c/2}\big)$$

Proof. To lighten the notation, for every $\boldsymbol{x}_{\geq k}$ and h let

$$Z^h_{\boldsymbol{x}_{\geq k}} := \tilde{T}^{-1}_{\boldsymbol{x}_{>k}}T^h_{x_k,\boldsymbol{x}_{>k}}\tilde{T}^{-1}_{\boldsymbol{x}_{>k}},$$

so that $V^g_{\boldsymbol{x}_{>k}} = \mathrm{E}_{x_k,y_k\neq x_k,y'_k\neq x_k}\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}Z^{g|x_k}_{\boldsymbol{x}_{\geq k}}\big(\hat{T}^{g|y'_k}_{y'_k,\boldsymbol{x}_{>k}}\big)^\dagger$. We first show the bound

$$\Big|\mathop{\mathrm{E}}_{\boldsymbol{x},y_k,y'_k\neq x_k}\sum_{g,a\neq g(\boldsymbol{x}_{\leq k})}\mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}Z^{g|x_k}_{\boldsymbol{x}_{\geq k}}\big(\hat{T}^{g|y'_k}_{y'_k,\boldsymbol{x}_{>k}}\big)^\dagger\otimes A^a_{\boldsymbol{x}}\otimes\big(\mathrm{Id}-A^{g(y'_k,\boldsymbol{x}_{<k})}_{y'_k,\boldsymbol{x}_{\neg k}}\big)\Big| = O\big(\varepsilon^{c/2}\big). \quad (5.16)$$

To prove (5.16), first bound

$$\begin{aligned}&\Big|\mathop{\mathrm{E}}_{\boldsymbol{x},y_k,y'_k\neq x_k}\sum_{g,a\neq g(\boldsymbol{x}_{\leq k})}\mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}Z^{g|x_k}_{\boldsymbol{x}_{\geq k}}\big(\hat{T}^{g|y'_k}_{'y_k,\boldsymbol{x}_{>k}}-\hat{S}^{g|y'_k}_{y'_k,\boldsymbol{x}_{\neg k}}\big)^\dagger\otimes A^a_{\boldsymbol{x}}\big)\Big|\\ &\qquad\leq \Big(\mathop{\mathrm{E}}_{\boldsymbol{x},y_k\neq x_k}\sum_g\mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}Z^{g|x_k}_{\boldsymbol{x}_{\geq k}}\big(Z^{g|x_k}_{\boldsymbol{x}_{\geq k}}\big)^\dagger\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\big)^\dagger\big)\Big)^{1/2}\Big(\mathrm{E}_{\boldsymbol{x},y_k\neq x_k}\sum_h\big\|\hat{T}^h_{y_k,\boldsymbol{x}_{>k}}-\hat{S}^h_{y_k,\boldsymbol{x}_{\neg k}}\big\|^2_\rho\Big)^{1/2}\\ &\qquad\leq O\big(\varepsilon^{c/2}\big), \qquad (5.17)\end{aligned}$$

where for the last inequality we bounded the second term using item 2 from Lemma 44. Next using item 1 and consistency of A,

$$\begin{aligned}
\Big| \mathop{\mathrm{E}}_{\boldsymbol{x},y_k,y'_k\neq x_k} & \sum_{g,a\neq g(\boldsymbol{x}_{\leq k})} \mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}} Z^{g|x_k}_{\boldsymbol{x}_{\geq k}} \big(\hat{S}^{g|y'_k}_{y'_k,\boldsymbol{x}_{\neg k}}\big)^\dagger \otimes A^a_{\boldsymbol{x}} \otimes \big(\mathrm{Id} - A^{g(y'_k,\boldsymbol{x}_{<k})}_{y'_k,\boldsymbol{x}_{\neg k}}\big)\big)\Big| \\
&\leq \Big(\mathop{\mathrm{E}}_{\boldsymbol{x},y_k\neq x_k} \sum_g \mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}} Z^{g|x_k}_{\boldsymbol{x}_{\geq k}} (Z^{g|x_k}_{\boldsymbol{x}_{\geq k}})^\dagger \big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}}\big)^\dagger\big)\Big)^{1/2} \\
&\qquad \cdot \Big(\mathop{\mathrm{E}}_{\boldsymbol{x},y_k\neq x_k} \sum_h \mathrm{Tr}_\rho\big(S^h_{y_k,\boldsymbol{x}_{>k}} \otimes \big(\mathrm{Id} - A^{h(\boldsymbol{x}_{<k})}_{y_k,\boldsymbol{x}_{\neg k}}\big)\big)\Big)^{1/2} \\
&\leq O\big(\varepsilon^{1/2}\big).
\end{aligned}$$

Finally, the proof of (5.16) can be concluded using a bound similar to (5.17). The same sequence of reasoning also shows that

$$\begin{aligned}
\mathop{\mathrm{E}}_{\boldsymbol{x},y_k,y'_k\neq x_k} & \sum_{g,a\neq g(\boldsymbol{x}_{\leq k})} \mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}} Z^{g|x_k}_{\boldsymbol{x}_{\geq k}} \big(\hat{T}^{g|y'_k}_{y'_k,\boldsymbol{x}_{>k}}\big)^\dagger \otimes A^a_{\boldsymbol{x}}\big) \\
&= \mathop{\mathrm{E}}_{\boldsymbol{x},y_k,y'_k\neq x_k} \sum_{g,a\neq g(\boldsymbol{x}_{\leq k})} \mathrm{Tr}_\rho\big(\hat{T}^{g|y_k}_{y_k,\boldsymbol{x}_{>k}} Z^{g|x_k}_{\boldsymbol{x}_{\geq k}} \big(\hat{T}^{g|y'_k}_{y'_k,\boldsymbol{x}_{>k}}\big) \otimes A^a_{\boldsymbol{x}} \otimes A^{g(y_k,\boldsymbol{x}_{<k})}_{y_k,\boldsymbol{x}_{\neg k}} \otimes A^{g(y'_k,\boldsymbol{x}_{<k})}_{y'_k,\boldsymbol{x}_{\neg k}}\big) + O\big(\varepsilon^{c/2}\big).
\end{aligned} \tag{5.18}$$

This last expression can be directly upper-bounded by $O(\sqrt{\varepsilon})$ using Cauchy-Schwarz and linearity of A (and of g) in the k-th direction. □

5.7.2 V is not too incomplete

Before showing that $\mathrm{Tr}_\rho(V)$ is not much smaller than $\mathrm{Tr}_\rho(T)$, we first prove the following useful preliminary bound.

Claim 51. *The following holds:*

$$\mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y'_k}} \sum_{\substack{(h,h',h'')\text{ not aligned}\\ (h(\boldsymbol{x}_{<k}),h'(\boldsymbol{x}_{<k}),h''(\boldsymbol{x}_{<k}))\\ \textit{aligned}}} \mathrm{Tr}_\rho\big(\hat{T}^h_{y_k,\boldsymbol{x}_{>k}} \tilde{T}^{-1} T^{h''}_{z_k,\boldsymbol{x}_{>k}} \tilde{T}^{-1} \big(\hat{T}^{h'}_{y'_k,\boldsymbol{x}_{>k}}\big)^\dagger \otimes A^{h'''(\boldsymbol{x}_{<k})}_{\boldsymbol{x}_{\neg k},y_k}\big) = O\big(n\sqrt{\mu(T,A)}\big), \tag{5.19}$$

where here h''' is the unique function aligned with (h',h'') in the k-th direction: $(h'''-h'')/(y_k-z_k) = (h''-h')/(z_k-y'_k)$.

The claim essentially follows from consistency of T and A, but the details are a bit cumbersome.

Proof. For any $\boldsymbol{x}_{>k}$ and h, let

$$Z^h_{\boldsymbol{x}_{>k}} := \mathrm{E}_{(z_k,y_k):\,x_k\neq z_k\neq y_k} \sum_{(h,h',h'')\ aligned} \tilde{T}^{-1} T^{h''}_{z_k,\boldsymbol{x}_{>k}} \tilde{T}^{-1} \big(\hat{T}^{h'}_{y_k,\boldsymbol{x}_{>k}}\big)^\dagger,$$

so that (5.19) can be expressed more succinctly as a bound on the quantity

$$\Delta \,:=\, \mathrm{E}_{\boldsymbol{x}} \sum_{\substack{h\neq h'\\ h(\boldsymbol{x})=h'(\boldsymbol{x})}} \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A_{\boldsymbol{x}}^{h'(\boldsymbol{x}_{<k})}\big).$$

For any two h, h' and $\boldsymbol{x} \in \mathbb{F}^n$ we let $\beta_0(h,h',\boldsymbol{x}) = 0$ and for $i = 1,\dots,k-1$ inductively define

$$\beta_i(h,h',\boldsymbol{x}) = \begin{cases} 1 & \text{if } h^c_{|\boldsymbol{x}_{<i}} \neq (h')^c_{|\boldsymbol{x}_{<i}} \text{ and } h^c_{|\boldsymbol{x}_{\leq i}} = (h')^c_{|\boldsymbol{x}_{\leq i}} \\ p/(p-1)\beta_{i-1}(h,h',\boldsymbol{x}) & \text{if } h^c_{|\boldsymbol{x}_{<i}} \neq (h')^c_{|\boldsymbol{x}_{<i}} \text{ and } h^c_{|\boldsymbol{x}_{\leq i}} \neq (h')^c_{|\boldsymbol{x}_{\leq i}} \\ \beta_{i-1}(h,h',\boldsymbol{x}) + 1/p & \text{if } h^c_{|\boldsymbol{x}_{<i}} = (h')^c_{|\boldsymbol{x}_{<i}}, \end{cases}$$

where for any function g in i variables $x_1,\dots,x_i$, g^c denotes the leading coefficient of g, i.e. its coefficient on the monomial $x_1\cdots x_i$. Using the fact that whenever $h^c_{|\boldsymbol{x}_{<i}} \neq (h')^c_{|\boldsymbol{x}_{<i}}$ there always exists a unique x_i such that $h^c_{|\boldsymbol{x}_{\leq i}} = (h')^c_{|\boldsymbol{x}_{\leq i}}$, one can verify that for any (h,h') it holds that

$$\sum_{\boldsymbol{x}} \beta_{k-1}(h,h',\boldsymbol{x}) \,=\, (k-1)\,p^{k-2}. \tag{5.20}$$

Moreover, for any $\boldsymbol{x}$ such that $h(\boldsymbol{x}_{<k}) = h'(\boldsymbol{x}_{<k})$ we have $\beta_{k-1}(h,h',\boldsymbol{x}) = 1 + (k-1-i)/p$, where $i \in [n]$ is the largest integer such that $h^c_{|\boldsymbol{x}_{<i}} \neq (h')^c_{|\boldsymbol{x}_{<i}}$ and $h^c_{|\boldsymbol{x}_{\leq i}} = (h')^c_{|\boldsymbol{x}_{\leq i}}$. We first bound

$$\Delta_1 := \mathrm{E}_{\boldsymbol{x}} \sum_{\substack{h\neq h'\\ h(\boldsymbol{x}_{<k})=h'(\boldsymbol{x}_{<k})}} \big(\beta_{k-1}(h,h',\boldsymbol{x})-1\big)\mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}}\big) \,=\, O\big(k^2/p\big). \tag{5.21}$$

Indeed, for any fixed $i < k$ we have

$$\begin{aligned} &\mathrm{E}_{\boldsymbol{x}} \sum_{\substack{h\neq h':\, h^c_{|\boldsymbol{x}_{<i}}\neq (h')^c_{|\boldsymbol{x}_{<i}}\\ h^c_{|\boldsymbol{x}_{\leq j}}=(h')^c_{|\boldsymbol{x}_{\leq j}}\ \forall j\geq i}} \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}}\big)\\ &\quad= \mathrm{E}_{\boldsymbol{x}} \sum_{h'} \mathrm{Tr}_\rho\Big(\Big(\sum_{\substack{h\neq h':\,\forall j\geq i,\\ h^c_{|\boldsymbol{x}_{\leq j}}=(h')^c_{|\boldsymbol{x}_{\leq j}}}} \hat{T}^h_{\boldsymbol{x}_{\geq k}} - \sum_{\substack{h\neq h':\,\forall j\geq i-1,\\ h^c_{|\boldsymbol{x}_{\leq j}}=(h')^c_{|\boldsymbol{x}_{\leq j}}}} \hat{T}^h_{\boldsymbol{x}_{\geq k}}\Big) Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}}\Big), \end{aligned}$$

which can be bounded by $O(1)$ by splitting it into two terms and applying the Cauchy-Schwarz inequality to each term. This proves (5.21). Using the recursive definition of β, one can prove the following bound by induction on $i = 0, \ldots, k-1$:

$$\Delta_i := \mathrm{E}_{\boldsymbol{x}} \sum_a \mathrm{Tr}_\rho\Big(\Big(\sum_{\substack{h'(\boldsymbol{x})=a \\ h \neq h'}} \beta_i(h,h',\boldsymbol{x}) \hat{T}^{h'}_{\boldsymbol{x}_{\geq k}} Z^h_{\boldsymbol{x}_{\geq k}}\Big)\Big(\sum_{\substack{h'(\boldsymbol{x})=a \\ h \neq h'}} \beta_i(h,h',\boldsymbol{x}) \hat{T}^{h'}_{\boldsymbol{x}_{\geq k}} Z^h_{\boldsymbol{x}_{\geq k}}\Big)^\dagger\Big) = O(i^2). \tag{5.22}$$

We can now re-write Δ as

$$\begin{aligned}
\Delta &= \mathrm{E}_{\boldsymbol{x}} \sum_{\substack{h \neq h' \\ h(\boldsymbol{x}_{<k}) = h'(\boldsymbol{x}_{<k})}} \beta_{k-1}(h,h',\boldsymbol{x})\, \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}}\big) + O(n^2/p) \\
&= \mathrm{E}_{\boldsymbol{x}} \sum_{\substack{h \neq h' \\ h(\boldsymbol{x}_{<k}) = h'(\boldsymbol{x}_{<k})}} \beta_{k-1}(h,h',\boldsymbol{x})\, \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}} A^{h'(\boldsymbol{x}_{\leq k})}_{\boldsymbol{x}}\big) + O\big(\sqrt{\mu(A)}\big) \\
&= \mathrm{E}_{\boldsymbol{x}} \sum_{h \neq h'} \beta_{k-1}(h,h',\boldsymbol{x})\, \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} A^{h'(\boldsymbol{x}_{<k})}_{\boldsymbol{x}}\big) + O\big(\sqrt{\mu(A)}\big),
\end{aligned}$$

where the first equality is by (5.21), the second uses self-consistency of A together with the Cauchy-Schwarz inequality, and in the last we added the terms for which $h(\boldsymbol{x}_{<k}) \neq h'(\boldsymbol{x}_{<k})$ by using that, in that case, $A^{h(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} A^{h'(\boldsymbol{x}_{<k})}_{\boldsymbol{x}} = 0$. Using consistency again, together with the Cauchy-Schwarz inequality and (5.22), we may remove the dependence of the last two A terms on $\boldsymbol{x}_{<k}$, leading to

$$\begin{aligned}
\Delta &:= \mathrm{E}_{\boldsymbol{x}} \sum_{h \neq h'} \beta_{k-1}(h,h',\boldsymbol{x})\, \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h}_{\boldsymbol{x}_{\geq k}} A^{h'}_{\boldsymbol{x}_{\geq k}}\big) + O\big(n\sqrt{\mu(T,A)}\big) \\
&= \frac{k-1}{p}\, \mathrm{E}_{\boldsymbol{x}} \sum_{h \neq h'} \mathrm{Tr}_\rho\big(\hat{T}^h_{\boldsymbol{x}_{\geq k}} Z^{h'}_{\boldsymbol{x}_{\geq k}} \otimes A^{h'}_{\boldsymbol{x}_{\geq k}}\big) + O\big(n\sqrt{\mu(T,A)}\big) \\
&= O\big(n\sqrt{\mu(T,A)}\big),
\end{aligned}$$

where the second equality follows from (5.20). □

Finally, we show that $\mathrm{Tr}_\rho(V)$ is not too small compared to $\mathrm{Tr}_\rho(T)$, proving item 2. in the conclusion of Lemma 43.

Claim 52. *Let $\{T^h_{\boldsymbol{x}_{\geq k}}\}$ be the (incomplete) POVM promised in the assumptions in Lemma 43, and $\{V^g_{\boldsymbol{x}_{>k}}\}$ the (incomplete) POVM defined in (5.15). Then*

$$\mathrm{E}_{\boldsymbol{x}} \sum_g \mathrm{Tr}_\rho\big(V^g_{\boldsymbol{x}_{>k}}\big) \geq \mathrm{Tr}_\rho(T) - O\big(n\sqrt{\mu(T,A)} + \varepsilon^{d/2}\big).$$

Proof. Proceeding as in the proof of Claim 50,

$$\begin{aligned}\mathrm{Tr}_\rho(V) \geq & \mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y_k'}} \sum_g \mathrm{Tr}_\rho\big(\hat{T}^{g_{|y_k}}_{y_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}T^{g_{|z_k}}_{z_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}\big(\hat{T}^{g_{|y_k'}}_{y_k',\boldsymbol{x}_{\geq k}}\big)\otimes A^{g(\boldsymbol{x}_{<k},y_k)}_{\boldsymbol{x}_{\neg k},y_k}\big)+O\big(\sqrt{\mu(T,A)}\big)\\
= & \mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y_k'}} \sum_g \mathrm{Tr}_\rho\big(\hat{T}_{y_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}T^{g_{|z_k}}_{z_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}\big(\hat{T}^{g_{|y_k'}}_{y_k',\boldsymbol{x}_{\geq k}}\big)^\dagger\otimes A^{g(\boldsymbol{x}_{<k},y_k)}_{\boldsymbol{x}_{\neg k},y_k}\big)+O\big(\sqrt{\mu(T,A)}\big)\\
& - \mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y_k'}} \sum_{\substack{g,h\neq g_{|y_k}\\ h(\boldsymbol{x}_{<k})=g(\boldsymbol{x}_{<k},y_k)}} \mathrm{Tr}_\rho\big(\hat{T}^h_{y_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}T^{g_{|z_k}}_{z_k,\boldsymbol{x}_{\geq k}}\tilde{T}^{-1}\big(\hat{T}^{g_{|y_k'}}_{y_k',\boldsymbol{x}_{\geq k}}\big)^\dagger\otimes A^{g(\boldsymbol{x}_{<k},y_k)}_{\boldsymbol{x}_{\neg k},y_k}\big),\end{aligned}$$

where the last equality uses again consistency of T and A to bound the terms for which $h(\boldsymbol{x}_{<k})\neq g(\boldsymbol{x}_{<k},y_k)$. The last expression above is bounded using Claim 51. Hence we obtain

$$\mathrm{Tr}_\rho(V) \geq \mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y_k'}} \sum_{(a,b,c)\text{ aligned}} \mathrm{Tr}_\rho\big(\hat{T}_{y_k,\boldsymbol{x}_{\neg k}}\tilde{T}^{-1}T^c_{z_k,\boldsymbol{x}_{\neg k}}\tilde{T}^{-1}\big(\hat{T}^b_{y_k',\boldsymbol{x}_{\neg k}}\big)^\dagger\otimes A^a_{\boldsymbol{x}_{\neg k},y_k}\big)+O\big(n\sqrt{\mu(T,A)}\big)$$

Using again consistency of T and A, and linearity of A in the k-th direction to remove the dependence of the last A term on y_k, we get

$$\mathrm{Tr}_\rho(V) \geq \mathop{\mathrm{E}}_{\substack{\boldsymbol{x}_{>k}\\ y_k\neq z_k\neq y_k'}} \sum_{(a,b,c)\text{ aligned}} \mathrm{Tr}_\rho\big(\hat{T}_{\boldsymbol{x}_{\neg k}}\tilde{T}^{-1}T^c_{z_k,\boldsymbol{x}_{\neg k}}\tilde{T}^{-1}\big(\hat{T}^b_{y_k',\boldsymbol{x}_{\neg k}}\big)^\dagger\otimes A^a_{\boldsymbol{x}_{\neg k},y_k}\big)+O\big(n\sqrt{\mu(T,A)}\big). \tag{5.23}$$

By definition, $(\mathrm{Id}-\hat{T}_{\boldsymbol{x}_{>k}}\tilde{T}^{-1}) = U(\mathrm{Id}-\Sigma\tilde{\Sigma}^{-1})U^\dagger$ is the projector on the singular values of $T_{\boldsymbol{x}_{>k}}=\hat{T}_{\boldsymbol{x}_{>k}}\big(\hat{T}_{\boldsymbol{x}_{>k}}\big)^\dagger$ that are at most ε^d. Hence

$$\mathrm{E}_{\boldsymbol{x}_{\neg k},z_k} \sum_{(a,b,c)\text{ aligned}} \mathrm{Tr}_\rho\big((\mathrm{Id}-\hat{T}_{\boldsymbol{x}_{>k}}\tilde{T}^{-1})T^c_{z_k,\boldsymbol{x}_{\geq k}}(\mathrm{Id}-\hat{T}_{\boldsymbol{x}_{>k}}\tilde{T}^{-1})^\dagger\big) \;=\; O\big(\varepsilon^d\big),$$

which from (5.23) gives

$$\mathrm{Tr}_\rho(V) \geq \mathrm{E}_{\boldsymbol{x}_{\neg k},z_k\neq y_k\neq y_k'} \sum_{(a,b,c)\text{ aligned}} \mathrm{Tr}_\rho\big(\hat{T}^c_{z_k,\boldsymbol{x}_{\neg k}}\tilde{T}^{-1}\big(\hat{T}^b_{y_k',\boldsymbol{x}_{\neg k}}\big)^\dagger\otimes A^a_{\boldsymbol{x}_{\neg k},y_k}\big)-O\big(\varepsilon^{d/2}+n\sqrt{\mu(T,A)}\big). \tag{5.24}$$

We may now conclude by consistency of T and A. □

Chapter 6

Immunizing games against entanglement

In this chapter we introduce a novel technique to design games in a way that limits the provers' ability to use their entanglement in order to collude against the verifier. To this end we design a new test, which can be added generically to *any* two-prover one-round game, and significantly limits the use of entanglement by the provers beyond its utility as shared randomness. We hope that this technique of "immunizing" a game against entanglement can be extracted to serve a wider purpose in other contexts where we want to limit the power of entanglement, possibly with cryptographic applications.

The goal in designing the new test lies is to prevent the provers from using entanglement to coordinate their replies, and hence increase their success probability. Our main idea is to modify the game by introducing a *third* prover. We use the extra prover to introduce a consistency test that forces two of the provers to give the *same* answer. As a result, to pass this test, the two original provers can only use an entangled state of a specific form; it must be (approximately) *extendable*, i.e., it must be the density matrix of a symmetric tripartite state. There are prior results pointing to the potential usefulness of a third prover to limit the cheating power of entanglement. For example, two entangled provers can cheat in the Odd Cycle game of Ref. [26], but if we add a third prover, then entangled provers can perform no better than classical ones [118]. Moreover, after the completion of the work presented in this chapter we learned from Andy Yao [131] of a way to add a third prover to the Magic Square game such that as a result the winning probability of entangled provers is nearly 0.94.

Our result has an important application to the computational complexity of multiplayer entangled games. We show that it is NP-hard to approximate the value of three-player games. This was the first hardness of approximation result for such games. Our main result can be stated as follows:

Theorem 53. *There exists a polynomial p such that it is* NP*-hard to decide, for an explicitly* *given three-prover one-round classical entangled game G, whether its value is* 1 *or at most* $1 - 1/p(|G|)$.[1]

This theorem implies that no polynomial-time algorithm can compute the value of an entangled game to within polynomial precision. Given the importance of semi-definite programs (SDPs) in results on entangled games, the following immediate corollary is of interest:

Corollary 54. *The success probability of classical entangled three-prover games cannot be computed by SDPs of polynomial size, unless* $\mathrm{P} = \mathrm{NP}$.

The proof of Theorem 53 is based on showing that by enforcing certain tests on the provers we obtain sets of projectors (which characterize their strategy) which pairwise "*almost commute*". We already introduced this condition in Chapter 4, and in this chapter we show how it can be exploited to derive a classical strategy for the original classical game.

Remark. *The results presented in this chapter are a subset of the results published in the paper [68]. In particular, that paper proves a similar result to the one discussed above for the case of* quantum *two-prover games, in which the verifier and players may exchange quantum messages (cf. Section 3.4.1 for a definition). In this chapter we chose to focus on the case of three-prover games.*

Related work. The techniques developed in this chapter were subsequently applied by Ito et al. [61] to show similar results for *binary* three-prover one-round classical entangled games. They also give a new upper-bound for the value of these games; or, as often called in this context, they give a family of new n-partite Tsirelson inequalities. Ito, Kobayashi and Matsumoto [59] extended our proof technique and showed how a certain form of oracularization could be used to prove a similar hardness result for the case of *two-prover* one-round entangled games with classical messages and a constant answer size.

6.1 Proof overview

Reduction

We prove our NP-hardness result by a reduction from the hardness of approximation result for classical (non-entangled) games, as implied by the PCP Theorem, which we state in the language of games:

Theorem (PCP Theorem [10, 9]). *There is a constant $s < 1$ such that it is* NP*-hard to decide, given a two-prover one-round game with a constant number of answers, whether its value is* 1 *or at most* s.

We start with an instance of such a classical two-prover one-round game and modify it to a a three-prover one-round classical entangled game with the property that the value of

[1] See Section 3.4.1 in Chapter 3 for a precise definition of the size $|G|$ of G.

the new entangled game is at least as big as the value of the original game. In other words, if the value of the original game is 1, the value of the new game is still 1.

We then need to show that if the value of the original game is at most s, then the value of the new entangled game is bounded away from 1 by an inverse-polynomial. In order to prove this we reason by contrapositive and use a successful strategy of the entangled provers to construct a strategy in the original game that achieves a large value (see *Rounding* below).

Rounding

The extra prover allows us to extract a mathematical condition on the operations of the entangled provers. More precisely it turns out that the projectors corresponding to the various questions of the verifier pairwise “almost commute” in some sense or “almost do not disturb” the entangled state. This means that the provers' actions are “almost classical”, in the sense that they allow us to take any strategy for the entangled game and convert it back to a strategy in the original classical game. We call this conversion *rounding* from a quantum solution to a classical solution, in analogy to the rounding schemes used to convert a solution to an SDP relaxation to a solution of the game. To explain the idea of our new rounding scheme, consider the case of two-prover one-round entangled games. Assume that the provers, when receiving a question from the verifier, perform a projective measurement on their share of the entangled state depending on the question, and answer with the outcome they get. In the *exact* case, when the value of the three-player entangled game is 1, the measurements corresponding to different questions *commute* exactly. Hence, there is a common basis in which the projectors corresponding to different answers are all diagonal for all questions. In other words, for each question, the projectors simply define a partition of the basis vectors. The probability that the provers give a certain pair of answers just corresponds to the size of the overlap of the supports of the two corresponding projectors, i.e., to the number of basis vectors that are contained in both of them. We can now construct a classical strategy for the original game, where the provers use shared randomness to sample a basis vector, check which projector/partition contains it, and output the corresponding answer. This classical strategy achieves exactly the same probability distribution on the answers, and hence the same value of the game.

Matters become more complicated in the case where the value of the entangled game is larger than $1-\varepsilon$. Now, the provers' measurements corresponding to different questions “almost commute”. To exploit this property in a rounding scheme, imagine the following pre-processing step to eliminate entanglement from the strategy: Before the game starts, the provers apply in sequence all possible measurements, corresponding to all possible questions, on a share of the entangled state, and write down a list of all the answers they obtain. Then, during the game, when they receive a question from the verifier, they respond with the corresponding answer in their list[2]. Because the measurements almost commute, the answer to any one particular question in this sequential measurement scheme is similarly distributed

[2] Obviously, the provers do not really need any entanglement to do this: all they have to do is sample from

to the scenario in the entangled game, where the prover only performs the measurement corresponding to that question. This can be seen by "commuting" the corresponding projectors through the list of projectors in the measurement, where each time we commute two operators we lose some small amount in precision. As a result, the success probability of this new unentangled strategy is similar to the one in the entangled game, or at least not too low.

A new mathematical challenge

As mentioned above, our tests enforce an almost-commuting condition on the operators of the provers. If they commuted exactly, they would be diagonal in a common basis, meaning that the strategy is essentially classical and does not use entanglement. If one could conclude that the operators are *nearly diagonal* in some basis, one could again extract a classical strategy as in the exact case. Hence we reduce proving *constant* hardness of approximation to the question whether one can approximate our operators by commuting ones. This touches upon a deep question in operator algebra: *Do almost commuting matrices nearly commute?* Here *almost commuting* means that the commutator is small in some norm, and nearly commuting means that the matrices can be approximated by matrices that are diagonal in a common basis. This famous question was asked for *two Hermitian* matrices by Halmos back in 1976 [50].[3] It was shown subsequently [123],[4] using methods from algebraic topology, that this conjecture is false for two *unitary* matrices. Then, Halmos' conjecture was disproved for the case of three Hermitian matrices [124]. Finally Halmos' conjecture was proved [78] by a "long tortuous argument" [32] using von Neumann algebras, almost twenty years after the conjecture had been publicized. In our case we reduce proving hardness of approximation of the value of an entangled game to the conjecture for a set of pairwise almost commuting *projectors*, where the norm is the univariate ρ-norm introduced in Chapter 4:

$$\big\|A\big\|_{\mathbf{u},\rho}^2 := \mathrm{Tr}\big(AA^\dagger\rho\big),$$

and ρ is the reduced density of the provers' shared entangled state on either prover's subsystem.

Conjecture. *Let $W_1,\dots,W_n$ be d-dimensional projectors such that for some $\varepsilon \geq 0$ and for all $i,j \in \{1,\dots,n\}$, $\big\|\, W_iW_j - W_jW_i\,\big\|_{\mathbf{u},\rho}^2 \leq \varepsilon$. Then there exists a $\delta \geq 0$, and pairwise commuting projectors $\tilde{W}_1,\dots\tilde{W}_n$ such that $\big\|W_i - \tilde{W}_i\big\|_{\mathbf{u},\rho}^2 \leq \delta$ for all $i \in \{1,\dots,n\}$.*

Our proof shows that the conjecture with a constant δ implies hardness of approximation of the value of entangled games to within a *constant*, i.e., the best possible result. Moreover,

the joint distribution that corresponds to the distribution of all the answers in this sequence of measurements.

[3] For the operator norm.

[4] For a simpler, elegant proof see Ref. [39].

when scaled up to the setting of interactive proofs, the conjecture with a constant δ implies inclusion of NEXP in MIP$^*(3,1)$ with completeness 1 and soundness bounded away from 1.

For two, three or a constant number of projectors the conjecture is easy to prove for a constant δ. We do not know if it is true in general.

6.2 Hardness of three-prover classical entangled games

In this section we prove Theorem 53, which we now state as:

Theorem 55. *There is a constant $s_2 > 0$ such that it is* NP*-hard to decide, given a three-prover classical entangled game with a constant number of answers, whether its value is* 1 *or at most* $1 - \varepsilon$ *for* $\varepsilon = \frac{s_2}{|Q|^2}$.

We will prove this by a reduction from the PCP Theorem. We begin by describing how to modify any two-prover classical game $G = (\pi, V)$ (which is assumed to be symmetric per Lemma 13) to a three-prover classical game G' of equal or higher value.

The modified three-prover game

In the constructed game G' the verifier chooses one of the provers uniformly at random. Rename the chosen prover Alice and call the other provers Bob and Cleve. The verifier samples questions q and q' according to $\pi(q, q')$. He sends question q to both Alice and Cleve, and question q' to Bob. He receives answers a, a', and a'', respectively, and accepts iff the following are both true:

Classical Test: The answers of Alice and Bob would win the game G, i.e., the answers a and a' satisfy $V(a, a' \mid q, q') = 1$.

Consistency Test: Alice and Cleve give the same answer, i.e., $a = a''$.

It is clear that the value of the constructed game is at least as large as the value of the original game G: if the provers reply according to an optimal classical strategy (which can be assumed to be symmetric per Lemma 13), they always pass the consistency test. Also, it is clear that the size of the description of the constructed game is linearly related to the size of the description of the original game, hence we have the same complexity parameter.

To prove Theorem 55, we need to show the following lemma.

Lemma 56. *If $\omega^*(G') > 1 - \varepsilon$, then $\omega(G) > s$.*

Here $\varepsilon = \frac{s_2}{|Q|^2}$ for the constant s_2 in Theorem 55 and the constant s is from the PCP Theorem.

Proof. Consider an entangled strategy for G' that succeeds with probability greater than $1-\varepsilon$.[5] Since the game G' is symmetric, we can assume that this strategy is symmetric, per Lemma 13. Suppose that the provers share a symmetric state $|\Psi\rangle \in \mathcal{H}_{\mathrm{A}} \otimes \mathcal{H}_{\mathrm{B}} \otimes \mathcal{H}_{\mathrm{C}}$, where each $\mathcal{H}_{\mathrm{A}}$, $\mathcal{H}_{\mathrm{B}}$, and $\mathcal{H}_{\mathrm{C}}$ is isomorphic to a same $\mathcal{H}$. Let $\rho_{\mathrm{AB}} = \mathrm{Tr}_{\mathcal{H}_{\mathrm{C}}} |\Psi\rangle\langle\Psi|$ be the reduced state of $|\Psi\rangle\langle\Psi|$ on Alice and Bob. When asked question q_i, each prover measures their part of $|\Psi\rangle$. Following standard arguments (extending the private space of the provers) we can assume that this measurement is projective. Let $W_{q_i}^{a_i}$ be the projector corresponding to question q_i and answer a_i. This defines the entangled strategy for G'; it passes the classical test with probability

$$\pi_1 = \sum_{a,a',q,q'} \pi(q,q') V(a,a' \mid q,q') p_{\mathrm{ent}}(a,a' \mid q,q'),$$

where

$$p_{\mathrm{ent}}(a,a' \mid q,q') = \mathrm{Tr}\left(W_q^a \otimes W_{q'}^{a'} \rho_{\mathrm{AB}}\right) = \langle\Psi| W_q^a \otimes W_{q'}^{a'} \otimes \mathrm{Id} |\Psi\rangle. \tag{6.1}$$

It passes the consistency test with probability $\pi_2 = \sum_q \pi(q)\pi_2(q)$, where $\pi(q)$ is the marginal of $\pi(q,q')$ and

$$\pi_2(q) = \sum_a \mathrm{Tr}\left(W_q^a \otimes W_q^a \rho_{\mathrm{AB}}\right) = \sum_a \langle\Psi| W_q^a \otimes W_q^a \otimes \mathrm{Id} |\Psi\rangle, \tag{6.2}$$

where we made use of the symmetry. Note that $\pi_1, \pi_2 > 1-\varepsilon$.

Eqs. (6.1) and (6.2) clarify the role of the third prover, Cleve. The main purpose of introducing the third prover is *not* to allow the two tests to be performed at the same time: Indeed, it is possible to modify the protocol so that the verifier chooses two of the provers at random (say Alice and Bob) and only sends questions to them, not interacting with the third prover at all.[6] The presence of the third prover would not be important if the provers were executing a classical strategy, but it can (and does) make a difference if their strategy requires entanglement. Indeed, if there were only two provers, then they could share any state ρ_{AB}, whereas here we require that ρ_{AB} be *extendable*, i.e., it must be the reduced density matrix of a symmetric tripartite state. To give a concrete example, it is not possible for ρ_{AB} to be the maximally entangled state $|\Psi^-\rangle\langle\Psi^-|$. This is termed *monogamy of entanglement* [127].

Rounding to a classical strategy

We construct a classical strategy for G from the entangled strategy for G' in a similar fashion as in the case of quantum entangled games, with

$$D(a_1,\ldots,a_n,a_1',\ldots,a_n') = \left\| W_{q_n}^{a_n} \cdots W_{q_1}^{a_1} \otimes W_{q_n}^{a_n'} \cdots W_{q_1}^{a_1'} \otimes \mathrm{Id} |\Psi\rangle \right\|^2,$$

[5] We in fact consider a strategy with finite entanglement that has success probability greater than $1-\varepsilon-\delta$ for some $\delta = O(\varepsilon)$, which we will not write.

[6] With probability p, he sends them different questions and performs the classical test; with probability $1-p$, he sends the same question and performs the consistency test—this modification does not materially change our conclusions, but it does weaken the bounds in Theorem 55.

where $q_1, \ldots, q_n$ is an ordering of the questions in Q such that $\pi(q_1) \geq \pi(q_2) \geq \ldots \geq \pi(q_n)$. As before, we define $p_{\text{class}}(a_i, a'_j \mid q_i, q_j)$ to be the marginal of D on a_i, a'_j.

Lemma 57. *The (weighted) statistical distance between* p_{class} *and* p_{ent} *is*

$$\Delta(p_{\text{class}}, p_{\text{ent}}) = \sum_{q,q'} \pi(q,q') \sum_{a,a'} \left| p_{\text{class}}(a, a' \mid q, q') - p_{\text{ent}}(a, a' \mid q, q') \right| < 12|Q|\sqrt{\varepsilon}.$$

We first show how this lemma proves Lemma 56. Since the strategy in the entangled game passes the classical test with probability greater than $1 - \varepsilon$, the classical strategy succeeds in the original game with probability greater than $1 - \varepsilon - \Delta(p_{\text{class}}, p_{\text{ent}}) > 1 - \varepsilon - 12|Q|\sqrt{\varepsilon}$. For $\varepsilon = \frac{s_2}{|Q|^2}$ and for sufficiently small constant s_2, this probability is larger than s. □

Why is Lemma 57 true? Rather than showing that the order of measurements is not important (although it will turn out in hindsight that this is true), we show that each measurement does not disturb ρ_{AB} very much. The key observation is as follows. Assume the provers pass the consistency test with high probability. If a particular measurement result occurs with certainty, the quantum state cannot be changed by the measurement. We use this fact in the following way: suppose Cleve were to perform the measurement corresponding to question q and assume he obtains an outcome a. Then, if Alice is asked question q, she must also give answer a with high probability. Thus her measurement does not change the quantum state much. However, since quantum theory is non-signaling, it cannot matter who measured first. It follows that Alice's measurement does not change ρ_{AB} much. Note that only the bipartite state ρ_{AB} is approximately unchanged — Alice's measurement can change the tri-partite state $|\Psi\rangle\langle\Psi|$ considerably. We then use a hybrid argument to show that performing all the measurements one after the other also leaves ρ_{AB} approximately unchanged.

Proof of Lemma 57. Let $\mathcal{W}_q$ be the superoperator corresponding to the projective measurement performed on question q, i.e., $\mathcal{W}_q(\sigma) = \sum_a W_q^a \sigma (W_q^a)^\dagger$ is the post-measurement state after performing $\{W_q^a\}$ on state σ.

To quantify how much a measurement changes a state we use Winter's Gentle Measurement Lemma, as state in Lemma 117 in Appendix A. The following simple corollary quantifies how much the measurement $\mathcal{W}_q \otimes \text{Id}$ changes ρ_{AB}:

Claim 58. *The trace distance between* $\mathcal{W}_q \otimes \text{Id}(\rho_{\text{AB}})$ *and* ρ_{AB} *is bounded by*

$$\|\mathcal{W}_q \otimes \text{Id}(\rho_{\text{AB}}) - \rho_{\text{AB}}\|_{\text{tr}} \leq 6\sqrt{1 - \pi_2(q)}.$$

Proof. Using the relations $\mathcal{W}_q \otimes \text{Id}(\rho_{\text{AB}}) = \text{Tr}_{\mathcal{H}_\text{C}}(\mathcal{W}_q \otimes \text{Id} \otimes \text{Id}(|\Psi\rangle\langle\Psi|))$ and $\rho_{\text{AB}} = \text{Tr}_{\mathcal{H}_\text{C}}(\text{Id} \otimes$

$\mathrm{Id} \otimes \mathcal{W}_q(|\Psi\rangle\langle\Psi|))$,

$$\begin{aligned}
&\|\mathcal{W}_q \otimes \mathrm{Id}(\rho_{\mathrm{AB}}) - \rho_{\mathrm{AB}}\|_{\mathrm{tr}} \\
&\quad\le \|\mathcal{W}_q \otimes \mathrm{Id} \otimes \mathrm{Id}(|\Psi\rangle\langle\Psi|) - \mathrm{Id} \otimes \mathrm{Id} \otimes \mathcal{W}_q(|\Psi\rangle\langle\Psi|)\|_{\mathrm{tr}} \\
&\quad\le \Big\|\mathcal{W}_q \otimes \mathrm{Id} \otimes \mathrm{Id}(|\Psi\rangle\langle\Psi|) - \sum_a W_q^a \otimes \mathrm{Id} \otimes W_q^a|\Psi\rangle\langle\Psi|W_q^a \otimes \mathrm{Id} \otimes W_q^a\Big\|_{\mathrm{tr}} \\
&\qquad + \Big\|\sum_a W_q^a \otimes \mathrm{Id} \otimes W_q^a|\Psi\rangle\langle\Psi|W_q^a \otimes \mathrm{Id} \otimes W_q^a - \mathrm{Id} \otimes \mathrm{Id} \otimes \mathcal{W}_q(|\Psi\rangle\langle\Psi|)\Big\|_{\mathrm{tr}} \\
&\quad\le 2\Big\|\sum_a W_q^a \otimes \mathrm{Id} \otimes W_q^a|\Psi\rangle\langle\Psi|W_q^a \otimes \mathrm{Id} \otimes W_q^a - \mathrm{Id} \otimes \mathrm{Id} \otimes \mathcal{W}_q(|\Psi\rangle\langle\Psi|)\Big\|_{\mathrm{tr}} \\
&\quad\le 6\sqrt{1-\pi_2(q)},
\end{aligned}$$

by monotonicity of the trace distance under partial trace, the triangle inequality, symmetry, and then taking $\rho = \bigoplus_a W_q^a \otimes \mathrm{Id} \otimes \mathrm{Id}|\Psi\rangle\langle\Psi|W_q^a \otimes \mathrm{Id} \otimes \mathrm{Id}$ and $X = \bigoplus_a \mathrm{Id} \otimes \mathrm{Id} \otimes W_q^a$, $Y = \mathrm{Id}$ in Lemma 117. □

For $1 \le i, j \le n$, let

$$\rho_{\mathrm{AB}}(i,j) = (\mathcal{W}_{q_{i-1}} \circ \cdots \circ \mathcal{W}_{q_1}) \otimes (\mathcal{W}_{q_{j-1}} \circ \cdots \circ \mathcal{W}_{q_1})\rho_{\mathrm{AB}}.$$

Then

$$p_{\mathrm{class}}(a_i, a'_j \mid q_i, q'_j) = \mathrm{Tr}\Big(\big(W_{q_i}^{a_i} \otimes W_{q'_j}^{a'_j}\big)\rho(i,j)\Big).$$

Hence we can bound $\sum_{a_i,a'_j} |p_{\mathrm{class}}(a_i, a'_j \mid q_i, q'_j) - p_{\mathrm{ent}}(a_i, a'_j \mid q_i, q'_j)|$ by bounding $\|\rho(i,j) - \rho\|_{\mathrm{tr}}$, since the trace distance between two states is an upper bound on the variation distance of the probability distribution resulting from making any measurement on those two states.

The following technique was introduced by Ambainis, Nayak, Ta-Shma and Vazirani [7] and has been used extensively by Aaronson [1, 2].

Claim 59. *The trace distance between $\rho_{\mathrm{AB}}(i,j)$ and ρ_{AB} is bounded by*

$$\|\rho_{\mathrm{AB}}(i,j) - \rho_{\mathrm{AB}}\|_{\mathrm{tr}} \le 6\sum_{i'=1}^{i-1} \sqrt{1-\pi_2(q_{i'})} + 6\sum_{j'=1}^{j-1} \sqrt{1-\pi_2(q_{j'})}.$$

Proof. By induction. The claim is clearly true for $(i,j) = (1,1)$. Given it is true for a particular value of (i,j), we show it is also true for $(i+1,j)$. In view of the symmetry, this is sufficient to establish the claim. We have

$$\begin{aligned}
&\|\rho_{\mathrm{AB}}(i+1,j) - \rho_{\mathrm{AB}}\|_{\mathrm{tr}} \\
&\quad\le \|\rho_{\mathrm{AB}}(i+1,j) - \mathcal{W}_{q_i} \otimes \mathrm{Id}(\rho_{\mathrm{AB}})\|_{\mathrm{tr}} + \|\mathcal{W}_{q_i} \otimes \mathrm{Id}(\rho_{\mathrm{AB}}) - \rho_{\mathrm{AB}}\|_{\mathrm{tr}} \\
&\quad\le \|\mathcal{W}_{q_i} \otimes \mathrm{Id}\,(\rho_{\mathrm{AB}}(i,j) - \rho_{\mathrm{AB}})\,\|_{\mathrm{tr}} + 6\sqrt{1-\pi_2(q_i)} \\
&\quad\le \|\rho_{\mathrm{AB}}(i,j) - \rho_{\mathrm{AB}}\|_{\mathrm{tr}} + 6\sqrt{1-\pi_2(q_i)},
\end{aligned}$$

where we used the triangle inequality, Claim 58, and monotonicity of the trace distance. □

Putting everything together, it follows that

$$\begin{aligned}
\Delta(p_{\text{class}}, p_{\text{ent}}) &\leq \sum_{i,j=1}^{n} \pi(q_i, q'_j) \|\rho_{\text{AB}}(i,j) - \rho_{\text{AB}}\|_{\text{tr}} \\
&\leq 6 \sum_{i,j=1}^{n} \pi(q_i, q'_j) \left(\sum_{i'=1}^{i-1} \sqrt{1 - \pi_2(q_{i'})} + \sum_{j'=1}^{j-1} \sqrt{1 - \pi_2(q_{j'})} \right) \\
&= 12 \sum_{i=1}^{n} \sum_{i'=1}^{i-1} \pi(q_i) \sqrt{1 - \pi_2(q_{i'})} \\
&\leq 12|Q| \sum_{i'=1}^{n} \pi(q_{i'}) \sqrt{1 - \pi_2(q_{i'})} \\
&\leq 12|Q| \sqrt{1 - \pi_2} \\
&< 12|Q| \sqrt{\varepsilon},
\end{aligned}$$

since $\sqrt{1-x}$ is concave and $\pi_2 = \sum_q \pi(q)\pi_2(q) > 1 - \varepsilon$.

Chapter 7

Parallel repetition of entangled games

One of the most interesting questions in the context of multiplayer games is the parallel repetition question. It is well known that one can reduce both the value and the entangled value of a game by repeating it sequentially, or alternatively, by repeating it in parallel with several independent pairs of players. However, for many applications (like hardness of approximation results or amplifications preserving zero-knowledge) we need a way to decrease the winning probability without increasing the number of rounds or the number of players, i.e., while staying in the model of two-player one-round games. Parallel repetition is designed to do just that: in its most basic form, in the ℓ-parallel repeated game, the referee simply chooses ℓ pairs of questions independently and sends to each player his corresponding ℓ-tuple of questions. Each player then replies with an ℓ-tuple of answers, which are accepted if and only if each of the ℓ answer pairs would have been accepted in the original game.

Clearly the value of an ℓ-parallel repeated game is *at least* the ℓ-th power of the value of the original game, since the players can just answer each of the ℓ questions independently as in the original protocol. However, contrary to what intuition might suggest and to the case of sequential repetition, parallel repetition does *not* necessarily decrease the value of a game in a straightforward exponential manner[1]. The parallel repetition question is that of finding *upper bounds* on the value of a repeated game, and for a long time no such upper bound, even very weak, could be proved. For the case of *classical* two-player games, first results date to Verbitsky [122] who showed that indeed the value goes to zero with the number of repetitions. Following this, Feige and Kilian [42] showed that the value decreases polynomially with the number of repetitions for the special case of so-called *projection games* (in which the second player's answer is uniquely determined by the first player's). They used a modified parallel repetition procedure in which a large fraction of the repetitions are made of *dummy* rounds, that is, rounds in which the questions are chosen independently at random for both players, and in which any answer is accepted. In this chapter we deviate somewhat

[1] See [41] for a classical example, and [27] for an example using entangled players due to Watrous. See also [65] for another example where parallel repetition does not reduce the value of a game at the exact rate one would expect if the players were playing independently.

from the common terminology, and use the term "parallel repetition" even when referring to such more general procedures. Finally, in a breakthrough result, Raz [96] showed that the value of a classical game repeated in parallel indeed decreases exponentially with the number of repetitions (albeit not exactly at the same rate as sequential repetition). There is still very active research in this area, mostly on simplifying the analysis, which, over a decade later, remains quite involved, and improving it for certain special cases of games [53, 94, 43, 95, 13, 14, 11, 99].

Previous work

The only two previous results regarding the parallel repetition of entangled games are for two special classes of games. First, Cleve et al. showed that for the class of *XOR games* (i.e., games with binary answers in which the referee's decision is based solely on the XOR of the two answers), *perfect* parallel repetition holds [27]. This means that the entangled value of an ℓ-parallel repeated game is exactly the ℓ-th power of the entangled value of the original game. Parallel repetition has also been shown to hold for the more general (but still quite restricted) class of *unique games* [66] (i.e., games where the referee applies some permutation to the answers of the second player and accepts if and only they match those from the first player). One might also add a third result by Holenstein [53], who proved a parallel repetition theorem for the so-called *no-signaling value*; since the no-signaling value is an upper bound on the entangled value, this can sometimes be used to upper bound the entangled value of repeated games. However, there is in general no guarantee regarding the quality of this upper bound, and in many cases (e.g., all unique games) the no-signaling value is always 1, making it useless as an upper bound on the entangled value.

It is important to note that in these results the entangled value of the parallel repeated game is never analyzed directly; instead, one uses a "proxy" such as a semidefinite program [27, 66] or the no-signaling value [53], whose behavior under parallel repetition is well understood. Moreover, in all these cases, the proxy's value is efficiently computable. This unfortunately gives a very strong indication that such techniques cannot be extended to deal with general games. Indeed, it is known that it is NP-hard to tell if the entangled value of a given game is 1 or not [68, 59]; hence, unless P=NP, for any efficiently computable upper bound on the entangled value, there are necessarily games whose entangled value is strictly less than 1 yet for which that upper bound is 1 (and such games can often be exhibited explicitly without relying on P$\neq$NP). We note that some of the early parallel repetition results for the classical value [44] followed the same route (of upper bounding the value by a semidefinite program) and were limited to special classes of games for the exact same reason.

In parallel to work on the parallel repetition problem, the related question of *product testing* arose in the context of error amplification for PCPs [37, 35, 54, 56]. Roughly speaking, the question here is to design tests by which a referee can check that the players play according to a *product strategy*, i.e., answer each question independently of the other questions (as one would expect from an honest behavior). Note that if the players are constrained to follow a product strategy, then their maximum winning probability must necessarily go down exponentially, hence the connection to the parallel repetition question. The result of Feige and Kilian [42] mentioned above in fact also shows that the strategy of the players must have some product structure, and recently there has been lots of renewed interest in this question leading to much stronger product testers [36]. In the case of entangled players, however, prior to this work nothing was known, raising the following question: Is there a way to test if the strategy of entangled players is in some sense close to a product strategy?

Results

The main result of this chapter can be informally stated as follows.

Theorem 60 (informal). *For any $s < 1$, $\delta > 0$, and entangled game G, there is a corresponding ℓ-parallel repeated game G', where $\ell = \mathrm{poly}((1-s)^{-1}, \delta^{-1})$, such that if the value of G is less than s then the value of G' is at most δ, whereas if the value of G is 1 then this also holds*[2] *for the repeated game.*

The dependency of ℓ on δ in Theorem 60 is polynomial, whereas as we already mentioned it is known that in some cases this dependence can be made poly-logarithmic (and this is certainly the case if the players are assumed to play independently). While a poly-logarithmic dependence is important in some applications for which one would like to perform amplification up to an exponentially small value, in many cases the main use of parallel repetition is to amplify a small "gap" between value 1 and value $1 - 1/poly(|G|)$ to a constant gap, say between 1 and 1/2. In this case the polynomial dependence of ℓ on $(1-s)^{-1}$ that we obtain is optimal (up to the exact value of the exponent).

In the course of the proof of this theorem we also establish that the player's strategies have a certain "serial" or "product" structure (more on this in the proof ideas and techniques section below). The informal statement above hides some details, which we now discuss. The kind of parallel repetition we perform depends on the structure of the game G, and we distinguish whether it is a projection game or not.

Repetition for projection games. If G is a projection game, then the repeated game is obtained by independently playing the original G on a subset of the repetitions, and playing dummy rounds in the other repetitions. We note that projection games form a wide class of games that captures most of the games one typically encounters in the classical literature (see [94]).

If, in addition, the game happens to be a *free* game (i.e., a game in which the referee's distribution on question pairs is a product distribution), then the dummy questions are no longer needed and hence our analysis applies to the *standard* ℓ-fold repetition.

[2] See the discussion following the theorem for some caveats.

Repetition for general games. If the game G does not have the projection property, then it is necessary to add a number of *consistency* rounds to the repetition. In those rounds the referee sends identical questions to the players, and expects identical answers. As before, the other rounds of the repetition are either the game G or dummy rounds. The consistency questions are added to play the role of the projection constraints.

This kind of repetition raises the following issue[3]: namely, it is not obvious that honest entangled players can answer the consistency questions correctly. This implies that, even if the original game had value 1, players might not be able to succeed in the consistency questions and hence the value of the repeated game might not equal 1 anymore. This may or may not be an issue depending on where the original game comes from. In many cases it is known that, if there is a perfect strategy, it does not require any entanglement at all, or it can be achieved using the maximally entangled state. In both cases it is not hard to see that players will be able to answer consistency questions perfectly, and hence our result holds. Because of this we regard this issue as a minor one; however it might be important in some contexts.

Proof idea and techniques

We focus on the case of projection games, as the proof of the other cases does not present additional challenges. The starting point of our proof is the work of Feige and Kilian [42], for which the following intuition can be given[4]. Our goal as the referee is to force the players to use a product strategy, preventing any elaborate cheating strategies. In other words, we want to make sure that the player chooses his answer to the ith question based only on that question and not on any of the other $\ell - 1$ questions. Towards this end, the referee chooses a certain (typically large) fraction of the ℓ question pairs to be independently distributed *dummy questions*, the answers to which are ignored. These dummy questions are meant to confuse the players: if they were indeed trying to carefully choose their answer to a certain question by looking at many other questions, now most of these other questions will be completely random and uncorrelated with the other player's questions, so that such a strategy cannot possibly be helpful.

In more detail, Feige and Kilian prove the following dichotomy theorem on the structure of single-player repeated strategies (that is, maps from ℓ-tuples of questions to ℓ-tuples of answers): either the strategy looks rather *random* (in which case the players cannot win the game with good probability — this is where the projection property is used) or it is almost a *serial* or *product* strategy, i.e., the answer to each question is chosen based on that question only (in which case the player is playing the rounds independently, and his success probability will suffer accordingly).

[3] This is why we treat the projection case separately, despite it leading to similar decay.

[4] We refer to Ryan O'Donnell's excellent lecture notes [89, 88] for a helpful exposition of Feige and Kilian's proof.

Our proof follows a similar structure. However, an important challenge immediately surfaces: the proof in [42], and indeed *all* proofs of parallel repetition theorems or direct product tests, make the important initial step of assuming that the player's strategies are deterministic (which is easily seen to hold without loss of generality). And indeed, it is not at all trivial to extend those proofs to even the randomized setting without making this initial simplifying assumption. To give a simple example, an important notion in Feige and Kilian's proof is that of a *dead* question — simply put, a question to which the player does not give any majority answer, when one goes over all possible ways of completing that specific question into a tuple of questions for the repeated game. It is easily seen that, in the case of a deterministic strategy, dead questions are harmful, as the players are unlikely to satisfy the projection property on them. However, it is just as easily seen that for most randomized strategies, good or bad, *all* questions are dead.

This illustrates the kinds of difficulties that one encounters while trying to show parallel repetition in the case of entangled players, when one cannot simply "fix the randomness". The issue we just raised is not too hard to solve, and others are more challenging. Indeed the main difficulty is to define a proper notion of *almost serial* for operators, which would in particular incorporate the inherent randomness of quantum strategies. It turns our that the right notion is the notion of consecutive measurements (rather than tensor products of measurements for each question, a tempting but excessively strong possibility). Based on a quantum analogue of Feige and Kilian's dichotomy theorem, we are able to show that the almost serial condition induces a condition of *almost orthogonality* on the player's operators. At this point we use a variant of the *orthogonalization lemma* introduced in Section 4.3.1 of Chapter 4, which lets us extract a *product* strategy from the almost-orthogonal condition. We obtain that the players approximately perform a series of consecutive measurements, each depending only on the current question. An upper bound on the value of the repeated game then follows.

Repeated games

We will consider two different types of repeated games. The first one, originally used by Feige and Kilian, applies to projection games, and we describe it in Definition 61. The second type of repetition applies to consistency games, and is closer to the direct product testing technique originally introduced by Dinur and Reingold [37]; we explain it in Definition 62.

Definition 61 (Feige-Kilian repetition)**.** *Let ℓ be any integer, and define $C_1 := \ell^{1/2}$ and $C_2 := \ell - C_1$. Given a two-player projection game $G = (\pi, V, Q, A)$, its ℓ-th Feige-Kilian repetition is the following game $G_{FK(\ell)}$:*

- *The referee picks a random partition $[\ell] = M \cup F$, where $|M| = C_1$ and $|F| = C_2 = \ell - C_1$. Indices in M will be called "game" indices, while indices in F will be called "confuse" indices.*

- *The referee picks* $(q'_M, q_M) \sim_{\pi^{C_1}} (Q \times Q)^{C_1}$.
- *He picks* $(q'_F, q_F) \sim_{(\pi_A \times \pi_B)^{C_2}} (Q \times Q)^{C_2}$*, where* π_A *is the marginal of* π *on the first player, and* π_B *the marginal on the second player.*
- *The referee sends the questions to the players (without specifying which questions are of which type). On game questions he verifies that the original game constraint is satisfied. He accepts any answers to confuse questions.*

Definition 62 (Dinur-Reingold repetition)**.** *Let* ℓ *be any integer, and define* $C'_1 := \ell^{1/2}$*,* $C_1 = 2C'_1$ *and* $C_2 := \ell - C_1$*. Given a two-player symmetric game* $G = (\pi, V, Q, A)$*, its* ℓ*-th Dinur-Reingold repetition is the following game* $G_{DR(\ell)}$*:*

- *The referee picks a random partition* $[\ell] = R \cup G \cup F$*, where* $|R| = C'_1$*,* $|G| = C'_1$*, and* $|F| = C_2$*. Indices in R will be called "consistency" indices, those in G will be called "game" indices, and those in F "confuse" indices.*
- *The referee picks* C'_1 *questions* $q_R \sim_{\pi_A^{C'_1}} Q^{C'_1}$ *and sets* $q'_R = q_R$*, where* π_A *is the marginal of* π *on the first player (since we assumed G was symmetric, this is the same as* π_B*, the marginal on the second player).*
- *The referee picks* C'_1 *pairs of questions* $(q'_G, q_G) \sim_{\pi^{C'_1}} (Q \times Q)^{C'_1}$.
- *He picks* $(q'_F, q_F) \sim_{(\pi_A \times \pi_B)^{C_2}} (Q \times Q)^{C_2}$.
- *The referee sends the questions to the players (without specifying which questions are of which type). On consistency questions he verifies that both answers, from Alice and from Bob, are identical. On game questions he verifies that the original game constraint is satisfied. He accepts any answers to confuse questions.*

Note that, if a game G has value 1, then its Dinur-Reingold repetition does not necessarily also have value 1, as the player's optimal strategy in G might not be *consistent.* A consistent strategy is one in which whenever the players are asked the same question they provide the same answer with certainty. This may not always hold of an optimal strategy; nevertheless Lemma 13 from Chapter 3 shows that we can assume it holds in some natural settings. That lemma shows that, if G is any game, then we may symmetrize it and assume that the provers are also playing according to a symmetric strategy. In particular, if G had value 1, and the optimal strategy used either no entanglement or a maximally entangled state, then this also holds of the optimal strategy in the symmetrized game. Such a strategy is automatically consistent.

7.1 Proof overview

We first give a formal account of our results in the next section, before proceeding to give a detailed overview of their proof in Section 7.1.2.

7.1.1 Results

We first state our main theorems. They refer to the two types of repetition of an entangled game G defined in the previous section, its ℓ-th *Feige-Kilian repetition* $G_{FK(\ell)}$, and its ℓ-th *Dinur-Reingold* repetition $G_{DR(\ell)}$. Both types of repeated games are made of ℓ independent rounds, played in parallel. Some of these rounds consist of independent repetitions of G, while others are either *confuse* or *consistency* rounds, containing simple tests independent of the original game (except for the distribution with which questions are chosen in those rounds). Our first result pertains to projection games.

Theorem 63. *There exists a constant $c \geq 1$ such that, for all $s < 1$ and $\delta > 0$ there is a $\ell = O((\delta^{-1}(1-s)^{-1})^c)$ such that, if G is a projection game with value $\omega^*(G) \leq s$, then the entangled value of the game $G_{FK(\ell)}$ is at most δ. Moreover, if the value of G is 1 then the value of $G_{FK(\ell)}$ is also 1.*

In the case of free projection games, questions to the players are chosen independently, so that the distribution on questions in the confuse rounds of the game $G_{FK(\ell)}$ is exactly the same as that in the original game. The only difference is that in such a round, all answers are accepted, which can only help the players. Hence the direct parallel repetition of G has a smaller value than its Feige-Kilian repetition, which implies the following.

Corollary 64. *Let $s < 1$ and $\delta > 0$. Then there is a $\ell = O((\delta^{-1}(1-s)^{-1})^c)$ such that, if G is a free projection game such that $\omega^*(G) \leq s$, then the (direct) ℓ-fold parallel repetition of G has value at most δ.*

Our second result is more general, as it applies to arbitrary games. It only comes with the mild caveat that, in order to preserve the fact that the original game had value 1 (whenever this indeed holds), it is required that in that case there also exists a perfect strategy which is consistent.

Theorem 65. *There exists a constant $c \geq 1$ such that, for all $s < 1$ and $\delta > 0$ there is a $\ell = O((\delta^{-1}(1-s)^{-1})^c)$ such that, if G is an arbitrary game with value $\omega^*(G) \leq s$, then the entangled value of the game $G_{DR(\ell)}$ is at most δ. Moreover, if G has a perfect consistent strategy then the value of $G_{DR(\ell)}$ is also 1.*

Lemma 13 shows that the requirement that G has a perfect consistent strategy (which is only a requirement in cases where we are interested in preserving the fact that G might have value 1) is satisfied for many examples of games, including those for which we know a priori that, if the value of G is 1, then there is an optimal strategy that either does not use any entanglement at all, or uses the maximally entangled state.

7.1.2 Proof overview

In the remainder of this section we describe the main ideas behind the proof of Theorem 63 and Theorem 65; full details can be found in Section 7.2. Our goal is to understand *repeated* quantum strategies, that is, maps $q \in Q^\ell \mapsto \{ X_q^a \}_{a\in A^\ell}$ which map tuples of questions $q = (q_1, \dots, q_\ell)$ to projective measurements $\{ X_q^a \}_{a\in A^\ell}$ in dimension d. The semantics are that, on receiving questions q, a player measures his share of the entangled state $|\Psi\rangle$ according to $\{ X_q^a \}_{a\in A^\ell}$, resulting in him sending back answer a with probability $\langle\Psi|\mathrm{Id} \otimes X_q^a|\Psi\rangle$. Interestingly, most of the proof will be directly concerned with the measurements $\{ X_q^a \}_{a\in A^\ell}$ themselves (together with the reduced density $\rho = \mathrm{Tr}_A|\Psi\rangle\langle\Psi|$), without reference to the other player's measurements or even the underlying game.

We will be interested in a strategy's *marginals*: given a fixed subset of indices $S \subseteq [\ell]$ and a set of questions q_S on the indices in S, one can define the marginalized measurement

$$\Big\{ X_{q_S}^{a_S} : \rho \mapsto \mathrm{E}_{q\sim\pi^{[\ell]\setminus S}} \Big[\sum_{a\in A^{[\ell]\setminus S}} \sqrt{X_{q_S q}^{a_S\, a}}\, \rho\, \sqrt{X_{q_S q}^{a_S\, a}} \Big] \Big\}_{a_S \in A^S}$$

which corresponds to choosing a tuple $q \in Q^{[\ell]\setminus S}$ by picking the question in each coordinate independently according to some fixed distribution π,[5] making the measurement corresponding to the POVM described by $\{ X_{q_S q}^{a_S a} \}_{(a_S,a)\in A^\ell}$, and marginalizing over those answers a corresponding to indices not in S.

Given that X was a projective measurement, the marginalized strategy is a POVM — it is not necessarily projective any more. Our main results will pertain to the structure of such marginalized strategies. We will show that they are either very random (this is formally called *dead* later on, and morally means that the marginalized strategy is very far from a projective measurement; rather its singular values tend to be small and spread out), or highly structured (this is called *serial* later on, and after some work we will show that it implies that the marginalized strategy has somewhat of a product form, i.e. it can be decomposed as a product $\Pi_{q_1}^{a_1} \cdots \Pi_{q_s}^{a_s}$ on a subset of the coordinates). The attentive reader might already see that once this is proven it will be possible to bound the success probability of both types of strategies in the repeated game; however we should warn that the exact statements, and their proofs, are quite technical and carry only a fair share of the intuition we have just given.

We proceed to give a few more details on the structure of the proof of our results. It can be divided into three main steps. The first two steps establish facts about the structure of repeated single-player strategies, and are independent of the game being played, as well as of the other player's strategy.

[5] We will often drop the reference to π and simply write $\mathrm{E}_q[\cdot]$. π will be fixed throughout, and later instantiated to the (marginal) distribution on questions from the original game G that is being repeated.

Step 1: A quantum dichotomy theorem. In the first step we prove an analogue of Feige and Kilian's dichotomy theorem [42]. The precise statement is given in Lemma 68, and its simple proof very closely follows that of Feige and Kilian's theorem. Informally, it states that there exists an integer $1 \leq r^* \ll \ell$, such that a tuple of questions (R, q_R), where $R \subseteq [\ell]$ denotes a subset of r^* indices, and q_R fixed questions in those positions, can be of two types only. Either it is *dead* (case 1 in the lemma), or it is $(1-\eta)$-*serial*, where $\eta > 0$ is a small parameter (case 2 in the lemma). Both types of strategies are precisely defined in Definition 67, and the meaning of dead is the easiest to grasp. The technical definition is simply that the (marginalized) measurement $\{X_{q_R}^{a_R}\}_{a_R \in A^R}$, when performed twice (sequentially) on the same half[6] of the state $|\Psi\rangle$, is unlikely to produce the same result. This kind of strategy is easily seen to be bad for the players, as is shown in step 3. of the proof.

Serial strategies are more subtle. In the case of a classical deterministic player, a serial strategy is such that, when one conditions on the player giving answers a_R to the questions q_R in R, the answers to most other questions (not in R) are for the most part determined by the player as a direct function of the corresponding question, i.e. he is playing an honest product strategy on those coordinates. In the quantum case, we will adopt a seemingly weaker definition, which is that a strategy is serial on (q_R, a_R) if, in expectation over the choice of an additional question q_i in position i, when the marginalized measurement $\{X_{q_R q_i}^{a_R a_i}\}_{(a_R, a_i) \in A^{R \cup \{i\}}}$ is performed twice on the same half of $|\Psi\rangle$, the probability that the same answer (a_R, a_i) is obtained twice is almost as large as the probability that just a_R is obtained twice: conditioned on being consistent on the answers to the questions in R, the strategy is also consistent in its answer on a random additional question q_i in position i.

Fleshing out the consequences of this definition to eventually show that it implies something close to the classical definition requires some work, and is the object of the second step of the proof.

Step 2: A product theorem for serial strategies. While for a classical deterministic player a serial strategy, as defined in the previous section, is one which decides on the answer a_i to most questions q_i not in R as a function of that question alone, in the quantum setting this is much less clear. The first task is to decide on what one expects from a serial strategy. For instance, one might ask for the measurements to take some "approximately-tensor" form; however we find that this is too strong a requirement. Instead, we first show that the serial property implies that the player's measurement operator $\{X_{q_R q_i}^{a_R a_i}\}_{(a_R, a_i) \in A^{R \cup \{i\}}}$ has a certain

[6] In fact we will also need to consider the outcome of performing the same measurement simultaneously on the two halves of $|\Psi\rangle$.

block-diagonal form, in the sense that[7]

$$X_{q_R q_i}^{a_R a_i} \approx \Pi_{q_i}^{a_i} X_{q_R q_i}^{a_R a_i} \Pi_{q_i}^{a_i}$$

where $\{\Pi_{q_i}^{a_i}\}_{a_i \in A}$ are *orthogonal* projectors; the precise statement is given in Claim 72. Its proof goes through a technical statement about sets of operators which are close to being pairwise orthogonal. That statement, proven in Lemma 127 in Appendix A, shows the natural fact that such operators are close to having a common block-diagonalization basis.

Once this is shown it is not hard to extend the approximation to a small number of additional questions $q_1, \ldots, q_g$, showing that the corresponding measurement also has a block-diagonal form, this time described by the product of the corresponding projectors $\Pi_{q_1}^{a_1} \cdots \Pi_{q_g}^{a_g}$; a precise statement is given in Lemma 73. It is in the precise sense described in that lemma that we can say that a serial strategy has a product form, based on which we can think of the player as playing sequentially on a subset of the coordinates.

Step 3: Both dead and serial strategies fail the repeated game. In the last step of the proof we show that both types of strategies, dead or serial, must fail in the repeated game with high probability (provided the value of the original game was bounded away from 1). For the case of dead strategies this is fairly intuitive: since a dead strategy does not assign consistent answers to a certain subset of the questions q_R, this implies that the player's answers in positions R will very much depend on the questions present in those indices not in R; not only that but it will be virtually impossible for the other player to correlate well with this player's answers on those indices. Here we crucially use the "projection", or "consistency" rounds of the repeated game in order to show that such strategies will fail in those rounds with high probability. This is proven in Claim 74.

The case of serial strategies is slightly harder to analyze, but it boils down to showing that the block-diagonal form we described earlier roughly implies that we can in fact see one of the players as making a sequential measurement governed by the $\Pi_{q_i}^{a_i}$. Since in this case the player's answer to question q_i is decided by applying a projective measurement depending on q_i alone, in case the original game had a value $s < 1$ such a strategy will fail in at least a fraction $s/2$ of the "game" rounds with high probability, and be caught by the referee provided there are enough such rounds. This is shown in Claim 75.

Finally note that the "confuse" rounds of the repeated game are not used in this stage (and indeed the referee accepts any answers in those rounds), but they are crucial to show the dichotomy lemma and the following claims, which only hold for strategies which have been marginalized over a sufficiently large number of questions; in order to be able to perform this marginalization it is important that questions to the players in the confuse rounds are picked independently.

[7] Note that this "approximation" should be taken with a grain of salt; in particular one cannot expect to extract any information about the measurement operators themselves simply by observing statistics of measurement outcomes. Rather, all our estimates will bear on the post-measurement state, resulting from applying the measurement corresponding to $X_{q_R q_i}^{a_R a_i}$ to one half of $|\Psi\rangle$.

7.2 Proof of the main theorem

In this section we give the proof our main results, Theorems 63 and 65. It is divided in three parts. The first, in Section 7.2.2, establishes our "quantum dichotomy theorem". The second, in Section 7.2.3, investigates the structure of serial strategies, and shows that they admit a certain block structure. The results in this section are based on our "orthogonalization lemma", which we introduced in Section 4.3.1 in Chapter 3, and is proved in Appendix A. Finally, in the third part, Section 7.2.4, we use the results from the first two parts to bound the success probability of the players in the repeated game.

Because of the nature of repeated strategies, which are indexed by large tuples of questions and answers, we are constrained to use rather heavy notation. We explain it in detail in the following section, which can also serve as a reading guide for the statements that are to follow.

7.2.1 Notation

Recall that for every $q \in Q^\ell$, $\{X_q^a\}_{a\in A^\ell}$ is an arbitrary projective measurement in d dimensions, that is, the X_q^a are projector matrices, and for any fixed q they sum to the identity over a. The position of the questions (or answers) in a tuple will always be fixed and usually clear from the context; for example when we write $q = (q_G, q_F)$, where $G, F \subseteq [\ell]$ are sets of indices, it is not necessary that the questions q_G are placed before the questions q_F in the tuple q; rather their position is determined by the indices in G, F. When precision is needed we shall write (i, q_i) to express the fact that question q_i is destined to appear in the i-th position of some tuple q. We also write $q_{\neg i}$ to denote $q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_\ell$.

We will often consider *marginalized* POVMs over a certain set $S \subseteq [\ell]$. Given questions q_S indexed by S, the marginalized POVM is the POVM indexed by answers a_S, which results from applying $\{X_{q_S q}^{a_S a}\}_{a_S a}$ for a random $q \in Q^{[\ell]\setminus S}$, and ignoring the answers a not in S. More precisely, given (S, q_S, a_S) it will be convenient to work with the Stinespring representation

$$\hat{X}_{q_S}^{a_S} := \sum_q \sum_a \sqrt{\pi(q)}\,\sqrt{X_{q_S q}^{a_S a}} \otimes \langle q, a|_E$$

where E is an extra register of the appropriate dimension, and π denotes an arbitrary distribution, fixed throughout (it will later be instantiated to the marginal distribution that arises from the original game G that is being repeated). This definition satisfies, for any $\rho \geq 0$,

$$\mathrm{E}_q\Big[\sum_a \sqrt{X_{q_S q}^{a_S a}}\,\rho\,\sqrt{X_{q_S q}^{a_S a}}\Big] = \hat{X}_{q_S}^{a_S}\,(\rho \otimes \mathrm{Id}_E)\,(\hat{X}_{q_S}^{a_S})^\dagger$$

where the identity Id_E was created on the additional register E introduced in the definition of $\hat{X}_{q_S}^{a_S}$, and the expectation is with respect to the distribution π. In order to make measurements corresponding to marginalization over different sets S, we will assume that the register E is always of large enough dimension, and if necessary $\hat{X}_{q_S}^{a_S}$ is tensored with $\frac{1}{\sqrt{|Q|^{|S|}|A|^{|S|}}}\sum_{q,a}\langle q,a|$ on the extra $2|S|$ registers. Note that there is nothing in the definitions above that require the questions and answers in $\hat{X}_{q_S}^{a_S}$ to be indexed to the same set, hence we extend them to define $\hat{X}_{q_S}^{a_T}$, for $T\subseteq S\subseteq[\ell]$, in the obvious way.

For any $\rho\geq 0$, we write $\mathrm{Tr}_\rho(A)$ for $\mathrm{Tr}(A(\rho\otimes\mathrm{Id}_E))$, so that in particular

$$\mathrm{Tr}_\rho\big((\hat{X}_{q_S}^{a_T})^\dagger\hat{X}_{q_S}^{a_T}\big) \;=\; \mathrm{E}_q\Big[\sum_a \mathrm{Tr}\big(\sqrt{X_{q_S q}^{a_T a}}\,\rho\,\sqrt{X_{q_S q}^{a_T a}}\big)\Big] \;=\; \mathrm{Tr}\big(X_{q_S}^{a_T}\rho\big)$$

where we define

$$X_{q_S}^{a_T} \;:=\; \hat{X}_{q_S}^{a_T}(\hat{X}_{q_S}^{a_T})^\dagger \;=\; \mathrm{E}_{q\in Q^{[\ell]\setminus S}}\Big[\sum_{a\in A^{[\ell]\setminus T}} X_{q_S q}^{a_T a}\Big]$$

Terms such as $\mathrm{Tr}\big(X_{q_S}^{a_T}\rho\big)$ will frequently appear on the right-hand side of our inequalities, and they should simply be considered as normalization factors, accounting for the (possibly unnormalized) underlying state ρ, and the conditioning on a fixed a_T. Finally, given $\rho\geq 0$ and a matrix A of appropriate dimension, we introduce the semi-norm

$$\|A\|_\rho^2 \;:=\; \mathrm{Tr}\big(A\rho^{1/2}A^\dagger\rho^{1/2}\big) \qquad (7.1)$$

Note that $\|\cdot\|_\rho$ is definite only if ρ has full rank. We will mostly use this norm for notational convenience. At this point it suffices to observe that it derives from a semi inner-product, so that it satisfies the Cauchy-Schwarz inequality.

At a first reading it may be helpful for the reader to consider the special case of the totally mixed state $\rho=d^{-1}\mathrm{Id}$; putting the notation in context this corresponds to the players sharing the maximally entangled state. In this case very little of the above is really needed, and in particular $\mathrm{Tr}_\rho\big((X_{q_S}^{a_T})^\dagger X_{q_S}^{a_T}\big)$ is simply the normalized trace $\mathrm{E}_q\left[\sum_a d^{-1}\mathrm{Tr}\big(X_{q_S q}^{a_T a}\big)\right]$. Many of our statements are easier to prove, and to understand, in this setting (the main cause of simplification being the commutation between ρ and the X operators), so that the reader may wish to consider it first.

7.2.2 A quantum dichotomy theorem

In this section we prove two important lemmas on the structure of any quantum strategy in a repeated game. The main lemma, Lemma 68, is the analogue of Lemma 11 in [42]. It establishes a dichotomy between two different types of strategies that a player can use, showing that either the strategy is very random, or it must have a relatively strong sequential structure. Its proof follows that of the classical setting without too much added difficulty, provided the definitions are made correctly — which we now proceed to do.

A crucial difficulty in adapting Feige and Kilian's argument is to define an appropriate measure of a strategy's *unpredictability*. In the classical case of a deterministic strategy, this can be measured through the entropy of the marginalized distribution on answers; however in the quantum or even the randomized setting such a measure is no longer helpful, as even honest product strategies can be very random, just by being convex combinations of distinct deterministic strategies. Instead, we measure unpredictability as follows.

Definition 66. *Given a strategy X_q^a, a state ρ, and a fixed set of questions q_R in positions $R \subseteq [\ell]$, define the collision probability of X on q_R as*

$$P_{col}(q_R|X,\rho) := \sum_{a_R} P_{col}(q_R, a_R|X,\rho) \tag{7.2}$$

where

$$P_{col}(q_R, a_R|X,\rho) \,:=\, \Big(\mathrm{Tr}_\rho\big((\hat{X}_{q_R}^{a_R})^\dagger\,\hat{X}_{q_R}^{a_R}(\hat{X}_{q_R}^{a_R})^\dagger\,\hat{X}_{q_R}^{a_R}\big) + \mathrm{Tr}\big(X_{q_R}^{a_R}\,\rho^{1/2}X_{q_R}^{a_R}\,\rho^{1/2}\big)\Big) \tag{7.3}$$

To understand this definition, first consider the case when ρ is the totally mixed state $d^{-1}\mathrm{Id}$. In this case both terms inside the summation are equal to the normalized squared Frobenius norm $d^{-1}\|X_{q_R}^{a_R}\|_F^2$. Expression (7.2) can be interpreted in two different ways. From an operational point of view, it corresponds to the probability that one obtains twice the same answers when one sequentially performs a measurement using the POVM with elements $\{X_{q_R}^{a_R}\}_{a_R}$. In this sense, P_{col} is a measure of the predictability of the strategy X_q^a: pick two completions q, q' at random and measure using first $\{X_{q_Rq}^{a_Ra}\}_{a_Ra}$ and then using $\{X_{q_Rq'}^{a_Ra'}\}_{a_Ra'}$; $P_{\mathrm{col}}(q_R|X,\rho)$ is the probability of getting twice the same result a_R (and ignoring the other answers a, a'). The analytic interpretation is that this is a measure of the entropy of the spectrum of $X_{q_R}^{a_R}$, which is maximized when $X_{q_R}^{a_R}$ is a projector (for a fixed value of the trace).

In case ρ is not the identity, and hence does not commute with the X_q^a, we need to adopt the more cumbersome definition (7.2) for technical reasons. However, note that the operational interpretation remains — the first term on the right-hand side of (7.3) is the probability of obtaining the same answer when performing the measurement twice on the *same half* of $|\Psi\rangle$, while the second term is the same, when the measurement is performed on the two *different halves* of $|\Psi\rangle$: indeed, note that $\mathrm{Tr}\big(X_{q_R}^{a_R}\,\rho^{1/2}X_{q_R}^{a_R}\,\rho^{1/2}\big) = \langle\Psi|X_{q_R}^{a_R} \otimes (X_{q_R}^{a_R})^T|\Psi\rangle$.[8]

The following lets us make the distinction between the two different types of strategies alluded to above.

Definition 67. *We will say that:*

- *A block (R, q_R) is ε-*dead *if $P_{col}(q_R|X,\rho) \leq \varepsilon$. If a block is not ε-dead it is ε-*alive. *Moreover, we say that the answer a_R is ε-*alive *if it satisfies*

$$P_{col}(q_R, a_R|X,\rho) \,\geq\, \varepsilon\,\mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big)$$

Note that any ε-alive block has at least one ε-alive answer. Sometimes we will simply say that a block or an answer are alive or dead, leaving the parameter ε implicit.

[8] Note the transpose sign, which indicates that our interpretation is only rigorously correct for the case of real symmetric X.

- *A block (R, q_R, a_R) is $(1-\eta)$-serial if a_R is alive and the following holds:*

$$\mathrm{E}_{(i,q_i)}\left[\, P_{col}(q_R, q_i|X,\rho)\,\right] \;\geq\; (1-\eta)P_{col}(q_R|X,\rho) \tag{7.4}$$

Lemma 68. *Assume that $\varepsilon, \eta > 0$ are chosen such that $\eta\,\varepsilon^3 > 16\,C_1^{-1/2}$.*[9] *Then one of the following holds*

1. *At least a $(1-\varepsilon)$ fraction of blocks (R, q_R) are ε-dead.*
2. *At least an ε fraction of blocks (R, q_R) are ε-alive, and moreover if (R, q_R) is an ε-alive block then*

$$\sum_{\substack{a_R:\, a_R \text{ alive but} \\ (q_R, a_R) \text{ is not } (1-\eta)\text{-serial}}} \mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big) \;\leq\; \varepsilon/2 \tag{7.5}$$

i.e. alive answers which are not $(1-\eta)$-serial have a small probability of occurring.

Proof. We extend the definition of the collision probability to measuring collisions over answers which are not necessarily on the same indices as the questions:

$$P_{\mathrm{col}}(q|R,X,\rho) \;:=\; \sum_{a_R}\Big(\mathrm{Tr}_\rho\big((\hat{X}_q^{a_R})^\dagger\,\hat{X}_q^{a_R}(\hat{X}_q^{a_R})^\dagger\,\hat{X}_q^{a_R}\big) + \mathrm{Tr}\big(X_q^{a_R}\,\rho^{1/2}X_q^{a_R}\,\rho^{1/2}\big)\Big)$$

where now q can be any subset of fixed questions, and R denotes the subset of answers on which we are measuring the collision probability.

Claim 69. *There exists an integer $1 \leq r^* \leq C_1$ such that*

$$\mathrm{E}_{R,q_R}\left[\, P_{col}(q_R|R,X,\rho)\,\right] - \mathrm{E}_{R,q_R,i,q_i}\left[\, P_{col}(q_R,q_i|R\cup\{i\},X,\rho)\,\right] \leq 8\,C_1^{-1/2}$$

where the expectation is taken over all subsets R of size $|R| = r^*$.

Proof. There is a similar statement in [42]. Here we closely follow the proof of Corollary 3.2 in the lecture notes [89]; since the argument is very similar (mostly replacing the use of Fact 1.3 in those notes by our Claims 130 and 132, proven in Appendix A) we only outline it here, leaving the details to the reader. The proof goes by considering what happens to the collision probability when one conditions on an additional question, resp. one considers collisions over an additional answer. First, note that if one extends R by an index i, then $P_{\mathrm{col}}(q|R\cup\{i\},X,\rho) \leq P_{\mathrm{col}}(q|R,X,\rho)$, since obtaining identical answers on R is a necessary condition to obtain identical answers on $R\cup\{i\}$. The following equation is the analogue of Fact 1.4 in [89]:

$$\big|\mathrm{E}_{(i,q_i)}\left[P_{\mathrm{col}}(q,q_i|R,X,\rho)\right] - P_{\mathrm{col}}(q|R,X,\rho)\big| \;\leq\; 4\,C_1^{-1/2} \tag{7.6}$$

[9] Recall that C_1, C_2 are chosen so that $C_1 + C_2 = \ell$: see Definitions 61 and 62 for more details.

The proof of (7.6) follows directly from Claims 130 and 132, and we omit it. It shows that the collision probability cannot increase by too much when one conditions on an additional question, in expectation. The proof of the claim is then concluded exactly as in the classical case: consider a sequence of steps in which one successively looks for collisions on an additional coordinate i, and conditions on an additional question q_i. In expectation over the choice of (i, q_i), P_{col} will never go up by more than $4C_1^{-1/2}$ when one performs this operation. Since P_{col} is always between 0 and 1, the fact that it never goes up by much implies that there must be a step in which it doesn't decrease by more than $8C_1^{-1/2}$: the total decrease cannot be larger than the total increase plus 1. r^* is chosen so that this step occurs when r^* indices (and questions) have already been fixed. □

Towards a contradiction, assume the negation of both 1. and 2. With probability at least ε a random block (R, q_R) is alive, and moreover if (R, q_R) is alive then alive answers which are not $(1-\eta)$-serial have a significant contribution. Fix such an answer a_R. Since (7.4) is not satisfied, summing over all a_R which are alive but not $(1-\eta)$-serial one can see that the collision probability, for this (R, q_R), must decrease by at least

$$\eta \cdot \sum_{\substack{a_R:\, a_R \text{ alive but} \\ (q_R, a_R) \text{ is not } (1-\eta)\text{-serial}}} P_{\text{col}}(q_R, a_R | X, \rho)$$

By the negation of (7.5) and the fact that the answers are alive, this quantity is at least $\eta\varepsilon^2/2$. Finally, taking the expectation over the choice of (R, q_R) gives a total decrease in P_{col} of at least $\eta\varepsilon^3/2$, contradicting Claim 69 if $\eta\varepsilon^3/2 > 8\,C_1^{-1/2}$. □

7.2.3 Serial strategies

The main result of this section is Lemma 73, which shows that serial strategies have a product structure. Given that most of the strategies that we consider in this section will have a fixed q_R and a_R, we introduce the useful notation $Y_{q_S}^{a_S} := X_{q_R q_S}^{a_R a_S}$ (resp. $\hat{Y}_{q_S}^{a_S} := \hat{X}_{q_R q_S}^{a_R a_S}$) for any $S \subseteq [\ell]\backslash R$; the value of q_R and a_R should always be clear from the context. We will also simply write Y for $X_{q_R}^{a_R}$ (resp. $\hat{Y}$ for $\hat{X}_{q_R}^{a_R}$). For the totality of this section $\eta > 0$ is a fixed parameter, which one can think of as polynomial in the soundness δ that we are aiming for in the repeated game.

We start with a simple fact which expands on the defining property of $(1-\eta)$-serial strategies.

Fact 70. *Let $q_R \in Q^R$. For every $a_R \in A^R$ there exists $\alpha_{a_R} \geq \mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big)$ such that $\sum_{a_R} \alpha_{a_R} \leq 3$ and the following holds. Suppose (R, q_R, a_R) is $(1-\eta)$-serial, and assume* *that $\eta \geq C_2^{-1/2}$. Then for a fraction at least $(1-\eta^{1/4})$ of all (i,q_i) for $i \notin R$ we have that*

$$0 \;\leq\; \mathrm{Tr}_\rho\big(\hat{Y}_{q_i}^\dagger\,\hat{Y}_{q_i}\hat{Y}_{q_i}^\dagger\,\hat{Y}_{q_i}\big) - \sum_{a_i}\mathrm{Tr}_\rho\big((\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}(\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}\big) \leq 4\,\eta^{3/4}\,\alpha_{a_R} \tag{7.7}$$

$$0 \;\leq\; \mathrm{Tr}\big(Y_{q_i}\,\rho^{1/2}Y_{q_i}\,\rho^{1/2}\big) - \sum_{a_i}\mathrm{Tr}\big(Y_{q_i}^{a_i}\,\rho^{1/2}Y_{q_i}^{a_i}\,\rho^{1/2}\big) \leq 4\,\eta^{3/4}\,\alpha_{a_R} \tag{7.8}$$

Proof. By condition (7.4) in the definition of $(1-\eta)$-serial, the $Y_{q_i}^{a_i}$ satisfy

$$\begin{aligned}
&E_{(i,q_i)}\Big[\mathrm{Tr}_\rho\big(\hat{Y}^\dagger\,\hat{Y}\hat{Y}^\dagger\,\hat{Y}\big) - \sum_{a_i}\mathrm{Tr}_\rho\big((\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}(\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}\big)\Big]\\
&\qquad\qquad + \;E_{(i,q_i)}\Big[\mathrm{Tr}\big(Y\,\rho^{1/2}Y\,\rho^{1/2}\big) - \sum_{a_i}\mathrm{Tr}\big(Y_{q_i}^{a_i}\,\rho^{1/2}Y_{q_i}^{a_i}\,\rho^{1/2}\big)\Big]\\
&\leq \eta\Big(\mathrm{Tr}_\rho\big(\hat{Y}^\dagger\,\hat{Y}\hat{Y}^\dagger\,\hat{Y}\big) + \mathrm{Tr}\big(Y\,\rho^{1/2}Y\,\rho^{1/2}\big)\Big)
\end{aligned} \tag{7.9}$$

For any $a'_R \in A^R$, let

$$\alpha_{a'_R} \;:=\; \max\Big(\mathrm{Tr}(X_{q_R}^{a'_R}\rho), \eta^{-1}\,\mathrm{E}_{(i,q_i)}\,\Big[\big|\mathrm{Tr}_\rho\big((\hat{X}_{q_R}^{a'_R})^\dagger X_{q_R}^{a'_R}\hat{X}_{q_R}^{a'_R}\big) - \mathrm{Tr}_\rho\big((\hat{X}_{q_Rq_i}^{a'_R})^\dagger X_{q_Rq_i}^{a'_R}\hat{X}_{q_Rq_i}^{a'_R}\big)\big|\Big]\Big) \tag{7.10}$$

By applying Claim 132 to the $\hat{X}_{q_Rq}^{a'_R}$ we obtain

$$\sum_{a'_R}\mathrm{E}_{(i,q_i)}\,\Big[\big|\mathrm{Tr}_\rho\big((\hat{X}_{q_R}^{a'_R})^\dagger X_{q_R}^{a'_R}\hat{X}_{q_R}^{a'_R}\big) - \mathrm{Tr}_\rho\big((\hat{X}_{q_Rq_i}^{a'_R})^\dagger X_{q_Rq_i}^{a'_R}\hat{X}_{q_Rq_i}^{a'_R}\big)\big|\Big] \;\leq\; 2C_2^{-1/2}\mathrm{Tr}(\rho)$$

which, by using our assumption that $C_2^{-1/2} \leq \eta$ and $\sum_{a'_R}\mathrm{Tr}(X_{q_R}^{a'_R}\rho) \leq \mathrm{Tr}(\rho)$, implies $\sum_{a'_R}\alpha_{a'_R} \leq 3\mathrm{Tr}(\rho) \leq 3$. Applying Claim 130 to the Y_q^a we also obtain

$$\mathrm{E}_{(i,q_i)}\,\big[\big|\mathrm{Tr}\big(Y\,\rho^{1/2}Y\,\rho^{1/2}\big) - \mathrm{Tr}\big(Y_{q_i}\,\rho^{1/2}Y_{q_i}\,\rho^{1/2}\big)\big|\big] \leq \eta\,\alpha_{a_R}$$

Hence (7.9), together with an application of Markov's inequality, implies that, for a fraction at least $(1-\eta^{1/4})$ of all (i,q_i),

$$\begin{aligned}
&\Big(\mathrm{Tr}_\rho\big(\hat{Y}_{q_i}^\dagger\,\hat{Y}_{q_i}\hat{Y}_{q_i}^\dagger\,\hat{Y}_{q_i}\big) - \sum_{a_i}\mathrm{Tr}_\rho\big((\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}(\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}\big)\Big) + \Big(\mathrm{Tr}\big(Y_{q_i}\,\rho^{1/2}Y_{q_i}\,\rho^{1/2}\big) - \sum_{a_i}\mathrm{Tr}\big(Y_{q_i}^{a_i}\,\rho^{1/2}Y_{q_i}^{a_i}\,\rho^{1/2}\big)\Big)\\
&\qquad \leq \eta^{3/4}\,\Big(\mathrm{Tr}_\rho\big(\hat{Y}^\dagger\,\hat{Y}\hat{Y}^\dagger\,\hat{Y}\big) + \mathrm{Tr}\big(Y\,\rho^{1/2}Y\,\rho^{1/2}\big) + 2\,\alpha_{a_R}\Big)
\end{aligned}$$

By expanding out the Y_{q_i} terms, one can verify that both terms on the left-hand-side of this equation are positive, hence each of them must be smaller than the right-hand-side, itself smaller than $4\,\eta^{3/4}\alpha_{a_R}$. This proves both (7.7) and (7.8). □

We now prove a simple claim which shows that $(1-\eta)$-serial strategies are close to being orthogonal; this is how we will subsequently exploit that property.

Claim 71. *Let $q_R \in Q^R$. For every $a_R \in A^R$ there exists $\alpha_{a_R} \geq \mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big)$ such that $\sum_{a_R}\alpha_{a_R} \leq 3$ and the following holds. Suppose that (R, q_R, a_R) is $(1-\eta)$-serial. Then for a fraction at least $(1-\eta^{1/4})$ of all (i, q_i) for $i \notin R$,*

$$\sum_{a_i \neq a_i'} \mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}_{q_i}^{a_i})^\dagger \hat{Y}_{q_i}^{a_i'}\,(\hat{Y}_{q_i}^{a_i'})^\dagger\,\hat{Y}_{q_i}^{a_i}\big) \;\leq\; 8\eta^{3/4}\,\alpha_{a_R} \tag{7.11}$$

where $\rho_{a_i} = \rho^{1/2}Y_{q_i}^{a_i}\rho^{1/2}$.

Proof. Define α_{a_R} as in (7.10). Letting $Z_i = \hat{Y}_{q_i}^\dagger(\hat{Y}_{q_i}\hat{Y}_{q_i}^\dagger)\hat{Y}_{q_i} - \sum_{a_i}(\hat{Y}_{q_i}^{a_i})^\dagger \hat{Y}_{q_i}^{a_i}\,(\hat{Y}_{q_i}^{a_i})^\dagger\,\hat{Y}_{q_i}^{a_i}$, Eq. (7.7) from Fact 70 can be re-written (for the (i, q_i) for which it holds) as

$$\mathrm{Tr}_\rho(Z_i) \;\leq\; 4\eta^{3/4}\,\alpha_{a_R}$$

Let $\rho_i := \sum_{a_i}\rho_{a_i}$, where $\rho_{a_i} = \rho^{1/2}Y_{q_i}^{a_i}\rho^{1/2}$. Since $\rho_i \leq \rho$ and $Z_i \geq 0$, we get

$$\mathrm{Tr}_{\rho_i}(Z_i) \;\leq\; \mathrm{Tr}_\rho(Z_i) \;\leq\; 4\eta^{3/4}\,\alpha_{a_R}$$

and hence, expanding out Z_i,

$$\sum_{a_i \neq a_i'} \mathrm{Tr}_{\rho_i}\big((\hat{Y}_{q_i}^{a_i})^\dagger \hat{Y}_{q_i}^{a_i'}\,(\hat{Y}_{q_i}^{a_i'})^\dagger\,\hat{Y}_{q_i}^{a_i}\big) \;\leq\; 4\eta^{3/4}\,\alpha_{a_R} \tag{7.12}$$

Finally, we can use (7.8) to upper-bound

$$\sum_{a_i \neq a_i'', a_i'} \mathrm{Tr}_{\rho_{a_i''}}\big((\hat{Y}_{q_i}^{a_i})^\dagger \hat{Y}_{q_i}^{a_i'}\,(\hat{Y}_{q_i}^{a_i'})^\dagger\,\hat{Y}_{q_i}^{a_i}\big) \;\leq\; 4\eta^{3/4}\,\alpha_{a_R}$$

where we used $\sum_{a_i'}\hat{Y}_{q_i}^{a_i'}(\hat{Y}_{q_i}^{a_i'})^\dagger \leq \mathrm{Id}$. Together with (7.12), this proves the claim. □

Claim 72. *Let $q_R \in Q^R$. For every $a_R \in A^R$ there exists $\alpha_{a_R} \geq \mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big)$ such that $\sum_{a_R}\alpha_{a_R} \leq 3$ and the following holds. Suppose that (R, q_R, a_R) is $(1-\eta)$-serial, let $1 \leq g \leq C_1/2$ be a fixed parameter, and (G, q_G) chosen at random under the constraint that $G \cap R = \emptyset$ and $|G| = g$. Then with probability at least $(1-\eta^{1/4}-e^{-2g})$ over the choice of (G, q_G), there is a partition $G = G' \cup G''$, where $g'' = |G''| \geq (1-4\eta^{c/4})\,g$, such that for every $i \in G''$*

$$\sum_{a_i} \mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_i})^\dagger(\mathrm{Id}-\Pi_{q_i}^{a_i})\hat{Y}_{q_G}^{a_i}\big) \leq O\big(g\,\eta^{1/c_2}\big)\,\alpha_{a_R} \tag{7.13}$$

where for $i \in G''$, $\{\Pi_{q_i}^{a_i}\}_{a_i}$ is an orthogonal measurement depending only on q_R, a_R and q_i (it is independent of the particular choice of (G, q_G)), $\rho_G = \rho^{1/2}Y_{q_G}\rho^{1/2}$, and $c > 0, c_2 \geq 1$ are universal constants.

Proof. Since (q_R, a_R) is $(1-\eta)$-serial, we can apply Claim 71 to obtain that a fraction $(1-\eta^{1/4})$ of (i, q_i) satisfy

$$\sum_{a_i\neq a_i'} \mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}^{a_i}_{q_i})^\dagger \hat{Y}^{a_i'}_{q_i}\,(\hat{Y}^{a_i'}_{q_i})^\dagger\,\hat{Y}^{a_i}_{q_i}\big) \;\leq\; 8\eta^{3/4}\,\alpha_{a_R} \tag{7.14}$$

where as before $\rho_{a_i} = \rho^{1/2}Y^{a_i}_{q_i}\rho^{1/2}$. We can now apply Lemma 127 (proven in Appendix A) to the $Y^{a_i}_{q_i}$ (with the states ρ_{a_i}) to obtain, for the fraction $(1-\eta^{1/4})$ of (i, q_i) considered above, orthogonal projectors $\{\Pi^{a_i}_{q_i}\}_{a_i}$ satisfying

$$\sum_{a_i} \mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}^{a_i}_{q_i})^\dagger(\mathrm{Id}-\Pi^{a_i}_{q_i})\hat{Y}^{a_i}_{q_i}\big) \leq O\big(\eta^{3c/4}\big)\alpha^c_{a_R}\Big(\sum_{a_i}\mathrm{Tr}\big(\rho_{a_i}\big)\Big)^{1-c} \tag{7.15}$$

Moreover, the $\Pi^{a_i}_{q_i}$ can easily be made into a projective measurement by enlarging one of them, so that they sum to identity; this will not harm the above bound. By Markov's inequality, with probability at least $(1-\eta^{c/4})$ over the choice of (i, q_i) it holds that $\mathrm{Tr}\big(Y_{q_i}\rho\big) \leq \eta^{-c/4}\mathrm{Tr}\big(Y\rho\big) \leq \eta^{-c/4}\alpha_{a_R}$. For any given (G, q_G), let $G'' \subseteq G$ denote those indices i in G for which this property holds for (i, q_i), and moreover (i, q_i) falls in the set of indices for which (7.15) holds. By the union bound and a Chernoff bound, the probability that $|G''| \leq (1-4\eta^{c/4})g$ is less than e^{-2g}, and for ever $i \in G''$ we have

$$\sum_{a_i} \mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}^{a_i}_{q_i})^\dagger(\mathrm{Id}-\Pi^{a_i}_{q_i})\hat{Y}^{a_i}_{q_i}\big) \leq O\big(\eta^{1/c_2}\big)\alpha_{a_R} \tag{7.16}$$

for some constant $c_2 > 0$. Applying Claim 130 to the $\hat{Y}^{a_i}_{q_i}$, and summing over a_i, we find that in expectation

$$E_{(G,q_G)}\Big[\sum_{a_i}\big|\mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}^{a_i}_{q_i})^\dagger\hat{Y}^{a_i}_{q_i}\big) - \mathrm{Tr}_{\rho_{G,a_i}}\big((\hat{Y}^{a_i}_{q_G})^\dagger\hat{Y}^{a_i}_{q_G}\big)\big|\Big] \leq g\,C_2^{-1}\,\mathrm{Tr}\big(Y_{q_i}\rho\big) \;\leq\; g\eta^{3/4}\alpha_{a_R}$$

where we used $C_2^{-1} \leq \eta$, $\rho_{G,a_i} := \rho^{1/2}Y^{a_i}_{q_G}\rho^{1/2}$, and we think of the choice of (G, q_G) as first picking (i, q_i) and then the remaining positions and questions. Another application of Claim 130 combined with (7.8) shows that for every $i \in G''$,

$$E_{(G,q_G)}\Big[\sum_{a_i\neq a_i'}\mathrm{Tr}_{\rho_{G,a_i'}}\big((\hat{Y}^{a_i}_{q_i})^\dagger\hat{Y}^{a_i}_{q_i}\big)\Big] \leq O(g\,\eta^{3/4})\,\alpha_{a_R}$$

Hence, letting $\rho_G := \rho^{1/2}Y_{q_G}\rho^{1/2} = \sum_{a_i}\rho_{G,a_i}$, combining the two previous equations we get

$$E_{(G,q_G)}\Big[\sum_{a_i}\big|\mathrm{Tr}_{\rho_{a_i}}\big((\hat{Y}^{a_i}_{q_i})^\dagger\hat{Y}^{a_i}_{q_i}\big) - \mathrm{Tr}_{\rho_G}\big((\hat{Y}^{a_i}_{q_G})^\dagger\hat{Y}^{a_i}_{q_G}\big)\big|\Big] \leq O(g\,\eta^{3/4})\,\alpha_{a_R}$$

Using Markov's inequality, his lets us replace $\hat{Y}_{q_i}^{a_i}$ by $\hat{Y}_{q_G}^{a_i}$ in (7.15) for a fraction $(1-\eta^{1/4})$ of (G, q_G), losing an additional factor $O(g\eta^{1/2})\alpha_{a_R}$. Hence

$$\sum_{a_i} \mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_i})^\dagger(\mathrm{Id}-\Pi_{q_i}^{a_i})\hat{Y}_{q_G}^{a_i}\big) \leq O\big(g\,\eta^{1/c_2}\big)\,\alpha_{a_R} \tag{7.17}$$

where we safely assumed that $c_2 \geq 2$. □

Lemma 73. *Let $q_R \in Q^R$. For every $a_R \in A^R$ there exists $\alpha_{a_R} \geq \mathrm{Tr}\big(X_{q_R}^{a_R}\rho\big)$ such that $\sum_{a_R}\alpha_{a_R} \leq 3$ and the following holds. Under the same conditions as in Claim 72, except for a lower fraction $(1-2\eta^{1/4c_2}-e^{-2g})$ of (G, q_G), it holds that*

$$\sum_{a_{G''}} \mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}}-\hat{Y}_{q_G})^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}(\hat{Y}_{q_G}^{a_{G''}}-\hat{Y}_{q_G})\big) \leq O\big(g^2\eta^{1/(4c_2)}\big)\,\alpha_{a_R} \tag{7.18}$$

$$\sum_{a_{G''}} \mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}})^\dagger\hat{Y}_{q_G}^{a_{G''}}-(\hat{Y}_{q_G}^{a_{G''}})^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}\hat{Y}_{q_G}^{a_{G''}}\big) \leq O\big(g\eta^{1/(8c_2)}\big)\,\alpha_{a_R} \tag{7.19}$$

Proof. Let $\{\Pi_{q_i}^{a_i}\}$ be the orthogonal projectors promised by Claim 72. Let $g'' = |G''|$, and assume for simplicity that the first g'' questions in G are those in G''. To prove the first inequality, we show the following by induction on $i = 1, \ldots, g''$: there exists a constant $C > 0$ such that, if we let $F_i = \{1, \ldots, i\}$, then

$$\sum_{a_{F_i}} \mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{F_i}}-\hat{Y}_{q_G})^\dagger\Pi_{q_i}^{a_i}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_i}^{a_i}(\hat{Y}_{q_G}^{a_{F_i}}-\hat{Y}_{q_G})\big) \;\leq\; C\,i\,g\,\eta^{1/(3c_2)}\,\alpha_{a_R} \tag{7.20}$$

The statement for $i = g''$ will imply (7.18). Let C_0 be the constant implicit in (7.13) from Claim 72. For $i = 1$, (7.20) is simply a re-statement of (7.13), provided C is chosen larger than C_0. Assume the inequality verified for $i-1$, and prove it for i. Write

$$\hat{Y}_{q_G} - \hat{Y}_{q_G}^{a_{F_i}} = (\hat{Y}_{q_G}-\hat{Y}_{q_G}^{a_i}) + (\hat{Y}_{q_G}^{a_i}-\hat{Y}_{q_G}^{a_{F_i}})$$

The first term on the right-hand side (when plugged back into (7.20)) can be bounded directly using (7.13) (and the fact that the projectors $\Pi_{q_j}^{a_j}$ sum to identity over a_j, for $j \in \{1, \ldots, i-1\}$). Regarding the second, we can use the Cauchy-Schwarz inequality together with (7.13) to bound

$$\sum_{a_{F_i}} \big|\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_i})^\dagger(\mathrm{Id}-\Pi_{q_i}^{a_i})\Pi_{q_{i-1}}^{a_{i-1}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_i}^{a_i}(\hat{Y}_{q_G}^{a_{F_i}}-\hat{Y}_{q_G}^{a_i})\big)\big| \;\leq\; 2\sqrt{C_0}\sqrt{g}\eta^{1/(2c_2)}\alpha_{a_R}^{1/2}\mathrm{Tr}\big(Y_{q_G}\rho\big)^{1/2}$$

By Markov's inequality, $\mathrm{Tr}\big(Y_{q_G}\rho\big) \leq \eta^{-1/4c_2}\mathrm{Tr}\big(Y\rho\big)$ for a fraction at least $(1-\eta^{1/4c_2})$ of (G, q_G), so that for those indices the bound above can be replaced by $2\sqrt{C_0}\sqrt{g}\eta^{1/(4c_2)}\alpha_{a_R}$.

For the rest of this proof we only consider questions (G, q_G) for which the bound $\mathrm{Tr}\big(Y_{q_G}\rho\big) \leq \eta^{-1/4c_2}\mathrm{Tr}\big(Y\rho\big)$ applies. We can similarly obtain

$$\sum_{a_{F_i}} \big|\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{F_i}})^\dagger(\mathrm{Id} - \Pi_{q_i}^{a_i})\Pi_{q_{i-1}}^{a_{i-1}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_i}^{a_i}(\hat{Y}_{q_G}^{a_{F_i}} - \hat{Y}_{q_G}^{a_i})\big)\big| \leq 2\sqrt{C_0}\sqrt{g}\eta^{1/(4c_2)}\alpha_{a_R}$$

so that

$$\begin{aligned}
&\sum_{a_{F_i}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{F_i}} - \hat{Y}_{q_G}^{a_i})^\dagger\Pi_{q_i}^{a_i}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_i}^{a_i}(\hat{Y}_{q_G}^{a_{F_i}} - \hat{Y}_{q_G}^{a_i})\big)\\
&\quad\leq \sum_{a_{F_i}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{F_i}} - \hat{Y}_{q_G}^{a_i})^\dagger\Pi_{q_{i-1}}^{a_{i-1}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{i-1}}^{a_{i-1}}(\hat{Y}_{q_G}^{a_{F_i}} - \hat{Y}_{q_G}^{a_i})\big) + 16\sqrt{C_0}\sqrt{g}\eta^{1/(4c_2)}\alpha_{a_R}\\
&\quad= \sum_{a_{F_i}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{F_{i-1}}} - \hat{Y}_{q_G})^\dagger\Pi_{q_{i-1}}^{a_{i-1}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{i-1}}^{a_{i-1}}(\hat{Y}_{q_G}^{a_{F_{i-1}}} - \hat{Y}_{q_G}^{)}\big) + 16\sqrt{C_0}\sqrt{g}\eta^{1/(4c_2)}\alpha_{a_R}
\end{aligned}$$

which can then be bounded using the induction hypothesis. This concludes the induction step, provided $C \geq C_0 + 16\sqrt{C_0}$, and proves (7.18).

We now prove (7.19). Use the Cauchy-Schwarz inequality to bound

$$\begin{aligned}
&\sum_{a_{G''}}\big|\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}} - \hat{Y}_{q_G})^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}\hat{Y}_{q_G}^{a_{G''}}\big)\big|\\
&\quad\leq \Big(\sum_{a_{G''}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}} - \hat{Y}_{q_G})^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}(\hat{Y}_{q_G}^{a_{G''}} - \hat{Y}_{q_G})\big)\Big)^{1/2}\\
&\qquad\cdot\Big(\sum_{a_{G''}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}})^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}\hat{Y}_{q_G}^{a_{G''}}\big)\Big)^{1/2}\\
&\quad\leq O\big(g\eta^{1/(8c_2)}\big)\,\alpha_{a_R}
\end{aligned}$$

by (7.18). We obtain (7.19) by noting that

$$\sum_{a_{G''}}\mathrm{Tr}_{\rho_G}\big((\hat{Y}_{q_G}^{a_{G''}})^\dagger\hat{Y}_{q_G}^{a_{G''}} - \hat{Y}_{q_G}^\dagger\Pi_{q_{g''}}^{a_{g''}}\cdots\Pi_{q_1}^{a_1}\cdots\Pi_{q_{g''}}^{a_{g''}}\hat{Y}_{q_G}\big) = 0$$

since the $\Pi_{q_i}^{a_i}$ sum to identity over a_i. □

7.2.4 Bounding the success of players in a repeated game

We proceed to show how the results from the previous section can be combined in order to prove Theorems 63 and 65. For the remainder of this section we fix a game G with question set Q and answer set A, and consider the ℓ-repeated games $G_{FK(\ell)}$ and $G_{DR(\ell)}$ for some fixed integer ℓ. Let s be the entangled value of the original game G, and $\{A_{q'}^{a'}\}_{a'}$ (resp. $\{(B_q^a)^T\}_a$)

be an arbitrary fixed projective strategy for Alice (resp. Bob), using entangled state $|\Psi\rangle$, in the ℓ-repeated game.[10] Let $\rho = \mathrm{Tr}_A|\Psi\rangle\langle\Psi|$ be the reduced density of $|\Psi\rangle$ on Bob's subsystem.

We note here that both types of ℓ-repeated games have the same overall structure, in that they consist of a set of C_1 "correlated" rounds, in which the referee sends either "game" or "consistency" questions, and C_2 "independent" rounds, in which he asks questions chosen independently from a product distribution (we refer to Definitions 61 and 62 for more details, including the definition of C_1 and C_2). In both cases, we can think of the referee as choosing the ℓ pairs of questions in the following order.

1. First, a subset $R \subseteq [\ell]$ of size $r^* \leq C_1/2$ is chosen, and designated as indices for either game rounds (in the case of a projection game), or otherwise consistency rounds. Pairs of questions (q'_R, q_R) are then picked according to the appropriate distribution.

2. A subset $G \subseteq [\ell]\backslash R$ of size $C_1 - r^*$ is chosen. In the case of a projection game, all the indices in G are designated as game rounds. In the other cases, $C_1/2$ of the indices in G are designated (at random) as game rounds, and the remaining indices are designated as consistency rounds. Pairs of questions (q'_G, q_G) are chosen accordingly. Note that the referee doesn't know the value of r^*, but he doesn't need to explicitly distinguish between the game and consistency rounds, since they use the same distribution on pairs of questions. The distinction is made only as a convenience for the analysis.

3. Finally, we let $F = [\ell]\backslash(R \cup G)$. F has size C_2, and the indices it contains are designated as confuse rounds, with corresponding pairs of questions (q'_F, q_F).

We will denote by $(q', q) := (q'_R q'_G q'_F, q_R q_G q_F)$ the ℓ-tuple of pairs of questions chosen by the referee. Since questions on the indices in R always correspond to cases where for every answer of Alice there is a unique possible valid answer for Bob, and since we will only perform consistency (as opposed to game) checks on questions in those indices, we may regroup Alice's tuples of answers a'_R when they induce the same a_R for Bob. Hence we re-define $A_{q_R q}^{a_R a} := \sum_{a'_R} A_{q_R q}^{a'_R a}$, where the summation runs over all a'_R such that (a'_R, a_R) are valid answers to the questions (q'_R, q_R).

Our first claim shows that the players have a low success probability on blocks (R, q_R) which are dead.

Claim 74. *Let $\varepsilon > 0$ be such that $\varepsilon \geq C_1 C_2^{-1}$, and suppose that (R, q_R) is an ε-dead block. Then the success probability of the players, conditioned on the referee picking questions (q', q) such that q includes q_R in the positions in R, is at most $\sqrt{2}\,\varepsilon$.*

[10] The transpose sign on Bob's operators is there for consistency of notation. For simplicity we will omit this transpose in the future whenever we consider expressions of the form $\langle\Psi|A \otimes B|\Psi\rangle$, which should be read as $\langle\Psi|A \otimes B^T|\Psi\rangle$.

Proof. The definition of (R, q_R) being ε-dead implies that

$$\sum_{a_R} \mathrm{Tr}\big(B_{q_R}^{a_R}\,\rho^{1/2} B_{q_R}^{a_R}\,\rho^{1/2}\big) \;\leq\; \varepsilon$$

By applying Claim 130 to the $B_{q_R q}^{a_R}$ together with Markov's inequality, we obtain that in expectation

$$E_{G,q_G}\Big[\sum_{a_R} \mathrm{Tr}\big(B_{q_R q_G}^{a_R}\,\rho^{1/2} B_{q_R q_G}^{a_R}\,\rho^{1/2}\big)\Big] \;\leq\; \varepsilon + C_1 C_2^{-1} \;\leq\; 2\varepsilon \tag{7.21}$$

where we used $|G| \leq C_1$ and our assumption on ε. Condition on (q'_R, q_R) being chosen as part of the referee's questions in the game, and assume that the referee only checks consistency of Alice and Bob's answers to the questions in R. This can only increase their success probability, which can then be bounded as

$$\begin{aligned} E_{(G,F),(q'_G q'_F, q_G q_F)}\Big[\sum_{a_R,a',a} \langle\Psi| A_{q'_R q'_G q'_F}^{a_R a'} \otimes B_{q_R q_G q_F}^{a_R a} |\Psi\rangle\Big] &\leq E_{G,(q'_G,q_G)}\Big[\Big(\sum_{a_R} \|A_{q'_R q'_G}^{a_R}\|_\rho^2\Big)^{1/2}\Big(\sum_{a_R} \|B_{q_R q_G}^{a_R}\|_\rho^2\Big)^{1/2}\Big] \\ &\leq \sqrt{2\,\varepsilon} \end{aligned}$$

where we used that (q'_F, q_F) are chosen according to a product distribution, the first inequality follows from Cauchy-Schwarz (recall the definition of $\|\cdot\|_\rho$ given in (7.1)), and for the second we upper-bounded $\sum_{a_R} \|A_{q'_R q'_G}^{a_R}\|_\rho^2$ by 1 and used Jensen's inequality together with (7.21) to bound the other term. □

We note informally that one can combine this claim with Lemma 73 to obtain a form of "direct product test" for entangled strategies. Indeed, if two entangled players Alice and Bob win the game with probability $s \gg \varepsilon$, then by the previous claim a fraction at least $s^2/2$ of blocks (R, q_R) should be alive; moreover a non-negligible fraction[11] of answers a_R to those blocks must be $(1-\eta)$-serial. Hence one can apply Lemma 73 to those blocks (R, q_R, a_R) and obtain a product form for the corresponding marginalized strategy.

The next claim shows that strategies which are product, even on a subset of the coordinates, also have a low success probability.

Claim 75. *Fix (R, q_R, a_R), and for every (i, q_i), where $i \in [\ell]\backslash R$ and $q_i \in Q$, let $\{\Pi_{q_i}^a\}_{a\in A}$ be a fixed projective measurement. Suppose that Bob's strategy is such that, with probability at least $1-\delta$ over the choice of (G, q_G) and $G_1 \subseteq G$ of size $|G_1| = g$, there is a partition $G_1 = G' \cup G''$ such that $g'' = |G''| \geq (1-\delta')g$ and Bob's POVM satisfies that for every $a_{G''}$*

$$B_{q_R q_G}^{a_R a_{G''}} = (\hat{B}_{q_R}^{a_R})^\dagger \Pi_{q_{g''}}^{a_{g''}} \cdots \Pi_{q_1}^{a_1} \cdots \Pi_{q_{g''}}^{a_{g''}} \hat{B}_{q_R}^{a_R}$$

[11] Note that one cannot hope to obtain any structural result on the strategies which would hold for more than a fraction s of questions or answers, as the player's strategy could be a mixture of a perfect winning strategy with probability s, and a random strategy with probability $(1-s)$.

where for simplicity we wrote $G'' = \{1, \dots, g''\}$.

Then the success probability of the players, conditioned on the referee asking questions (q', q) such that q includes q_R in the positions in R, and summed over all valid answers which include a_R for Bob, is at most

$$\big(\delta + e^{-(1-s-\delta')^2 g}\big)\,\mathrm{Tr}\big(B^{a_R}_{q_R}\rho\big)$$

Proof. Fixing the questions in R and G, and conditioning on the players consistently answering a_R to (q'_R, q_R), their probability of being accepted is at most

$$\begin{aligned}\sum_{a'_{G''},a_{G''}} \langle\Psi|A^{a_R a'_{G''}}_{q'_R q'_G} \otimes B^{a_R a_{G''}}_{q_R q_G}|\Psi\rangle &= \sum_{a'_{G''},a_{G''}} \langle\Psi|A^{a_R a'_{G''}}_{q'_R q'_G} \otimes (\hat{B}^{a_R}_{q_R})^\dagger \Pi^{a_{g''}}_{q_{g''}} \cdots \Pi^{a_1}_{q_1} \cdots \Pi^{a_{g''}}_{q_{g''}} \hat{B}^{a_R}_{q_R}|\Psi\rangle \\ &= \sum_{a'_{G''},a_{G''}} \big(\langle\Psi|\mathrm{Id} \otimes (\hat{B}^{a_R}_{q_R})^\dagger\big) \cdot A^{a_R a'_{G''}}_{q'_R q'_G} \otimes \Pi^{a_{g''}}_{q_{g''}} \cdots \Pi^{a_1}_{q_1} \cdots \Pi^{a_{g''}}_{q_{g''}} \cdot \big(\mathrm{Id} \otimes \hat{B}^{a_R}_{q_R}|\Psi\end{aligned} \tag{7.22}$$

The fact that sequential strategies cannot succeed in many rounds of the repeated game implies that

$$\Big\|E_{(G,q'_G,q_G)}\Big[\sum_{a'_{G''},a_{G''}} A^{a_R a'_{G''}}_{q'_R q'_G} \otimes \Pi^{a_{g''}}_{q_{g''}} \cdots \Pi^{a_1}_{q_1} \cdots \Pi^{a_{g''}}_{q_{g''}}\Big]\Big\|_\infty \leq \exp(-(1-s-\delta')^2 g)$$

Indeed, the expression on the left-hand side can be upper-bounded by the maximum success probability of an Alice playing an arbitrary strategy and Bob a sequential strategy described by the measurements $\Pi^{a_i}_{q_i}$, provided the referee only checks the answers to those questions in $G'' \subseteq G_1$, where G_1 is a random subset of G of size g chosen by the referee. But this success probability is even lower than the success probability that Alice and Bob would have if Bob played his sequential strategy on *all* questions in G_1, but the referee was to accept as long as at least g'' out of Alice and Bob's g answers were correct. Since the probability of such a serial strategy succeeding in any round is at most the value s of the original game, and $g'' \geq (1-\delta')g$, by a Chernoff bound the probability that the players succeed in g'' out of the g rounds is at most $\exp(-(1-s-\delta')^2 g)$. Hence the expression in (7.22) can be upper-bounded, in expectation, by

$$e^{-(1-s-\delta')^2 g}\,\langle\Psi|\mathrm{Id} \otimes (\hat{B}^{a_R}_{q_R})^\dagger \hat{B}^{a_R}_{q_R}|\Psi\rangle \;=\; e^{-(1-s-\delta')^2 g}\,\mathrm{Tr}\big(B^{a_R}_{q_R}\rho\big)$$

Finally, we must account for the small probability δ that the serial property does not hold; for those sets G we can trivially bound the success probability, conditioned on Bob answering a_R to q_R, by $\mathrm{Tr}\big(B^{a_R}_{q_R}\rho\big)$. □

We finally turn to the proof of our main theorem.

Proof of Theorem 63. We first set parameters: let C_0 be a large enough constant, $\varepsilon = C_0^{-1}\delta^2$ (recall that δ is the target value for the repeated game $G_{FK}(\ell)$), $\eta = C_0^{-1}\delta^{24c_2}(1-s)$ (where c_2 is the constant which appears in Claim 72), $g = C_0 \log(1/\delta)(1-s)^{-1}$, and $\ell \geq C_0^{15}\delta^{-125c_2}(1-s)^{-4}$. Recall also that C_1 was defined as $C_1 = \sqrt{\ell}$, and $C_2 = \ell - C_1$. This choice of parameters satisfies the following constraints:

- $\eta\,\varepsilon^3 > 16\,C_1^{-1/2}$, which is used in Lemma 68.
- $\eta \geq C_2^{-1/2}$, which is used in Fact 70 and subsequent claims.
- $\varepsilon \geq C_1\,C_2^{-1}$, which is used in Claim 74.

As before, in game $G_{FK(\ell)}$, we can think of the referee as first picking $r^* \leq C_1/2$ pairs of questions $(R, (q'_R, q_R))$ for the players, then picking g pairs $(G_1, (q'_{G_1}, q_{G_1}))$, then $C_1 - r^* - g$ pairs $(G_2, (q'_{G_2}, q_{G_2}))$ and finally C_2 independent pairs of confuse questions $(F, (q'_F, q_F))$. Let $G = G_1 \cup G_2$ and $(q', q) = (q'_R q'_G q'_F, q_R q_G q_F)$. Let $\{A^{a'}_{q'}\}_{a'}$ be Alice's POVM on questions q', and $\{B^a_q\}_a$ Bob's POVM on questions q.

By Lemma 68, one of two cases hold. Either a $(1-\varepsilon)$ fraction of blocks (R, q_R) are ε-dead, in which case the player's success probability is readily bounded by $\varepsilon + \sqrt{2\varepsilon}$ by Claim 74. Otherwise, it must be that we are in case 2 of the lemma, so that ε-alive blocks are for the most part serial. Note that any dead blocks contribute at most $\sqrt{2\varepsilon}$ to the success probability, by Claim 74. A similar argument to that in Claim 74 shows that alive blocks which are not $(1-\eta)$-serial also contribute at most $\sqrt{2\varepsilon}$, given the fact that we are in the case 2. of Lemma 68, and there can only be few such blocks by (7.5).

Suppose (R, q_R, a_R) is $(1-\eta)$-serial. By Lemma 73, for every (i, q_i) there exists a projective measurement $\{\Pi^{a_i}_{q_i}\}_{a_i}$, depending only on q_R, a_R, q_i, a_i, such that with probability at least $(1-2\eta^{1/4c_2}-e^{-2g})$ over the choice of (G, q_G) such that $|G| = g$ there is a partition $G_1 = G' \cup G''$ such that $g'' = |G''| \geq (1-4\eta^{c/4})g$ such that Eqs. (7.18) and (7.19) from Lemma 73 are satisfied, where $\rho_G = \rho^{1/2} B^{a_R}_{q_R q_G} \rho^{1/2}$. To alleviate notation we let $\Pi = \Pi^{a_1}_{q_1} \cdots \Pi^{a_{g''}}_{q_{g''}}$, and we first use Cauchy-Schwarz to bound

$$\begin{aligned}
\sum_{a'_{G''}, a_{G''}} & \langle \Psi | A^{a_R a'_{G''}}_{q'_R q'_G} \otimes (\hat{B}^{a_R a_{G''}}_{q_R q_G})^\dagger (\mathrm{Id} - \Pi^\dagger \Pi) \hat{B}^{a_R a_{G''}}_{q_R q_G} | \Psi \rangle \\
& \leq \| A^{a_R}_{q'_R q_G} \|_\rho \, \Big\| \sum_{a_{G''}} (\hat{B}^{a_R a_{G''}}_{q_R q_G})^\dagger (\mathrm{Id} - \Pi^\dagger \Pi) \hat{B}^{a_R a_{G''}}_{q_R q_G} \Big\|_\rho \\
& \leq \| A^{a_R}_{q'_R q_G} \|_\rho \Big(\sum_{a_{G''}} \mathrm{Tr}_{\rho_G}\big((\hat{B}^{a_R a_{G''}}_{q_R q_G})^\dagger (\mathrm{Id} - \Pi^\dagger \Pi) \hat{B}^{a_R a_{G''}}_{q_R q_G} \big) \Big)^{1/2} \\
& \leq O\big(\sqrt{g}\eta^{1/(16c_2)}\big) \| A^{a_R}_{q'_R q_G} \|_\rho \, \alpha^{1/2}_{a_R} \qquad (7.23)
\end{aligned}$$

where $\rho_G = \rho^{1/2} B^{a_R}_{q_R q_G} \rho^{1/2}$, the first inequality is by Cauchy-Schwarz, the second uses $(\mathrm{Id} - \Pi^\dagger \Pi) \leq \mathrm{Id}$, the last is by Eq. (7.19) from Lemma 73, and α_{a_R} was defined in Eq. (7.10) (where here we substitute $\hat{B}^{a_R}_{q_R}$ for $\hat{X}^{a_R}_{q_R}$). A similar argument, using this time Eq. (7.18), lets us bound

$$\sum_{a'_{G''},a_{G''}} \langle\Psi| A^{a_R a'_{G''}}_{q'_R q'_G} \otimes \big(\hat{B}^{a_R a_{G''}}_{q_R q_G} - \hat{B}^{a_R}_{q_R q_G}\big)^\dagger \Pi^\dagger \Pi \big(\hat{B}^{a_R a_{G''}}_{q_R q_G} - \hat{B}^{a_R}_{q_R q_G}\big)|\Psi\rangle \leq O\big(g\eta^{1/(8c_2)}\big)\|A^{a_R}_{q'_R q_G}\|_\rho\, \alpha^{1/2}_{a_R} \tag{7.24}$$

and hence combining (7.23) and (7.24) we get

$$\sum_{a'_{G''},a_{G''}} \big|\langle\Psi| A^{a_R a'_{G''}}_{q'_R q'_G} \otimes \big(B^{a_R a_{G''}}_{q_R q_G} - (\hat{B}^{a_R}_{q_R q_G})^\dagger \Pi^\dagger \Pi \hat{B}^{a_R}_{q_R q_G}\big)|\Psi\rangle\big| \leq O\big(\sqrt{g}\eta^{1/(16c_2)}\big)\|A^{a_R}_{q'_R q_G}\|_\rho\, \alpha^{1/2}_{a_R}$$

Finally, by Claim 130 we have

$$\begin{aligned}
\mathrm{E}_{(G,q_G)}\Big[\sum_{a'_{G''},a_{G''}} \big|\langle\Psi| A^{a_R a'_{G''}}_{q'_R q'_G} \otimes \big((\hat{B}^{a_R}_{q_R})^\dagger \Pi^\dagger \Pi \hat{B}^{a_R}_{q_R} - (\hat{B}^{a_R}_{q_R q_G})^\dagger \Pi^\dagger \Pi \hat{B}^{a_R}_{q_R q_G}\big)|\Psi\rangle\big|\Big]& \\
\leq 4\|A^{a_R}_{q'_R q_G}\|_\rho\, \mathrm{E}_{(G,q_G)}\left[\Big|\big\|B^{a_R}_{q_R}\big\|^2_\rho - \big\|B^{a_R}_{q_R q_G}\big\|^2_\rho\Big]^{1/2}& \\
\leq 4\eta\|A^{a_R}_{q'_R q_G}\|_\rho\, \alpha^{1/2}_{a_R}&
\end{aligned}$$

where for the first inequality we used $\sum_{a''_G} \Pi^\dagger\Pi = \mathrm{Id}$, and for the second that $\eta \geq C_2^{-1}$. Hence the statistical distribution of outcomes produced by Alice and Bob (conditioned on answering a_R to q_R) is close to that which would be obtained if Bob was to use the operators $(B^{a_R}_{q_R})^\dagger \Pi^\dagger \Pi B^{a_R}_{q_R}$ as his POVM on questions q_G. But the success probability of the latter, when summed over all valid answers to the pair of questions $(q'_{G''}, q_{G''})$, can be bounded by Claim 75. Hence summing over all a_R (and using $\sum_{a_R} \|A^{a_R}_{q'_R q_G}\|_\rho\, \alpha^{1/2}_{a_R} \leq 3$) and taking the expectation over q_R, the average winning probability of the players for all $(1-\eta)$-serial blocks (R, q_R, a_R) is at most

$$O\big(\sqrt{g}\,\eta^{1/(16c_2)} + 2\eta^{c/4} + e^{-2g} + e^{-(1-s-4\eta^{1/4c_2})^2 g}\big)$$

where we also accounted for those (rare) choices of (G, q'_G, q_G) for which the previous bounds do not hold. Given our choice of parameters ε, η, g and ℓ, it can be checked that this expression is $\ll \delta$. Combining this bound with the one resulting from dead blocks shows that the winning probability of the players is at most δ, which proves the theorem as long as $\ell = \mathrm{poly}(\delta^{-1}, (1-s)^{-1})$ is large enough. □

We conclude this section by briefly explaining how the proof of Theorem 63 can be adapted to prove Theorem 65. The main reason the proof carries over is that, in the proof of Theorem 63, we only used the projection property for a subset of the game questions (to bound the success over dead blocks), while for $(1-\eta)$-serial blocks the game questions were only used in conjunction with the fact that the value of the game was at most s. Here, consistency rounds will play the role of the game questions previously in R, and game rounds will play the role of those game questions previously in G (or rather its small subset G_1).

Proof of Theorem 65. In game $G_{DR(\ell)}$, we think of the referee as first picking $r^* \leq C_1/2$ pairs of consistency questions $(R, (q'_R, q_R))$ for the players, then picking $C_1/2 - r^*$ additional consistency pairs $(R', (q'_{R'}, q_{R'}))$, $C_1/2$ pairs of game questions $(G, (q'_G, q_G))$ and finally C_2 independent pairs of confuse questions $(F, (q'_F, q_F))$. Let $(q', q) = (q'_R q'_{R'} q'_G q'_F, q_R q_{R'} q_G q_F)$.

Assume a choice of parameters made that is similar to the one in the proof of Theorem 63. As before, we can apply Lemma 68 to Bob's strategy B^a_q, distinguishing between two cases.

In the first case, a fraction $(1-\varepsilon)$ of blocks (R, q_R) are dead, for $|R| = r^*$. Then Claim 74 again applies, as the only property we used in its proof was that any answer of Alice induced a fixed answer for Bob, which is the case for consistency questions.

In the second case, a fraction ε of blocks (R, q_R) are alive. Those blocks which are dead can be dealt with as in the previous case, and we can focus on blocks (R, q_R, a_R) which are $(1-\eta)$-serial. Here we can reason exactly as in Theorem 63, using Claim 75 with G_1 chosen as a subset of the questions in G, and the remaining consistency questions playing the role of the remaining game questions before. □

7.3 Discussion and open questions

The work presented in this chapter shows for the first time that the entangled value of games can be decreased through parallel repetition. Even though we framed and proved our results in the context of 2-player games, it should not be hard to extend them in some cases to multiple players, depending on the kind of projection or consistency constraints that one can assume on the game. On the other hand, extending the result to either many-round games, or games with quantum messages, is an interesting open question.

One implication of our result is the following. The celebrated PCP theorem says that given a game, it is NP-hard to tell if its value is 1 or less than, say, 0.99. Combined with Raz's parallel repetition result, one obtains that it is also hard to tell if the value is 1 or less than, say, 0.01. The latter statement led to an enormous body of work on strong hardness of approximation results [52]. It is currently a major open question whether an analogue of the PCP theorem holds for the entangled value. *If* such a result was proved, our results would allow to amplify the hardness to 1 vs. 0.01, as in the classical case, possibly leading to further surprising implications. (While we show such a hardness result in Chapter 5, the protocol used in that result requires a large (constant) number of provers, together with a polynomial number of rounds of interaction: it is an open question whether the results in this chapter can be applied to improve the gap of that protocol, while keeping the number of rounds of interaction as small as possible.)

The main open question left by the results presented in this chapter is whether it is possible to show a better rate of decay, in particular an exponential rate as Raz obtained from direct parallel repetition, or [56] first obtained in the setting of direct product testers. Another open question is whether our statement can be extended to hold for simple parallel repetition for arbitrary entangled games (i.e. without adding dummy or consistency questions).

We believe that our main conceptual contributions are the extension of the notion of "approximately serial" to the setting of measurements, and our subsequent *orthogonalization lemma*. We hope that these techniques might prove useful elsewhere. Lastly, product testers are very useful in the area of property testing, and it remains to be seen if our result can be applied similarly.

Chapter 8

Trevisan's extractor in the presence of quantum side information

In this chapter we show that a family of extractor constructions originally introduced by Trevisan [119] in the classical setting is secure against quantum adversaries. This construction, and its proof of security, will be crucial to the results in Chapter 9. We first introduce the task of randomness extraction, emphasizing the importance of taking into account any "side information", classical or quantum, that an adversary to the extractor may have about its random source. In §8.2 we give formal definitions of extractors and discuss briefly how much randomness can be extracted from a given source. Section 8.3 contains the description of Trevisan's extractor construction paradigm and a proof that it is still sound in the presence of quantum side information, in the cases of both uniform and weakly random seeds. Then in §8.4 we plug in various one-bit extractors and pseudo-random seed constructions, resulting in, amongst others, a construction which is nearly optimal in the amount of randomness extracted in §8.4.1 (which is identical to the best known bound in the classical case [98] for Trevisan's extractor), and a construction which is still sound if there is a small linear entropy loss in the seed in §8.4.4. Finally, in §8.5, we mention a few classical results which modify and improve Trevisan's extractor, but for which the correctness in the presence of quantum side information does not seem to follow immediately from our work.

8.1 Introduction

Randomness extraction is the art of generating (almost) uniform randomness from any weakly random source X. More precisely, a *randomness extractor* (or, simply *extractor*) is a function Ext that takes as input X together with a uniformly distributed (and usually short) string Y, called the *seed*, and outputs a string Z. One then requires Z to be almost uniformly distributed whenever the min-entropy of X is larger than some threshold k, i.e.,

$$H_{\min}(X) \geq k \implies Z := \mathrm{Ext}(X, Y) \text{ statistically close to uniform.} \tag{8.1}$$

The min-entropy of a random variable X is directly related to the probability of correctly guessing the value of X using an optimal strategy: $2^{-H_{\min}(X)} = \max_x P_X(x)$. Hence Criterion (8.1) can be interpreted operationally: if the maximum probability of successfully guessing the input of the extractor, X, is sufficiently low then its output is statistically close to uniform.

The randomness of a value X always depends on the information one has about it, in the following called *side information.* In cryptography, for instance, a key is supposed to be uniformly random from the point of view of an adversary, who may have access to messages exchanged by the honest parties, which we would therefore consider as side information. Here, extractors are typically used for *privacy amplification* [19, 20] , i.e., to turn a partially secure raw key (about which the adversary may have non-trivial information) into a perfectly secure key. We thus demand that the extractor output be uniform with respect to the side information held by the adversary. Another example is *randomness recycling* in a computation, which can be done using extractors [58]. The aim is that the recycled randomness is independent of the outputs of previous computations, which are therefore considered as side information.

In the following, we make side information explicit and denote it by E. The notions of randomness we are going to use, such as the *guessing probability*, *min-entropy* or the *uniformity* of a random variable, must then be defined with respect to E. We can naturally reformulate Criterion (8.1) as

$$H_{\min}(X|E) \geq k \implies Z := \text{Ext}(X,Y) \begin{array}{l}\text{statistically close to uniform}\\ \text{conditioned on } E,\end{array} \tag{8.2}$$

where $H_{\min}(X|E)$ is the conditional min-entropy, formally defined in §3.3. This conditioning naturally extends the operational interpretation of the min-entropy to scenarios with side information, i.e., $2^{-H_{\min}(X|E)}$ is the maximum probability of correctly guessing X, given access to side information E [75].

Interestingly, the relationship between the two Criteria (8.1) and (8.2) depends on the physical nature of the side information E, i.e., whether E is represented by the state of a classical or a quantum system. In the case of purely classical side information, E may be modeled as a random variable and it is known that the two criteria are essentially equivalent (see Lemma 78 for a precise statement). But in the general case where E is a quantum system, Criterion (8.2) is *strictly stronger* than (8.1): it was shown in [46] that there exist extractors that fulfill (8.1) but for which (8.2) fails (see also [74] for a discussion).

Since our world is inherently non-classical, it is of particular importance that (8.2) rather than the weaker Criterion (8.1) be taken as the relevant criterion for the definition of extractors. In cryptography, for instance, there is generally nothing that prevents an adversary from holding quantum side information. In fact, even if a cryptographic scheme is purely classical, an adversary may acquire information using a non-classical attack strategy. Hence, when using extractors for privacy amplification, Criterion (8.1) does not generally imply security. A similar situation may arise in the context of randomness recycling. If we run a (simulation of) a quantum system E using randomness X, approximately $H_{\min}(X|E)$ bits of X can be reused. If we now, in an attempt to recycle the randomness, apply a function Ext which fulfills (8.1) but not (8.2), the output Z may still be correlated to the system E.

It is known that the conditional min-entropy accurately characterizes the maximum amount of uniform randomness that can be extracted from X while being independent from E. (More precisely, the *smooth min-entropy*, an entropy measure derived from $H_{\min}(X|E)$ by maximizing the latter over all states in an ε-neighborhood, is an upper bound on the amount of uniform randomness that can be extracted; see §3.3 in Chapter 3 and [100] for details). In other words, the characterization of extractors in terms of $H_{\min}(X|E)$ is essentially optimal, and one may thus argue that Criterion (8.2) is indeed the correct definition for randomness extraction (see also [100, 74, 76]). In this chapter, we follow this line of argument and call an extractor *quantum-proof* if it satisfies Criterion (8.2) (see §8.2.1).

We note that there have been alternatives proposals in the literature for defining extractors in the context of quantum side information, which do however not satisfy the above optimality condition. One prominent example is the bounded storage model (see §8.4.3), where the (quantum) side information E is characterized by the number of qubits, $H_0(E)$, required to store it. In this case, the entropy $H_{\min}(X|E)$ of a source X conditioned on E is generally lower-bounded by $H_{\min}(X) - H_0(E)$. However, this characterization of side information is strictly weaker than that using $H_{\min}(X|E)$: there are sources X and nontrivial side information E such that $H_{\min}(X) - H_0(E) \ll H_{\min}(X|E)$.[1] In particular, even if an extractor can provably extract $H_{\min}(X) - H_0(E)$ bits of uniform randomness from a source X, we do not know whether the same extractor can attain the optimal $H_{\min}(X|E)$ bits. Note also that the same considerations apply to the purely classical case. In fact, no recent work defines classical extractors for randomness sources with side information stored in bounded classical memories.[2]

Finally we remark that the increased generality attained by the notion of quantum-proof extractors used here is crucial for applications. For example in quantum key distribution, where extractors are used for privacy amplification [100], it is generally impossible to bound the adversary's memory size.

[1] This can easily be seen by considering the following example. Let X be uniformly distributed on $\{0,1\}^n$ and E be X with each bit flipped with constant probability $\varepsilon < 1/2$. Then $H_{\min}(X|E) = \Theta(n)$, but $H_{\min}(X) - H_0(E) = 0$.

[2] Restricting the class of randomness sources further than by bounding their min-entropy can have advantages, e.g., if we consider only bit-fixing sources, or sources generated by a random walk on a Markov chain, then the extractor can be deterministic. (See [104] for a brief overview of restricted families of sources studied in the literature.) There is however no known advantage (e.g., in terms of seed length) in considering only input sources with side information stored in memory of bounded size, whether it is classical or quantum memory.

8.1.1 Related results

In the standard literature on randomness extraction, constructions of extractors are usually shown to fulfill Criterion (8.1), for certain values of the threshold k (see [31] as well as [104] for an overview). However, only a few constructions have been shown to fulfill Criterion (8.2) with arbitrary quantum side information E. Among them is two-universal hashing [100, 117], constructions based on the sample-and-hash approach [74], as well as all extractors with one-bit output [76].

Recently, Ta-Shma [112] studied Trevisan's construction of extractors [119] in the bounded quantum storage model. The result was a breakthrough because it, for the first time, implied the existence of quantum-proof extractors requiring only short seeds (logarithmic in the input length). Unfortunately, his proof technique requires the output length to be much smaller than the min-entropy of the original data: it scales as $(H_{\min}(X)/H_0(E))^{1/c}$, where $c > 1$ is a constant. Furthermore, Ta-Shma's result is proved in the bounded quantum storage model, which, as discussed previously, only allows the extractor to output at most $H_{\min}(X) - H_0(E)$ bits. This expression can in general be arbitrarily smaller than $H_{\min}(X|E)$, and in some cases may even become 0 (or negative) for n-bit sources for which it is possible to extract $\Omega(n)$ bits of randomness.[1]

Subsequent to this work, Ben-Aroya and Ta-Shma [16] showed how two versions of Trevisan's extractor, shown quantum-proof in this paper, can be combined to extract a constant fraction of the min-entropy of an n-bit source with a seed of length $O(\log n)$, when $H_{\min}(X|E) > n/2$. This is better than the straightforward application of Trevisan's extractor analyzed here, which requires $O(\log^2 n)$ bits of seed for the same output size (but works for any $H_{\min}(X|E)$).

8.1.2 Contributions

We show that the performance of Trevisan's extractor does not suffer in the presence of quantum side information. This improves on the best previously known result [112] in two major ways. First, we prove our results in the most general model, where the min-entropy of the source is measured relative to quantum side information (Criterion (8.2)). Second, we show that the output length of the extractor can be close to the optimal conditional min-entropy $H_{\min}()$. This provides the first proof of soundness for an extractor with poly-logarithmic seed meeting Criterion (8.2) in the presence of arbitrary quantum side information.

More generally, we show that a whole class of extractors is quantum-proof. It has been observed, by, e.g., Lu and Vadhan [79, 120], that Trevisan's extractor [119] (and variations of it, such as [98]) is a concatenation of the outputs of a one-bit extractor with different pseudo-random seeds. Since the proof of the extractor property is independent of the type of the underlying one-bit extractor (and to some extent the construction of the pseudo-random seeds), our result is valid for a generic scheme (defined in §8.3.1, Definition 83). We find that the performance of this generic scheme in the context of quantum side information is roughly

	Min-entropy	Output length	Seed length	Note
Corollary 89	any k	$m = k - 4\log 1/\varepsilon$	$d = O(\log^3 n)$	optimized output length
Corollary 90	$k = n^\alpha$	$m = n^{\alpha-\gamma}$	$d = O(\log n)$	optimized seed length
Corollary 91	$k = \alpha n$	$m = (\alpha - \gamma)n$	$d = O(\log^2 n)$	local extractor
Corollary 92	$k = n^\alpha$	$m = n^{\alpha-\gamma}$	$d = O(\log n)$	seed with min-entropy βd

Table 8.1: Plugging various weak designs and 1-bit extractors in Trevisan's construction, we obtain these concrete extractors. Here n is the input length, $\varepsilon = \text{poly}1/n$ the error, α and γ are arbitrary constants such that $0 < \gamma < \alpha \leq 1$, and $\frac{1}{2} < \beta < 1$ is a specific constant.

equivalent to the (known) case of purely classical side information (§8.3.2, Theorem 86).

In practical situations where quantum-proof extractors are used, e.g., privacy amplification in quantum key distribution [100], the players do not necessarily have access to a uniform source of randomness. We therefore analyze separately the situation where the seed is only weakly random, and show that Trevisan's extractor is quantum-proof in that setting as well (§8.3.2, Theorem 87).

By "plugging" various one-bit extractors and pseudo-random seeds into the generic scheme, we obtain different final constructions, optimized for different needs, e.g., maximizing the output length, minimizing the seed, or using a non-uniform seed. In Table 8.1 we give a brief overview of the final constructions proposed.

8.1.3 Proof technique

The proof proceeds by contradiction. We first assume that a player holding the side information E can distinguish the output from uniform with probability greater than ε. We then show that such a player can reconstruct the input X with high probability, which means that X must have low min-entropy ($H_{\min}(X|E) < k$). Taking the contrapositive proves that the extractor is sound.

Trevisan [119] originally proved the correctness of his extractor this way. His construction starts by encoding the source X using a list-decodable code C [111]. The output of the extractor then consists of certain bits of $C(X)$, which are specified by the seed and a construction called a (weak) design [86, 98]. (See §8.3.1 for a precise description of Trevisan's extractor.) His proof can then be broken down in two steps. He first shows that a player who can distinguish the output from uniform can guess a random bit of $C(X)$. In the second step, he shows that such a player can reconstruct X.

Proving the soundness of Trevisan's extractor in the quantum min-entropy framework requires some important changes. In order to better explain these new elements, it will be useful to first give a brief overview of the main steps that go into Ta-Shma's proof [112]. For the sake of contradiction, assume that there is a test T, which performs a measurement on the side information E in order to distinguish the output from uniform with advantage

ε. Using a standard hybrid argument, along with properties of the (weak) design, one can then construct a new test T' (using a little extra classical advice about X) which predicts a random bit of $C(X)$ with probability $\frac{1}{2} + \frac{\varepsilon}{m}$, where m is the number of output bits. Further, T' makes exactly *one* query to T.

The proof in [112] proceeds by showing how from such a test, one can construct another test T'' which predicts any bit of X with probability 0.99 and queries T' at most $q = (m/\varepsilon)^c$ times ($c = 15$ for the code in [112]). This gives a random access code (RAC) [6] for X; however, since it requires q queries to the side information E, the no-cloning theorem forces us to see it as querying a single system of length $qH_0(E)$ (recall that Ta-Shma's result was proved in the bounded storage model, where one bounds the information provided by E by its number of qubits $H_0(E)$). Finally, using a new bound on the dimension of RACs [112], one finds that $H_{\min}(X) \gtrsim m^c H_0(E)$, hence $m \lesssim (H_{\min}(X)/H_0(E))^{1/c}$.

Our proof improves upon Ta-Shma's through two major changes. First, we model the side information E explicitly, instead of viewing it as an oracle which one queries. Indeed, the measurement performed by the test T' to predict the bits of $C(X)$ will be different from the measurement performed by T'' to reconstruct X, and this cannot be captured by the "oracle-side-information" model of Ta-Shma. We also generalize previous versions of this step by considering non-uniform seeds. We thus show (in §8.3.2, Proposition 1) that if the output of the extractor can be distinguished from uniform with probability $\frac{1}{2} + \varepsilon$ by a player holding the side information E, then the bits of $C(X)$ can be guessed with probability $\frac{1}{2} + \frac{\varepsilon}{m}$ by a player holding E and some extra small classical information G.

Second, we depart from the reconstruction paradigm at the heart of the second half of the proof of both Trevisan's and Ta-Shma's results. Instead of defining explicitly the measurement and computation necessary to reconstruct X, we use the fact that for any list-decodable code $C : \{0,1\}^n \to \{0,1\}^{\bar{n}}$, the function

$$\begin{aligned} C' : \{0,1\}^n \times [\bar{n}] &\to \{0,1\} \\ (x,i) &\mapsto C(x)_i \end{aligned}$$

is a one-bit extractor (see §A.5.5 for more details). The second part of the (classical) reconstruction paradigm can be understood as a proof that these codes are one-bit extractors according to Criterion (8.1). It was however proved by König and Terhal [76], that in the one-bit setting the more general Criterion (8.2) is essentially equivalent to the usual Criterion (8.1). This result lets us conclude the proof directly, without having needed to explicitly show that the source X could be reconstructed from the extractor's output and the side information.

This proof structure results in a very modular extractor construction paradigm, which allows arbitrary one-bit extractors and pseudo-random seeds to be "plugged in," producing many different final constructions, some of which are given in Table 8.1 and detailed in §8.4.

8.2 Extractors

8.2.1 Extractors, side information, and privacy amplification

An extractor $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a function which takes a weak source of randomness X and a uniformly random, short seed Y, and produces some output $\mathrm{Ext}(X,Y)$, which is almost uniform. The extractor is said to be strong, if the output is approximately independent from the seed.

Definition 76 (strong extractor [87]). *A function* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* (k,ε)-strong extractor with uniform seed*, if for all distributions* X *with min-entropy* $H_{\min}(X) \geq k$ *and a uniform seed* Y*, we have*[3]

$$\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)Y} - \rho_{U_m} \otimes \rho_Y\|_{\mathrm{tr}} \leq \varepsilon,$$

where ρ_{U_m} *is the fully mixed state on a system of dimension* 2^m*.*

Using the connection between min-entropy and guessing probability (Eq. (3.1)), a (k,ε)-strong extractor can be seen as a function which guarantees that if the guessing probability of X is not too high ($p_{\mathrm{guess}}(X) \leq 2^{-k}$), then it produces a random variable which is approximately uniform and independent from the seed Y.

As discussed in the introduction, we consider here a more general situation involving side information, denoted by E, which may be represented by the state of a quantum system. We then want to find some function Ext such that, if the probability of guessing X *given* E is not too high, Ext can produce a random variable $\mathrm{Ext}(X,Y)$ which is approximately uniform and independent from the seed Y and the side information E. Equivalently, one may think of a *privacy amplification* scenario [19, 20], where E is the information available to an adversary and where the goal is to turn weakly secret data X into a *secret* key $\mathrm{Ext}(X,Y)$, where the seed Y is assumed to be public. (In typical key agreement protocols, the seed is chosen by the legitimate parties and exchanged over public channels.)

The following definition covers the general situation where the side information E may be represented quantum-mechanically. The case of purely classical side information is then formulated as a restriction on the nature of E.

Definition 77 (quantum-proof strong extractor [74]). *A function* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* quantum-proof *(or simply* quantum*)* (k,ε)-strong extractor with uniform seed*, if for all states* ρ_{XE} *classical on* X *with* $H_{\min}(X|E)_\rho \geq k$*, and for a uniform seed* Y*, we have*

$$\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} \leq \varepsilon,$$

[3] A more standard classical notation would be $\frac{1}{2}\|\mathrm{Ext}(X,Y) \circ Y - U_m \circ Y\| \leq \varepsilon$, where the distance metric is the variational distance. However, since classical random variables can be represented by quantum states diagonal in the computational basis, and the trace distance reduces to the variational distance, we use the quantum notation for compatibility with the rest of this work.

where ρ_{U_m} is the fully mixed state on a system of dimension 2^m.

The function Ext *is a* classical-proof (k, ε)-strong extractor with uniform seed *if the same holds with the system E restricted to classical states.*

It turns out that if the system E is restricted to classical information about X, then this definition is essentially equivalent to the conventional Definition 76.

Lemma 78 ([76, Proposition 1])**.** *Any (k, ε)-strong extractor is a classical-proof $(k+\log 1/\varepsilon, 2\varepsilon)$-strong extractor.*

However, if the system E is quantum, this does not necessarily hold. Gavinsky et al. [46] give an example of a (k, ε)-strong extractor, which breaks down in the presence of quantum side information, even when $H_{\min}(X|E)$ is significantly larger than k.

Remark 79. *In this section we defined extractors with a uniform seed, as this is the most common way of defining them. Instead one could use a seed which is only weakly random, but require it to have a min-entropy larger than a given threshold, $H_{\min}(Y) \geq s$. The seed must still be independent from the input and the side information. Since having access to a uniform seed is often an unrealistic assumption, it is much more useful for practical applications to define and prove the soundness of extractors with a weakly random seed. We redefine extractors formally this way in §A.5.1, and show in §8.3.2 that Trevisan's extractor is still quantum-proof in this setting.*

All the considerations of this section, in particular Lemma 78 and the gap between classical and quantum side-information, also apply if the seed is only weakly random. In the following, when we simply talk about a strong extractor, without specifying the nature of the seed, we are referring to both uniform seeded and weakly random seeded extractors.

8.2.2 Extracting more randomness

Radhakrishnan and Ta-Shma [93] have shown that a (k, ε)-strong extractor Ext : $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ will necessarily have

$$m \leq k - 2\log 1/\varepsilon + O(1). \tag{8.3}$$

However, in some situations we can extract much more randomness than the min-entropy. For example, let X be distributed on $\{0,1\}^n$ with $\Pr[X = x_0] = 1/n$ and for all $x \neq x_0$, $\Pr[X = x] = \frac{n-1}{n(2^n-1)}$. We have $H_{\min}(X) = \log n$, so using a $(\log n, 1/n)$-strong extractor we could obtain at most $\log n$ bits of randomness. But X is already $1/n$-close to uniform, since $\frac{1}{2}\|\rho_X - \rho_{U_n}\|_{\text{tr}} \leq \frac{1}{n}$. So we already have n bits of nearly uniform randomness, exponentially more than by using a $(\log n, 1/n)$-strong extractor.

In the case of quantum extractors, similar examples can be found, e.g., in [116, Remark 22]. However, an upper bound on the extractable randomness can be obtained by replacing the min-entropy by the *smooth* min-entropy (Definition 10). More precisely, the total number of ε-uniform bits that can be extracted in the presence of side information E can never exceed $H^{\varepsilon}_{\min}(X|E)$ [100, Section 5.6].

Conversely, the next lemma implies that an extractor which is known to extract m bits from any source such that $H_{\min}(X|E) \geq k$ can in fact extract the same number of bits, albeit with a slightly larger error, from sources which only satisfy $H^{\varepsilon'}_{\min}(X|E) \geq k$, a much weaker requirement in some cases.

Lemma 80. *If* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a quantum-proof* (k, ε)*-strong extractor, then for any state* ρ_{XE} *and any* $\varepsilon' > 0$ *with* $H^{\varepsilon'}_{\min}(X|E)_\rho \geq k$*,*

$$\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} \leq \varepsilon + 2\varepsilon'.$$

Proof. Let $\tilde{\rho}_{XE}$ be the state ε'-close to ρ_{XE} for which $H_{\min}(X|E)_{\tilde{\rho}}$ reaches its maximum. Then

$$\begin{aligned}
&\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} \\
&\quad\leq \frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YE} - \tilde{\rho}_{\mathrm{Ext}(X,Y)YE}\|_{\mathrm{tr}} + \frac{1}{2}\|\tilde{\rho}_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \tilde{\rho}_E\|_{\mathrm{tr}} \\
&\qquad + \frac{1}{2}\|\rho_{U_m} \otimes \rho_Y \otimes \tilde{\rho}_E - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} \\
&\quad\leq \frac{1}{2}\|\tilde{\rho}_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \tilde{\rho}_E\|_{\mathrm{tr}} + \|\rho_{XE} - \tilde{\rho}_{XE}\|_{\mathrm{tr}} \\
&\quad\leq \varepsilon + 2\varepsilon'.
\end{aligned}$$

In the second inequality above we used (twice) the fact that a trace-preserving quantum operation can only decrease the trace distance. And in the last line we used the fact that the purified distance — used to measure the distance between two states (see Definition 10) — is larger than the trace distance. □

Remark 81. *Since a* (k, ε)*-strong extractor can be applied to any source with smooth min-entropy* $H^{\varepsilon'}_{\min}(X|E) \geq k$*, we can measure the entropy loss of the extractor — namely how much entropy was not extracted — with*

$$\Delta := k - m,$$

where m *is the size of the output. From Eq. (8.3) we have that an extractor has optimal entropy loss if* $\Delta = 2 \log 1/\varepsilon + O(1)$*.*

8.3 Constructing m-bit extractors from one-bit extractors and weak designs

In this section we show how to construct a quantum m-bit extractor from any (classical) 1-bit strong extractor.

This can be seen as a derandomization of a result by König and Terhal [76], who also extract m bits in the presence of quantum side information by concatenating m times a 1-bit extractor. They however choose a different seed for each bit, thus having a seed of total length $d = mt$, where t is the length of the seed of the 1-bit extractor. In the case of classical side information, this derandomization was done by Trevisan [119], who shows how to concatenate m times a 1-bit extractor using only $d = \operatorname{poly} t, \log m$ bits of seed.[4] We combine the weak designs from Raz et al. [98], which they use to improve Trevisan's extractor, and a previous observation by two of the authors [33], that since 1-bit extractors were shown to be quantum-proof in [76], Trevisan's extractor is also quantum-proof.

This results in a generic scheme, which can be based on any weak design and 1-bit strong extractor. We define it in §8.3.1, then prove bounds on the min-entropy and error in §8.3.2.

8.3.1 Description of Trevisan's construction

In order to shorten the seed while still outputting m bits, in Trevisan's extractor construction paradigm the seed is treated as a string of length $d < mt$, which is then split in m overlapping blocks of t bits, each of which is used as a (different) seed for the 1-bit extractor. Let $y \in \{0,1\}^d$ be the total seed. To specify the seeds for each application of the 1-bit extractor we need m sets $S_1, \cdots, S_m \subset [d]$ of size $|S_i| = t$ for all i. The seeds for the different runs of the 1-bit extractor are then given by y_{S_i}, namely the bits of y at the positions specified by the elements of S_i.

The seeds for the different outputs of the 1-bit extractor must however be nearly independent. To achieve this, Nisan and Wigderson [86] proposed to minimize the overlap $|S_i \cap S_j|$ between the sets, and Trevisan used this idea in his original work [119]. Raz et al. [98] improved this, showing that it is sufficient for these sets to meet the conditions of a *weak design*.[5]

Definition 82 (weak design [98, Definition 5]). *A family of sets $S_1, \ldots, S_m \subset [d]$ is a* weak (t,r)-design *if*

1. *For all i, $|S_i| = t$.*
2. *For all i, $\sum_{j=1}^{i-1} 2^{|S_j \cap S_i|} \leq rm$.*

We can now describe Trevisan's generic extractor construction.

[4] Trevisan's original paper does not explicitly define his extractor as a pseudo-random concatenation of a 1-bit extractor. It has however been noted in, e.g., [79, 120], that this is basically what Trevisan's extractor does.

[5] The second condition of the weak design was originally defined as $\sum_{j=1}^{i-1} 2^{|S_j \cap S_i|} \leq r(m-1)$. We prefer to use the version of [51], since it simplifies the notation without changing the design constructions.

Definition 83 (Trevisan's extractor [119]). *For a one-bit extractor $C : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}$, which uses a (not necessarily uniform) seed of length t, and for a weak (t,r)-design $S_1, \dots, S_m \subset [d]$, we define the m-bit extractor $\mathrm{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ as*

$$\mathrm{Ext}_C(x,y) := C(x, y_{S_1}) \cdots C(x, y_{S_m}).$$

Remark 84. *The length of the seed of the extractor Ext_C is d, one of the parameters of the weak design, which in turn depends on t, the size of the seed of the 1-bit extractor C. In §8.4 we will give concrete instantiations of weak designs and 1-bit extractors, achieving various entropy losses and seed sizes. The size of the seed will always be $d = \mathrm{polylog}\, n$, if the error is $\varepsilon = \mathrm{poly} 1/n$. For example, to achieve a near optimal entropy loss (§8.4.1), we need $d = O(t^2 \log m)$ and $t = O(\log n)$, hence $d = O(\log^3 n)$.*

8.3.2 Analysis

We now prove that the extractor defined in the previous section is a quantum-proof strong extractor. The first step follows the structure of the classical proof [119, 98]. We show that a player holding the side information and who can distinguish the output of the extractor Ext_C from uniform can — given a little extra information — distinguish the output of the underlying 1-bit extractor C from uniform. This is summed up in the following proposition:

Proposition 1. *Let X be a classical random variable correlated to some quantum system E, let Y be a (not necessarily uniform) seed, independent from XE, and let*

$$\|\rho_{\mathrm{Ext}_C(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} > \varepsilon, \tag{8.4}$$

where Ext_C is the extractor from Definition 83. Then there exists a fixed partition of the seed Y in two substrings V and W, and a classical random variable G, such that G has size $H_0(G) \leq rm$, where r is one of the parameters of the weak design (Definition 82), $V \leftrightarrow W \leftrightarrow G$ form a Markov chain,[6] and

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_{VWGE}\|_{\mathrm{tr}} > \frac{\varepsilon}{m}. \tag{8.5}$$

We provide a proof of Proposition 1 in §A.5.3, where it is restated as Proposition 3.[7]

For readers familiar with Trevisan's scheme [119, 98], we briefly sketch the correspondence between the variables of Proposition 1 and quantities analyzed in Trevisan's construction.

[6] Three random variables are said to form a Markov chain $X \leftrightarrow Y \leftrightarrow Z$ if for all x, y, z we have $P_{Z|YX}(z|y,x) = P_{Z|Y}(z|y)$, or equivalently $P_{ZX|Y}(z,x|y) = P_{Z|Y}(z|y)P_{X|Y}(x|y)$.

[7] Note that Ta-Shma [112] has already implicitly proved that this proposition must hold in the presence of quantum side information, by arguing that the side information can be viewed as an oracle. The present statement is a strict generalization of that reasoning, which allows conditional min-entropy as well as non-uniform seeds to be used.

Trevisan's proof proceeds by assuming by contradiction that there exists a player, holding E, who can distinguish between the output of the extractor and the uniform distribution (Eq. (8.4)). Part of the seed is then fixed (this corresponds to W in the above statement) and some classical advice is taken (this corresponds to G in the above statement) to construct another player who can distinguish a specific bit of the output from uniform. But since a specific bit of Trevisan's extractor is just the underlying 1-bit extractor applied to a substring of the seed (V in the above statement), this new player (who holds WGE) can distinguish the output of the 1-bit extractor from uniform (Eq. (8.5)).

In the classical case Proposition 1 would be sufficient to prove the correctness of Trevisan's scheme, since it shows that if a player can distinguish Ext_C from uniform, then he can distinguish C from uniform given a few extra advice bits, which contradicts the assumption that C is an extractor.[8] But since our assumption is that the underlying 1-bit extractor is only classical-proof, we still need to show that the quantum player who can distinguish $C(X,V)$ from uniform is not more powerful than a classical player, and so if he can distinguish the output of C form uniform, so can a classical player. This has already been done by König and Terhal [76], who show that 1-bit extractors are quantum-proof.

Theorem 85 ([76, Theorem III.1]). *Let $C : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}$ be a (k,ε)-strong extractor. Then C is a quantum-proof $(k + \log 1/\varepsilon, 3\sqrt{\varepsilon})$-strong extractor.*[9]

We now need to put Proposition 1 and Theorem 85 together to prove that Trevisan's extractor is quantum-proof. The cases of uniform and weak random seeds differ somewhat in the details. We therefore give two separate proofs for these two cases in §8.3.2 and §8.3.2.

Uniform seed

We show that Trevisan's extractor is a quantum-proof strong extractor with uniform seed with the following parameters.

Theorem 86. *Let $C : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}$ be a (k,ε)-strong extractor with uniform seed and $S_1, \dots, S_m \subset [d]$ a weak (t,r)-design. Then the extractor given in Definition 83, $\mathrm{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$, is a quantum-proof $(k + rm + \log 1/\varepsilon, 3m\sqrt{\varepsilon})$-strong extractor.*

Proof. In Proposition 1, if the seed Y is uniform, then V is independent from W and hence by the Markov chain property from G as well, so Eq. (8.5) can be rewritten as

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_V \otimes \rho_{WGE}\|_{\mathrm{tr}} > \frac{\varepsilon}{m},$$

which corresponds to the exact criterion of the definition of a quantum-proof extractor.

[8] In the classical case, [119, 98] still show that a player who can distinguish $C(X,V)$ from uniform can reconstruct X with high probability. But this is nothing else than proving that C is an extractor.

[9] This result holds whether the seed is uniform or not.

Let C be a (k,ε)-strong extractor with uniform seed, and assume that a player holds a system E such that

$$\|\rho_{\text{Ext}_C(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\text{tr}} > 3m\sqrt{\varepsilon}.$$

Then by Proposition 1 and because Y is uniform, we know that there exists a classical system G with $H_0(G) \leq rm$, and a partition of Y in V and W, such that,

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_V \otimes \rho_{WGE}\|_{\text{tr}} > 3\sqrt{\varepsilon}. \tag{8.6}$$

Since C is a (k,ε)-strong extractor, we know from Theorem 85 that we must have $H_{\min}(X|WGE) < k + \log 1/\varepsilon$ for Eq. (8.6) to hold. Hence by Lemma 139, $H_{\min}(X|E) = H_{\min}(X|WE) \leq H_{\min}(X|WGE) + H_0(G) < k + rm + \log 1/\varepsilon$. □

Weak random seed

We show that Trevisan's extractor is a quantum-proof strong extractor with weak random seed, with the following parameters.

Theorem 87. *Let $C : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}$ be a (k,ε)-strong extractor with an s-bit seed — i.e., the seed needs at least s bits of min-entropy — and $S_1,\ldots,S_m \subset [d]$ a weak (t,r)-design. Then the extractor given in Definition 83,* $\text{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$, *is a quantum-proof $(k+rm+\log 1/\varepsilon, 6m\sqrt{\varepsilon})$-strong extractor for any seed with min-entropy $d-(t-s-\log\frac{1}{3\sqrt{\varepsilon}})$.*

The main difference between this proof and that of Theorem 86, is that since the seed Y is not uniform in Proposition 1, the substring W of the seed not used by the 1-bit extractor C is correlated to the seed V of C, and acts as classical side information about the seed. To handle this, we show in Lemma 135 that with probability $1-\varepsilon$ over the values of W, V still contains a lot of min-entropy, roughly $s'-d'$, where d' is the length of W and s' the min-entropy of Y. And hence a player holding WGE can distinguish the output of C from uniform, even though the seed has enough min-entropy.

Proof. Let C be a (k,ε)-strong extractor with s bits of min-entropy in the seed, and assume that a player holds a system E such that

$$\|\rho_{\text{Ext}_C(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\text{tr}} > 6m\sqrt{\varepsilon}.$$

Then by Proposition 1 we have

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_{VWGE}\|_{\text{tr}} > 6\sqrt{\varepsilon}. \tag{8.7}$$

Since this player has classical side-information W about the seed V, we need an extra step to handle it. Lemma 135 tells us that from Eq. (8.7) and because by Theorem 85, C is a quantum $(k+\log 1/\varepsilon, 3\sqrt{\varepsilon})$-strong extractor, we must have either for some w, $H_{\min}(X|GEW = w) < k + \log 1/\varepsilon$ and hence

$$\begin{aligned} H_{\min}(X|E) = H_{\min}(X|EW = w) \\ \leq H_{\min}(X|GEW = w) + H_0(G) < k + rm + \log 1/\varepsilon, \end{aligned}$$

or $H_{\min}(V|W) < s + \log \frac{1}{3\sqrt{\varepsilon}}$, from which we obtain using Lemma 137,

$$H_{\min}(Y) \leq H_{\min}(V|W) + H_0(W) < s + \log \frac{1}{3\sqrt{\varepsilon}} + d - t.$$

8.4 Concrete constructions

Depending on what goal has been set — e.g., maximize the output, minimize the seed length — different 1-bit extractors and weak designs will be needed. In this section we give a few examples of what can be done, by taking various classical extractors and designs, and plugging them into Theorem 86 (or Theorem 87), to obtain bounds on the seed size and entropy loss in the presence of quantum side information.

The results are usually given using the O-notation. This is always meant with respect to all the free variables, e.g., $O(1)$ is a constant independent of the input length n, the output length m, and the error ε. Likewise, $o(1)$ goes to 0 for both n and m large.

We first consider the problem of extracting all the min-entropy of the source in §8.4.1. This was achieved in the classical case by Raz et al. [98], so we use the same 1-bit extractor and weak design as them.

In §8.4.2 we give a scheme which uses a seed of length $d = O(\log n)$, but can only extract part of the entropy. This is also based on Raz et al. [98] in the classical case.

In §8.4.3 we combine an extractor and design which are locally computable (from Vadhan [120] and Hartman and Raz [51] respectively), to produce a quantum m-bit extractor, such that each bit of the output depends on only $O(\log(m/\varepsilon))$ bits of the input.

And finally in §8.4.4 we use a 1-bit extractor from Raz [97], which only requires a weakly random seed, resulting in a quantum m-bit extractor, which also works with a weakly random seed.

These constructions are summarized in Table 8.1 on page 126.

8.4.1 Near optimal entropy loss

To achieve a near optimal entropy loss we need to combine a 1-bit extractor with near optimal entropy loss and a weak $(t, 1)$-design. We use the same extractor and design as Raz et al. [98] to do so, namely Lemma 141 for the design and Proposition 4 for the 1-bit extractor. Plugging this into Theorem 86 we get an extractor against quantum adversaries with parameters similar to Raz et al. [98].

Corollary 88. *Let C be the extractor from Proposition 4 with error $\varepsilon' = \frac{\varepsilon^2}{9m^2}$ and let us use the weak design from Lemma 141. Then Trevisan's extractor* $\mathrm{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a $(m + 8\log m + 8\log 1/\varepsilon + O(1), \varepsilon)$-strong extractor with uniform seed against quantum adversaries, with $d = O(\log^2(n/\varepsilon)\log m)$.*

For $\varepsilon = \mathrm{poly}1/n$ the seed has length $d = O(\log^3 n)$. The entropy loss is $\Delta = 8\log m + 8\log 1/\varepsilon + O(1)$, which means that the input still has this much randomness left in it (conditioned on the output). We can extract a bit more by now applying a second extractor to the input. For this we will use the extractor by Tomamichel et al [117], which is a quantum (k', ε')-strong extractor with seed length $d' = O(m' + \log n' + \log 1/\varepsilon')$ and entropy loss $\Delta' = 4\log 1/\varepsilon' + O(1)$, where n' and m' are the input and output string lengths. Since we will use it for $m' = 8\log m + 4\log 1/\varepsilon' + O(1)$, we immediately get the following corollary from Lemma 136.

Corollary 89. *By applying the extractors from Corollary 88 and [117, Theorem 10] in succession, we get a new function* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$*, which is a quantum-proof $(m + 4\log 1/\varepsilon + O(1), \varepsilon)$-strong extractor with uniform seed, with $d = O(\log^2(n/\varepsilon)\log m)$.*

For $\varepsilon = \mathrm{poly}1/n$ the seed has length $d = O(\log^3 n)$.

The entropy loss is $\Delta = 4\log 1/\varepsilon + O(1)$, which is only a factor 2 times larger than the optimal entropy loss. By Lemma 80 this extractor can produce $m = H^{\varepsilon}_{\min}(X|E) - 4\log 1/\varepsilon - O(1)$ bits of randomness with an error 3ε.

8.4.2 Seed of logarithmic size

The weak design used in §8.4.1 requires the seed to be of size $d = \Theta(t^2 \log m)$, where t is the size of the seed of the 1-bit extractor. Since t cannot be less than $\Omega(\log n)$ [93], a scheme using this design will always have $d = \Omega(\log^2 n \log m)$. If we want to use a seed of size $d = O(\log n)$ we need a different weak design, e.g., Lemma 142, at the cost of extracting less randomness from the source.

For the 1-bit extractor we can use the same as in the previous section, Proposition 4. Plugging this into Theorem 86 we get an extractor against quantum adversaries with logarithmic seed length.

Corollary 90. *If for any constant $0 < \alpha \leq 1$, the source has min-entropy $H_{\min}(X|E) = n^{\alpha}$, and the desired error is $\varepsilon = \mathrm{poly}1/n$, then using the extractor C from Proposition 4 with error $\varepsilon' = \frac{\varepsilon^2}{9m^2}$ and the weak design from Lemma 142 with $r = n^{\gamma}$ for any $0 < \gamma < \alpha$, we have that Trevisan's extractor* $\mathrm{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a $(n^{\gamma}m + 8\log m + 8\log 1/\varepsilon + O(1), \varepsilon)$-strong extractor with uniform seed against quantum adversaries, with $d = O\left(\frac{1}{\gamma}\log n\right)$.*

Choosing γ to be a constant results in a seed of length $d = O(\log n)$. The output length is $m = n^{\alpha-\gamma} - o(1) = H_{\min}(X|E)^{1-\frac{\gamma}{\alpha}} - o(1)$. By Lemma 80 this can be increased to $m = H^{\varepsilon}_{\min}(X|E)^{1-\frac{\gamma}{\alpha}} - o(1)$ with an error of 3ε.

8.4.3 Locally computable extractor

Another interesting feature of extractors is *locality*, that is, the m-bit output depends only a small subset of the n input bits. This is useful in, e.g., the bounded storage model (see [82, 79, 120] for the case of a classical adversary and [74] for a general quantum treatment), where we assume a huge source of random bits, say n, are available, and the adversary's storage is bounded by αn for some constant $\alpha < 1$. Legitimate parties are also assumed to have bounded workspace for computation. In particular, for the model to be meaningful, the bound is stricter than that on the adversary. So to extract a secret key from the large source of randomness, they need an extractor which only reads $\ell \ll n$ bits. An extractor with such a property is called ℓ-local. We will use a construction of an ℓ-local extractor by Vadhan [120], stated in Lemma 146.

Since we assume that the available memory is limited, we also want the construction of the weak design to be particularly efficient. For this we can use a construction by Hartman and Raz [51], given in Lemma 143. Plugging this into Theorem 86 we get a local extractor against quantum adversaries.

Corollary 91. *If for any constant* $0 < \alpha \leq 1$*, the source has min-entropy* $H_{\min}(X|E) = \alpha n$*, then using the weak design from Lemma 143 for any constant* $r > 1$*, and the extractor* C *from Lemma 146 with error* $\epsilon' = \frac{\varepsilon^2}{9m^2}$ *and any constant* $\gamma < \alpha$*, we have that Trevisan's extractor* $\text{Ext}_C : \{0,1\}^n \times \{0,1\}^d \rightarrow \{0,1\}^m$ *is an* ℓ*-local* $(\gamma n + rm + 2\log m + 2\log 1/\varepsilon + O(1), \varepsilon)$*-strong extractor with uniform seed against quantum adversaries, with* $d = O(\log^2(n/\varepsilon))$ *and* $\ell = O(m \log(m/\varepsilon))$*. Furthermore, each bit of the output depends on only* $O(\log(m/\varepsilon))$ *bits of the input.*

With these parameters the extractor can produce up to $m = (\alpha - \gamma)n/r - O(\log 1/\varepsilon) = (H_{\min}(X|E) - \gamma n)/r - O(\log 1/\varepsilon)$ bits of randomness, with an error of $\varepsilon = \text{poly}1/n$. By Lemma 80 this can be increased to $m = (H^{\varepsilon}_{\min}(X|E) - \gamma n)/r - O(\log 1/\varepsilon)$ with an error of 3ε.

8.4.4 Weak random seed

Extractors with weak random seeds typically require the seed to have a min-entropy linear in its length. Theorem 87 says that the difference between the length and the min-entropy of the seed needed in Trevisan's extractor is roughly the same as the difference between the length and min-entropy of the seed of the underlying 1-bit extractor. So we will describe in detail how to modify the construction from §8.4.2 to use a weakly random seed. As that extractor uses a seed of length $O(\log n)$, this new construction allows us to preserve the linear loss in the min-entropy of the seed. Any other version of Trevisan's extractor can be modified in the same way to use a weakly random seed, albeit with weaker parameters.

For this we need a 1-bit extractor which uses a weakly random seed. We will use a result by Raz [97] (Lemma 147), which allows us to construct the extractor from Corollary 148.

Plugging this and the weak design from Lemma 142 in Theorem 87, we get the following extractor with weak random seed.

Corollary 92. *Let $\alpha > 0$ be a constant such that the source has min-entropy $H_{\min}(X|E) = n^\alpha$, and the desired error is $\varepsilon = \operatorname{poly} 1/n$. Using the extractor C from Corollary 148 with error $\varepsilon' = \frac{\varepsilon^2}{9m^2}$ and the weak design from Lemma 142 with $r = n^\gamma$ for any $0 < \gamma < \alpha$, we have that Trevisan's extractor $\operatorname{Ext}_C : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a $(n^\gamma m + 8 \log m + 8 \log 1/\varepsilon + O(1), \varepsilon)$-strong extractor with an s-bit weak random seed against quantum adversaries, where the seed has length $d = O\left(\frac{1}{\beta^2\gamma} \log n\right)$ and min-entropy $s = \left(1 - \frac{\frac{1}{2}-\beta}{c}\right) d$, for some constant c.*[10]

Choosing β and γ to be constants results in a seed of length $d = O(\log n)$ with a possible entropy-loss linear in d. The output length is the same as in §8.4.2, $m = n^{\alpha-\gamma} - o(1) = H_{\min}(X|E)^{1-\frac{\gamma}{\alpha}} - o(1)$.

If we are interested in extracting all the min-entropy of the source, we can combine Lemma 147 with the extractor from §8.4.1. The results in a new extractor with seed length $d = O(\log^3 n)$ and seed min-entropy $s = d - O(\sqrt[3]{d})$.

8.5 Other variations of Trevisan's scheme

There exist many results modifying and improving Trevisan's extractor. We briefly describe a few of them here, and refer to [104] for a more extensive review.

Some of these constructions still follow the "design and 1-bit extractor" pattern — hence our work implies that they are immediately quantum-proof with roughly the same parameters — e.g., the work of Raz et al. [98] and Lu [79], which were mentioned in §8.4 and correspond to modifications of the design and 1-bit extractor respectively. Other results such as [98, 114, 105] replace the binary list-decoding codes with multivariate codes over a field F. Raz et al. [98] thus reduce the dependence of the seed on the error from $O(\log^2 1/\varepsilon)$ to $O(\log 1/\varepsilon)$. Ta-Shma et al. [114] and Shaltiel and Umans [105] reduce the size of the seed to $d \leq 2 \log n$ in several constructions with different parameters for the min-entropy. In these constructions the connection to 1-bit extractors is not clear anymore, and it is therefore not guaranteed that the construction is quantum-proof.

Raz et al. [98] extract a little more randomness than we do in §8.4.1. They achieve this by composing (in the sense described in §A.5.2) the scheme of Corollary 88 with an extractor by Srinivasan and Zuckerman [110], which has an optimal entropy loss of $\Delta = 2 \log 1/\varepsilon + O(1)$. In the presence of quantum side information this extractor has only been proven to have an entropy loss of $\Delta = 4 \log 1/\varepsilon + O(1)$ in [117], hence our slightly weaker result in Corollary 89. This still leaves room for a small improvement.

[10] If we work out the exact constant, we find that $c \approx d/t \approx \frac{8(1+4a)}{\beta\gamma \ln 2}$, for $\varepsilon = n^{-a}$.

Impagliazzo et al. [57] and then Ta-Shma et al. [113] modify Trevisan's extractor to work for a sub-polynomial entropy source, still using a seed of size $d = O(\log n)$. Ta-Shma et al. [113] achieve a construction which can extract all the min-entropy k of the source with such a seed length, for some $k = o(n)$. While it is unclear whether these modifications preserve the "design and 1-bit extractor" structure, it is an interesting open problem to analyze them in the context of quantum side information.

Another research direction consists in making these constructions practically implementable. Whether the extractor is used for privacy amplification [19, 20], randomness recycling [58], or for generating true randomness [130], the extractor has to be efficiently computable. This does not seem to be the case of Trevisan's construction [108]. An important open problem is thus to find variations which are efficient to execute.

Chapter 9

Certifiable Quantum Dice

In Chapter 8 we studied how random bits could be manipulated through the use of extractors. In this chapter we introduce a method by which one may generate *certifiably random* bits through an interaction with two *untrusted* provers. The only assumption we will need in order to guarantee the bits' randomness is that the provers obey the no-signaling condition.

We first motivate the problem of generating certifiable randomness, and introduce our results, in the following section. In Section 9.2 we define the *guessing game*, an important conceptual tool in the proofs of our main results, stated in Theorem 93 and Theorem 94. In Section 9.3 we prove Theorem 93, while Theorem 94 is proven in Section 9.4.

9.1 Introduction

A source of independent random bits is a basic resource in many modern-day computational tasks, such as cryptography, game theoretic protocols, algorithms and physical simulations. Moreover, these tasks place different demands on the quality of the randomness (e.g. the need for privacy in cryptographic applications). It is of great interest, therefore, to construct a physical device for reliably and provably outputting a stream of random bits. Testing such a device poses a fundamental problem — since all outputs should be output with equal probabilitythere is no basis for rejecting any particular output of the device.

Starting in the mid-80's, computer scientists considered the question of extracting truly random bits from adversarially controlled physical sources of randomness, such as the semi-random source [101], and weak random sources [31]. This sequence of papers has culminated in sophisticated algorithms called randomness extractors that are guaranteed to output a sequence of truly random bits from physical sources of low-quality randomness (see [104] for a survey). It was clear, in a classical world, that these results were the best one could hope for — while it was necessary to assume that the physical device outputs randomness (since that could not be tested), minimal assumptions were made about the quality of randomness output.

Quantum mechanics provides a surprising path around this fundamental barrier — it provides a way of testing that the output of a certain kind of device is truly random. Recall the famous CHSH game, illustrated in Figure 9.1. In this game two *non-communicating* parties, represented by spatially separated boxes A, B, are given inputs $x, y \in \{0, 1\}$ respectively. Their task is to produce outputs $a, b \in \{0, 1\}$ such that the *CHSH condition* $a \oplus b = x \wedge y$ holds. Let $p_{\texttt{CHSH}}$ be the probability that a certain pair of boxes produces outputs satisfying this condition, when the inputs x, y are chosen uniformly at random.

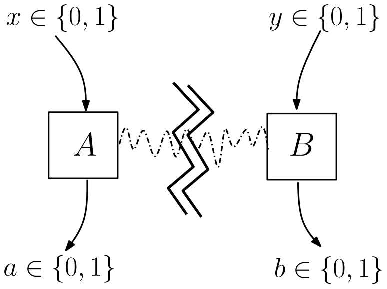

Check: $a \oplus b = x \wedge y$

Figure 9.1: The CHSH game. Any pair of boxes A, B is characterized by a distribution $p(a, b|x, y)$ which is required to be *no-signaling*: the marginal distribution of b is independent of x, and that of a is independent of y.

Classical players can achieve a success probability at most $p_{\texttt{CHSH}} \leq \frac{3}{4}$, but there is a quantum strategy that succeeds with $p_{\texttt{CHSH}} = \cos^2 \pi/8 \approx 0.85$. Indeed, we may define the *quantum regime* corresponding to success probability $3/4 < p_{\texttt{CHSH}} \leq \cos^2 \pi/8 \approx 0.85$. For any value in that range there is a simple quantum-mechanical pair of boxes, still obeying the no-signaling condition, which achieves that success probability.

These well-known facts have a striking consequence: any boxes producing correlations that fall in the quantum regime *must be randomized*! Indeed, deterministic boxes are inherently classical, so that their success probability must fall in the classical regime $p_{\texttt{CHSH}} \leq 3/4$. Hence a simple *statistical test* guaranteeing the presence of randomness, under a single assumption on the process that produced the bits: that it obeys the no-signaling condition. This powerful observation was first made in Colbeck's Ph.D. thesis [28] (see also [29] for an expanded version). The idea was then developed in a paper by Pironio et. al. [92], where the first quantitative bounds on the amount of randomness produced were shown.

An efficient and testable randomness-generation protocol

This method of generating randomness is not very efficient. Choosing a pair of inputs for the boxes requires 2 bits of randomness, so the 2 bits that are output certainly do not contain

Protocol A

1. Let ℓ, Δ be two integers given as input. Set $k = \lceil 10 \log^2 \ell \rceil$ and $m = \Delta\,\ell$.
2. Choose $T \subseteq [m]$ uniformly at random by selecting each position independently with probability $1/\ell$.
3. Repeat, for $i = 1, \ldots, m$:
    - 3.1 If $i \notin T$, then
        - 3.1.1 Set $x = y = 0$ and choose x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0,1\}^k$.
        - 3.1.2 If $a \oplus b$ has more than $\lceil 0.16k \rceil$ 1's then reject and abort the protocol. Otherwise, continue.
    - 3.2 If $i \in T$,
        - 3.2.1 Pick $x, y \in \{0,1\}$ uniformly at random, and set x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0,1\}^k$.
        - 3.2.2 If $a \oplus b$ differs from $x \wedge y$ in more than $\lceil 0.16k \rceil$ positions then reject and abort the protocol. Otherwise, continue.
4. If all steps accepted, then accept.

Figure 9.2: Protocol A uses $O(\Delta \log \ell)$ bits of randomness and makes $O(\ell \log^2 \ell)$ uses of the boxes. Theorem 93 shows that $\Omega(\ell)$ bits of randomness are produced, with security $\varepsilon = \exp(-\Omega(\Delta))$.

more randomness than was used.[1]

Instead, consider the following randomness-efficient protocol. Let n be the target number of random bits to be generated, and ε a "security" parameter. Inputs in the protocol are grouped in $m = C\,\lceil n \log(1/\varepsilon) \rceil$ successive blocks of $k = 10\lceil \log^2 n \rceil$ pairs of inputs each, where C is a large constant. Inputs in a given block consist of a fixed pair (x, y) repeated k times. Most blocks use the $(0,0)$ input, but approximately $10^3 \lceil \log(1/\varepsilon) \rceil$ of them are selected at random and marked as "Bell" blocks. In those blocks a random pair of inputs $(x, y) \in \{0,1\}^2$ is chosen, and used as inputs throughout the block. Finally, the sequence of outputs produced by the boxes is accepted if, in every block, the CHSH constraint is satisfied by at least $0.84k$ of the blocks's input/output pairs.[2]

[1] In fact, one may show that boxes having a probability of success in the CHSH game that is close to the optimal quantum value produce at most 1.25 random bits per use, on average [92].

[2] Note that honest boxes, playing each round independently, will indeed satisfy the CHSH condition in each block on average with probability $1 - 2^{-\Omega(\log^2 n)}$, so that by a union bound it is very unlikely that they

The following theorem shows that this protocol (formally described as Protocol A in Figure 9.2) can be used to generate certifiably random bits.

Theorem 93. *There exists a constant $C > 1$ such that the following holds. Let $\varepsilon > 0$ be given, and n an integer. Set $\Delta = 10^3 \lceil \log(1/\varepsilon) \rceil$, and $\ell = C\,n$. Let $(\mathcal{A}, \mathcal{B})$ be an arbitrary pair of no-signaling boxes used to execute Protocol A, B the random variable describing the bits output by $\mathcal{B}$ in protocol A, and CHSH the event that the boxes' outputs are accepted in the protocol. Then for all large enough n at least one of the following holds :*

- *Either* $H_\infty^\varepsilon(B|\mathit{CHSH}) \geq n$,
- *Or* $\Pr\big(\mathit{CHSH}\big) \leq \varepsilon$.

Moreover, Protocol A requires $O(\log n \log(1/\varepsilon))$ bits of randomness, and makes $O(n \log^2 n \log(1/\varepsilon))$ uses of the boxes.

We note that the second condition in the theorem is necessary, as there is always an unavoidable chance that the boxes successfully guess their whole inputs, and deterministically produce matching outputs. The theorem guarantees that the probability of this happening can be bounded by an inverse-exponential in the number of random bits used.

The theorem as stated only guarantees that the bits output by the device have large (smooth) min-entropy. In order to obtain bits that are (close to) uniformly random, one may apply an extractor. There exists efficient constructions of such devices which will convert B into roughly $H_\infty^\varepsilon(B|\text{CHSH})$ bits that are ε-close, in statistical distance, to uniform. In order to do so, the best extractors will require an additional $O(\log n)$ many uniformly random bits to be used as seed [48].

Compared to the basic procedure outlined earlier, Protocol A uses two main ideas in order to save on the randomness required. The first idea is to restrict the inputs to $(0, 0)$ most of the time. Only a few randomly placed checks (the Bell blocks) are performed in order to verify that the boxes are generating their inputs honestly. This idea was already used in [92], and led to a protocol with a quadratic $\sqrt{n} \to n$ expansion of randomness.

The second idea is to systematically group inputs to the boxes into blocks of k successive, identical pairs and check that the CHSH correlations are satisfied on average *in every block*. This is necessary: if one was to only check that the CHSH condition is satisfied on average over the whole protocol, then boxes systematically producing the outputs $(0, 0)$ would lead to a large — close to 100% — violation. Hence the more robust checking that we perform forces the boxes to play "honestly" and produce randomness in essentially every block.

Moreover, the block structure of the inputs also plays a key role in the analysis of the protocol, which is based on the definition of a simple "guessing game", explained in Section 9.2. The main point is that if box $\mathcal{B}$'s output in a certain block is likely to be a particular string, then Alice, given access to $\mathcal{A}$, can guess $\mathcal{B}$'s input $y \in \{0, 1\}$ based on whether $\mathcal{A}$'s output is

will fail the CHSH condition in any of the blocks.

"close" or "far" in Hamming distance from that particular string. This provides a way for Alice to guess $\mathcal{B}$'s input with probability greater than 1/2, violating the no-signaling condition placed on the boxes. This style of reasoning can be used to establish that $\mathcal{B}$'s output must have high min-entropy, thus yielding Theorem 93. The proof is given in Section 9.3.

To understand the significance of Theorem 93, it may be instructive to recall the common paraphrasing of Einstein's quote from his 1926 letter to Max Born expressing his unhappiness with quantum mechanics as "God does not play dice with the Universe." Clearly a device based on quantum mechanics can be used to generate randomness — simply prepare a qubit in the $|0\rangle$ state, apply a Hadamard gate, and measure the resulting state in the computational basis: the outcome is a uniformly random bit. However, in addition to believing the correctness of quantum mechanics, to trust that such a device produces random bits one must believe that the manufacturer is trustworthy, experimentally skilled, and that the device is always well calibrated. These difficulties are compounded by the fact that the postulates of quantum physics forbid any classical observer from getting more than a small probabilistic digest of the internal quantum state of the system. The randomness generation protocol presented above has the property that the output is guaranteed to be random based only on the observed correlations in the output (violations of Bell inequalities), and on the relativistic assumption that information does not travel faster than light. In this sense it might be appropriate to deem that it is "Einstein certifiable"!

Quantum adversaries

We have described a simple protocol that guarantees the production of bits that are statistically close to uniform. Suppose these random bits were used later in an interactive cryptographic protocol. In that case it is crucial that the bits generated appear close to uniform not only to the (honest) user of the protocol, but also to any adversary in the cryptographic protocol.

For concreteness, consider the following catastrophic scenario: the maker of the boxes, call her Eve, inserted an undetectable "back-door" by not only entangling $\mathcal{A}$ and $\mathcal{B}$ together, but extending this entanglement to reach into her own, private, laboratory. Eve knows that the protocol mostly uses 0's as inputs to $\mathcal{B}$. Betting on this she repeatedly makes a specific measurement on her system, which reliably produces the same output as $\mathcal{B}$ in case its input was a 0. If we assume that $\mathcal{B}$'s outputs are uniformly distributed then such a strategy does not obviously violate the no-signaling constraint between $\mathcal{B}$ and Eve. But Eve learns most of $\mathcal{B}$'s output: while in isolation it may be random, it is totally insecure!

We rule out this scenario by showing an analogue to Theorem 93 which also holds in the presence of a quantum adversary. The theorem applies to a slight variant of the protocol used in the previous section, described as Protocol B in Figure 9.3. The main differences are that the number of random bits used in that protocol is slightly larger, $O(\log^3 n)$ instead of $O(\log n)$, and the protocol is based on using an "extended" version of the CHSH game, which will be introduced in Section 9.4.

Theorem 94. *Let $\alpha, \gamma > 0$ be such that $\gamma \leq 1/(10 + 8\alpha)$, and n an integer. Set $C = \lceil 100\,\alpha \rceil$, and $\ell = n^{1/\gamma}$. Let $(\mathcal{A}, \mathcal{B})$ be an arbitrary pair of no-signaling boxes used to execute Protocol B, CHSH the event that the boxes' outputs are accepted in the protocol, and B' the random variable describing the bits output by $\mathcal{B}$, conditioned on CHSH. Let E be an arbitrary quantum system, possibly entangled with $\mathcal{A}$ and $\mathcal{B}$, but such that no communication occurs between $\mathcal{A}, \mathcal{B}$ and E once the protocol starts. Then for all large enough n at least one of the following holds:*

- *Either $H_\infty^\varepsilon(B'|E) \geq n$,*
- *Or $\Pr\big(\mathit{CHSH}\big) \leq \varepsilon$,*

where $\varepsilon = n^{-\alpha}$. Moreover, Protocol B uses only $O(\gamma^{-3}\log^3 n)$ bits of randomness.

Indication that dealing with quantum, rather than classical, adversaries may present substantial new difficulties may be found in the area of strong extractor constructions. There are examples of such constructions, secure against classical adversaries, that dramatically fail in the presence of quantum adversaries with even smaller prior information [46]. Luckily, other constructions, such as a very efficient construction due to Trevisan [119], have been shown secure even against quantum adversaries [112, 34]. One may use such a "quantum-proof" extractor in order to efficiently transform the bits output in Protocol B into ones that are statistically close to uniform even from the point of view of the adversary at the cost of an additional $O(\log^2 n)$ bits of fresh randomness.

A reason to think that the power of a quantum adversary in learning $\mathcal{B}$'s output may be limited comes from a delicate property of entanglement, its *monogamy* [115]. Informally, monogamy states that a tripartite entangled state $|\Psi\rangle_{ABE}$ cannot be maximally entangled both between A and B and between B and E. Since Protocol B enforces very strict correlations between the outputs of $\mathcal{A}$ and $\mathcal{B}$, one may hope that these correlations will pre-empt any strong correlation between $\mathcal{B}$ and an arbitrary E.

Interestingly, the proof of Theorem 94 makes crucial use of the properties of a specific construction of a quantum-proof extractor, based on Trevisan's construction and the t-XOR code, that was first outlined in [33]. This construction is used to prove the following information-theoretic lemma. The lemma gives an operational interpretation to a random variable having small smooth min-entropy conditioned on a quantum system, and may be of independent interest.

Lemma 95. *Let ρ_{XE} be a state such that X is a classical random variable distributed over m-bit strings, and E is an arbitrarily correlated quantum system. Let $\varepsilon, \delta > 0$, and $K = H_\infty^\varepsilon(X|E)$. Then there exists a subset $V \subseteq [m]$ of size $v = |V| = O(K\log^2 m)$, and for every v-bit string z a measurement M_z on E such that, with probability at least $\Omega(\varepsilon^6/m^6)$, M_{X_V} produces a string Y that agrees with X in a fraction at least $1 - \frac{1}{\log m}$ of positions.*

In essence Lemma 95 states that, given access to some of the bits of X (the ones indexed by V), and to the quantum system E, one can predict the remainder of the string X with inverse-polynomial success probability. In the range of large K (at least inverse-polynomial in m), this is much higher than the inverse-exponential probability that one would get by measuring E directly, without using any "advice" bits.

The proof of lemma 95 mostly follows from the proof of security of Trevisan's extractor against quantum adversaries presented in [34]. Since however it does not follow as a black-box, we give a detailed outline of the proof of the lemma in Appendix A.6.2.

Related work. In independent recent work, Fehr, Gelles and Schaffner [40] showed that security of a randomness-generation scheme against quantum adversaries could be automatically deduced from its security against classical adversaries, *provided* the result proved in the non-adversarial setting takes on a very specific form. Namely, the bound on the randomness generated in that case should be a function of the average violation of a Bell inequality by the devices' outputs, and should hold even conditioned on any inputs to the boxes in the protocol. They use their results to present an exponential randomness-expansion scheme based on the use of four, instead of two, no-signaling devices.

Neither condition required for the reduction in [40] to hold seems to be satisfied in our setting: our bound does not rely solely on the average violation of the CHSH inequality by the devices[3]. Moreover, we prove a bound on the min-entropy that holds on average over the choice of inputs in the protocol, rather than for all inputs, as needed for the reduction in [40].

Recent work by Colbeck and Renner [30] studies a related question, that of improving the quality of a given source of weak randomness. Specifically, they show that if one is given access to a so-called Santha-Vazirani source then one can produce bits that are guaranteed to be statistically close to uniform by using the violation of a specific Bell inequality by a pair of untrusted no-signaling devices.

9.2 The guessing game

Consider the following simple guessing game. In this game, there are two cooperating players, Alice and Bob. At the start of the game Bob receives a single bit $y \in \{0,1\}$ chosen uniformly at random. The players are then allowed to perform arbitrary computations, but are not allowed to communicate. At the end of the game Alice outputs a bit a, and the players win if $a = y$.

Clearly, any strategy with success probability larger than $\frac{1}{2}$ indicates a violation of the no-communication assumption between Alice and Bob. At the heart of the proofs of both Theorem 93 and Theorem 94 is a reduction to the guessing game. Assuming there existed a

[3] In fact, we argued that merely observing a large average violation would not suffice to guarantee an exponential expansion of randomness.

pair of boxes violating the conclusions of either theorem, we will show how these boxes may be used to devise a successful strategy in the guessing game, contradicting the no-signaling assumption placed on the boxes.

To illustrate the main features of the strategies we will design later, consider the following simplified setting. Let $\mathcal{A}, \mathcal{B}$ be a given pair of boxes taking inputs $X, Y \in \{0,1\}$ and producing outputs $A, B \in \{0,1\}^k$ respectively. Assume the following two properties hold. First, if the input to $\mathcal{B}$ is $Y = 0$ then its output B is essentially deterministic, in the sense that $B = b_0$ with high probability. Second, whatever their inputs, the boxes' outputs satisfy the CHSH constraint on average: at least 84% of $i \in [k]$ are such that $A_i \oplus B_i = X \wedge Y$. Then we claim that there is a strategy for Alice and Bob in the guessing game, using $\mathcal{A}$ and $\mathcal{B}$, that succeeds with probability strictly larger than 1/2, demonstrating that the boxes must be signaling.

Alice and Bob's strategy is the following. Alice is given access to $\mathcal{A}$ and Bob to $\mathcal{B}$. Upon receiving his secret bit y, Bob inputs it to $\mathcal{B}$, collecting outputs $b \in \{0,1\}^k$. Alice chooses an $x \in \{0,1\}$ uniformly at random, and inputs it to $\mathcal{A}$, collecting outputs $a \in \{0,1\}^k$. Let b_0 be the k-bit string with the highest probability of being output by $\mathcal{B}$, conditioned on $y = 0$. Alice makes a decision as follows: she computes the relative Hamming distance $d = d_H(a, b_0)$. If $d < 0.2$ she claims "Bob's input was 0". Otherwise, she claims "Bob's input was 1".

By assumption, if Bob's secret bit was $y = 0$, then his output is almost certainly b_0. By the CHSH constraint, independently of her input Alice's output a lies in a Hamming ball of radius 0.16 around b_0. So in this case she correctly decides to claim "Bob's input was 0".

In the case that Bob's secret bit was $y = 1$, the analysis is more interesting. Let b be the actual output of $\mathcal{B}$. Let a_0 and a_1 be $\mathcal{A}$'s output in the two cases $x = 0$ and $x = 1$ respectively. We claim that the Hamming distance $d_H(a_0, a_1) \geq 0.68$. This is because by the CHSH constraint, $d_H(a_0, b) \leq 0.16$, while $d_H(a_1, b) \geq 0.84$. Applying the triangle inequality gives the lower bound on the distance between a_0 and a_1. This lower bound is large enough that both a_0 and a_1 cannot lie in the Hamming ball of radius 0.16 around b_0 (observe that this argument makes no use of the actual location of b!). Thus in the case $y = 1$, Alice correctly outputs "Bob's input was 1" with probability at least 1/2.

Overall Alice and Bob succeed in the guessing game with probability 3/4, which contradicts no-signaling.

Clearly there is a lot of slack in the above reasoning, since for contradiction it suffices to succeed in the guessing game with any probability strictly greater than 1/2. By being more careful it is possible to allow Bob's output on $y = 0$ to have more min entropy, as well as allow for a small probability that the boxes' outputs may not satisfy the CHSH constraint:

Lemma 96. *Let $\beta, \gamma > 0$ be such that $\gamma + 2\beta < 1/4$, and k an integer. Suppose given a pair of boxes $\mathcal{A}, \mathcal{B}$, taking inputs $X, Y \in \{0,1\}$ and producing outputs $A, B \in \{0,1\}^k$ each. Suppose the following conditions hold:*

1. *When given input* 0*, the distribution of outputs of* $\mathcal{B}$ *has low min-entropy: there exists a* $b_0 \in \{0,1\}^k$ *such that* $\Pr(B = b_0|Y = 0) \geq 1 - \gamma$*,*

2. *The boxes' outputs satisfy the CHSH condition, on average:*

$$\Pr\big(\#\big\{i \in [k],\, A_i \oplus B_i \neq X \wedge Y\big\} \,>\, 0.16\,k\big) \,\leq\, \beta.$$

Then there is a strategy for Alice and Bob, using $\mathcal{A}$ *and* $\mathcal{B}$*, with gives them success probability strictly greater than* $1/2$ *in the guessing game.*

Proof. Alice and Bob's strategy in the guessing game is as described above. Let b_0 be the k-bit string that is most likely to be output by $\mathcal{B}$, conditioned on $y = 0$.

We first show that, if Bob's input was $y = 0$, then Alice claims that Bob had a 0 with probability at least $1 - \gamma - 2\beta$. By the first condition in the lemma, Bob obtains the output b_0 with probability at least $1 - \gamma$. Moreover, by the second condition the CHSH constraint will be satisfied with probability at least $1 - 2\beta$ on average over Alice's choice of input, given that Bob's input was $y = 0$. Given $y = 0$, whatever the input to $\mathcal{A}$ the CHSH constraint states that $d_H(a, b) < 0.16$. Hence by a union bound Alice will obtain an output string a at relative Hamming distance at most 0.16 from b_0 with probability at least $1 - \gamma - 2\beta$.

Next we show that, in case Bob's input in the guessing game is $y = 1$, Alice claims that Bob had a 1 with probability at least $\frac{1}{2}\big(1 - 8\beta\big)$. Let b' the actual output produced by Bob. By the second condition in the lemma and Markov's inequality, with probability at least $1 - 4\beta$ the output b' is such that the CHSH constraint will be satisfied with probability at least $1 - 4\beta$ simultaneously for both of Alice's possible choices of input.

Suppose this holds. If Alice chooses $x = 0$ then the CHSH constraint indicates that the corresponding a_0 should be such that $d_H(a_0, b') \leq 0.16$, while in case she chooses $x = 1$ her output a_1 should satisfy $d_H(a_1, b') \geq 0.84$. By the triangle inequality, $d_H(a_0, a_1) \geq 0.68$: whatever the value of b', only one of a_0 or a_1 can be at distance less than 0.2 from b_0. By a union bound, with probability at least $1 - 8\beta$ there is a choice of input for Alice that will make her claim Bob had a 1, and she chooses that input with probability 1/2.

The two bounds proven above together show that Alice's probability of correctly guessing Bob's input in the guessing game is at least

$$p_{succ} \,\geq\, \frac{1}{2}\big(1 - 2\gamma\big) + \frac{1}{2}\frac{1 - 8\beta}{2} \,=\, \frac{1}{2} + \Big(\frac{1}{4} - 2\beta - \gamma\Big),$$

which is greater than 1/2 whenever $2\beta + \gamma < 1/4$, proving Lemma 96. □

9.3 Proof of the main result

Theorem 93 asserts that, given any pair $(\mathcal{A}, \mathcal{B})$ of non-signaling boxes, if the outputs of $\mathcal{B}$ do not contain much min-entropy (when its inputs are chosen as in Protocol A, described in

Figure 9.2), then the boxes can only satisfy the CHSH constraints imposed in the protocol with small probability.

We prove Theorem 93 by a reduction to the guessing game introduced in Section 9.2. Suppose that there existed a pair of boxes such that neither of the theorem's conclusions was satisfied. Recall that Protocol A calls for a total of mk uses of the boxes, divided into m blocks of k pairs of identical inputs each. We show that, provided the CHSH constraints are satisfied in all blocks with non-negligible probability, there must exist a special block $i_0 \in [m]$ in which the boxes' outputs, conditioned on specific past values, have properties close to those required in Lemma 96. This lets us carry out a reduction to the guessing game, leading to a contradiction of the no-signaling assumption. The exact properties of the special block that we obtain are described in Claim 97 below.

Modeling events in the protocol. To model the situation, we introduce four sequences of random variables $X = (X_i), Y = (Y_i), A = (A_i), B = (B_i) \in \big(\{0,1\}^k\big)^m$, where m is the number of blocks of the protocol. X and Y are distributed as in Protocol A, and A, B are random variables describing the boxes' respective outputs when their inputs are X and Y. For $i \in [m]$, let CHSH_i be the event that $d_H(A_i \oplus B_i, X_i \wedge Y_i) \leq 0.16$, and $\mathrm{CHSH} = \bigwedge_i \mathrm{CHSH}_i$. We will also use the shorthand $\mathrm{CHSH}_{<i} = \bigwedge_{j<i} \mathrm{CHSH}_j$. Finally, we let T_j be a random variable denoting the j-th Bell block, chosen jointly by Alice and Bob at the start of Protocol A.

Claim 97. *There exists a constant $C > 1$ such that the following holds. Let $2^{-Cn} < \varepsilon < 1/5$ and $\Delta = 10^3\lceil\log(1/\varepsilon)\rceil$. Suppose that (i) $H_\infty^\varepsilon(B|CHSH) \leq n$, and (ii) $\Pr(CHSH) \geq \varepsilon$. Let $m = C\,\Delta n$. Then for all large enough n there exists an index j_0 and a set G satisfying $\Pr(G) \geq \varepsilon^5$ such that the following hold.*

- *$\mathcal{B}$'s output in the j_0-th Bell block T_{j_0} is essentially deterministic:*

$$\forall b \in G, \qquad \Pr(B_{T_{j_0}} = b_{T_{j_0}} | CHSH_{<T_{j_0}}, B_{<T_{j_0}} = b_{<T_{j_0}}) \geq 0.99, \tag{9.1}$$

- *Irrespective of the input $Y_{T_{j_0}} = y_{T_{j_0}} \in \{0,1\}$ to $\mathcal{B}$ in the T_{j_0}-th block, the CHSH condition is satisfied with high probability:*

$$\forall b \in G, \qquad \Pr(CHSH_{T_{j_0}} | CHSH_{<T_{j_0}}, B_{<T_{j_0}} = b_{<T_{j_0}}, Y_{T_{j_0}} = y_{T_{j_0}}) \geq 0.9. \tag{9.2}$$

The proof of Claim 97 mostly follows from an appropriate chained application of Baye's rule, and is given in Appendix A.6.1. In order to conclude the proof of Theorem 93 it remains to show how the special block identified in Claim 97 can be used to show that boxes $\mathcal{A}$ and $\mathcal{B}$ satisfying the claim's assumptions may be used successfully in the guessing game.

Consider the following strategy for Alice and Bob in the guessing game. In a preparatory phase (before Bob receives his secret bit y), Alice and Bob run protocol A with the boxes $\mathcal{A}$ and $\mathcal{B}$, up to the i_0-th block (excluded). Bob communicates $\mathcal{B}$'s outputs up till that block to Alice. Together they check that the CHSH constraint is satisfied in all blocks preceding the i_0-th; if not they abort. They also verify that Bob's outputs are the prefix of a string $b \in G$; if not they abort. The guessing game can now start: Alice and Bob are separated and Bob is given his secret input y.

Given the conditioning that Alice and Bob have performed before the game started, once it starts boxes $\mathcal{A}$ and $\mathcal{B}$ can be seen to satisfy both conditions of Lemma 96. Indeed, since under the input distribution specified in Protocol A $\mathcal{B}$ receives a 0 as input in block i_0 with probability at least $1/2$, condition 1. in Lemma 96 holds with $\gamma = 1/50$ as a consequence of item 1 in Claim 97. Condition 2 in Lemma 96 puts a bound on the probability of the CHSH condition being satisfied under the uniform input distribution. Hence item 2 from Claim 97 implies that condition 2 holds with $\beta = 1/10$. Since $\gamma + 2\beta = 0.22 < 1/4$, Lemma 96 concludes that the boxes $\mathcal{A}$ and $\mathcal{B}$ must be signaling in the i_0-th block, a contradiction. This finishes the proof of Theorem 93.

9.4 Producing random bits secure in the presence of a quantum adversary

In this section we prove Theorem 94. We first give an overview of the proof, describing the main steps, in the next section. The formal proof is given in Section 9.4.2

9.4.1 Proof overview

Theorem 94 is based on Protocol B, a variant of Protocol A which replaces the use of the CHSH game by the following "extended" variant. In this game each box may receive one of four possible inputs, labeled $(A,0),(A,1),(B,0),(B,1)$. An input such as "$(A,1)$" to either box means: "perform the measurement that $\mathcal{A}$ would have performed in the honest CHSH strategy, in case its input had been a 1". The advantage of working with this game is that there exists an optimal strategy (the one directly derived from the honest CHSH strategy) in which both players always output identical answers when their inputs are equal.

Protocol B follows the same structure as Protocol A. Inputs are divided into groups of $k = \lceil 10 \log^2 n \rceil$ identical inputs. There are $m = O(n^{1/\delta} \log^2 n)$ successive blocks, where $\delta > 0$ is a small parameter. Most blocks use the same input $(A,0)$ to both boxes. A random subset $T \subseteq [m]$ of approximately $\log^2 n$ blocks are designated as Bell blocks. In such blocks $\mathcal{A}$ is given an input at random in $\{(A,0),(A,1)\}$, while $\mathcal{B}$ is given an input at random in $\{(A,0),(B,0)\}$.

As in the proof of Theorem 93 we will prove Theorem 94 by contradiction, through a reduction to the guessing game. In the non-adversarial case the crux of the reduction consisted in identifying a special block $i_0 \in [m]$ in which $\mathcal{B}$'s output B was essentially deterministic, conditioned on past outputs. In the adversarial setting, however, B may

Protocol B

1. Let ℓ, C be two integers given as input. Set $k = \lceil 10 \log^2 \ell \rceil$ and $m = \lceil C\ell \log^2 \ell \rceil$.
2. Choose $T \subseteq [m]$ uniformly at random by selecting each position independently with probability $1/\ell$.
3. Repeat, for $i = 1, \ldots, m$:
   - 3.1 If $i \notin T$, then
     - 3.1.1 Set $x = y = (A, 0)$ and choose x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0, 1\}^k$.
     - 3.1.2 If $a \neq b$ then reject and abort the protocol. Otherwise, continue.
   - 3.2 If $i \in T$,
     - 3.2.1 Pick $x \in \{(A, 0), (A, 1)\}$ and $y \in \{(A, 0), (B, 0)\}$ uniformly at random, and set x, y as inputs for k consecutive steps. Collect outputs $a, b \in \{0, 1\}^k$.
     - 3.2.2 If either $a = b$ and $x = y$, or $d_H(a, b) \leq 0.16$ and $y = (B, 0)$, or $d_H(a, b) \in [0.49, 0.51]$ and $x = (A, 1)$ and $y = (A, 0)$ then continue. Otherwise reject and abort the protocol.
4. If all steps accepted, then accept.

Figure 9.3: Protocol B uses $O(\log^3 \ell)$ bits of randomness and makes $O(\ell \log^4 \ell)$ uses of the boxes. Theorem 94 shows that $\Omega(\ell^\gamma)$ bits of randomness are produced, where $\gamma > 0$ is a constant depending on the security parameter ε one wants to achieve.

be perfectly uniform, and such a block may not exist. Instead, we start by assuming for contradiction that the min-entropy of Bob's output conditioned on Eve's information is small: $H_\infty^\varepsilon(B|E) \leq n$.

Previously in the guessing game Alice tried to guess Bob's secret input $y \in \{0,1\}$. She did so by using her prediction for $\mathcal{B}$'s outputs, together with the CHSH constraint and her own box $\mathcal{A}$'s outputs. Here we team up Alice and Eve. Alice will provide Eve with some information she obtained in previous blocks of the protocol, and based on that information Eve will attempt to make an accurate prediction for $\mathcal{B}$'s outputs in the special block. Alice will then use that prediction to guess y, using as before the CHSH constraint and her own box $\mathcal{A}$'s outputs.

The reconstruction paradigm. We would like to show that, under our assumption on $H_\infty^\varepsilon(B|E)$, Eve can perform the following task: accurately predict (part of) B, given auxiliary information provided by Alice. We accomplish this by using the "reconstruction" property of certain extractor constructions originally introduced by Trevisan [119]. Recall that an extractor is a function which maps a string B with large min-entropy (conditioned on side information contained in E) to a (shorter) string Z that is statistically close to uniform even from the point of view of an adversary holding E. The reconstruction proof technique proceeds as follows: Suppose an adversary breaks the extractor. Then there exists another adversary who, given a small subset of the bits of the extractor's input as "advice", can reconstruct the *whole* input. Hence the input's entropy must have been at most the number of advice bits given.

For the purposes of constructing extractors, one would then take the contrapositive to conclude that, provided the input has large enough entropy, the extractor's output must be indistinguishable from uniform, thereby proving security. Here we work *directly* with the reconstruction procedure. Suppose that B has low min-entropy, conditioned on Eve's side information. If we were to apply an extractor to B in order to extract *more* bits than its conditional min-entropy, then certainly the output would not be secure: Eve would be able to distinguish it from a uniformly random string. The reconstruction paradigm states that, as a consequence, there is a strategy for Eve that successfully predicts the *entire* string B, given a subset of its bits as advice — exactly what is needed from Eve to facilitate Alice's task in the guessing game.

The t-XOR extractor. At this stage we are faced with two difficulties. The first is that the reconstruction paradigm was developed in the context of classical adversaries, who can repeat predictive measurements at will. Quantum information is more delicate, and may be modified by the act of measuring. The second has to do with the role of the advice bits: since they come from $\mathcal{B}$'s output B we need to ensure that, in the guessing game, Alice can indeed provide this auxiliary information to Eve, *without* communicating with Bob.

In order to solve both problems we focus on a specific extractor construction, the t-XOR extractor E_t (here t is an integer such that $t = O(\log^2 n)$). For our purposes it will suffice to think of E_t as mapping the mk-bit string B to a string of $r \ll n$ bits, each of which is the parity of a certain subset of t out of B's mk bits. Which parities is dictated by an extra argument to the extractor, its seed, based on the use of combinatorial designs. Formally,

$$\begin{aligned} E_t \,:\, \{0,1\}^{mk} \times \{0,1\}^s &\to \{0,1\}^r \\ (b,y) \qquad &\mapsto \big(C_t^1(b,y),\dots,C_t^r(b,y)\big), \end{aligned}$$

where $C_t^i(b,y)$ is the parity of a specific subset of t bits of x, depending on both i and y.

Suppose that Eve can distinguish the output of the extractor $Z = E_t(B,Y)$ from a uniformly random string with success probability ε. In the first step of the reconstruction proof, a hybrid argument is used to show that Eve can predict the parity of t bits of B chosen at random with success ε/r, given access to the parities of $O(r)$ other subsets of t bits of B as advice. This step uses specific properties of the combinatorial designs.

The next step is the most critical. One would like to argue that, since Eve can predict the parity of a random subset of t of B's bits, she can recover a string that agrees with *most* of the t-XORs of B. One could then appeal to the approximate list-decoding properties of the t-XOR code in order to conclude that Eve may deduce a list of guesses for the string B itself. Since, however, Eve is quantum, the fact that she has a measurement predicting *any* t-XOR does not imply she has one predicting *every* t-XOR: measurements are destructive and distinct measurements need not be compatible. This is a fundamental difficulty, which arises e.g. in the analysis of random access codes [7]. To overcome it one has to appeal to a subtle argument due to Koenig and Terhal [76]. They show that without loss of generality one may assume that Eve's measurement has a specific form, called the *pretty-good measurement.* One can then argue that this specific measurement may be refined into one that predicts a guess for the whole list of t-XORs of B, from which a guess for B can be deduced by list-decoding the t-XOR code.

The security of the t-XOR extractor against quantum adversaries was first shown by Ta-Shma [112], and later improved in [33, 34]. As such, the argument above is not new. Rather, our contribution is to observe that it proves *more* than just the extractor's security. Indeed, summarizing the discussion so far we have shown that, if $H_\infty^\varepsilon(B|E) \leq n$, then there is a measurement on E which, given a small amount of information about B as advice, reconstructs a good approximation to the *whole* string B with success probability $\mathrm{poly}(\varepsilon/r)$. (This is essentially the statement that is made in Lemma 95.) Most crucially, the bits of information required as advice are localized to a small subset of bits of B, of the order of the number of bits of information Eve initially has about that string. This property holds thanks to the specific extractor we are using, which is *local*: every bit of the output only depends on few bits of the input.

Completing the reduction to the guessing game. In the guessing game it is Alice who needs to hand the advice bits to Eve. Indeed, if Bob, holding box $\mathcal{B}$, was to hand them over, they could leak information about his secret input y: some of the advice bits may fall in blocks of the protocol that occur *after* the special block i_0 in which Bob is planning to use his secret y as input. This leak of information defeats the purpose of the guessing game, which is to demonstrate signaling between $\mathcal{A}$ and $\mathcal{B}$.

Hence the "extended" variant of the CHSH game introduced in Protocol B: since in most blocks the inputs to both $\mathcal{A}$ and $\mathcal{B}$ are identical, by the extended CHSH constraint enforced in the protocol their outputs should be identical. The relatively few advice bits needed by Eve occupy a fixed set of positions, and with good probability all Bell blocks will fall outside of these positions, in which case Alice can obtain the advice bits required by Eve directly from $\mathcal{A}$'s outputs.

The proof of Theorem 94 is now almost complete, and one may argue as in Lemma 96 that Alice and Eve together will be able to successfully predict Bob's secret input in the guessing game, contradicting the no-signaling assumption placed on $\mathcal{A}$ and $\mathcal{B}$. A more detailed proof of the theorem is given in the next section.

9.4.2 Proof of Theorem 94

We proceed to formally prove Theorem 94, using Lemma 95 to perform a reduction to the guessing game (Lemma 95 is proved in Appendix A.6.2). Protocol B is described in Figure 9.3. It consists of $m = \lceil C\ell \log^2 \ell \rceil$ blocks of $k = \lceil 10 \log^2 n \rceil$ rounds each, where C is a large constant, $\ell = n^{1/\gamma}$ and n is the target amount of min-entropy. Each round of the protocol selects inputs to the boxes coming from the "extended CHSH" game. That game has four questions per party: $(A,0),(A,1),(B,0),(B,1)$. We expect honest boxes to apply the following strategy. They share a single EPR pair, and perform the same measurement if provided the same input. On input $(A,0)$ the measurement is in the computational basis, and on input $(A,1)$ it is in the Hadamard basis $\{|+\rangle,|-\rangle\}$, with the outcome $|+\rangle$ being associated with the output '0'. On input $(B,0)$ the measurement is in the basis $\{\cos^2(\pi/8)|0\rangle + \sin^2(\pi/8)|1\rangle, \sin^2(\pi/8)|0\rangle - \cos^2(\pi/8)|1\rangle\}$, with the first vector being associated with the outcome '0'.

Modeling. To model the situation, introduce four sequences of random variables $X = (X_i), Y = (Y_i), A = (A_i), B = (B_i) \in \left(\{0,1\}^k\right)^m$. X and Y are distributed as in protocol B, while A, B are random variables describing the boxes' respective outputs when their inputs are X and Y. For $i \in [m]$, let CHSH_i be the following event:

$$\mathrm{CHSH}_i = \begin{cases} A_i = B_i & \text{if } X_i = Y_i, \\ d_H(A_i,B_i) \leq 0.16 & \text{if } Y_i = (B,0), \\ d_H(A_i,B_i) \in [0.49, 0.51] & \text{if } X_i = (A,1) \text{ and } Y_i = (A,0). \end{cases}$$

Honest CHSH boxes as described above satisfy CHSH_i with probability $1 - 2^{-\Omega(k)}$. Let $\mathrm{CHSH} = \bigwedge_i \mathrm{CHSH}_i$.

We introduce two new random variables to model the adversary Eve's behavior, when she performs the measurement promised by Lemma 95. We use $E^A = (E_i^A) \in \big(\{0,1\}^k\big)^m$ to denote the outcome of that measurement when the required advice bits are the bits A_V taken from $\mathcal{A}$'s outputs, and $E^B = (E_i^B) \in \big(\{0,1\}^k\big)^m$ to denote its outcome when they are the bits B_V taken from $\mathcal{B}$'s output (here V is a fixed subset of $[km]$ that will be specified later). Let G^A be the event that $d_H(E^A, B) < f_e$, and G^B the event that $d_H(E^B, B) < f_e$, where $f_e > 0$ is a parameter to be specified later. Let $j \in T$ be an index that runs over the blocks that have been designated as Bell blocks in the protocol (where T itself is a random variable). Given a Bell block j, let G_j^A be a boolean random variable such that $G_j^A = 1$ if and only if either $d_H(E_j^A, B) < 0.01$ and $Y_j = (A, 0)$, or $d_H(E_j^A, B) < 0.17$ and $Y_j = (B, 0)$. Define G_j^B symmetrically with respect to E^B instead of E^A.

We prove Theorem 94 by contradiction. Assume that both the theorem's conclusions are violated, so that (i) $H_\infty^\varepsilon(B'|E) \geq n$, where B' is a random variable describing the distribution of $\mathcal{B}$'s outputs conditioned on CHSH, and $\Pr\big(\text{CHSH}\big) \leq \varepsilon$. Here $\varepsilon = n^{-\alpha}$, where $\alpha > 0$ is a parameter.

The first step is to apply Lemma 95 with $X = B'$. The conclusion of the lemma is that there exists a subset $V \subseteq [km]$ of size $|V| = O(m^\gamma \log^2 m)$ such that, letting $f_e = 1/(\log mk)$, we have $p_s := \Pr(G^B|\text{CHSH}) = \Omega(\varepsilon^7/n^6) = \Omega(n^{-7(\alpha+\gamma)})$.

G^B denotes the event that Eve correctly predicts B on a fraction at least $1 - f_e$ of positions. Since in Protocol B the Bell blocks form only a very small fraction of the total, a priori it could still be that Eve's prediction is systematically wrong on all Bell blocks, preventing us from successfully using them in the guessing game.

The following claim shows Eve's errors cannot be concentrated in the Bell blocks. The intuition is the following. If $\mathcal{B}$'s input in a Bell block is $(A, 0)$ then nothing distinguishes this block from most others, so that Eve's prediction has no reason of being less correct than average. However, blocks in which its input is $(B, 0)$ are distinguished. We rule out the possibility that Eve's errors are concentrated in such blocks by appealing to the no-signaling condition between Eve and $\mathcal{A}$. Indeed, about half of Bell blocks in which $\mathcal{B}$'s input is $(B, 0)$ are such that $\mathcal{A}$'s input for the same block is $(A, 0)$: looking only at $\mathcal{A}$'s inputs they are indistinguishable from most other blocks. We will argue that, as long as the CHSH constraint is satisfied, Eve might as well have been given the advice bits by Alice, in which case there is no reason for her to make more errors than average in those blocks.

Claim 98. *Let T be the set of Bell blocks selected in Protocol B. Then there exists a constant $c_e < 10^{-3}$ such that the following holds.*

$$\Pr\Big(\mathrm{E}_{j\in T}\big[\,G_j^A\,\big] \;>\; 1 - \frac{c_e}{\log n},\; \mathit{CHSH}\Big) \;=\; \Omega(p_s\varepsilon) \;=\; \Omega\big(n^{-8(\alpha+\gamma)}\big).$$

The proof of Claim 98 is given in Appendix A.6.1. Based on this claim we can show an analogue of Claim 97 which will let us complete the reduction to the guessing game.

Claim 98 shows that with probability $\Omega(p_s\varepsilon)$ Eve's prediction will be correct on a fraction at least $1-c_e/\log n$ of Bell blocks. Since there are $O(\log^2 n)$ such blocks in Protocol B, with the same probability Eve only makes errors on a total number $w_e = O(\log n)$ of Bell blocks. Group the Bell blocks in groups of $20w_e$ successive blocks, and let k be an index that runs over such groups; there are $O(\log n)$ of them. Let G_k^A be the event that Eve's prediction is correct in at least 99% of the Bell blocks in group k: $G_k^A = 1$ if and only if $\mathrm{E}_{j\sim k} G_j^A \geq 0.99$, where the average is taken over the Bell blocks comprising group k. By Markov's inequality, it follows from Claim 98 that $\Pr(\wedge_k G_k^A, \text{ CHSH}) = \Omega(p_s\varepsilon)$.

Claim 99. *For all large enough n there exists a Bell block $j_0 \in T$ such that, in that block, it is highly likely that both Eve's prediction (when given advice bits from $\mathcal{A}$'s output) is correct and the CHSH constraint is satisfied, conditioned on this being so in past iterations:*

$$\Pr(G_{j_0}^A, CHSH_{j_0} | CHSH_{j<j_0}, G_{k<k_0}^A) \geq 0.98, \tag{9.3}$$

where k_0 is the index of the group containing the j_0-th Bell block.

Proof. By the chain rule, since there are $O(\log n)$ groups there will exist a group k_0 in which Eve's prediction is correct, and the CHSH condition is satisfied, with probability at least 0.99, when conditioned on the same holding of all previous groups. Since by definition Eve being correct in the group means that she is correct in 99% of that group's blocks, there is a specific block j_0 in which she is correct with probability at least 0.98. □

The reduction to the guessing game should now be clear, and follows along the same lines as the proof of Theorem 93 given in Section 9.3. Alice and Bob run protocol B, including the selection of all Bell blocks T, with the boxes $\mathcal{A}$ and $\mathcal{B}$, up to the j_0-th Bell block (excluded). Bob communicates $\mathcal{B}$'s outputs up till that block to Alice. They check that the CHSH constraint is satisfied in all blocks previous to the j_0-th; if not they abort. The guessing game can now start: Alice and Bob are separated and Bob is given his secret input y. If $y = 0$ then he chooses $(A, 0)$ as input to $\mathcal{B}$ in the j_0-th block; otherwise he chooses $(B, 0)$. He then completes the protocol honestly. Alice chooses an input $x \in \{(A, 0), (A, 1)\}$ at random for the j_0-th block, and then completes the protocol honestly.

In order to help her guess Bob's input, Alice has access to the eavesdropper Eve. Alice gives the bits a_V taken from $\mathcal{A}$'s output string a as advice bits to Eve. Eve makes a prediction e for Bob's output. Alice checks that the event $G_{<k_0}^A$ is satisfied. If not she aborts. If so, by Claim 99 we know that both CHSH_{j_0} and $G_{j_0}^A$ are satisfied with probability at least 0.98, so this must be so with probability at least 0.92 for each of the four possible pair of inputs (x, y) given to $\mathcal{A}$ and $\mathcal{B}$ in the j_0-th block.

Alice makes her prediction as follows: if either $\mathcal{A}$'s input was $(A, 0)$ and its output agrees with Eve's prediction on at least a 0.99 fraction of positions (in the j_0-th block), or $\mathcal{A}$'s input was $(A, 1)$ and its output agrees with Eve's prediction on a fraction of positions that is between 0.48 and 0.52 she claims "Bod had a 0". Otherwise she claims "Bob had a 1".

Clearly if Bob is using $(A,0)$ as his input then Alice will predict correctly with probability at least 0.92, since in that case $G_{j_0}^A$ implies that Eve predicts $\mathcal{B}$'s output with at most 1% of error. If he is using $(B,1)$ then $G_{j_0}^A$ implies that Eve's prediction will be within 0.17 relative Hamming distance of $\mathcal{B}$'s output in block j_0. By the CHSH constraint $\mathcal{A}$'s output must also be within 0.16 of $\mathcal{B}$'s output, whatever input Alice chooses. Hence $\mathcal{A}$'s output is always within $0.43 < 0.49$ of $\mathcal{B}$'s, meaning Alice will correctly claim Bob had a 1 whenever her input is $(A,1)$. Hence in that case she correctly predicts Bob's input with probability at least $0.92/2$.

Overall, conditioned on Alice not aborting her prediction is correct with probability at least 0.69 over the choice of a random input for Bob, indicating a violation of the no-signaling assumption on the boxes and proving Theorem 94.

Chapter 10

Multiprover interactive proofs with quantum messages

In this chapter we prove some general structural properties of multiprover interactive proof systems in which the verifier is quantum and exchanges quantum messages with a polynomial number of entangled provers, i.e. QMIP^* systems. A novelty of our approach is that our results make extensive use of prior shared entanglement, which is required even for *honest* provers. To the best of our knowledge all previous results in this area have focused on studying the *negative* effects of entanglement, i.e., whether or not *dishonest* entangled provers can break proof systems that are sound for any dishonest *unentangled* provers. Our work is the first to focus on the *positive* aspects of entanglement, where shared entanglement may be advantageous to *honest* provers.

The main result of this chapter is that any quantum k-prover interactive proof system that may involve polynomially many rounds can be parallelized to a *one-round* quantum $(k+1)$-prover interactive proof system of *perfect* completeness and such that the gap between completeness and soundness is still bounded by an inverse-polynomial.

To state our results more precisely, let $\mathrm{QMIP}^*(k, m, c, s)$ denote the class of languages having m-turn quantum k-prover interactive proof systems with completeness at least c and soundness at most s, where provers are allowed to share an arbitrary amount of entanglement, as defined in Chapter 3. We will call the difference $c - s$ the “gap”. As commonly used in classical multi-prover interactive proofs we use the term “round” to describe an interaction consisting of questions from the verifier followed by answers from the provers. We use the term “turn” for messages sent in one direction. One round consists of two turns: a turn for the verifier and a turn for the provers. Throughout this chapter we assume that the number m of turns and the number k of provers are functions in poly, and that completeness c and soundness s are polynomial-time computable functions of the input size, with values in $[0, 1]$. We show the following main theorem.

Theorem 100. *For any $k, m \in \mathrm{poly}$ and c, s satisfying $c - s \in \mathrm{poly}^{-1}$ there exists a function*

$p \in \mathrm{poly}$ such that

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*\left(k+1, 2, 1, 1-\frac{1}{p}\right).$$

Since it is easy to amplify the success probability without increasing the number of rounds by running multiple instances of a proof system in parallel using a different set of provers for every instance, the above theorem shows that one-round (i.e., two-turn) QMIP* systems are as powerful as general QMIP* systems.

Corollary 101. *For any $k, m, p \in \mathrm{poly}$ and c, s satisfying $c - s \in \mathrm{poly}^{-1}$, there exists a function $k' = O(k\,p\,m^2/(c-s)^2)$ such that*

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*(k', 2, 1, 2^{-p}).$$

The proof of our main theorem comes in three parts, corresponding to Section 10.2, Section 10.3, and Section 10.4. The first part shows how to convert any QMIP* system with two-sided bounded error into one with one-sided bounded error of perfect completeness without changing the number of provers. The second part shows that any QMIP* system with polynomially many turns can be parallelized to one with only three turns in which the gap between completeness and soundness is still bounded by an inverse-polynomial. Again the number of provers remains the same in this transformation. Finally, the third part shows that any three-turn QMIP* system with sufficiently large gap can be converted into a two-turn (i.e., one-round) QMIP* system with inverse-polynomial gap, by adding an extra prover.

Similar statements to our first and second parts have already been shown in [70] for *single-prover* quantum interactive proofs. Their proofs, however, heavily rely on the fact that a single quantum prover can apply arbitrary operators over all the space except for the private space of the verifier. This is not the case any more for quantum multi-prover interactive proofs, since now a quantum prover cannot access the qubits in the private spaces of the other quantum provers, in addition to those in the private space of the verifier. Hence new methods are required for the multi-prover case.

10.1 Proof overview

To transform proof systems so that they have perfect completeness, our basic idea is to adapt the quantum rewinding technique developed for quantum zero-knowledge proofs by [126] to our setting. We show how the main idea behind this technique can be used to "rewind" an unsuccessful computation that would result in rejection into a successful one. To this end, we first modify the proof system so that the honest provers can convince the verifier with probability exactly $\frac{1}{2}$ using some initial shared state and moreover no other initial shared state achieves a higher acceptance probability. This initial shared state corresponds to the auxiliary state in the case of quantum zero-knowledge proofs, and as in that scenario we can prove that the sequence of forward, backward, and forward executions of the protocol achieves perfect completeness. The obvious problem of this construction lies in proving soundness, as the dishonest provers may not use the same strategies for all of the three executions of the proof system. To settle this, we design a simple protocol that tests if the second backward execution is indeed a backward simulation of the first forward execution. The verifier performs with equal probability either the original rewinding protocol or this invertibility test without revealing which test the provers are undergoing. This forces the provers to use essentially the same strategies for the first two executions of the protocol, which is sufficient to bound the soundness. As a result we prove the following.

Theorem 102. *For any $k, m, p \in \mathrm{poly}$ and c, s satisfying $c - s \in \mathrm{poly}^{-1}$, there exists a function $m' \in \mathrm{poly}$ such that*

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*(k, m', 1, 2^{-p}).$$

For the parallelization to three turns, our approach is to first show that any QMIP* system with sufficiently large gap can be converted into another QMIP* system with the same number of provers, in which the number of rounds (turns) becomes almost half of that in the original proof system. The proof idea is that the verifier in the first turn receives the snapshot state from the original system after (almost) half of turns have been executed, and then with equal probability executes either a forward-simulation or a backward-simulation of the original system from that turn on. Thus, honest provers have to share the snapshot state of the original system, but only have to simulate the original system to convince the verifier after that. In contrast, any strategy of dishonest provers with unallowable high success probability would lead to a strategy of dishonest provers in the original system that contradicts the soundness condition. By repeatedly applying this modification, together with Theorem 102 as preprocessing, we can convert any QMIP* system into a three-turn QMIP* system with the same number of provers that still has an inverse-polynomial gap.

Theorem 103. *For any $k, m \in \mathrm{poly}$ and c, s satisfying $c - s \in \mathrm{poly}^{-1}$, there exists a function $p \in \mathrm{poly}$ such that*

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*\left(k, 3, 1, 1 - \frac{1}{p}\right).$$

For $k = 1$, this gives an alternative proof of the parallelization theorem due to [70] for single-prover quantum interactive proofs. It is interesting to note that our parallelization method does not need the controlled-swap test at all, while it *is* the key test in the Kitaev-Watrous parallelization method. Another point worth mentioning in our method is that, at every time step of our parallelized protocol, the whole system has only one snapshot state of the original system. This is in contrast to the fact that the whole system has to simultaneously treat many snapshot states in the Kitaev-Watrous method. The merit of our method is, thus, that we do not need to treat the possible entanglement among different snapshot states when analyzing soundness, which may be a main reason why our method works well even for the multi-prover case. Moreover, our method is more space-efficient than the Kitaev-Watrous method, in particular when we parallelize a system with polynomially many rounds.

To prove the third part, we will take a detour by proving that

(i) any three-turn QMIP^* system with sufficiently large gap can be modified to a three-turn *public-coin* QMIP^* system with the same number of provers and a gap of roughly similar order of magnitude, and

(ii) any three-turn public-coin QMIP^* system can be converted into a two-turn QMIP^* system without changing completeness and soundness, by adding one extra prover.

The notion of public-coin QMIP^* systems we use is a natural generalization of public-coin quantum interactive proofs in the single-prover case introduced by [81]. The corresponding complexity class is denoted by $\mathrm{QMIP}^*_{\mathrm{pub}}(k, m, c, s)$. Intuitively, at every round, a public-coin quantum verifier flips a fair classical coin at most polynomially many times, and then simply broadcasts the result of these coin-flips to all the provers. Property (i) is a generalization of the result by [81] to the multi-prover case, whereas property (ii) is completely new. We note that the protocol that arises in the proof of property (ii) is no longer public-coin. It is not hard to see that this cannot be avoided unless $\mathrm{BP} \cdot \mathrm{PP} = \mathrm{PSPACE}$: two-turn public-coin systems with any number of provers are in fact equivalent to single-prover two-turn public-coin systems (i.e., QAM systems), which are at most as powerful as $\mathrm{BP} \cdot \mathrm{PP}$ [81]. A simple proof of this fact is given in Theorem 112 at the end of Section 10.4.1.

The idea to prove (ii), assuming that the number of provers in the original proof system is k, is to send questions only to the first k provers in the new $(k+1)$-prover system, requesting the original second messages from the k provers in the original system. The verifier expects to receive from the $(k+1)$-st prover the original first messages from the k provers in the original system without asking any question to that prover. The public-coin property of the original system implies the nonadaptiveness of the messages from the verifier, which is essential to prove (ii). In fact, there is a way of directly proving the third part, but our detour enables us to show another two important properties of QMIP^* systems. Specifically, property (i) essentially proves the equivalence of *public-coin* quantum k-prover interactive proofs and general quantum k-prover interactive proofs, for any k.

Theorem 104. *For any* $k, m, p \in \mathrm{poly}$ *and* c, s *satisfying* $c - s \in \mathrm{poly}^{-1}$*, there exists a func-*

tion $m' \in \mathrm{poly}$ *such that*

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*_{\mathrm{pub}}(k, m', 1, 2^{-p}).$$

Note that in the classical case, public-coin multi-prover interactive proofs are only as powerful as single-prover interactive proofs — since every prover receives the same question from the verifier, every prover knows how other provers will behave and the joint strategy of the provers can therefore simulate any strategy of a single prover. Hence, these systems cannot be as powerful as general classical multi-prover interactive proofs unless NEXP = PSPACE. In contrast, our result shows that in the quantum case, public-coin QMIP* systems *are* as powerful as general QMIP* systems. The non-triviality of public-coin QMIP systems may be explained as follows: even if every quantum prover knows how other quantum provers will behave, still each quantum prover can apply only local transformations over a part of some state that may be entangled among the provers, which is not enough to simulate every possible strategy a single quantum prover could follow.

Property (ii) for the case $k = 1$ implies that any language in QIP (and thus in PSPACE) has a *two-prover one-round* quantum interactive proof system of perfect completeness with exponentially small error in soundness, since any language in QIP has a three-turn public-coin quantum interactive proof system of perfect completeness with exponentially small error in soundness [81].

Corollary 105. *For any* $p \in \mathrm{poly}$*,*

$$\mathrm{QIP} \subseteq \mathrm{QMIP}^*(2, 2, 1, 2^{-p}).$$

In the classical case a similar statement to the last corollary was shown by [24] (and the stronger statement that two-prover one-round interactive proofs are as powerful as general multi-prover interactive proofs was shown later by [44]). All these results are, however, not known to hold under the existence of prior entanglement among the provers. Before our result, it has even been open if PSPACE has two-prover one-round quantum interactive proof systems. Very recently, [68] succeeded in proving that the classical two-prover one-round interactive proof system for PSPACE by [24] is sound in a weak sense against any pair of dishonest prior-entangled provers: soundness is bounded away from one by an inverse-polynomial. After the completion of the present work, [59] improved the result by [68] to show that the same system for PSPACE has exponentially small soundness even against no-signaling provers (hence against entangled provers), and thus, PSPACE is now known to have classical two-prover one-round interactive proof systems even with entangled provers. These results are incomparable to ours since on one hand we have the inclusion even for QIP, and on the other hand both the verifier and the honest provers must be quantum. In contrast, in [68] and [59] both of them just follow a classical protocol.

10.2 Achieving Perfect Completeness

In this section we prove Theorem 102, showing that any QMIP* system with two-sided bounded error can be transformed into a one with one-sided bounded error of perfect completeness without changing the number of provers. For the case of a single prover, this was shown by [70], but their proof relies on the single prover performing a global unitary on the whole system, and therefore does not carry over to the multi-prover case (no prover has access to the private spaces of other provers and the private space of each prover might be arbitrarily large, and thus, we cannot use the verifier to transfer those spaces from one prover to any other).

When proving statements that involve the perfect-completeness property, we assume that our universal gate set satisfies some conditions, which may not hold with an arbitrary universal gate set. Specifically, we assume that the Hadamard transformation and any classical reversible transformations are exactly implementable in our gate set. Note that this condition is satisfied by most of the standard gate sets including the Shor basis [107] consisting of the Hadamard gate, the controlled-i-phase-shift gate, and the Toffoli gate, and thus, we believe that this condition is not restrictive. We stress that most of our main statements do hold with an arbitrary choice of universal gate set (the only thing that would change is that the completeness and soundness conditions may become worse by negligible amounts in some of the claims, which does not affect the final main statements).

First, we introduce the notion of *perfectly rewindable* QMIP* systems.

Definition 106. *Let $s < \frac{1}{2}$. A language L has a perfectly rewindable m-turn quantum k-prover interactive proof system with soundness at most s iff there exists an m-turn polynomial-time quantum verifier V, such that, for every input x:*

(Perfect Rewindability) *if $x \in L$, there exists a set of m-turn quantum provers $P_1, \ldots, P_k$ such that $\max_{|\Phi\rangle} p_{\mathrm{acc}}(x, V, P_1, \ldots, P_k, |\Phi\rangle) = \frac{1}{2}$, where the maximum is taken over all a priori shared states $|\Phi\rangle$ prepared by $P_1, \ldots, P_k$.*

(Soundness) *if $x \notin L$, for any set of m-turn quantum provers $P'_1, \ldots, P'_k$ and any a priori shared state $|\Phi'\rangle$, $p_{\mathrm{acc}}(x, V, P'_1, \ldots, P'_k, |\Phi'\rangle) \leq s$.*

Note that in the perfect rewindability property we first fix the provers' transformations and then maximize over all a priori shared states, which hence have a fixed dimension. We first show how to modify any general QMIP* system (with some appropriate conditions on completeness and soundness) to a perfectly rewindable one with the same k and m.

Lemma 107. *Let $c \geq \frac{1}{2} > s$. Then, any language L in* QMIP*(k, m, c, s) *has a perfectly rewindable m-turn quantum k-prover interactive proof system with soundness at most s.*

Proof. Let L be a language in QMIP*(k, m, c, s) and V be the corresponding m-turn quantum verifier. We slightly modify V to construct another m-turn quantum verifier W for a perfectly rewindable proof system for L. The new verifier W, in addition to the registers of V, prepares another single-qubit register B, initialized to $|0\rangle$. For the first $m-2$ turns, W simply simulates V. In the $(m-1)$-st turn, a turn for the verifier, W proceeds like V would, but sends B to the first prover in addition to the qubits V would send in the original proof system. In the m-th turn the first prover is requested to send B back to W, in addition to the qubits sent to V in the original proof system. Then W proceeds for the final decision procedure like V would, but accepts iff V would have accepted *and* B is in the state $|1\rangle$. Notice that W accepts only if V would have accepted. Hence the soundness is obviously at most s in the constructed proof system.

For perfect rewindability we slightly modify the protocol for honest provers in the case $x \in L$. Let $|\Phi^*\rangle$ be the a priori shared state in the original proof system that maximizes the acceptance probability for the original honest provers and let $p_{\max}$ be that maximal acceptance probability. The new provers use $|\Phi^*\rangle$ as the a priori shared state and simulate the original provers except for the last turn. The only difference is that in the last turn the first prover proceeds as P_1 would, *and* applies a one-qubit unitary T to the qubit in B,

$$T : |0\rangle \mapsto \sqrt{1 - \frac{1}{2p_{\max}}}|0\rangle + \sqrt{\frac{1}{2p_{\max}}}|1\rangle.$$

¿From the construction it is obvious that the maximum accepting probability is exactly equal to $\frac{1}{2}$ and that this maximum is achieved when the provers use the a priori shared state $|\Phi^*\rangle$. □

Now, we are ready to show the following lemma.

Lemma 108. *Let $c \geq \frac{1}{2}$ and $s < \frac{1}{25}$. Then,*

$$\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*\Big(k, 3m, 1, \frac{1}{2} + 2\sqrt{s} + \frac{5s}{2}\Big).$$

Proof. The intuitive idea behind the proof of this lemma, using the "quantum rewinding technique" by [126], has already been explained in the introduction. We add some more intuition before proceeding to the technical proof. Using Lemma 107 we can assume that in the case of honest provers (i.e., $x \in L$) the acceptance probability with shared state $|\Phi^*\rangle$ is exactly $\frac{1}{2}$ and furthermore that no other a priori shared state achieves higher acceptance probability. The acceptance probability when the provers use any shared state $|\Phi\rangle$ can be written as $p_{\mathrm{acc}} = \|\Pi_{\mathrm{acc}} Q|\Psi\rangle\|^2 = \|\Pi_{\mathrm{acc}} Q \Pi_{\mathrm{init}}|\Psi\rangle\|^2$, where $|\Psi\rangle = |0\cdots 0\rangle_{(\mathsf{V},\mathsf{M}_1,\ldots,\mathsf{M}_k)} \otimes |\Phi\rangle$, Q is the unitary transformation induced by the QMIP* system just before the verifier's final measurement, Π_{init} is the projection onto states in which all the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$ are in state $|0\rangle$ and Π_{acc} is the projection onto accepting states in the original proof system (i.e., states with the designated output qubit being $|1\rangle$). In other words the state $|\Psi^*\rangle = |0\cdots 0\rangle_{(\mathsf{V},\mathsf{M}_1,\dots,\mathsf{M}_k)} \otimes |\Phi^*\rangle$ maximizes the expression

$$\max_{|\Psi\rangle} \langle\Psi|\Pi_{\rm init} Q^\dagger \Pi_{\rm acc} Q \Pi_{\rm init}|\Psi\rangle,$$

meaning that the matrix $M = \Pi_{\rm init} Q^\dagger \Pi_{\rm acc} Q \Pi_{\rm init}$ has maximum eigenvalue $\frac{1}{2}$ with corresponding eigenvector $|\Psi^*\rangle$. Now we apply the quantum rewinding technique by performing forward, backward, and forward executions of the proof system in sequence. Perfect completeness follows from the fact that the initial state is an eigenvector of M with the corresponding eigenvalue exactly $\frac{1}{2}$, and the proof is similar to that of the zero-knowledge scenario [126].

The challenge of this construction lies in the proof of soundness. If the input is a no-instance, the maximum eigenvalue of any matrix M corresponding to our proof system is small. This shows that if the dishonest provers are actually "not so dishonest", i.e., if they use the same strategies for all of the three (forward, backward, and forward) executions of the original proof system, the acceptance probability is still small. However, the problem arises when the dishonest provers change their strategies for some of the three executions. To settle this, we design a simple protocol that tests if the backward execution is indeed a backward simulation of the first forward execution. The verifier performs the original rewinding protocol or this invertibility test uniformly at random without revealing which test the provers are undergoing. Honest provers always pass this invertibility test, and thus perfect completeness is preserved. When the input is a no-instance, this forces the provers to use approximately the same strategies for the first two executions of the proof system, which is sufficient to bound the soundness. We note that, as is shown by the example at the end of this section, the invertibility test *is* necessary — without it, in some proof systems, the provers can apply a backwards execution that is different from the inverse of their forward execution, and are accepted with certainty even if their maximum acceptance probability in the original proof system was zero.

We now proceed with the technical details of the proof. Let L be a language in QMIP$^*(k, m, c, s)$ and let V be the verifier in the perfectly rewindable m-turn quantum k-prover interactive proof system for L as per Lemma 107. We construct a $3m$-turn quantum verifier W of a new quantum k-prover interactive proof system for L. W has the same registers as V in the original proof system, and performs one of the two tests, which we call "Rewinding Test" and "Invertibility Test". The exact protocol is described in Figure 10.1, where for simplicity it is assumed that m is even (the case in which m is odd can be proved in a similar manner).

Completeness: Assume the input x is in L. From the original provers $P_1, \dots, P_k$ we design honest provers $R_1, \dots, R_k$ for the constructed $3m$-turn system. Each new prover R_i has the same quantum register P_i as P_i has, and the new provers initially share $|\Phi^*\rangle$. For the first m turns each R_i simulates P_i. At the $(m+2j)$-th turn for $1 \leq j \leq \frac{m}{2}$, R_i applies $(P_i^{\frac{m}{2}-j+1})^\dagger$ (i.e., the inverse of P_i's transformation at the $(m-2j+2)$-nd turn in the original system). Finally, for the $(2m+2j)$-th turn for $1 \leq j \leq \frac{m}{2}$, R_i again applies P_i^j.

Verifier's Protocol for Achieving Perfect Completeness

1. Simulate the original verifier for the first m turns.

2. Choose $b \in \{0,1\}$ uniformly at random. If $b = 0$, move to the Rewinding Test described in Step 3, while if $b = 1$, move to the Invertibility Test described in Step 4.

3. (Rewinding Test)

   3.1 Apply $V^{\frac{m}{2}+1}$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if the content of $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ corresponds to an accepting state in the original proof system. Otherwise apply $(V^{\frac{m}{2}+1})^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$, and send M_i to the i-th prover, for $1 \leq i \leq k$.

   3.2 For $j = \frac{m}{2}$ down to 2, do the following:
   Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^j)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$, and send M_i to the i-th prover, for $1 \leq i \leq k$.

   3.3 Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Perform a controlled-phase-flip: multiply the phase by -1 if all the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ are in state $|0\rangle$. Apply V^1 to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$, and send M_i to the i-th prover, for $1 \leq i \leq k$.

   3.4 For $j = 2$ to $\frac{m}{2}$, do the following:
   Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply V^j to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$, and send M_i to the i-th prover, for $1 \leq i \leq k$.

   3.5 Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $V^{\frac{m}{2}+1}$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if the content of $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ corresponds to an accepting state in the original proof system, and reject otherwise.

4. (Invertibility Test)

   4.1 Send M_i to the i-th prover, for $1 \leq i \leq k$.

   4.2 For $j = \frac{m}{2}$ down to 2, do the following:
   Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^j)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$, and send M_i to the i-th prover, for $1 \leq i \leq k$.

   4.3 Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if all the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ are in state $|0\rangle$, and reject otherwise.

Figure 10.1: Verifier's protocol for achieving perfect completeness

It is obvious from this construction that the provers $R_1, \ldots, R_k$ can convince W with certainty when W performs the Invertibility Test. We show that $R_1, \ldots, R_k$ can convince W with certainty even when W performs the Rewinding Test. In short, this holds for essentially the same reason that the quantum rewinding technique works well in the case of quantum zero-knowledge proofs, and we will closely follow that proof.

For notational convenience, let

$$\widetilde{P}^j = P_1^j \otimes \cdots \otimes P_k^j$$

for $1 \le j \le \frac{m}{2}$, and let

$$Q = V^{\frac{m}{2}+1} \widetilde{P}^{\frac{m}{2}} V^{\frac{m}{2}} \cdots \widetilde{P}^1 V^1.$$

Recall that $M|\Psi^*\rangle = \frac{1}{2}|\Psi^*\rangle$ where $M = \Pi_{\mathrm{init}} Q^\dagger \Pi_{\mathrm{acc}} Q \Pi_{\mathrm{init}}$. Define the unnormalized states $|\phi_0\rangle$, $|\phi_1\rangle$, $|\psi_0\rangle$, and $|\psi_1\rangle$ by

$$\begin{aligned} |\phi_0\rangle &= \Pi_{\mathrm{acc}} Q |\Psi^*\rangle, & |\phi_1\rangle &= \Pi_{\mathrm{rej}} Q |\Psi^*\rangle, \\ |\psi_0\rangle &= \Pi_{\mathrm{init}} Q^\dagger |\phi_0\rangle, & |\psi_1\rangle &= \Pi_{\mathrm{illegal}} Q^\dagger |\phi_0\rangle, \end{aligned}$$

where $\Pi_{\mathrm{illegal}} = I_{(\mathsf{V},\mathsf{M}_1,\ldots,\mathsf{M}_k)} - \Pi_{\mathrm{init}}$ is the projection onto states orthogonal to $|0 \cdots 0\rangle_{(\mathsf{V},\mathsf{M}_1,\ldots,\mathsf{M}_k)}$ and $\Pi_{\mathrm{rej}} = I_{(\mathsf{V},\mathsf{M}_1,\ldots,\mathsf{M}_k)} - \Pi_{\mathrm{acc}}$ is the projection onto rejecting states. Then, noticing that $|\Psi^*\rangle = \Pi_{\mathrm{init}} |\Psi^*\rangle$, we have

$$|\psi_0\rangle = \Pi_{\mathrm{init}} Q^\dagger \Pi_{\mathrm{acc}} Q |\Psi^*\rangle = \Pi_{\mathrm{init}} Q^\dagger \Pi_{\mathrm{acc}} Q \Pi_{\mathrm{init}} |\Psi^*\rangle = M|\Psi^*\rangle = \frac{1}{2}|\Psi^*\rangle,$$

and thus,

$$\begin{aligned} Q^\dagger |\phi_1\rangle &= Q^\dagger \Pi_{\mathrm{rej}} Q |\Psi^*\rangle = |\Psi^*\rangle - Q^\dagger \Pi_{\mathrm{acc}} Q |\Psi^*\rangle \\ &= |\Psi^*\rangle - Q^\dagger |\phi_0\rangle = 2|\psi_0\rangle - (|\psi_0\rangle + |\psi_1\rangle) = |\psi_0\rangle - |\psi_1\rangle. \end{aligned}$$

Hence, the state just before the controlled-phase-flip in Step 3.3 when entering the Rewinding Test is exactly

$$\frac{1}{\| |\phi_1\rangle \|} Q^\dagger |\phi_1\rangle = \frac{1}{\| |\phi_1\rangle \|} (|\psi_0\rangle - |\psi_1\rangle).$$

Since $\Pi_{\mathrm{init}} |\psi_0\rangle = |\psi_0\rangle$ and $\Pi_{\mathrm{init}} |\psi_1\rangle = 0$, the controlled-phase-flip changes the state to

$$-\frac{1}{\| |\phi_1\rangle \|} (|\psi_0\rangle + |\psi_1\rangle) = -\frac{1}{\| |\phi_1\rangle \|} Q^\dagger |\phi_0\rangle.$$

Therefore, the state just after $V^{\frac{m}{2}+1}$ is applied in Step 3.5 is exactly

$$-\frac{1}{\| |\phi_1\rangle \|} Q Q^\dagger |\phi_0\rangle = -\frac{1}{\| |\phi_1\rangle \|} |\phi_0\rangle,$$

and thus, the fact that $\Pi_{\mathrm{acc}}|\phi_0\rangle = |\phi_0\rangle$ implies that the verifier W always accepts in Step 3.5.

Soundness: Now suppose that the input x is not in L. Let $R'_1, \dots, R'_k$ be any k provers for the constructed $3m$-turn proof system, and let $|\psi\rangle$ be any a priori shared state. Let R_i^j be the transformation that R'_i applies at his $2j$-th turn, for $1 \le i \le k$ and $1 \le j \le \frac{3m}{2}$ and let Z denote the controlled-phase-flip operator in Step 3.3. Let

$$\widetilde{R}^j = R_1^j \otimes \cdots \otimes R_k^j$$

for $1 \le j \le \frac{3m}{2}$, and define

$$\begin{aligned}
U_1 &= \widetilde{R}^{\frac{m}{2}} V^{\frac{m}{2}} \cdots \widetilde{R}^2 V^2 \widetilde{R}^1 V^1,\\
U_2 &= (V^1)^\dagger \widetilde{R}^m \cdots (V^{\frac{m}{2}-1})^\dagger \widetilde{R}^{\frac{m}{2}+2} (V^{\frac{m}{2}})^\dagger \widetilde{R}^{\frac{m}{2}+1},\\
U_3 &= \widetilde{R}^{\frac{3m}{2}} V^{\frac{m}{2}} \cdots \widetilde{R}^{m+2} V^2 \widetilde{R}^{m+1} V^1.
\end{aligned}$$

There are three cases of acceptance in the constructed proof system. In the first case, the verifier W performs the Rewinding Test and accepts in Step 3.1. This happens with probability $\frac{p_1}{2}$, where

$$p_1 = \| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_1 |\psi\rangle \|^2.$$

In the second case, the verifier W performs the Rewinding Test and accepts in Step 3.5. This happens with probability $\frac{p_2}{2}$, where

$$p_2 = \| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle \|^2.$$

Finally, in the third case, the verifier W performs the Invertibility Test and accepts in Step 4.3. This happens with probability $\frac{p_3}{2}$, where

$$p_3 = \| \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle \|^2.$$

Hence, the total probability p_{acc} that W accepts x when communicating with $R'_1, \dots, R'_k$ is given by $p_{\mathrm{acc}} = \frac{1}{2}(p_1 + p_2 + p_3)$. From the soundness condition of the original proof system, it is obvious that $p_1 \le s$. We shall show that $p_2 \le 1 + 4\sqrt{s} + 4s - p_3$. This implies that $p_{\mathrm{acc}} \le \frac{1}{2} + 2\sqrt{s} + \frac{5s}{2}$, and the soundness condition follows.

Using the triangle inequality, we have that

$$\begin{aligned}
&\| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle \| \\
&\quad \le \| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle \| \\
&\qquad + \| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle \| \\
&\qquad + \| \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle \|.
\end{aligned} \tag{10.1}$$

The first term of the right-hand side of inequality (10.1) can be bounded from above as follows:

$$\begin{aligned}
&\|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle\| \\
&\quad \le \|V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle - V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle\| \\
&\quad = \|(V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle - U_1 |\psi\rangle\| \\
&\quad = \|\Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle - V^{\frac{m}{2}+1} U_1 |\psi\rangle\| \\
&\quad = \| - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_1 |\psi\rangle\| = \|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_1 |\psi\rangle\| = \sqrt{p_1} \le \sqrt{s}.
\end{aligned}$$

The second term of the right-hand side of inequality (10.1) can be bounded from above as follows:

$$\begin{aligned}
&\|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| \\
&\quad \le \|V^{\frac{m}{2}+1} U_3 Z U_2 U_1 |\psi\rangle - V^{\frac{m}{2}+1} U_3 Z \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| \\
&\quad = \|U_2 U_1 |\psi\rangle - \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| = \|\Pi_{\mathrm{illegal}} U_2 U_1 |\psi\rangle\| = \sqrt{1 - p_3}.
\end{aligned}$$

Here the last equality follows from the facts that $U_2 U_1 |\psi\rangle = \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle + \Pi_{\mathrm{illegal}} U_2 U_1 |\psi\rangle$ is a unit vector, that $\Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle$ and $\Pi_{\mathrm{illegal}} U_2 U_1 |\psi\rangle$ are orthogonal, and that $\|\Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\|^2 = p_3$.

Finally, since $\Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle$ is an unnormalized state parallel to some legal initial state and $Z\Pi_{\mathrm{init}} = -\Pi_{\mathrm{init}}$ from the definitions of Z and Π_{init}, the third term of the right-hand side of inequality (10.1) can be bounded as follows by using the soundness condition of the original proof system:

$$\begin{aligned}
\|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| &= \| - \Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| \\
&= \|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 \Pi_{\mathrm{init}} U_2 U_1 |\psi\rangle\| \le \sqrt{s}.
\end{aligned}$$

Putting everything together, we have

$$\begin{aligned}
p_2 &= \|\Pi_{\mathrm{acc}} V^{\frac{m}{2}+1} U_3 Z U_2 (V^{\frac{m}{2}+1})^\dagger \Pi_{\mathrm{rej}} V^{\frac{m}{2}+1} U_1 |\psi\rangle\|^2 \\
&\le \left(2\sqrt{s} + \sqrt{1 - p_3}\right)^2 = 1 + 4\sqrt{s(1 - p_3)} + 4s - p_3 \le 1 + 4\sqrt{s} + 4s - p_3,
\end{aligned}$$

as desired. □

To finish the proof of Theorem 102 it suffices to repeat sequentially the proof system obtained in Lemma 108 an appropriate number of times (and accept if and only if all the original verifiers would have accepted every time). To see that this reduces soundness exponentially with the number of repetitions, imagine by contradiction that there exists a set of provers that succeeds in the k-th repetition of the proof system with probability s' greater than the soundness s in the original proof system. Then we can construct provers for the original proof system by letting them initially share the state of the provers at the end of the $(k - 1)$-st repetition of the new proof system. These provers would be accepted with probability $s' > s$, a contradiction.

Necessity of the invertibility test. Consider the following single-prover two-turn quantum interactive proof system. The verifier initially has a state $|000\rangle$ of three qubits, where the first two qubits are in his private space and the last one is the message register. He sends the message register to the prover, receives it back, and then applies the transformation U which maps $|001\rangle \mapsto |010\rangle$, $|011\rangle \mapsto |110\rangle$, and conversely $|010\rangle \mapsto |001\rangle$, $|110\rangle \mapsto |011\rangle$, and leaves all the other basis states unchanged. After his transformation, the verifier again sends the message register to the prover, receives it back, and finally accepts if anf only if the final state of the system is $|111\rangle$.

It is easy to see that the prover's maximum winning probability in this protocol is zero. However, we can design a prover that passes with certainty the Rewinding Test induced from this protocol, by having him apply the transformation $|0\rangle \mapsto |1\rangle$ in both of his actions during the first forward phase, applying the identity and then $|0\rangle \mapsto |1\rangle$ in the backward phase, and applying twice the identity in the last forward phase. It is easy to check that this prover succeeds in the Rewinding Test with certainty, and thus, we cannot achieve perfect completeness with the Rewinding Test only. Note that this prover fails the Invertibility Test with certainty, and the Invertibility Test is indeed helpful in this case.

10.3 Parallelizing to Three Turns

In this section we prove Theorem 103, which reduces the number of turns to three without changing the number of provers. This is done by repeatedly converting any $(2^l + 1)$-turn QMIP* system into a $(2^{l-1} + 1)$-turn QMIP* system where the gap decreases, but is still bounded by an inverse polynomial. We first show the following lemma.

Lemma 109. *For any c, s satisfying $c^2 > s$,*

$$\mathrm{QMIP}^*(k, 4m+1, c, s) \subseteq \mathrm{QMIP}^*\left(k, 2m+1, \frac{1+c}{2}, \frac{1+\sqrt{s}}{2}\right).$$

Proof. Let L be a language in $\mathrm{QMIP}^*(k, 4m+1, c, s)$ and let V be the corresponding $(4m+1)$-turn quantum verifier. We construct a $(2m+1)$-turn quantum verifier W for the new quantum k-prover interactive proof system for L. The idea is that W first receives the snapshot state that V would have in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$ just after the $(2m+1)$-st turn of the original system. W then executes with equal probability either a forward-simulation of the original system from the $(2m+1)$-st turn or a backward-simulation of the original system from the $(2m+1)$-st turn. In the former case W accepts if and only if the simulation results in acceptance in the original proof system, while in the latter case W accepts if and only if all the qubits in V are in state $|0\rangle$ (recall that in the original proof system the first turn was done by the provers, hence we do not measure the qubits in each M_i here). The details are given in Figure 10.2.

Verifier's Protocol to Half the Number of Turns

1. Receive V and M_1 from the first prover and M_i from the i-th prover for $2 \leq i \leq k$.
2. Choose $b \in \{0, 1\}$ uniformly at random.
3. If $b = 0$, execute a forward-simulation of the original proof system as follows:

   3.1 Apply V^{m+1} to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Send b and M_i to the i-th prover, for $1 \leq i \leq k$.

   3.2 For $j = m + 2$ to $2m$, do the following:
   Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply V^j to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Send M_i to the i-th prover, for $1 \leq i \leq k$.

   3.3 Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply V^{2m+1} to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if the content of $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ is an accepting state of the original proof system, and reject otherwise.

4. If $b = 1$, execute a backward-simulation of the original proof system as follows:

   4.1 Send b and M_i to the i-th prover, for $1 \leq i \leq k$.

   4.2 For $j = m$ down to 2, do the following:
   Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^j)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Send M_i to the i-th prover, for $1 \leq i \leq k$.

   4.3 Receive M_i from the i-th prover, for $1 \leq i \leq k$. Apply $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if all the qubits in V are in state $|0\rangle$, and reject otherwise.

Figure 10.2: Verifier's protocol to reduce the number of turns by half.

Completeness: Assume the input x is in L. Let $P_1, \dots, P_k$ be the honest quantum provers in the original proof system with a priori shared state $|\Phi\rangle$. Let $|\psi_{2m+1}\rangle$ be the quantum state in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k, \mathsf{P}_1, \dots, \mathsf{P}_k)$ just after the $(2m + 1)$-st turn in the original proof system. We construct honest provers $R_1, \dots, R_k$ for the new $(2m + 1)$-turn system. In addition to V and M_1, R_1 prepares P_1 in his private space. Similarly, in addition to M_i, R_i prepares P_i in his private space for $2 \leq i \leq k$. $R_1, \dots, R_k$ initially share $|\psi_{2m+1}\rangle$ in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k, \mathsf{P}_1, \dots, \mathsf{P}_k)$. At the first turn of the constructed proof system, R_1 sends V and M_1 to W, while each R_i, for $2 \leq i \leq k$, sends M_i to W. At the $(2j - 1)$-st turn for $2 \leq j \leq m + 1$, if $b = 0$, each R_i applies P_i^{m+j} (i.e., P_i's transformation at the $(2m + 2j - 1)$-st turn in the original system) while if $b = 1$, each R_i applies $(P_i^{m-j+3})^\dagger$ (i.e., the inverse of P_i's transformation at the $(2m - 2j + 5)$-th turn in the original system) to the qubits in $(\mathsf{P}_i, \mathsf{M}_i)$, for $1 \leq i \leq k$. The provers $R_1, \dots, R_k$ can then clearly convince W with probability at least c if $b = 0$, and with certainty if $b = 1$. Hence, W accepts every input $x \in L$ with probability at least $\frac{1+c}{2}$.

Soundness: Now suppose that x is not in L. Let $R'_1, \ldots, R'_k$ be arbitrary provers for the constructed proof system, and let $|\psi\rangle$ be an arbitrary quantum state that represents the state just after the first turn in the constructed system. Suppose that, at the $(2j-1)$-st turn for $2 \leq j \leq m+1$, each R'_i applies X_i^j if $b=0$ and Y_i^j if $b=1$, for $1 \leq i \leq k$ and write

$$\widetilde{X}^j = X_1^j \otimes \cdots \otimes X_k^j, \qquad \widetilde{Y}^j = Y_1^j \otimes \cdots \otimes Y_k^j.$$

Define unitary transformations U_0 and U_1 by

$$\begin{aligned} U_0 &= V^{2m+1}\widetilde{X}^{m+1}V^{2m} \cdots \widetilde{X}^2 V^{m+1}, \\ U_1 &= (V^1)^\dagger \widetilde{Y}^{m+1} \cdots (V^m)^\dagger \widetilde{Y}^2, \end{aligned}$$

and let

$$|\alpha\rangle = \frac{1}{\|\Pi_{\mathrm{acc}} U_0 |\psi\rangle\|} \Pi_{\mathrm{acc}} U_0 |\psi\rangle, \qquad |\beta\rangle = \frac{1}{\|\Pi_{\mathrm{init}} U_1 |\psi\rangle\|} \Pi_{\mathrm{init}} U_1 |\psi\rangle,$$

where Π_{acc} is the projection onto accepting states in the original proof system and Π_{init} is the projection onto states in which all the qubits in V are in state $|0\rangle$. Then

$$\begin{aligned} \|\Pi_{\mathrm{acc}} U_0 |\psi\rangle\| &= \frac{1}{\|\Pi_{\mathrm{acc}} U_0 |\psi\rangle\|} \bigl|\langle\psi| U_0^\dagger \Pi_{\mathrm{acc}} U_0 |\psi\rangle\bigr| \\ &= F\bigl(|\alpha\rangle\langle\alpha|, U_0|\psi\rangle\langle\psi|U_0^\dagger\bigr) = F\bigl(U_0^\dagger|\alpha\rangle\langle\alpha|U_0, |\psi\rangle\langle\psi|\bigr), \end{aligned}$$

and thus, the probability p_0 of acceptance when $b=0$ is given by

$$p_0 = F\bigl(U_0^\dagger|\alpha\rangle\langle\alpha|U_0, |\psi\rangle\langle\psi|\bigr)^2.$$

Similarly, the probability p_1 of acceptance when $b=1$ is given by

$$p_1 = F\bigl(U_1^\dagger|\beta\rangle\langle\beta|U_1, |\psi\rangle\langle\psi|\bigr)^2.$$

Hence the probability p_{acc} of acceptance when W communicates with $R'_1, \ldots, R'_k$ is given by

$$p_{\mathrm{acc}} = \frac{1}{2}(p_0 + p_1) = \frac{1}{2}\Bigl(F\bigl(U_0^\dagger|\alpha\rangle\langle\alpha|U_0, |\psi\rangle\langle\psi|\bigr)^2 + F\bigl(U_1^\dagger|\beta\rangle\langle\beta|U_1, |\psi\rangle\langle\psi|\bigr)^2\Bigr).$$

Therefore, from Lemma 7, we have

$$\begin{aligned} p_{\mathrm{acc}} &\leq \frac{1}{2}\left(1 + F\bigl(U_0^\dagger|\alpha\rangle\langle\alpha|U_0, U_1^\dagger|\beta\rangle\langle\beta|U_1\bigr)\right) \\ &= \frac{1}{2}\left(1 + F\bigl(|\alpha\rangle\langle\alpha|, U_0U_1^\dagger|\beta\rangle\langle\beta|U_1U_0^\dagger\bigr)\right). \end{aligned}$$

Note that $\Pi_{\mathrm{init}}|\beta\rangle = |\beta\rangle$ and thus $|\beta\rangle$ is a legal quantum state which could appear in the original proof system just after the first turn. Hence, from the soundness property of the original proof system,

$$\begin{aligned} &\bigl\|\Pi_{\mathrm{acc}} U_0 U_1^\dagger |\beta\rangle\bigr\|^2 \\ &\quad = \bigl\|\Pi_{\mathrm{acc}} V^{2m+1}\widetilde{X}^{m+1}V^{2m} \cdots \widetilde{X}^2 V^{m+1} (\widetilde{Y}^2)^\dagger V^m \cdots (\widetilde{Y}^{m+1})^\dagger V^1 |\beta\rangle\bigr\|^2 \leq s, \end{aligned}$$

since V^1, $(\widetilde{Y}^{m+1})^\dagger, \cdots, V^m$, $(\widetilde{Y}^2)^\dagger$, V^{m+1}, $\widetilde{X}^2, \cdots, V^{2m}$, $\widetilde{X}^{m+1}, V^{2m+1}$ form a legal sequence of transformations in the original proof system.

Now, from the fact that $\Pi_{\rm acc}|\alpha\rangle = |\alpha\rangle$, we have

$$\begin{aligned} F\big(|\alpha\rangle\langle\alpha|, U_0U_1^\dagger|\beta\rangle\langle\beta|U_1U_0^\dagger\big) &= \big|\langle\alpha|U_0U_1^\dagger|\beta\rangle\big| = \big|\langle\alpha|\Pi_{\rm acc}U_0U_1^\dagger|\beta\rangle\big| \\ &\leq \|\Pi_{\rm acc}U_0U_1^\dagger|\beta\rangle\| \leq \sqrt{s}. \end{aligned}$$

Hence the probability $p_{\rm acc}$ that W accepts x is bounded by $p_{\rm acc} \leq \frac{1}{2} + \frac{\sqrt{s}}{2}$, which completes the proof. □

Now, by repeatedly applying the construction in the proof of Lemma 109, we can reduce the number of turns to three. The proof is straightforward, but we need to carefully keep track of the efficiency of the constructed verifiers in each application, since the construction is sequentially applied a logarithmic number of times.

Lemma 110. *For any $m \geq 4$ and any c, s such that $\varepsilon = 1 - c$ and $\delta = 1 - s$ satisfy $\delta > 2(m-1)\varepsilon$,*

$$\mathrm{QMIP}^*(k, m, 1-\varepsilon, 1-\delta) \subseteq \mathrm{QMIP}^*\left(k, 3, 1 - \frac{2\varepsilon}{m-1}, 1 - \frac{\delta}{(m-1)^2}\right).$$

Proof. Let l be such that $2^l + 1 \leq m \leq 2^{l+1} + 1$. Trivially, the inclusion $\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*(k, 2$ holds, and we show the inclusion $\mathrm{QMIP}^*(k, 2^{l+1} + 1, 1 - \varepsilon, 1 - \delta) \subseteq \mathrm{QMIP}(k, 3, 1 - \frac{2\varepsilon}{m-1}, 1 - \frac{\delta}{(m-1)^2})$.

Let L be a language in $\mathrm{QMIP}^*(k, 2^{l+1} + 1, 1 - \varepsilon, 1 - \delta)$ and let $V^{(0)}$ be the corresponding $(2^{l+1} + 1)$-turn quantum verifier. Given a description of $V^{(0)}$ one can compute in polynomial time a description of a $(2^l + 1)$-turn quantum verifier $V^{(1)}$ following the proof of Lemma 109. The resulting proof system has completeness at least $1 - \frac{\varepsilon}{2}$ and soundness at most $\frac{1}{2} + \frac{\sqrt{1-\delta}}{2} \leq 1 - \frac{\delta}{4}$. Crucially, the description of $V^{(1)}$ is at most some constant times the size of the description of $V^{(0)}$ plus an amount bounded by a polynomial in the input length. Hence it is obvious that, given a description of $V^{(0)}$, one can compute in polynomial time a description of a three-turn quantum verifier $V^{(l)}$ by repeatedly applying the construction in the proof of Lemma 109 l times. The resulting proof system has completeness at least $1 - \frac{\varepsilon}{2^l} \geq 1 - \frac{2\varepsilon}{m-1}$ and soundness at most $1 - \frac{\delta}{4^l} \leq 1 - \frac{\delta}{(m-1)^2}$, as desired. □

Theorem 103 now follows immediately from Theorem 102 and Theorem 110: For every $p \in$ poly there exist a function $m' \in$ poly such that the inclusions $\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*(k, m', 1, 2^{-p})$ $\subseteq \mathrm{QMIP}^*\big(k, 3, 1, 1 - \frac{1-2^{-p}}{(m'-1)^2}\big)$ hold. Now it suffices to observe that $\frac{1-2^{-p}}{(m'-1)^2} \in \mathrm{poly}^{-1}$.

10.4 Public-Coin Systems

In this section we present the last part to complete the proof of Theorem 100. We show how any three-turn QMIP* system with sufficiently large gap can be converted into a two-turn QMIP system with one extra prover, in which the gap is bounded by an inverse-polynomial. Although there is also a direct proof for this, given in Section 10.4.3, we first give another proof that takes a detour by showing how (i) any three-turn QMIP* system with sufficiently large gap can be modified to a three-turn *public-coin* QMIP* system with inverse-polynomial gap without changing the number of provers, and (ii) any three-turn public-coin QMIP* system can be converted into a two-turn QMIP* system without changing completeness and soundness, by adding an extra prover. By parallelizing to two turns we lose the public-coin property, and in Theorem 112 we show that this is unavoidable unless BP · PP = PSPACE: any two-turn public-coin proof system with a polynomial number of provers can be converted to a *single-prover* two-turn public-coin proof system (i.e., a quantum Arthur-Merlin proof system).

The added benefits of our detour are a proof of the equivalence of public-coin QMIP* systems and general QMIP* systems (Theorem 104) and a proof that QIP and hence PSPACE has two-prover one-round quantum interactive proof systems of perfect completeness and exponentially small soundness (Corollary 105).

Remark. The direct proof in Section 10.4.3 would only give the weaker corollary that QIP has a two-prover one-round quantum interactive proof system of perfect completeness, but with soundness only exponentially close to $\frac{1}{2}$. This is indeed weaker than what we can show with the detour, since it is not known how to amplify the success probability of QMIP* systems without increasing either the number of provers or the number of turns.

10.4.1 Converting to Public-Coin Systems

In this subsection we prove Theorem 104 showing that any language that has a quantum k-prover interactive proof system with two-sided bounded error also has a *public-coin* quantum k-prover interactive proof system of perfect completeness and exponentially small soundness.

We first show that any three-turn QMIP* system with sufficiently large gap can be modified to a three-turn public-coin QMIP* system with the same number of provers and inverse-polynomial gap. In the single-prover case, [81] proved a similar statement. Our proof is a generalization of their proof (Theorem 5.4 in [81]) to the multi-prover case.

Lemma 111. *For any c, s satisfying $c^2 > s$,*

$$\mathrm{QMIP}^*(k, 3, c, s) \subseteq \mathrm{QMIP}^*_{\mathrm{pub}}\left(k, 3, \frac{1+c}{2}, \frac{1+\sqrt{s}}{2}\right).$$

Moreover, the message from the verifier to each prover in the public-coin system consists of only one classical bit.

Verifier's Protocol in the Three-Turn Public-Coin System

1. Receive V from the first prover and receive nothing from the i-th prover, for $2 \leq i \leq k$.
2. Choose $b \in \{0, 1\}$ uniformly at random. Send b to each prover.
3. Receive M_i from the i-th prover for $1 \leq i \leq k$.

   3.1 If $b = 0$, apply V^2 to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if the content of $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ is an accepting state of the original proof system, and reject otherwise.

   3.2 If $b = 1$, apply $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$. Accept if all the qubits in V are in state $|0\rangle$, and reject otherwise.

Figure 10.3: Verifier's protocol in the three-turn public-coin system.

Proof. Let L be a language in $\mathrm{QMIP}^*(k, 3, c, s)$ and let V be the corresponding three-turn quantum verifier. We construct a new verifier W for the public-coin system. The idea is that in the first turn W receives the reduced state in the register V (corresponding to the private space of the original verifier) of the snapshot state just after the second turn (i.e., just after the first transformation of V) in the original proof system. W then flips a fair classical coin $b \in \{0, 1\}$ and broadcasts b to the provers. At the third turn the i-th prover is requested to send the message register M_i of the original proof system, for $1 \leq i \leq k$. If $b = 0$ the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ should form the quantum state the original verifier V would possess just after the third turn of the original proof system. Now W applies V^2 to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ and accepts if and only if the content of $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ is an accepting state of the original proof system. On the other hand, if $b = 1$, the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ should form the quantum state the original verifier V would possess just after the second turn of the original proof system. Now W applies $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k)$ and accepts if and only if all the qubits in V are in state $|0\rangle$. The detailed description of the protocol of W is given in Figure 10.3.

Completeness: Assume the input x is in L. Let $P_1, \dots, P_k$ be the honest quantum provers in the original proof system with a priori shared state $|\Phi\rangle$ in $(\mathsf{P}_1, \dots, \mathsf{P}_k)$. Let $|\psi_2\rangle$ be the quantum state in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k, \mathsf{P}_1, \dots, \mathsf{P}_k)$ just after the second turn in the original proof system. We construct honest provers $R_1, \dots, R_k$ for the public-coin system. In addition to V and M_1, R_1 prepares P_1 in his private space. Similarly, in addition to M_i, R_i prepares P_i in his private space, for $2 \leq i \leq k$. $R_1, \dots, R_k$ initially share $|\psi_2\rangle$ in $(\mathsf{V}, \mathsf{M}_1, \dots, \mathsf{M}_k, \mathsf{P}_1, \dots, \mathsf{P}_k)$. At the first turn of the constructed proof system, R_1 sends V to W, while each R_i, $2 \leq i \leq k$ send nothing to W. At the third turn, if $b = 0$ each R_i applies P_i^2 to the qubits in $(\mathsf{M}_i, \mathsf{P}_i)$ and then sends M_i to W, while if $b = 1$, each R_i does nothing and sends M_i to W. It is obvious that the provers $R_1, \dots, R_k$ can convince W with probability at least c if $b = 0$, and with certainty if $b = 1$. Hence, W accepts every input $x \in L$ with probability at least $\frac{1+c}{2}$.

Soundness: Now suppose that x is not in L. Let $R'_1, \dots, R'_k$ be arbitrary provers for the constructed proof system, and let $|\psi\rangle$ be an arbitrary quantum state that represents the state just after the first turn in the constructed system. Suppose that at the third turn each R'_i applies X_i if $b = 0$ and Y_i if $b = 1$, for $1 \leq i \leq k$ and write

$$\widetilde{X} = X_1 \otimes \cdots \otimes X_k, \qquad \widetilde{Y} = Y_1 \otimes \cdots \otimes Y_k.$$

Note that $\widetilde{X}$ and $\widetilde{Y}$ are unitary transformations that do not act over the qubits in V. Let

$$|\alpha\rangle = \frac{1}{\|\Pi_{\mathrm{acc}} V^2 \widetilde{X} |\psi\rangle\|} \Pi_{\mathrm{acc}} V^2 \widetilde{X} |\psi\rangle, \qquad |\beta\rangle = \frac{1}{\|\Pi_{\mathrm{init}} (V^1)^\dagger \widetilde{Y} |\psi\rangle\|} \Pi_{\mathrm{init}} (V^1)^\dagger \widetilde{Y} |\psi\rangle,$$

where Π_{acc} is the projection onto accepting states in the original proof system and Π_{init} is the projection onto states in which all the qubits in V are in state $|0\rangle$.

Then, with a similar argument to that in the proof of Lemma 109, the probability p_{acc} that W accepts x when communicating with $R'_1, \dots, R'_{k+1}$ is bounded by

$$\begin{aligned} p_{\mathrm{acc}} &\leq \frac{1}{2}\Big(1 + F\big(\widetilde{X}^\dagger (V^2)^\dagger |\alpha\rangle\langle\alpha| V^2 \widetilde{X}, \widetilde{Y}^\dagger V^1 |\beta\rangle\langle\beta| (V^1)^\dagger \widetilde{Y}\big)\Big) \\ &= \frac{1}{2}\Big(1 + F\big(|\alpha\rangle\langle\alpha|, V^2 \widetilde{X}\widetilde{Y}^\dagger V^1 |\beta\rangle\langle\beta| (V^1)^\dagger \widetilde{Y}\widetilde{X}^\dagger (V^2)^\dagger\big)\Big). \end{aligned}$$

Since $\Pi_{\mathrm{init}}|\beta\rangle = |\beta\rangle$ is a legal quantum state which could appear just after the first turn in the original proof system, V^1, $\big(\widetilde{X}\widetilde{Y}^\dagger\big), V^2$ form a legal sequence of transformations in the original proof system, and $\Pi_{\mathrm{acc}}|\alpha\rangle = |\alpha\rangle$, again a similar argument to that in the proof of Lemma 109 shows that

$$F\big(|\alpha\rangle\langle\alpha|, V^2 \widetilde{X}\widetilde{Y}^\dagger V^1 |\beta\rangle\langle\beta| (V^1)^\dagger \widetilde{Y}\widetilde{X}^\dagger (V^2)^\dagger\big) \leq \sqrt{s}.$$

Hence the probability p_{acc} that W accepts x is bounded by $p_{\mathrm{acc}} \leq \frac{1}{2} + \frac{\sqrt{s}}{2}$, as desired. □

Theorem 104 now follows directly from Theorem 103 and Theorem 111 together with sequential repetition: Theorem 103 and Theorem 111 imply that there exists a function $p' \in$ poly such that $\mathrm{QMIP}^*(k, m, c, s) \subseteq \mathrm{QMIP}^*\big(k, 3, 1, 1 - \frac{1}{p'}\big) \subseteq \mathrm{QMIP}^*_{\mathrm{pub}}\big(k, 3, 1, 1 - \frac{1}{4p'}\big)$, since $\frac{1}{2}\left(1 + \sqrt{1 - \frac{1}{p'}}\right) \leq 1 - \frac{1}{4p'}$. Finally, sequential repetition gives that for all $p \in$ poly there exists a function $m' \in$ poly such that $\mathrm{QMIP}^*_{\mathrm{pub}}\big(k, 3, 1, 1 - \frac{1}{4p'}\big) \subseteq \mathrm{QMIP}^*_{\mathrm{pub}}(k, m', 1, 2^{-p})$.

We end this section by proving the following theorem, which shows that the parallelization to three turns in Lemma 111 is optimal when considering public-coin systems: Theorem 112 below implies that it cannot be brought down to two turns unless BP · PP = PSPACE, which would imply a collapse of the counting hierarchy to the second level, since BP·PP $\subseteq$ BPP$^{\mathrm{PP}}$ $\subseteq$ PP$^{\mathrm{PP}}$. Here BP·PP is the complexity class obtained by applying the BP-operator to the class

PP. Indeed, on one hand the inclusion QAM $\subseteq$ BP $\cdot$ PP holds for the class QAM of languages having two-turn public-coin quantum single-prover interactive proof systems [81], and on the other hand the inclusions PSPACE $\subseteq$ QIP $\subseteq$ QMIP are obvious. We write QAM(c, s) to specify completeness c and soundness s.

Theorem 112. *For any $k \in$ poly and for any c, s,*

$$\mathrm{QMIP}^*_{\mathrm{pub}}(k, 2, c, s) = \mathrm{QMIP}^*_{\mathrm{pub}}(1, 2, c, s) = \mathrm{QAM}(c, s).$$

Proof (sketch). The inclusion QAM$(c, s) \subseteq \mathrm{QMIP}^*_{\mathrm{pub}}(k, 2, c, s)$ is clear. To show that $\mathrm{QMIP}^*_{\mathrm{pub}}(k, 2, c, s) \subseteq \mathrm{QAM}(c, s)$, we transform a $\mathrm{QMIP}^*_{\mathrm{pub}}$ protocol into a QAM protocol in the most straightforward manner: the verifier receives all the k proofs from the single prover, after having sent him the results of public coin-flips. Completeness is obvious, and we only need to check for soundness. Suppose the prover in the QAM system answers ρ_r when asked r, a string of n random bits. Then the k provers in the $\mathrm{QMIP}^*_{\mathrm{pub}}$ system could share the 2^n possible quantum states $\rho_1, \ldots, \rho_{2^n}$ among themselves before the protocol starts, and could simply answer ρ_r to the question r (which they all received). These provers are accepted with the same probability as the original QAM prover, and the claim follows. □

10.4.2 Parallelizing to Two Turns

Finally, we prove the last piece of Theorem 100 by showing that any three-turn public-coin quantum k-prover interactive proof system can be converted into a two-turn (i.e., one-round) $(k + 1)$-prover system without changing completeness and soundness. The idea of the proof is to send questions only to the first k provers to request the original second messages from the k provers in the original system and to receive from the $(k + 1)$-st prover the original first messages of the k provers in the original system without asking him any question.

Lemma 113. *For any $k \in$ poly and for any c, s,*

$$\mathrm{QMIP}^*_{\mathrm{pub}}(k, 3, c, s) \subseteq \mathrm{QMIP}^*(k + 1, 2, c, s).$$

Proof. Let L be a language in $\mathrm{QMIP}^*_{\mathrm{pub}}(k, 3, c, s)$ and let V be the corresponding verifier.

The original three-turn system can be viewed as follows: At the first turn, V first receives a quantum register M_i from the ith prover, for each $1 \leq i \leq k$. Then V flips a fair classical coin q_{M} times to generate a random string r of length q_{M}, and broadcasts r to all the provers. V also stores r in a quantum register Q in his private space. Finally, at the third turn, V receives a quantum register N_i from the i-th prover, for each $1 \leq i \leq k$. V then prepares a quantum register V for his work space, where all the qubits in V are initialized to state $|0\rangle$, applies the transformation V^{final} to the qubits in $(\mathsf{Q}, \mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{N}_1, \ldots, \mathsf{N}_k)$, and performs

Verifier's Protocol in the One-Round System

1. Prepare a quantum register V, and initialize all the qubits in V to state $|0\rangle$. Flip a fair classical coin q_{M} times to generate a random string r of length q_{M}. Store r in a quantum register Q, and send r to the i-th prover for $1 \leq i \leq k$. Send nothing to the $(k+1)$-st prover.
2. Receive a quantum register N_i from the i-th prover, for $1 \leq i \leq k$, and k quantum registers $\mathsf{M}_1, \ldots, \mathsf{M}_k$ from the $(k+1)$-st prover. Apply V^{final} to the qubits in $(\mathsf{Q}, \mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{N}_1, \ldots, \mathsf{N}_k)$ and accept if and only if the content of $(\mathsf{Q}, \mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{N}_1, \ldots, \mathsf{N}_k)$ is an accepting state of the original proof system.

Figure 10.4: Verifier's protocol to reduce the number of turns to two.

the measurement $\Pi = \{\Pi_{\text{acc}}, \Pi_{\text{rej}}\}$ to decide acceptance or rejection. We construct a two-turn quantum verifier W for the new quantum $(k+1)$-prover interactive proof system for L.

The constructed prover W starts with generating a random string r of length q_{M} in the first turn, and sends r to the first k provers. To the last prover, W does not send any question. In the second turn W receives N_i from the i-th prover expecting the original second message from the original i-th prover, for $1 \leq i \leq k$. From the $(k+1)$-st prover W receives k quantum registers $\mathsf{M}_1, \ldots, \mathsf{M}_k$, expecting the original first messages of the original k provers. Then W proceeds like V would. A detailed description of the protocol of W is given in Figure 10.4.

Completeness: Assume the input x is in L. Let $P_1, \ldots, P_k$ be the honest provers in the original proof system. Let $|\psi_1\rangle$ be the quantum state in $(\mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{P}_1, \ldots, \mathsf{P}_k)$ in the original proof system just after the first turn. We construct honest provers $R_1, \ldots, R_{k+1}$ for the two-turn system. For $1 \leq i \leq k$, R_i prepares quantum register P_i in his private space, where some of the qubits in P_i form the quantum register N_i, while R_{k+1} prepares the quantum registers $\mathsf{M}_1, \ldots, \mathsf{M}_k$ in his private space. Initially, $R_1, \ldots, R_{k+1}$ share $|\psi_1\rangle$ in $(\mathsf{M}_1, \ldots, \mathsf{M}_k, \mathsf{P}_1, \ldots, \mathsf{P}_k)$. At the second turn, R_{k+1} just sends the qubits in $(\mathsf{M}_1, \ldots, \mathsf{M}_k)$ to W, while each R_i, after receiving r, just behaves like P_i would at the third turn of the original system, and then sends N_i, which is a part of P_i, to W, for $1 \leq i \leq k$. It is obvious from the construction that the provers $R_1, \ldots, R_{k+1}$ can convince W with the same probability with which $P_1, \ldots, P_k$ could convince V, which is at least c.

Soundness: Now assume the input x is not in L. Let $R'_1, \ldots, R'_{k+1}$ be any provers for the constructed proof system and let R'_i be the quantum register consisting of all the qubits in the private space of R'_i, for $1 \leq i \leq k+1$. For R'_{k+1}, some of the qubits in R'_{k+1} form the register $\mathsf{M} = (\mathsf{M}_1, \ldots, \mathsf{M}_k)$. Let $|\psi\rangle$ be an arbitrary quantum state in $(\mathsf{R}'_1, \ldots, \mathsf{R}'_{k+1})$ that is initially shared by $R'_1, \ldots, R'_{k+1}$. Suppose that, at the second turn, each R'_i applies $X_i^{(r)}$, for $1 \leq i \leq k$, if the message from W is r. Without loss of generality, we assume that R'_{k+1} does nothing, and just sends the qubits in $(\mathsf{M}_1, \dots, \mathsf{M}_k)$ at the second turn, since R'_{k+1} receives nothing from W (that R'_{k+1} applies some transformation Z is equivalent to sharing $Z|\psi\rangle$ at the beginning).

Consider three-turn quantum provers $P'_1, \dots, P'_k$ for the original proof system with the following properties: (1) each P'_i prepares the quantum register R'_i in his private space, for $1 \leq i \leq k$, (2) $P'_1, \dots, P'_k$ initially share $|\psi\rangle$ in $(\mathsf{R}'_1, \dots, \mathsf{R}'_{k+1})$, where all the qubits in R'_{k+1} except for those in $(\mathsf{M}_1, \dots, \mathsf{M}_k)$ are shared arbitrarily, (3) at the first turn, each P'_i sends M_i to V, for $1 \leq i \leq k$, and (4) if the message from V is r, at the third turn, each P'_i applies $X_i^{(r)}$ to the qubits in R'_i, for $1 \leq i \leq k$. It is obvious that these provers $P'_1, \dots, P'_k$ can convince the original verifier V with the same probability that $R'_1, \dots, R'_{k+1}$ can convince W. Hence, the probability W accepts x is at most s, as desired. □

Now Theorem 100 follows from Theorem 103, Theorem 111, and Theorem 113. Corollary 105, claiming $\mathrm{QIP} \subseteq \mathrm{QMIP}^*(2, 2, 1, 2^{-p})$ for any $p \in \mathrm{poly}$ follows directly from Theorem 113 and the fact shown by [81] that any language in QIP can be verified by a three-turn public-coin quantum interactive proof system of perfect completeness with exponentially small error in soundness (i.e., $\mathrm{QIP} \subseteq \mathrm{QMAM}(1, 2^{-p})$ for any $p \in \mathrm{poly}$).

10.4.3 Directly Modifying Three-Turn Systems to Two-Turn Systems

For completeness, here we give a direct proof of the fact that any k-prover three-turn system can be converted into a $(k+1)$-prover two-turn system.

Lemma 114. *For any c, s satisfying $c^2 > s$,*

$$\mathrm{QMIP}^*(k, 3, c, s) \subseteq \mathrm{QMIP}^*\left(k+1, 2, \frac{1+c}{2}, \frac{1+\sqrt{s}}{2}\right).$$

Proof. The proof is very similar to that of Theorem 111. Indeed, our starting point is the same, but this time we move to a two-turn proof system, instead of a three-turn public-coin system, by adding an extra prover. As in Theorem 113, the verifier first broadcasts a random bit $b \in \{0, 1\}$ to all but the extra prover, and ask the extra prover to send him a register V and the other provers to send him registers M_i. We then proceed as in Step 3 of the proof system given in Theorem 111: a detailed description is given in Figure 10.5.

Completeness: This follows immediately from the completeness of the proof system in Theorem 111: in Theorem 111 the first prover sends both V (before receiving the bit b) and M_1 (after); here we can imagine that before the protocol starts the first prover gives register V to the extra $(k+1)$-st prover, who sends it to the verifier.

Soundness: This also follows from the soundness of the proof system in Theorem 111: by combining the actions of the first prover and the extra $(k+1)$-st prover (and thus making

Verifier's Protocol in the One-Round System (Direct Construction)

1. Choose $b \in \{0, 1\}$ uniformly at random. Send b only to the first k provers, and send nothing to the $(k+1)$-st prover.

2. Receive M_i from the i-th prover, for $1 \le i \le k$, and V from the $(k+1)$-st prover.

    2.1 If $b = 0$, apply V^2 to the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$. Accept if the content of $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$ is an accepting state in the original proof system, and reject otherwise.

    2.2 If $b = 1$, apply $(V^1)^\dagger$ to the qubits in $(\mathsf{V}, \mathsf{M}_1, \ldots, \mathsf{M}_k)$. Accept if all the qubits in V are in state $|0\rangle$, and reject otherwise.

Figure 10.5: Verifier's protocol to reduce the number of turns to two (direct construction).

the provers only stronger), we can construct a set of provers that would succeed in the proof system of Theorem 111 with the same probability as they succeed here. □

10.5 Conclusion

We showed that restricting the number of turns to three and requiring the verifier to be public coin, or even restricting to two turns (without the public-coin property) does not affect the class QMIP. Moreover, we showed that any QMIP* system can be made to have perfect completeness.

An obvious drawback of some of our results is that they require an increase in the number of provers. This happens in two cases. First, when we want to improve soundness while keeping the number of rounds constant, we are forced to use new sets of prover to perform independent parallel repetitions. An interesting open problem is to show a parallel repetition theorem for QMIP, which would in particular allow us to improve soundness without increasing the number of provers. Second, we use an additional prover in the reduction from three turns to two turns. Finding a way to avoid this in general would in particular show that QIP(2) = QIP(3)(= QIP), which is an open question. Perhaps it might be easier to show that the additional prover can be avoided when there are at least two provers originally.

A related, more general open problem would be to study to what extent the number of provers can be reduced in a QMIP proof system. This study has been initiated by [73] for the case of multi-prover QMA (see also [3]), but nothing is known for multi-round proof systems with entangled provers.

Bibliography

[1] S. Aaronson. “Limitations of Quantum Advice and One-Way Communication”. In: *Theory Comput.* 1 (2005), pp. 1–28.

[2] S. Aaronson. “QMA/qpoly $\subseteq$ PSPACE/poly: De-Merlinizing Quantum Protocols”. In: *Proceedings of the 21st IEEE Conference on Computational Complexity,* Prague Czech Republic. IEEE Computer Society, pp. 261–273.

[3] S. Aaronson et al. “The Power of Unentanglement”. In: *Proceedings of the 23rd IEEE Conference on Computational Complexity,* College Park MD. 2008, pp. 223–236.

[4] D. Aharonov, A. Kitaev, and N. Nisan. “Quantum Circuits with Mixed States”. In: *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing,* Dallas TX. 1998, pp. 20–30.

[5] N. Alon et al. “Simple Constructions of Almost k-wise Independent Random Variables”. In: *Random Structures and Algorithms* 3.3 (1992), pp. 289–304.

[6] A. Ambainis et al. “Dense Quantum Coding and a Lower Bound for 1-Way Quantum Automata”. In: *Proceedings of the Thirty-first Annual ACM Symposium on the Theory of Computing,* Atlanta GA. ACM, 1999, pp. 376–383.

[7] A. Ambainis et al. “Dense Quantum Coding and Quantum Finite Automata”. In: *Journal of the ACM* 49.4 (2002), pp. 496–511.

[8] P K Aravind. *The Magic Squares and Bell’s Theorem.* Tech. rep. arXiv:quant-ph/0206070, 2002.

[9] S. Arora and S. Safra. “Probabilistic Checking of Proofs: A New Characterization of NP”. In: *Journal of the ACM* 45.1 (1998), pp. 70–122.

[10] S. Arora et al. “Proof Verification and the Hardness of Approximation Problems”. In: *Journal of the ACM* 45.3 (1998), pp. 501–555.

[11] S. Arora et al. “Unique Games on Expanding Constraint Graphs are Easy”. In: *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing,* Victoria BC. New York, NY, USA, 2008, pp. 21–28.

[12] L. Babai, L. Fortnow, and C. Lund. “Non-Deterministic Exponential Time has Two-Prover Interactive Protocols”. In: *Computational Complexity* 1.1 (1991), pp. 3–40.

[13] B. Barak et al. “Rounding Parallel Repetitions of Unique Games”. In: *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science,* Philadelphia PA. 2008, pp. 374–383.

[14] B. Barak et al. “Strong Parallel Repetition Theorem for Free Projection Games”. In: *Proc. 13th RANDOM.* 2009, pp. 352–365.

[15] J. S. Bell. “On the Einstein-Podolsky-Rosen Paradox”. In: *Physics* 1 (1964), pp. 195–200.

[16] A. Ben-Aroya and A. Ta-Shma. *Better Short-Seed Extractors Against Quantum Knowledge.* Tech. rep. arXiv:1004.3737, 2010.

[17] M. Ben-Or et al. “Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions”. In: *Proceedings of the Twentieth Annual ACM Symposium on the Theory of Computing,* Chicago IL. 1988, pp. 113–131.

[18] C. Bennett and G. Brassard. “Quantum Cryptography: Public Key Distribution and Coin Tossing”. In: *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing.* 1984, pp. 175–179.

[19] C. Bennett, G. Brassard, and J.-M. Robert. “Privacy Amplification by Public Discussion”. In: *SIAM Journal on Computing* 17.2 (1988), pp. 210–229.

[20] C. Bennett et al. “Generalized Privacy Amplification”. In: *IEEE Transactions on Information Theory* 41.6 (Nov. 1995), pp. 1915–1923.

[21] R. Bhatia. “Perturbation Inequalities for the Absolute Value Map in Norm Ideals of Operators”. In: *Journal of Operator Theory* 19.1 (1988), pp. 129–136.

[22] R. Bhatia and C. Davis. “A Cauchy–Schwarz Inequality for Operators with Applications”. In: *Linear Algebra and its Applications* 223–224 (July 1995), pp. 119–129.

[23] M. Blum, M. Luby, and R. Rubinfeld. “Self-Testing/Correcting with Applications to Numerical Problems”. In: *Journal of Computer and System Sciences* 47.3 (1993), pp. 549–595.

[24] J.-Y. Cai, A. Condon, and R. J. Lipton. “PSPACE is Provable by Two Provers in One Round”. In: *Journal of Computer and System Sciences* 48.1 (1994), pp. 183–193.

[25] J. F. Clauser et al. “Proposed Experiment to Test Local Hidden-Variable Theories”. In: *Phys. Rev. Lett.* 23 (1969), pp. 880–884.

[26] R. Cleve et al. “Consequences and Limits of Nonlocal Strategies”. In: *Proceedings of the 19th IEEE Conference on Computational Complexity,* Amherst MA. 2004, pp. 236–249.

[27] R. Cleve et al. “Perfect Parallel Repetition Theorem for Quantum XOR Proof Systems”. In: *Computational Complexity* 17 (2008), pp. 282–299.

[28] R. Colbeck. “Quantum And Relativistic Protocols For Secure Multi-Party Computation”. arXiv:0911.3814. PhD thesis. Trinity College, University of Cambridge, Nov. 2009.

[29] R. Colbeck and A. Kent. “Private Randomness Expansion with Untrusted Devices”. In: *Journal of Physics A: Mathematical and Theoretical* 44.9 (2011), p. 095305.

[30] R. Colbeck and R. Renner. *Free Randomness Amplification.* Tech. rep. arXiv:1105.3195, 2011.

[31] Zuckerman D. “General Weak Random Sources”. In: *Proceedings of the 31st Annual IEEE Symposium on Foundations of Computer Science,* St. Louis MO. IEEE, 1990, pp. 534–543.

[32] K. R. Davidson and S. J. Szarek. “Local Operator Theory, Random Matrices and Banach Spaces”. In: *Handbook of the Geometry of Banach Spaces, volume 1.* Ed. by W. B. Johnson and J. Lindenstrauss. North-Holland Publishing Company, 2001, pp. 317–366.

[33] A. De and T. Vidick. “Near-Optimal Extractors Against Quantum Storage”. In: *Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing,* Cambridge MA. Cambridge, Massachusetts, USA: ACM, 2010, pp. 161–170.

[34] A. De et al. *Trevisan's Extractor in the Presence of Quantum Side Information.* Tech. rep. arXiv:0912.5514, 2009.

[35] I. Dinur and E. Goldenberg. “Locally Testing Direct Product in the Low Error Range”. In: *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science,* Philadelphia PA. 2008, pp. 613–622.

[36] I. Dinur and O. Meir. “Derandomized Parallel Repetition of Structured PCPs”. In: *Proceedings of the 25th IEEE Conference on Computational Complexity,* Cambridge MA. 2010, pp. 16–27.

[37] I. Dinur and O. Reingold. “Assignment Testers: Towards a Combinatorial Proof of the PCP Theorem”. In: *SIAM Journal on Computing* 36.4 (2006), pp. 975–1024.

[38] A. C. Doherty et al. “The Quantum Moment Problem and Bounds on Entangled Multi-prover Games”. In: *Proceedings of the 23rd IEEE Conference on Computational Complexity,* College Park MD (2008), pp. 199–210.

[39] R. Exel and T. Loring. “Almost Commuting Unitary Matrices”. In: *Proc. Amer. Math. Soc.* 106.4 (1989), pp. 913–915.

[40] S. Fehr, R. Gelles, and C. Schaffner. “Security and Composability of Randomness Expansion from Bell Inequalities”. Manuscript. 2011.

[41] U. Feige. “On the Success Probability of Two Provers in One-Round Proof Systems”. In: *Proc. 6th IEEE Structure in Complexity Theory.* 1991, pp. 116–123.

[42] U. Feige and J. Kilian. “Two-Prover Protocols—Low Error at Affordable Rates”. In: *SIAM Journal on Computing* 30.1 (2000), pp. 324–346.

[43] U. Feige, G. Kindler, and R. O’Donnell. “Understanding Parallel Repetition Requires Understanding Foams”. In: *Proceedings of the 21st IEEE Conference on Computational Complexity,* San Diego CA. 2007, pp. 179–192.

[44] U. Feige and L. Lovász. “Two-Prover One-Round Proof Systems: Their Power and Their Problems”. In: *Proceedings of the Twenty-fourth Annual ACM Symposium on the Theory of Computing,* Victoria BC, Canada. 1992, pp. 733–744.

[45] U. Feige et al. “Interactive Proofs and the Hardness of Approximating Cliques”. In: *Journal of the ACM* 43.2 (1996), pp. 268–292.

[46] D. Gavinsky et al. “Exponential Separation for One-Way Quantum Communication Complexity, with Applications to Cryptography”. In: *SIAM Journal on Computing* 38.5 (2008), pp. 1695–1708.

[47] M. Genovese. “Research on hidden variable theories: A review of recent progresses”. In: *Physics Reports* 413.6 (2005), pp. 319 –396.

[48] V. Guruswami, C. Umans, and S. Vadhan. “Unbalanced Expanders and Randomness Extractors from Parvaresh-Vardy Codes”. In: *Proceedings of the 21st IEEE Conference on Computational Complexity,* San Diego CA. Washington, DC, USA: IEEE Computer Society, 2007, pp. 96–108.

[49] V. Guruswami et al. “Combinatorial Bounds for List Decoding”. In: *IEEE Transactions on Information Theory* 48.5 (2002), pp. 1021–1034.

[50] P. Halmos. “Some Unknown Problems of Unknown Depth About Operators on Hilbert Space”. In: *Proc. R. Soc. Edinburgh A* 76 (1976), pp. 67–76.

[51] T. Hartman and R. Raz. “On the Distribution of the Number of Roots of Polynomials and Explicit Weak Designs”. In: *Random Structures and Algorithms* 23.3 (2003), pp. 235–263.

[52] J. Håstad. “Some Optimal Inapproximability Results.” In: *Journal of the ACM* 48 (2001), pp. 798–859.

[53] T. Holenstein. “Parallel Repetition: Simplification and the No-Signaling Case”. In: *Theory of Computing* 5.1 (2009), pp. 141–172.

[54] R. Impagliazzo. “Uniform Direct Product Theorems: Simplified, Optimized, and Derandomized”. In: *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing,* Victoria BC. 2008, pp. 579–588.

[55] R. Impagliazzo, R. Jaiswal, and V. Kabanets. “Approximately List-Decoding Direct Product Codes and Uniform Hardness Amplification”. In: *Foundations of Computer Science, 2006. FOCS ’06. 47th Annual IEEE Symposium on.* 2006, pp. 187 –196.

[56] R. Impagliazzo, V. Kabanets, and A. Wigderson. "New Direct-Product Testers and 2-Query PCPs". In: *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing,* Bethesda MA. 2009, pp. 131–140.

[57] R. Impagliazzo, R. Shaltiel, and A. Wigderson. "Extractors and Pseudo-Random Generators with Optimal Seed Length". In: *Proceedings of the Thirty-second Annual ACM Symposium on the Theory of Computing,* Portland OR. ACM, 2000, pp. 1–10.

[58] R. Impagliazzo and D. Zuckerman. "How to Recycle Random Bits". In: *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science,* Research Triangle Park NC. 1989, pp. 248–253.

[59] T. Ito, H. Kobayashi, and K. Matsumoto. "Oracularization and Two-Prover One-Round Interactive Proofs against Nonlocal Strategies". In: *Proceedings of the 24th IEEE Conference on Computational Complexity,* Paris France. 2009, pp. 217–228.

[60] T. Ito and T. Vidick. "NEXP $\subseteq$ MIP*". Manuscript. 2011.

[61] T. Ito et al. "Generalized Tsirelson Inequalities, Commuting-Operator Provers, and Multi-Prover Interactive Proof Systems". In: *Proceedings of the 23rd IEEE Conference on Computational Complexity,* College Park MD. 2008, pp. 187–198.

[62] R. Jain, S. Upadhyay, and J. Watrous. "Two-Message Quantum Interactive Proofs Are in PSPACE". In: *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science,* Atlanta GA. 2009, pp. 534–543.

[63] R. Jain and J. Watrous. "Parallel Approximation of Non-interactive Zero-sum Quantum Games". In: *Proceedings of the 24th IEEE Conference on Computational Complexity,* Paris France. Washington, DC, USA, 2009, pp. 243–253.

[64] R. Jain et al. "QIP = PSPACE". In: *Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing,* Cambridge MA. 2010, pp. 573–582.

[65] J. Kempe and O. Regev. "No Strong Parallel Repetition with Entangled and Non-signaling Provers". In: *Proceedings of the 25th IEEE Conference on Computational Complexity,* Cambridge MA. 2010, pp. 7–15.

[66] J. Kempe, O. Regev, and B. Toner. "Unique Games with Entangled Provers Are Easy". In: *SIAM Journal on Computing* 39.7 (2010), pp. 3207–3229.

[67] J. Kempe and T. Vidick. "Parallel Repetition of Entangled Games". In: *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing,* San Jose CA. 2011, pp. 353–362.

[68] J. Kempe et al. "Entangled Games Are Hard to Approximate". In: *SIAM Journal on Computing* 40.3 (2011), pp. 848–877.

[69] J. Kempe et al. "Using Entanglement in Quantum Multi-Prover Interactive Proofs". In: *Computational Complexity* 18 (2 2009), pp. 273–307.

[70] A. Kitaev and J. Watrous. "Parallelization, Amplification, and Exponential Time Simulation of Quantum Interactive Proof Systems". In: *Proceedings of the Thirty-second Annual ACM Symposium on the Theory of Computing,* Portland OR. 2000, pp. 608–617.

[71] M. Koashi and A. Winter. "Monogamy of Quantum Entanglement and Other Correlations". In: *Physical Review A* 69.2 (2004), p. 022309.

[72] H. Kobayashi and K. Matsumoto. "Quantum Multi-Prover Interactive Proof Systems with Limited Prior Entanglement". In: *Journal of Computer and System Sciences* 66.3 (2003), pp. 429–450.

[73] H. Kobayashi, K. Matsumoto, and T. Yamakami. "Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur?" In: *Proceedings of 14th International Symposium on Algorithms and Computation ISAAC 2003,* Kyoto, Japan. Vol. 2906. Lecture Notes in Computer Science. 2003, pp. 189–198.

[74] R. König and R. Renner. *Sampling of Min-Entropy Relative to Quantum Knowledge.* Tech. rep. arXiv:0712.4291, 2007.

[75] R. König, R. Renner, and C. Schaffner. "The Operational Meaning of Min- and Max-Entropy". In: *IEEE Transactions on Information Theory* 55.9 (2009), pp. 4337–4347.

[76] R. König and B. M. Terhal. "The Bounded-Storage Model in the Presence of a Quantum Adversary". In: *IEEE Transactions on Information Theory* 54.2 (2008), pp. 749–762.

[77] D. Leung, B. Toner, and J. Watrous. *Coherent State Exchange in Multi-Prover Quantum Interactive Proof Systems.* Tech. rep. arXiv:0804.4118, 2008.

[78] H. Lin. "Almost Commuting Selfadjoint Matrices and Applications". In: *Operator Algebra and Their Applications.* Ed. by Peter A. Fillmore and James A. Mingo. Vol. 13. Fields Inst. Comm. American Mathematical Society, 1997, pp. 193–233.

[79] C.-J. Lu. "Encryption against Storage-Bounded Adversaries from On-Line Strong Extractors". In: *Journal of Cryptology* 17.1 (2004), pp. 27–42.

[80] C. Lund et al. "Algebraic Methods for Interactive Proof Systems". In: *Journal of the ACM* 39.4 (1992), pp. 859–868.

[81] C. Marriott and J. Watrous. "Quantum Arthur-Merlin games". In: *Computational Complexity* 14.2 (2005), pp. 122–152.

[82] U. Maurer. "Conditionally-Perfect Secrecy and a Provably-Secure Randomized Cipher". In: *Journal of Cryptology* 5.1 (1992), pp. 53–66.

[83] N. D. Mermin. "Simple Unified Form for the Major No-Hidden-Variables Theorems". In: *Physical Review Letters* 65.27 (1990), pp. 3373–3376.

[84] A. Nayak and P. Shor. “Bit-Commitment-Based Quantum Coin Flipping”. In: *Physical Review A* 67.1 (2003), p. 012304.

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

[86] N. Nisan and A. Wigderson. “Hardness vs Randomness”. In: *Journal of Computer and System Sciences* 49.2 (1994), pp. 149 –167.

[87] N. Nisan and D. Zuckerman. “Randomness is Linear in Space”. In: *Journal of Computer and System Sciences* 52.1 (1996), pp. 43–52.

[88] R. O’Donnell. *Lecture 12 : ”Confuse / Match” Games (I).* Available at `http://www.cs.washington.edu/education/courses/cse533/05au/`. 2005.

[89] R. O’Donnell. *Lecture 13: ”Confuse / Match” Games (II).* Available at `http://www.cs.washington.edu/education/courses/cse533/05au/`. 2005.

[90] T. Ogawa and H. Nagaoka. “Making Good Codes for Classical-Quantum Channel Coding via Quantum Hypothesis Testing”. In: *IEEE Transactions on Information Theory* 53.6 (2007), pp. 2261 –2266.

[91] A. Peres. “Incompatible Results of Quantum Measurements”. In: *Physical Review A* 151.3–4 (1990), pp. 107 –108.

[92] S. Pironio et al. “Random Numbers Certified by Bell’s Theorem.” In: *Nature* 464.7291 (2009), p. 10.

[93] J. Radhakrishnan and A. Ta-Shma. “Bounds for Dispersers, Extractors, and Depth-Two Superconcentrators”. In: *SIAM Journal on Discrete Mathematics* 13.1 (2000), pp. 2–24.

[94] A. Rao. “Parallel Repetition in Projection Games and a Concentration Bound”. In: *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing,* Victoria BC. 2008, pp. 1–10.

[95] R. Raz. “A Counterexample to Strong Parallel Repetition”. In: *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science,* Philadelphia PA. 2008, pp. 369–373.

[96] R. Raz. “A Parallel Repetition Theorem”. In: *SIAM Journal on Computing* 27 (1998), pp. 763–803.

[97] R. Raz. “Extractors with Weak Random Seeds”. In: *Proceedings of the 37th Annual ACM Symposium on the Theory of Computing,* Baltimore MD. ACM, 2005, pp. 11–20.

[98] R. Raz, O. Reingold, and S. Vadhan. “Extracting all the Randomness and Reducing the Error in Trevisan’s Extractors”. In: *Journal of Computer and System Sciences* 65.1 (2002), pp. 97–128.

[99] R. Raz and R. Rosen. "A Strong Parallel Repetition Theorem for Projection Games on Expanders". In: *Technical report ECCC TR10-142* (2010).

[100] R. Renner. "Security of Quantum Key Distribution". PhD thesis. Swiss Federal Institute of Technology Zurich, Sept. 2005. eprint: `quant-ph/0512258`.

[101] M. Santha and U. V. Vazirani. "Generating Quasi-Random Sequences From Slightly-Random Sources". In: *Proceedings of the 25th Annual IEEE Symposium on Foundations of Computer Science,* Singer Island FL. Washington, DC, USA: IEEE Computer Society, 1984, pp. 434–440.

[102] P. H. Schönemann. "A Generalized Solution of the Orthogonal Procrustes Problem". In: *Psychometrika* 31.1 (1966), pp. 1–10.

[103] J. T. Schwartz. "Fast Probabilistic Algorithms for Verification of Polynomial Identities". In: *Journal of the ACM* 27.4 (Oct. 1980), pp. 707–717.

[104] R. Shaltiel. "Recent Developments in Explicit Constructions of Extractors". In: *Bulletin of the European Association for Theoretical Computer Science* 77 (June 2002), pp. 67–95.

[105] R. Shaltiel and C. Umans. "Simple Extractors for all Min-Entropies and a New Pseudorandom Generator". In: *Journal of the ACM* 52.2 (2005), pp. 172–216.

[106] A. Shamir. "IP = PSPACE". In: *Journal of the ACM* 39.4 (1992), pp. 869–877.

[107] P. W. Shor. "Fault-Tolerant Quantum Computation". In: *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science,* Burlington VT. 1996, pp. 56–65.

[108] R. Solcà. "Efficient Simulation of Random Quantum States and Operators". MA thesis. Swiss Federal Institute of Technology, Zurich, 2010.

[109] R. W. Spekkens and T. Rudolph. "Degrees of Concealment and Bindingness in Quantum Bit-Commitment Protocols". In: *Physical Review A* 65.1 (2002), p. 012310.

[110] A. Srinivasan and D. Zuckerman. "Computing with Very Weak Random Sources". In: *SIAM Journal on Computing* 28.4 (1999), pp. 1433–1459.

[111] M. Sudan, L. Trevisan, and S. Vadhan. "Pseudorandom Generators Without the XOR Lemma". In: *Proceedings of the Thirty-first Annual ACM Symposium on the Theory of Computing,* Atlanta GA. ACM, 1999, pp. 537–546.

[112] A. Ta-Shma. "Short Seed Extractors Against Quantum Storage". In: *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing,* Bethesda MA. ACM, 2009, pp. 401–408.

[113] A. Ta-Shma, C. Umans, and D. Zuckerman. "Loss-Less Condensers, Unbalanced Expanders, and Extractors". In: *Proceedings of the Thirty-third Annual ACM Symposium on the Theory of Computing,* Hersonissos, Crete, Greece. ACM, 2001, pp. 143–152.

[114] A. Ta-Shma, D. Zuckerman, and S. Safra. “Extractors from Reed-Muller codes”. In: *Journal of Computer and System Sciences* 72.5 (2006), pp. 786–812.

[115] B. Terhal. “Is Entanglement Monogamous?” In: *IBM J. Res. Dev.* 48 (1 2004), pp. 71–78.

[116] M. Tomamichel, R. Colbeck, and R. Renner. “Duality Between Smooth Min- and Max-Entropies”. In: *IEEE Transactions on Information Theory* 56.9 (2010), pp. 4674–4681.

[117] M. Tomamichel et al. “Leftover Hashing Against Quantum Side Information”. In: *Proceedings of 2010 international symposium on information theory, ISIT*. IEEE, 2010, pp. 2703–2707.

[118] B. Toner. “Monogamy of Non-Local Quantum Correlations”. In: *Proceedings of the Royal Society A* 465.2101 (2009), pp. 59–69.

[119] L. Trevisan. “Extractors and Pseudorandom Generators”. In: *Journal of the ACM* 48.4 (2001), pp. 860–879.

[120] S. Vadhan. “Constructing Locally Computable Extractors and Cryptosystems in the Bounded-Storage Model”. In: *Journal of Cryptology* 17.1 (2004), pp. 43–77.

[121] U. V. Vazirani and T. Vidick. “Certifiable Quantum Dice”. Manuscript. 2011.

[122] O. Verbitsky. “Towards the Parallel Repetition Conjecture”. In: *Proc. 9th IEEE Conference on Structure in Complexity Theory.* 1994, pp. 304–307.

[123] D. Voiculescu. “Asymptotically Commuting Finite Rank Unitary Operators Without Commuting Approximants”. In: *Acta Sci. Math.* 45 (1983), pp. 429–431.

[124] D. Voiculescu. “Remarks on the Singular Extension in the C^*-Algebra of the Heisenberg Group”. In: *J. Operator Theory* 5.2 (1981), pp. 147–170.

[125] J. Watrous. “PSPACE has Constant-Round Quantum Interactive Proof Systems”. In: *Theoretical Computer Science* 292.3 (2003), pp. 575–588.

[126] J. Watrous. “Zero-Knowledge Against Quantum Attacks”. In: *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing,* Seattle WA. 2006, pp. 296–305.

[127] R. F. Werner. “An Application of Bell’s Inequalities to a Quantum State Extension Problem”. In: *Lett. Math. Phys.* 17.4 (1989), pp. 359–363.

[128] A. Wigderson and D. Zuckerman. “Expanders That Beat the Eigenvalue Bound: Explicit Construction and Applications”. In: *Combinatorica* 19.1 (1999), pp. 125–138.

[129] A. Winter. “Coding Theorem and Strong Converse for Quantum Channels”. In: *IEEE Transactions on Information Theory* 45.7 (1999), pp. 2481 –2485.

[130] F. Xu et al. “An ultrafast quantum random number generator based on quantum phase fluctuations”. arXiv:1109.0643. 2011.

[131] A. C.-C. Yao. Personal Communication. Feb. 2007.

[132] A. C.-C. Yao. “Theory and Applications of Trapdoor Functions”. In: *Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science,* Chicago IL. IEEE, 1982, pp. 80–91.

Appendix A

Auxiliary results

This appendix collects a few useful inequalities, together with proofs that were omitted from the main chapters of this dissertation.

A.1 Matrix inequalities

We first state a useful operator version of the Cauchy-Schwarz inequality.

Claim 115. *Let A, B be (possibly rectangular) matrices such that $A^\dagger B$ exists, and $B^\dagger B$ is invertible. Then*

$$(A^\dagger B)(B^\dagger B)^{-1}(B^\dagger A) \leq A^\dagger A$$

Proof. Let $\Delta = (B^\dagger B)^{-1}(B^\dagger A)$. Then the matrix $(A - B\Delta)^\dagger (A - B\Delta)$ is positive, which gives the result. □

The following is another version of a matrix Cauchy-Schwarz inequality. It follows from Eq. (3) of Bhatia and Davis [22] (see also [21]), substituting the norm $|||\cdot|||$ by $\|\cdot\|_1$,

Theorem 116. *Let A and B be arbitrary matrices. Then,*

$$\left\|A^\dagger B\right\|_1 \leq \left\|A\right\|_F \left\|B\right\|_F.$$

Winter's gentle measurement lemma [129, Lemma 9] is an key lemma formalizing the intuitive fact that if a measurement produces a certain outcome with near-certainty when performed on a specific state, then the post-measurement state is close to the original state. The following is a variant of that lemma.

Lemma 117. *Let ρ be a density operator on a Hilbert space $\mathcal{H}$, and X and Y be linear operators from $\mathcal{H}$ to a Hilbert space $\mathcal{K}$ such that $X^*X \preceq I$ and $Y^*Y \preceq I$. Then,*

$$\|X\rho X^* - Y\rho Y^*\|_1 \leq 2\sqrt{\operatorname{Tr}(X-Y)\rho(X-Y)^*}.$$

Proof. The inequality follows from a calculation similar to Ogawa and Nagaoka's proof [90, Appendix C] of Winter's "gentle measurement lemma" [129, Lemma 9]. By the triangle inequality,

$$\|X\rho X^* - Y\rho Y^*\|_1 \le \|(X-Y)\rho X^*\|_1 + \|Y\rho(X-Y)^*\|_1.$$

By Theorem 116,

$$\begin{aligned}\|(X-Y)\rho X^*\|_1 &\le \|(X-Y)\sqrt{\rho}\|_2 \|\sqrt{\rho}\, X^*\|_2 \\ &= \sqrt{\mathrm{Tr}(X-Y)\rho(X-Y)^*}\sqrt{\mathrm{Tr} X\rho X^*} \\ &\le \sqrt{\mathrm{Tr}(X-Y)\rho(X-Y)^*}.\end{aligned}$$

Similarly, $\|Y\rho(X-Y)^*\|_1 \le \sqrt{\mathrm{Tr}(X-Y)\rho(X-Y)^*}$, and the lemma follows. □

We state the following two corollaries of Lemma 117, that will be used in Chapter 5.

Claim 118. *Let $\{A_i\}$ and $\{B_i\}$ be two sets of positive matrices of the same dimension, and $\rho \ge 0$. Then*

$$\Big\| \sum_i \sqrt{A_i}\,\rho\,\sqrt{A_i} - \sqrt{B_i}\,\rho\,\sqrt{B_i} \Big\|_1 \le 2\Big(\sum_i \mathrm{Tr}\big(\big(\sqrt{A_i} - \sqrt{B_i}\big)^2 \rho\big)\Big)^{1/2}.$$

Proof. Let X be a block-column matrix with blocks the $\sqrt{A_i}$, and similarly for Y and the $\sqrt{B_i}$. Then

$$\Big\| \sum_i \sqrt{A_i}\,\rho\,\sqrt{A_i} - \sqrt{B_i}\,\rho\,\sqrt{B_i} \Big\|_1 \le \sum_i \Big\| \sqrt{A_i}\,\rho\,\sqrt{A_i} - \sqrt{B_i}\,\rho\,\sqrt{B_i} \Big\|_1 \le \big\|X\rho X^\dagger - Y\rho Y^\dagger\big\|_1,$$

and

$$\mathrm{Tr}\big((X-Y)\rho(X-Y)^\dagger\big) = \sum_i \mathrm{Tr}\big(\big(\sqrt{A_i} - \sqrt{B_i}\big)^2 \rho\big),$$

so that the claim follows from Lemma 117. □

Claim 119. *Let $\sigma \ge 0$ be symmetric, $\rho = \mathrm{Tr}_1\sigma = \mathrm{Tr}_2\sigma$ and $\{A_i\}_i$ a POVM, and let*

$$\delta := \sum_{i \ne j} \mathrm{Tr}\big(A_i \otimes A_j \sigma\big).$$

Then

$$\big\| \sum_i \sqrt{A_i}\rho\sqrt{A_i} - \rho \big\|_1 = O(\sqrt{\delta}).$$

Proof. First note that, $\{A_i\}_i$ being a POVM, $\mathrm{Tr}_2\big(\sum_i \mathrm{Id}\otimes\sqrt{A_i}\sigma\mathrm{Id}\otimes\sqrt{A_i}\big) = \mathrm{Tr}_2\big(\sigma\big) = \rho$. Hence by monotonicity of the trace norm

$$\begin{aligned}\big\|\sum_i \sqrt{A_i}\rho\sqrt{A_i}-\rho\big\|_1 &\leq \big\|\sum_{i,j}\sqrt{A_i}\otimes\sqrt{A_j}\sigma\sqrt{A_i}\otimes\sqrt{A_j}-\sum_j \mathrm{Id}\otimes\sqrt{A_j}\sigma\mathrm{Id}\otimes\sqrt{A_j}\big\|_1\\ &\leq \big\|\sum_i \sqrt{A_i}\otimes\sqrt{A_i}\sigma\sqrt{A_i}\otimes\sqrt{A_i}-\sum_i \mathrm{Id}\otimes\sqrt{A_i}\sigma\mathrm{Id}\otimes\sqrt{A_i}\big\|_1\\ &\qquad+\sum_{i\neq j}\mathrm{Tr}\big(A_i\otimes A_j\sigma\big)\\ &\leq 2\sqrt{\sum_i \mathrm{Tr}\big((\sqrt{A_i}\otimes\sqrt{A_i}-\mathrm{Id}\otimes\sqrt{A_i})^2\sigma\big)}\sqrt{\mathrm{Tr}\big(\mathrm{Id}\otimes A_i\sigma\big)}+\delta\\ &\leq 2\sqrt{\delta}+\delta\end{aligned}$$

where the second inequality is the triangle inequality, the third is by Claim 118, and the last uses the definition of δ. □

A.2 Omitted proofs from Chapter 5

The following useful lemma relates the consistency of a measurement when performed on two separate subsystems of a permutation-invariant state with the possibility of exchanging the sub-system on which the measurement is performed. Here ρ is the reduced density of a permutation-invariant state, and $\mu(A) = \sum_{i\neq j}\mathrm{Tr}_\rho\big(A_i\otimes A_j\big)$.

Lemma 120. *Let $\{A_i\}$ be a POVM, Z such that $ZZ^\dagger \leq \mathrm{Id}$, and $\big\{Z_i\big\}$ such that $\sum_i Z_iZ_i^\dagger \leq \mathrm{Id}$. Then*

$$\Big|\sum_i \mathrm{Tr}_\rho(Z_iA_i)-\sum_i \mathrm{Tr}_\rho(Z_i\otimes A_i)\Big| \leq \sqrt{\mu(A)} \tag{A.1}$$

$$\Big|\sum_{i\neq j}\mathrm{Tr}_\rho(A_jZA_i)\Big| \leq 2\sqrt{\mu(A)} \tag{A.2}$$

Proof. We first prove (A.1). We have

$$\begin{aligned}\Big|\sum_i \mathrm{Tr}_\rho(Z_iA_i)-\sum_i \mathrm{Tr}_\rho(Z_i\otimes A_i)\Big| &= \Big|\sum_i \mathrm{Tr}_\rho\big(Z_i(A_i\otimes\mathrm{Id}-\mathrm{Id}\otimes A_i)\big)\Big|\\ &\leq \Big(\sum_i \mathrm{Tr}_\rho\big(Z_iZ_i^\dagger\big)\Big)^{1/2}\Big(\sum_{i\neq j}\mathrm{Tr}_\rho\big((A_i\otimes A_j)^2\big)\Big)^{1/2}\\ &\leq \sqrt{\mu(A,A)},\end{aligned}$$

where the second inequality follows from Cauchy-Schwarz. Regarding (A.2), we have

$$\Big|\sum_{i\neq j}\mathrm{Tr}_\rho(A_iZA_j)\Big| = \Big|\mathrm{Tr}_\rho\big(Z\big) - \sum_i \mathrm{Tr}_\rho\big(A_iZA_i\big)\Big|$$

From (A.1) we know that

$$\Big|\sum_i \mathrm{Tr}_\rho\big(A_iZA_i\big) - \sum_i \mathrm{Tr}_\rho\big(A_iZ\otimes A_i\big)\Big| \leq \sqrt{\mu(A)}.$$

The second term on the left-hand side satisfies

$$\begin{aligned}\Big|\sum_i \mathrm{Tr}_\rho\big(A_iZ\otimes A_i\big) - \sum_i \mathrm{Tr}_\rho\big(Z\otimes A_i\big)\Big| &\leq \Big(\sum_i \mathrm{Tr}_\rho\big(Z^\dagger Z\otimes A_i\big)\Big)^{1/2}\Big(\sum_i \mathrm{Tr}_\rho\big((\mathrm{Id}-A_i)^2\otimes A_i\big)\Big)^{1/2}\\ &\leq \sqrt{\mu(A)},\end{aligned}$$

and this concludes the proof. □

The following lemma follows from the standard expansion properties of the hypercube. Recall that for $\rho \geq 0$ and any A, $\|A\|_\rho^2 = \mathrm{Tr}\big(AA^\dagger\rho\big)$.

Claim 121 (Expansion lemma)**.** *Let $A : \mathbb{F}^n \to \mathbb{C}^{d\times d}$ such that for every $\boldsymbol{x}$, $0 \leq A_{\boldsymbol{x}} \leq \mathrm{Id}$, and*

$$\mathrm{E}_{i,\boldsymbol{x}_{\neg i},x_i,x_i'}\big\|A_{\boldsymbol{x}} - A_{\boldsymbol{x}'}\big\|_\rho^2 \;\leq\; \varepsilon.$$

Then there exists $A_0 = \mathrm{E}_{\boldsymbol{x}}A_{\boldsymbol{x}} \geq 0$ such that

$$\mathrm{E}_{\boldsymbol{x}}\big\|A_{\boldsymbol{x}} - A_0\big\|_\rho^2 \;\leq\; 2n\varepsilon.$$

Proof. Let $M := \sum_{\boldsymbol{x},i,x_i'}|\boldsymbol{x}\rangle\langle\boldsymbol{x}'|$ be the adjacency matrix of the hypercube $\mathbb{F}^n$, $L := np\mathrm{Id}-M$ the Laplacian, and $\tilde{L} = L\otimes\rho^{1/2}$. Let $A = \sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle\otimes A_x$. Then

$$A^\dagger\tilde{L}\cdot A = \frac{1}{2}\sum_{\boldsymbol{x},i,x_i'}(A_{\boldsymbol{x}} - A_{\boldsymbol{x}'})\rho^{1/2}(A_{\boldsymbol{x}} - A_{\boldsymbol{x}'}). \tag{A.3}$$

The normalized Laplacian $L/(np)$ has smallest eigenvalue 0, and second smallest $\lambda_1 \geq 1/(2n)$. Let the smallest eigenvector of L be $|v_0\rangle = p^{-n/2}\sum_{\boldsymbol{x}}|\boldsymbol{x}\rangle$, and write $A = |v_0\rangle\otimes A_0 + |v_1\rangle\otimes A_1$, where $|v_1\rangle$ is orthogonal to $|v_0\rangle$, and $A_0 = p^{-n/2}\sum_{\boldsymbol{x}}A_{\boldsymbol{x}}$. Then

$$A^\dagger\tilde{L}A \;=\; \lambda_1 A_1\rho^{1/2}A_1 \;\geq\; \frac{1}{2n}A_1\rho^{1/2}A_1.$$

Using the assumption made in the claim's statement together with (A.3), we get $\|A_1\|_\rho^2 \leq 2n\varepsilon p^n$, and hence by definition of A,

$$\mathrm{Tr}\big((A - |v_0\rangle\otimes A_0)^\dagger(\mathrm{Id}\otimes\rho^{1/2})(A - |v_0\rangle\otimes A_0)\rho^{1/2}\big) \;=\; \|A_1\|_\rho^2 \;\leq\; 2n\varepsilon p^n,$$

which proves the claim. □

A.2.1 Proof of Corollary 37

In this section we give the proof of Corollary 37, used in Chapter 5. A standard method to convert multiple constraints to a single constraint involving an exponential sum is by using small-bias probability spaces.

Definition 122 (Small-bias probability space)**.** *Let $n \in \mathbb{N}$. A set $S \subseteq \mathbb{F}_2^n$ is called an ε-*bias probability space *if for every $\boldsymbol{c} \in \mathbb{F}_2^n \setminus \{0\}$, it holds that*

$$\big|\Pr_{\boldsymbol{\zeta} \in S}[\boldsymbol{c} \cdot \boldsymbol{\zeta} = 0] - \Pr_{\boldsymbol{\zeta} \in S}[\boldsymbol{c} \cdot \boldsymbol{\zeta} = 1]\big| \le \varepsilon.$$

Proposition 2. *Let $n \in \mathbb{N}$, and let $S \subset \mathbb{F}_2^n$ be an ε-bias probability space. Let $\mathbb{F}$ be a finite field of characteristic two. If $\boldsymbol{c} \in \mathbb{F}^n \setminus \{0\}$, then*

$$\Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} \zeta_i c_i = 0\right] \le \frac{1+\varepsilon}{2}.$$

Proof. If $\mathbb{F} = \mathbb{F}_2$, then the proposition holds because

$$\begin{aligned}\Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i\zeta_i = 0\right] &= \frac{1}{2} + \frac{1}{2}\left(\Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i\zeta_i = 0\right] - \Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i\zeta_i = 1\right]\right) \\ &\le \frac{1+\varepsilon}{2}.\end{aligned}$$

For general $\mathbb{F}$, regard $\mathbb{F}$ as a vector space over $\mathbb{F}_2$, and let $\{\alpha_1, \dots, \alpha_k\}$ be a basis of $\mathbb{F}$ over $\mathbb{F}_2$. Write $\boldsymbol{c}$ as $\boldsymbol{c} = \alpha_1\boldsymbol{c}^{(1)} + \cdots + \alpha_k\boldsymbol{c}^{(k)}$, where $\boldsymbol{c}^{(1)}, \dots, \boldsymbol{c}^{(k)} \in \mathbb{F}_2^n$. Because $\boldsymbol{c} \ne 0$, we have that $\boldsymbol{c}^{(j^*)} \ne 0$ for some j^*. By using the case of $\mathbb{F}_2$, it holds that

$$\Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i^{(j^*)}\zeta_i = 0\right] \le \frac{1+\varepsilon}{2}.$$

Since $\alpha_1, \dots, \alpha_k$ are linearly independent over $\mathbb{F}_2$, $\sum_{i=1}^{n} c_i\zeta_i = 0$ implies $\sum_{i=1}^{n} c_i^{(j)}\zeta_i = 0$ for all j, and therefore in particular $\sum_{i=1}^{n} c_i^{(j^*)}\zeta_i = 0$. Therefore,

$$\Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i\zeta_i = 0\right] \le \Pr_{\boldsymbol{\zeta} \in S}\left[\sum_{i=1}^{n} c_i^{(j^*)}\zeta_i = 0\right] \le \frac{1+\varepsilon}{2}. \qquad \square$$

Theorem 123 (Alon, Goldreich, Håstad, and Peralta [5])**.** *There exist a constant $c > 0$ and a polynomial-time algorithm C which, given $K, M \in \mathbb{N}$, $i \in \{1, \dots, K\}$ and $j \in \{1, \dots, M\}$, outputs $C(K, M, i, j) \in \mathbb{F}_2$ such that the set $\{\boldsymbol{\zeta}^{(j)} : 1 \le j \le M\}$ defined by $\boldsymbol{\zeta}^{(j)} = (C(K, M, 1, j), \dots, C(K, M, K, j))$ is an (K/M^c)-bias probability space in $\mathbb{F}_2^K$.*

By arithmetizing the Boolean circuit for C by using a similar idea to the proof of Proposition 4.2 of Ref. [12], we obtain the following corollary.

Corollary 124. *There exist a constant $c > 0$ and a polynomial-time algorithm A which, given 1^k and 1^m, outputs 1^t and an arithmetic expression $f(\boldsymbol{i}, \boldsymbol{j}, \boldsymbol{l})$ in $k + m + t$ variables such that the set $\{\boldsymbol{\zeta}^{(\boldsymbol{j})} \colon \boldsymbol{j} \in \{0,1\}^m\}$ defined by $\boldsymbol{\zeta}^{(\boldsymbol{j})} = (\sum_{\boldsymbol{l} \in \{0,1\}^t} f(\boldsymbol{i}, \boldsymbol{j}, \boldsymbol{l}))_{\boldsymbol{i} \in \{0,1\}^k}$ is an 2^{k-cm}-bias probability space in $\mathbb{F}_2^{2^k}$.*

Proof of Corollary 37. The protocol works as follows. The verifier first computes $m = \lceil (k + 2)/c \rceil$, where c is the constant in Corollary 124. He runs the algorithm of Corollary 124 with parameters k and m to obtain $t \in \mathbb{N}$ and an arithmetic expression $f(\boldsymbol{i}, \boldsymbol{j}, \boldsymbol{l})$ in $k + m + t$ variables. Let d' be the maximum degree of f in single variables. He chooses $\boldsymbol{j} \in \{0,1\}^m$ uniformly at random, and sends $\boldsymbol{j}$ to the prover. Then he simulates the protocol in Lemma 36 with explicit inputs $k + t$ and $d + d'$ and implicit input $h_{\boldsymbol{j}}(\boldsymbol{i}, \boldsymbol{l}) := f(\boldsymbol{i}, \boldsymbol{j}, \boldsymbol{l}) h(\boldsymbol{i})$.

For $\boldsymbol{i} \in \mathbb{F}^k$, $\boldsymbol{j} \in \mathbb{F}^m$, and $\boldsymbol{l} \in \mathbb{F}^t$, let $\zeta_{\boldsymbol{i}}^{(\boldsymbol{j})} = \sum_{\boldsymbol{l}} f(\boldsymbol{i}, \boldsymbol{j}, \boldsymbol{l}) \in \mathbb{F}$ and $\boldsymbol{\zeta}^{(\boldsymbol{j})} = (\zeta_{\boldsymbol{i}}^{(\boldsymbol{j})})_{\boldsymbol{i} \in \{0,1\}^k} \in \mathbb{F}^{2^k}$. Because $m \geq (k + 2)/c$, Corollary 124 guarantees that $\{\boldsymbol{\zeta}^{(\boldsymbol{j})} \colon \boldsymbol{j} \in \{0,1\}^m\}$ is a 1/4-bias probability space.

Let $c_{\boldsymbol{i}} = h(\boldsymbol{i})$. Then for all $\boldsymbol{j} \in \{0,1\}^m$, it holds that

$$\sum_{\boldsymbol{i} \in \{0,1\}^k, \boldsymbol{l} \in \{0,1\}^t} h_{\boldsymbol{j}}(\boldsymbol{i}, \boldsymbol{l}) = \sum_{\boldsymbol{i} \in \{0,1\}^k} \zeta_{\boldsymbol{i}}^{(\boldsymbol{j})} c_{\boldsymbol{i}}. \tag{A.4}$$

Completeness: Suppose that $c_{\boldsymbol{i}} = 0$ for all $\boldsymbol{i} \in \{0,1\}^k$. Then, by Eq. (A.4), it holds that

$$\sum_{\boldsymbol{i} \in \{0,1\}^k, \boldsymbol{l} \in \{0,1\}^t} h_{\boldsymbol{j}}(\boldsymbol{i}, \boldsymbol{l}) = 0$$

for all $\boldsymbol{j} \in \{0,1\}^m$. Therefore, the completeness of the protocol in Lemma 36 implies that the protocol constructed above also has perfect completeness.

Soundness: Suppose that $\boldsymbol{c} \neq 0$. By Proposition 2, it holds that

$$\Pr_{\boldsymbol{j} \in \{0,1\}^m} \Bigl[\sum_{\boldsymbol{i} \in \{0,1\}^k} \zeta_{\boldsymbol{i}}^{(\boldsymbol{j})} c_{\boldsymbol{i}} = 0 \Bigr] \leq \frac{1 + 1/4}{2} = \frac{5}{8}.$$

Eq. (A.4) and the soundness in Lemma 36 imply that for the choices of $\boldsymbol{j} \in \{0,1\}^m$ such that $\sum_{\boldsymbol{i} \in \{0,1\}^k} \zeta_{\boldsymbol{i}}^{(\boldsymbol{j})} c_{\boldsymbol{i}} \neq 0$, the acceptance probability conditioned on the choice of $\boldsymbol{j}$ is at most $(d + d')(k + t)/|\mathbb{F}|$. Therefore, the overall acceptance probability is at most $5/8 + (d + d')(k + t)/|\mathbb{F}|$. The corollary follows because d' and t are polynomially bounded in k. □

A.3 The orthogonalization lemma

In this section we prove different variants of our orthogonalization lemma. We first give a general statement of the lemma.

Lemma 125. *Let ρ_i, $i = 1, \dots, k$ be positive matrices, and $\rho := \sum_i \rho_i$. Let $P_1, \dots, P_k$ be d-dimensional projectors such that*

$$\sum_{i \neq j} \operatorname{Tr}(P_i P_j P_i \, \rho_i) \leq \varepsilon \qquad \text{and} \qquad \sum_{i \neq j} \operatorname{Tr}(P_i \, \rho_j) \leq \varepsilon$$

for some $0 < \varepsilon \leq \operatorname{Tr}(\rho)$. Then there exists orthogonal projectors $Q_1, \dots, Q_k$ such that

$$\sum_{i=1}^{k} \operatorname{Tr}\big((P_i - Q_i)^2 \, \rho_i\big) = O\big(\varepsilon^{1/2}\big) \operatorname{Tr}(\rho)^{1/2}$$

We first prove Lemma 125. We will then show how Lemma 31 stated in Chapter 4 follows from it. Finally, we will give a different corollary of Lemma 125, Lemma 127 below, which is adapted to our work on parallel repetition presented in Chapter 7.

Proof of Lemma 125. For every i write $P_i = \sum_l |x_{i,l}\rangle\langle x_{i,l}|$, where the $\{|x_{i,l}\rangle\}_l$ are orthonormal, and let $X_i := \sum_l |x_{i,l}\rangle\langle e_{i,l}|$, $X := \sum_i X_i$, where $|e_{i,l}\rangle$ is the canonical basis: X has the $|x_{i,l}\rangle$ as its columns. In order for X to be a square matrix, if necessary we extend the space in which the $|x_{i,l}\rangle$ vectors live, so as to make it the same dimension as $\operatorname{Span}\{|e_{i,l}\rangle\}$. The inner-product condition on the P_i implies that

$$\sum_{i \neq j} \operatorname{Tr}\big(P_i P_j P_i \, \rho_i\big) \;=\; \sum_{i \neq j} \sum_{l, l', l''} \langle x_{i,l} | x_{j,l'} \rangle \langle x_{j,l'} | x_{i,l''} \rangle \langle x_{i,l''} | \rho_i | x_{i,l} \rangle \;\leq\; \varepsilon \tag{A.5}$$

Write $X^\dagger X = \sum_{i,j,l,l'} \langle x_{i,l} | x_{j,l'} \rangle \, |e_{i,l}\rangle\langle e_{j,l'}|$, so that

$$\sum_i \operatorname{Tr}\big((X^\dagger X - \mathrm{Id})^2 \, X_i^\dagger \rho_i X_i\big) \;=\; \sum_{i,l,l''} \; \sum_{(j,l') \neq (i,l),(i,l'')} \langle x_{i,l} | x_{j,l'} \rangle \langle x_{j,l'} | x_{i,l''} \rangle \langle x_{i,l''} | \rho_i | x_{i,l} \rangle \;\leq\; \varepsilon \tag{A.6}$$

where we used (A.5) to upper-bound the expression in the middle by ε. Indeed, in the second summation, if $i = j$ then either $l' \neq l$ or $l' \neq l''$, so that one of the inner products $\langle x_{i,l} | x_{i,l'} \rangle$ or $\langle x_{i,l'} | x_{i,l''} \rangle$ is 0, since the $\{|x_{i,l}\rangle\}_l$ are orthogonal.

Let $X = U \Sigma V^\dagger$, where Σ is diagonal positive and U, V unitary, be the polar decomposition of X. By an appropriate choice of the basis $|e_{i,l}\rangle$ we can assume that $V = \mathrm{Id}$ (if not, re-define $X_i := X_i V$; this corresponds to changing $|e_{i,l}\rangle \to V^\dagger |e_{i,l}\rangle$). Let Π be the projector on the span of the eigenvectors of Σ with corresponding eigenvalue at least $1/2$ and at most 2. Π is needed to control eigenvalues of Σ which may be too small or too large.

Let $\tilde{U} = U\Pi$ and $\tilde{X} = X\Pi$. Let $|\tilde{u}_{i,l}\rangle$ (resp. $|\tilde{x}_{i,l}\rangle$) be the column vectors of $\tilde{U}$ (resp. $\tilde{X}$), so that $\tilde{U} = \sum_{i,l} |\tilde{u}_{i,l}\rangle\langle e_{i,l}|$. We will show that the projectors $Q_i := \sum_l |\tilde{u}_{i,l}\rangle\langle \tilde{u}_{i,l}|$ are close to the projectors P_i, in the sense claimed in the lemma (note that since U is unitary and Π a diagonal projector the Q_i are orthogonal projectors, which do not necessarily sum to identity). We first state some consequences of (A.6).

Fact 126. *The following inequalities holds*

$$\sum_{i,l,l'} \langle \tilde{u}_{i,l} - \tilde{x}_{i,l} | \tilde{u}_{i,l'} - \tilde{x}_{i,l'} \rangle \langle \tilde{x}_{i,l'} | \rho_i | \tilde{x}_{i,l} \rangle \;\leq\; \varepsilon \tag{A.7}$$

$$\sum_{i,l} |\langle \tilde{u}_{i,l} | \rho | \tilde{u}_{i,l} \rangle - \langle \tilde{x}_{i,l} | \rho | \tilde{x}_{i,l} \rangle| \;\leq\; 2\sqrt{2}\,\varepsilon^{1/2}\mathrm{Tr}(\rho)^{1/2} \tag{A.8}$$

Proof. We start with proving (A.7). Since Σ is diagonal, one can immediately check that $X^\dagger X - \mathrm{Id} = (X-U)^\dagger(X+U)$. Note also that $(X+U)(X+U)^\dagger = U(\mathrm{Id}+\Sigma)^2U^\dagger \geq \mathrm{Id}$. Hence

$$\begin{aligned}\sum_i \mathrm{Tr}\big((\Sigma - \mathrm{Id})^2 X_i^\dagger \rho_i X_i\big) &= \sum_i \mathrm{Tr}\big((X-U)^\dagger(X-U)X_i^\dagger \rho_i X_i\big) \\ &\leq \sum_i \mathrm{Tr}\big((X-U)^\dagger(X+U)(X+U)^\dagger(X-U)X_i^\dagger \rho_i X_i\big) \\ &\leq \varepsilon \end{aligned} \tag{A.9}$$

where the last inequality is by (A.6). This implies that $\sum_i \mathrm{Tr}((\Sigma - \mathrm{Id})^2(X_i\Pi)^\dagger\rho_i(X_i\Pi) \leq \varepsilon$ (note that Π commutes with Σ by definition), which is just (A.7).

Before turning to the proof of (A.8), first observe that

$$\begin{aligned}\mathrm{Tr}((\Sigma - \mathrm{Id})^2\Pi X^\dagger \rho X) &= \sum_{i,j} \mathrm{Tr}((\Sigma - \mathrm{Id})^2 \Pi X_i^\dagger \rho_j X_i\big) \\ &\leq 2\varepsilon \end{aligned} \tag{A.10}$$

where the equality uses that $(\Sigma - \mathrm{Id})^2\Pi)$ is diagonal, and the inequality is by (A.7) for the terms $i = j$ and uses $(\Sigma - \mathrm{Id})^2\Pi \leq \mathrm{Id}$ and the second condition in the lemma for the terms $i \neq j$. From (A.10) we get

$$\begin{aligned}\sum_{i,l} \langle \tilde{u}_{i,l} - \tilde{x}_{i,l} | \rho | \tilde{u}_{i,l} - \tilde{x}_{i,l} \rangle &= \mathrm{Tr}(\Pi(X-U)^\dagger \rho(X-U)) \\ &\leq 4\,\mathrm{Tr}\big(\Sigma\Pi\Sigma(X-U)^\dagger\rho(X-U)\big) \\ &= 4\,\mathrm{Tr}\big((\mathrm{Id}-\Sigma)\Pi(\mathrm{Id}-\Sigma)X^\dagger\rho X\big) \\ &\leq 8\varepsilon \end{aligned} \tag{A.11}$$

where the first inequality uses $\Pi\Sigma \geq 1/2\Pi$, by definition of Π, and the last is by (A.10).

We now prove (A.8). By Cauchy-Schwarz, for every (i,l)

$$\langle \tilde{u}_{i,l} - \tilde{x}_{i,l} | \rho | \tilde{u}_{i,l} \rangle \leq \langle \tilde{u}_{i,l} - \tilde{x}_{i,l} | \rho | \tilde{u}_{i,l} - \tilde{x}_{i,l} \rangle^{1/2} \langle \tilde{u}_{i,l} | \rho | \tilde{u}_{i,l} \rangle^{1/2}$$

hence by (A.11) we see that

$$\sum_{i,l} |\langle \tilde{u}_{i,l} - \tilde{x}_{i,l} | \rho | \tilde{u}_{i,l} \rangle| \leq 2\sqrt{2}\,\varepsilon^{1/2}\mathrm{Tr}(\rho)^{1/2}$$

A symmetric inequality can be obtained, and (A.8) follows by the triangle inequality. □

As a consequence of Fact 126, note that

$$\Big|\sum_{i,l,l'}\langle\tilde{u}_{i,l}|\tilde{x}_{i,l'}\rangle\,\langle\tilde{x}_{i,l'}|\rho_i|\tilde{u}_{i,l}-\tilde{x}_{i,l}\rangle\Big| \leq \Big(\sum_{i,l,l'}\langle\tilde{x}_{i,l}|\tilde{x}_{i,l'}\rangle\langle\tilde{x}_{i,l'}|\rho_i|\tilde{x}_{i,l}\rangle\Big)^{1/2}\Big(\sum_{i,l}\langle\tilde{u}_{i,l}-\tilde{x}_{i,l}|\rho_i|\tilde{u}_{i,l}-\tilde{x}_{i,l}\rangle\Big)^{1/2}$$

$$\leq \mathrm{Tr}(\rho)^{1/2}\cdot(8\varepsilon)^{1/2} \;=\; O(\varepsilon^{1/2})\mathrm{Tr}(\rho)^{1/2} \tag{A.12}$$

where the first inequality is by Cauchy-Schwarz (and the $|\tilde{u}_{i,l}\rangle$ being orthonormal) and the second uses $\tilde{X}_i\tilde{X}_i^\dagger \leq \mathrm{Id}$, and (A.11) (with $\rho_i \leq \rho$).

In order to bound the distance between $Q_i = \sum_l |\tilde{u}_{i,l}\rangle\langle\tilde{u}_{i,l}|$ and P_i, we first bound the distance between Q_i and $\tilde{P}_i := \tilde{X}_i\tilde{X}_i^\dagger$:

$$\begin{aligned}\sum_i \mathrm{Tr}\big((\tilde{P}_i-Q_i)^2\,\rho_i\big) &= \sum_{i,l}\big(\langle\tilde{x}_{i,l}|\rho_i|\tilde{x}_{i,l}\rangle+\langle\tilde{u}_{i,l}|\rho_i|\tilde{u}_{i,l}\rangle\big) - 2\sum_{i,l,l'}\Re\big(\langle\tilde{u}_{i,l}|\tilde{x}_{i,l'}\rangle\,\langle\tilde{x}_{i,l'}|\rho_i|\tilde{u}_{i,l}\rangle\big)\\ &\leq 2\sum_{i,l}\langle\tilde{x}_{i,l}|\rho_i|\tilde{x}_{i,l}\rangle - 2\sum_{i,l,l'}\Re\big(\langle\tilde{u}_{i,l}|\tilde{x}_{i,l'}\rangle\,\langle\tilde{x}_{i,l'}|\rho_i|\tilde{x}_{i,l}\rangle\big) + O(\varepsilon^{1/2}\mathrm{Tr}(\rho)^{1/2})\\ &\leq O(\varepsilon^{1/2}\mathrm{Tr}(\rho)^{1/2}) \end{aligned}\tag{A.13}$$

where the first inequality is by (A.8) and (A.12) and the second by (A.7). It remains to bound the distance between the $\tilde{P}_i$ and the P_i:

$$\begin{aligned}\sum_i \mathrm{Tr}\big((\tilde{P}_i-P_i)^2\,\rho_i\big) &= \sum_i \mathrm{Tr}\big((\mathrm{Id}-\Pi)X_i^\dagger\rho_iX_i\big)\\ &\leq 2\sum_i \mathrm{Tr}\big(|\mathrm{Id}-\Sigma|X_i^\dagger\rho_iX_i\big)\\ &\leq 2\Big(\sum_i \mathrm{Tr}\big((\mathrm{Id}-\Sigma)^2X_i^\dagger\rho_iX_i\big)\Big)^{1/2}\Big(\sum_i \mathrm{Tr}\big(X_i^\dagger\rho_iX_i\big)\Big)^{1/2}\\ &\leq 2\varepsilon^{1/2}\mathrm{Tr}(\rho)^{1/2}\end{aligned}\tag{A.14}$$

where the first inequality uses $(\mathrm{Id}-\Pi) \leq 2|\Sigma-\mathrm{Id}|$ by definition of Π, the second is Cauchy-Schwarz and the last is by (A.9). Combining (A.13) and (A.14) finishes the proof of the lemma. □

We now show how Lemma 31 follows from Lemma 125.

Proof of Lemma 31. Given that $\sum_a A^a = \mathrm{Id}$, the assumption (4.12) is equivalent to

$$\sum_{a\neq a'}\mathrm{Tr}\big(\big(\sqrt{A^a}(\mathrm{Id}-A^a)\sqrt{A^a}\big)\,\rho\big) \;\leq\; \varepsilon.$$

For every a, let P_a be the projection on the eigenvalues of A^a larger than $1/2$, and $\rho_a := P_a\rho P_a$. By definition, $A_a \geq (1/2)P_a$. Hence

$$\begin{aligned}\sum_a \mathrm{Tr}\big((P^a-\sqrt{A^a})^2\rho_a\big) &= \sum_a \Big(\mathrm{Tr}\big(A^a\rho_a\big)+\mathrm{Tr}\big(P^a\rho_a\big)-2\mathrm{Tr}\big(P^a\sqrt{A^a}\rho_a\big)\Big)\\ &\leq 2\sum_a \mathrm{Tr}\big(P^a(\mathrm{Id}-A^a)P^a\rho\big)\\ &\leq \sum_a \mathrm{Tr}\big(\sqrt{A^a}(\mathrm{Id}-A^a)\sqrt{A^a}\rho\big)\\ &\leq \varepsilon, \end{aligned}\tag{A.15}$$

where for the first inequality we used that the P^a were projectors, and $A^a \leq \sqrt{A^a}$. Next observe that

$$\begin{aligned}\sum_{a\neq a'} \mathrm{Tr}\big(P^aP^{a'}P^a\rho_a\big) &\leq 2\sum_{a\neq a'} \mathrm{Tr}\big(P^aA^{a'}P^a\rho_a\big)\\ &= 2\sum_{a\neq a'} \mathrm{Tr}\big(P^a(\mathrm{Id}-A^a)P^a\rho\big)\\ &\leq 4\sum_{a\neq a'} \mathrm{Tr}\big(\sqrt{A^a}(\mathrm{Id}-A^a)\sqrt{A^a}\rho\big)\\ &\leq 4\,\varepsilon. \end{aligned}\tag{A.16}$$

Hence the P^a, together with ρ^a, satisfy both assumptions of Lemma 125. The lemma gives us orthogonal projectors Q^a such that

$$\sum_a \mathrm{Tr}\big((P^a-Q^a)^2\rho_a\big) \;=\; O\big(\varepsilon^{1/2}\big),$$

which, combined with (A.15) and the triangle inequality, also shows that

$$\sum_a \mathrm{Tr}\big((\sqrt{A^a}-Q^a)^2\rho_a\big) \;=\; O\big(\varepsilon^{1/2}\big).\tag{A.17}$$

From (A.16) we know that $\sum_a \mathrm{Tr}(A^a\rho_a) \geq 1-O(\varepsilon)$, so that by the Cauchy-Schwarz inequality we also get $\sum_a \mathrm{Tr}(Q^a\rho_a) \geq 1-O\big(\varepsilon^{1/4}\big)$, and since the Q^a are orthogonal and $\sum_a \mathrm{Tr}(\rho_a)=1$, this implies

$$\sum_{a\neq a'} \mathrm{Tr}\big(Q^a\rho_{a'}\big) = O\big(\varepsilon^{1/4}\big).$$

Together with the same equation for A^a instead of Q^a, obtained from (A.16), and (A.17), we obtain

$$\sum_a \mathrm{Tr}\big((\sqrt{A^a}-Q^a)^2\rho\big) \;=\; O\big(\varepsilon^{1/4}\big).\tag{A.18}$$

To conclude it suffices to make the $\{Q^a\}$ into a projective measurement $\{B^a\}$, by setting $B^{a_0} := Q^{a_0} + (\mathrm{Id} - \sum_a Q^a)$ for some arbitrary a_0, and $B^a := Q^a$ for $a \neq a_0$. From the fact that $\sum_a A^a = \mathrm{Id}$ and (A.18) it is not hard to see that $\mathrm{Tr}((\mathrm{Id} - \sum_a Q^a)\rho) = O(\varepsilon^{1/4})$, hence (A.18) still holds with the B^a in place of the Q^a, and the lemma is proved. □

In Lemma 127 stated below, one can think of the $\hat{Y}_i$ as operators in the Stinespring representation of a measurement $\mathcal{M}_i : \rho \mapsto \hat{Y}_i(\rho \otimes \mathrm{Id})\hat{Y}_i^\dagger$, where i refers to the i-th outcome of the measurement. In that setting the hypothesis of the lemma is that, when $\mathcal{M}$ is performed twice sequentially on a specific state ρ, it is likely that identical answers will be obtained. The conclusion is that the operators $\hat{Y}_i$ have an approximate joint block-diagonal form, as described by the orthogonal projectors Π_i.

Lemma 127. *There is a $c > 0$ such that the following holds. Let ρ_i, $i = 1, \ldots, k$ be positive, ρ such that $\sum_i \rho_i \leq \rho$ and $\hat{Y}_i$, $i = 1, \ldots, k$ (possibly rectangular) matrices, be such that*

$$\sum_{i \neq j} \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger\, (\hat{Y}_j\, \hat{Y}_j^\dagger)\, \hat{Y}_i\big) \leq \alpha\, \mathrm{Tr}(\rho) \tag{A.19}$$

and $\sum_i \hat{Y}_i \hat{Y}_i^\dagger \leq \mathrm{Id}$. Then there exists orthogonal projectors $\{\Pi_i\}$ such that

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger(\mathrm{Id} - \Pi_i)\hat{Y}_i\big) \leq O\big(\alpha^c\big)\mathrm{Tr}(\rho).$$

Proof. The idea of the proof is simple. Let $\beta_1, \beta_2 > 0$ be parameters to be chosen later. For every i, let P_i be the projector on the eigenvectors of $\hat{Y}_i\, \hat{Y}_i^\dagger$ with corresponding eigenvalue at least β_1. Since P_i contains all the large eigenvalues, $P_i\hat{Y}_i \approx \hat{Y}_i$. Moreover, by definition $P_i \leq \beta_1^{-1}\hat{Y}_i\, \hat{Y}_i^\dagger$. These two properties together with (A.19) *almost* imply that $\sum_{i\neq j} \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger P_i\, P_j\, P_i \hat{Y}_i\big) \lesssim \beta^{-1}\alpha \mathrm{Tr}(\rho)$. Choosing $\beta_1 \approx \sqrt{\alpha}$, we could then apply Lemma 125 to the P_i and states $\sigma_i := \hat{Y}_i\rho_i\hat{Y}_i^\dagger$, recovering close orthogonal projectors Π_i which would satisfy the required condition. Carrying out this intuition precisely is a bit tedious, and we now proceed to the details. We will use the following simple fact.

Fact 128. *Let $A \geq 0$, $\rho \geq 0$, and Π a projection. Let*

$$a = \mathrm{Tr}_\rho(A), \qquad b = |\mathrm{Tr}_\rho((\mathrm{Id} - \Pi)A\Pi)| \qquad \textit{and} \qquad c = \mathrm{Tr}_\rho\big((\mathrm{Id} - \Pi)A(\mathrm{Id} - \Pi)\big).$$

Then both the following hold

$$\mathrm{Tr}_\rho\big(\Pi A \Pi\big) \leq \big(\sqrt{a} + \sqrt{c}\big)^2 \leq 2(a + c),$$
$$\mathrm{Tr}_\rho\big(\Pi A \Pi\big) \leq \Big(\frac{\sqrt{a} + \sqrt{a + 4b}}{2}\Big)^2 \leq a + 2b.$$

Proof. Write $\Pi = (\Pi - \mathrm{Id}) + \mathrm{Id}$, so $\mathrm{Tr}_\rho(\Pi A \Pi) \le |\mathrm{Tr}_\rho((\Pi - \mathrm{Id})A\Pi)| + |\mathrm{Tr}_\rho(A\Pi)|$. The second term can be bounded by $a^{1/2}\mathrm{Tr}_\rho(\Pi A\Pi)^{1/2}$ by Cauchy-Schwarz. Similarly bounding the first term by $c^{1/2}\mathrm{Tr}_\rho(\Pi A\Pi)^{1/2}$ yields the first equation. To get the second, let $X = \mathrm{Tr}_\rho(\Pi A \Pi)^{1/2}$ to obtain the equation

$$X^2 - a^{1/2}X - b \le 0.$$

Solving and using $X \ge 0$, one finds that this is equivalent to $X \le (\sqrt{a} + \sqrt{a+4b})/2$. □

Let $Y_{-i} := \sum_{j\neq i} \hat{Y}_j \hat{Y}_j^\dagger \le \mathrm{Id}$, and Q_i be the projector on the eigenvectors of $P_i Y_{-i} P_i$ with eigenvalue at most β_2. Note that, by definition, $Q_i \le P_i \le \beta_1^{-1}\hat{Y}_i\hat{Y}_i^\dagger$ (and in particular Q_i commutes with P_i). We first bound the distance between $\hat{Y}_i^\dagger$ and $\hat{Y}_i^\dagger Q_i$: since $\hat{Y}_i^\dagger(\mathrm{Id} - Q_i) = \hat{Y}_i^\dagger(\mathrm{Id} - P_i) + \hat{Y}_i^\dagger P_i(\mathrm{Id} - Q_i)P_i$,

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger(\mathrm{Id} - Q_i)\hat{Y}_i\big) = \sum_i \Big(\mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger(Id - P_i)\hat{Y}_i\big) + \mathrm{Tr}_{\rho_i}(\hat{Y}_i^\dagger P_i(Id - Q_i)P_i\hat{Y}_i)\Big) \quad \text{(A.20)}$$

The first term is easily bounded by $\beta_1 \,\mathrm{Tr}(\rho)$. For the second, note that $P_i(Id - Q_i)P_i \le \beta_2^{-1}P_iY_{-i}P_i$. Using Fact 128 with $A^i = Y_{-i}$, $\Pi^i = P_i$, and $\rho^i = \hat{Y}_i\rho_i(\hat{Y}_i)^\dagger$ we get $\sum_i a^i \le \alpha\mathrm{Tr}(\rho)$ and $\sum_i c^i \le \beta_1\mathrm{Tr}(\rho)$, so that

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger P_i Y_{-i} P_i \hat{Y}_i\big) \le 2(\alpha + \beta_1)\mathrm{Tr}(\rho)$$

Assuming $\alpha \le \beta_1$ (which will hold for our choice of parameters), from (A.20) we get

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger(\mathrm{Id} - Q_i)\hat{Y}_i\big) \;\le\; O(\beta_2^{-1}\beta_1)\mathrm{Tr}(\rho). \quad \text{(A.21)}$$

Next observe that, by definition of Q_i, followed by an application of the Cauchy-Schwarz inequality,

$$\begin{aligned}\sum_i \big|\mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger Q_i Y_{-i}(\mathrm{Id} - Q_i)\hat{Y}_i\big)\big| &= \sum_i \big|\mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger Q_i Y_{-i}(\mathrm{Id} - P_i)\hat{Y}_i\big)\big| \\ &\le \Big(\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger(\mathrm{Id} - P_i)\hat{Y}_i\big)\Big)^{1/2}\Big(\sum_i \mathrm{Tr}_{\rho_i}(\hat{Y}_i^\dagger Q_i Y_{-i}^2 Q_i \hat{Y}_i)\Big)^{1/2} \\ &\le \beta_1^{1/2}\beta_2\mathrm{Tr}(\rho), \qquad \text{(A.22)}\end{aligned}$$

where we used $Q_iY_{-i}^2Q_i \le \beta_2^2\mathrm{Id}$, which holds by definition of Q_i, to bound the second term in the last inequality. Using the second bound in Fact 128 with $A^i = Y_{-i}$, $\Pi^i = Q_i$, $\rho^i = \hat{Y}_i\rho_i(\hat{Y}_i)^\dagger$, we get $\sum_i a^i \le \alpha\mathrm{Tr}(\rho)$ and $\sum_i b^i \le \beta_1^{1/2}\beta_2\mathrm{Tr}(\rho)$ by (A.22), so that

$$\begin{aligned}\sum_{i\neq j} \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger Q_i Q_j Q_i \hat{Y}_i\big) &\le \beta_1^{-1}\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger Q_i Y_{-i} Q_i \hat{Y}_i\big) \\ &\le \beta_1^{-1}\big(\alpha + 2\beta_1^{1/2}\beta_2\big)\mathrm{Tr}(\rho).\end{aligned}$$

Set $\beta_2 = \beta_1^{3/4}$ and $\beta_1 = \alpha^{4/5}$ to obtain

$$\sum_{i \neq j} \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger Q_i Q_j Q_i \hat{Y}_i\big) \leq O(\alpha^{1/5})\, \mathrm{Tr}(\rho). \tag{A.23}$$

Let $\sigma_i := \hat{Y}_i \rho_i \hat{Y}_i^\dagger$. We are now ready to apply Lemma 125 to the Q_i and σ_i: the first condition holds by (A.23), and the second is a direct consequence of (A.19) and $Q_j \leq \beta_1^{-1} \hat{Y}_j \hat{Y}_j^\dagger$ for every j. The lemma then gives us pairwise orthogonal Π_i such that

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger (Q_i - \Pi_i)^2 \hat{Y}_i\big) \leq O(\alpha^{1/10}) \mathrm{Tr}(\rho).$$

Combined with (A.21) and the triangle inequality, this leads to

$$\sum_i \mathrm{Tr}_{\rho_i}\big(\hat{Y}_i^\dagger (\mathrm{Id} - \Pi_i) \hat{Y}_i\big) \leq O(\alpha^{1/10})\, \mathrm{Tr}(\rho).$$

□

A.4 Omitted proofs from Chapter 7

In this section we give a series of useful claims showing that, in a strategy which has been marginalized over a large number of indices, fixing a particular coordinate (i, q_i) does not have much influence on average. Throughout this question we fix a question set Q and a distribution μ on Q. Whenever an expectation over tuples of questions $q \in Q^C$ is taken, it will be over the product distribution μ^C.

Our claims will rely essentially on the following, which applies to *any* matrix semi-norm $\| \cdot \|$, provided it is derived from a semi-inner product $\langle \cdot, \cdot \rangle$.

Claim 129. *Let C be an integer, and $f : Q^C \to \{\, X \in \mathbb{C}^{d \times d} \,\}$. Let $M = \mathrm{E}_q\left[f(q)\right]$ and for any (i, q_i), $M_{i,q_i} = \mathrm{E}_{q_{\neg i}}\left[f(q)\right]$. Suppose that $\mathrm{E}_q\left[\|f(q)\|^2\right] \leq 1$. Then*

1. $0 \leq \mathrm{E}_{i,q_i}\left[\|M - M_{i,q_i}\|^2\right] \leq \frac{\mathrm{E}_q\left[\|f(q)\|^2\right]}{C} \leq \frac{1}{C}$.
2. $\mathrm{E}_{i,q_i}\left[\|M - M_{i,q_i}\|^2\right] = \mathrm{E}_{i,q_i}\left[\|M_{i,q_i}\|^2\right] - \|M\|^2$.
3. $\Pr_{i,q_i}(|\mathrm{Tr}(M) - \mathrm{Tr}(M_{i,q_i})| \geq C^{-1/3}) \leq C^{-1/3}$.

Proof. The proof of all three parts is in close analogy to that of Lemma 2.1 in [88], which shows similar statements for a *Boolean* function f. For part 1 note that $\mathrm{E}_{i,q_i}\left[\|M - M_{i,q_i}\|^2\right] =$

$\frac{1}{C}\sum_{i=1}^{C} \mathrm{E}_{q_i}\left[\|M - M_{i,q_i}\|^2\right]$ and hence it suffices to show that $\sum_{i=1}^{C} \mathrm{E}_{q_i}\left[\|M - M_{i,q_i}\|^2\right] \leq \mathrm{Tr}(M)$. Observe that

$$
\begin{aligned}
0 &\leq \mathrm{E}_q\left[\|f(q) - \sum_i (M_{i,q_i} - M)\|^2\right] \\
&= \mathrm{E}_q\left[\|f(q)\|^2\right] - \sum_i \mathrm{E}_{q_i}\left[\langle M_{i,q_i} - M, M_{i,q_i}\rangle + \langle M_{i,q_i}, M_{i,q_i} - M\rangle\right] + \sum_{i,j} \mathrm{E}_{q_i,q_j}\Big[\langle M - M_{i,q_i}, M - M_{j,q_j}\rangle \\
&= \mathrm{E}_q\left[\|f(q)\|^2\right] - \sum_i \mathrm{E}_{q_i}\left[\|M - M_{i,q_i}\|^2\right],
\end{aligned}
$$

where for the last equality we have used that $\mathrm{E}_{q_i}\left[M_{i,q_i} - M\right] = 0$ and hence $\mathrm{E}_{q_i}\left[\langle M_{i,q_i} - M, M_{i,q_i}\rangle\right] = \mathrm{E}_{q_i}\left[\langle M_{i,q_i} - M, M_{i,q_i} - M\rangle\right]$ and, for $i \neq j$,

$$
\mathrm{E}_{q_i,q_j}\left[\langle M - M_{i,q_i}, M - M_{j,q_j}\rangle\right] = \langle \mathrm{E}_{q_i}\left[M - M_{i,q_i}\right], \mathrm{E}_{q_j}\left[M - M_{j,q_j}\right]\rangle = 0
$$

Part 1. now follows, and the second inequality is simply the assumption that $\mathrm{E}_q\left[\|f(q)\|^2\right] \leq 1$.

Part 2 is trivial from the expansion of $\|M - M_{i,q_i}\|^2$. Part 3 follows from part 1 using Markov's inequality, which gives $\Pr_{i,q_i}((\mathrm{Tr}(M - M_{i,q_i}))^2 \geq C^{-2/3}) \leq C^{2/3}\mathrm{E}_{i,q_i}\left[(\mathrm{Tr}(M - M_{i,q_i}))^2\right]$. Observing that for $A := M - M_{i,q_i}$ we have $(\mathrm{Tr}(A))^2 = \langle A, \mathrm{Id}\rangle^2 \leq \|A\|^2 \cdot \|\mathrm{Id}\|^2 = \|A\|^2$ gives the desired bound. □

The following is a direct corollary of Claim 129, obtained for a specific instantiation of the norm $\|\cdot\|$.

Claim 130. *Let Y_q^a, for $q \in Q^C$ and $a \in A^C$, be positive matrices such that $Y_q := \sum_a Y_q^a \leq \mathrm{Id}$, and $\rho \geq 0$. Let $Y = \mathrm{E}_q\left[Y_q\right]$. Then*

$$
\mathrm{E}_{(i,q_i)}\left[\left|\mathrm{Tr}\big(Y\rho^{1/2}Y\rho^{1/2}\big) - \mathrm{Tr}\big(Y_{q_i}\rho^{1/2}Y_{q_i}\rho^{1/2}\big)\right|\right] \leq C^{-1}\mathrm{E}_q\left[\mathrm{Tr}\big(Y_q\rho^{1/2}Y_q\rho^{1/2}\big)\right] \leq \mathrm{Tr}_\rho(Y).
$$

Proof. The statement follows from Claim 129, applied to $f(q) = Y_q$ and the (semi)-norm $\|A\|^2 = \mathrm{Tr}\big(A\rho^{1/2}A^\dagger\rho^{1/2}\big)$, which is derived from the inner-product $(A, B) \mapsto \mathrm{Tr}\big(A\rho^{1/2}B^\dagger\rho^{1/2}\big)$. The second inequality holds since $0 \leq Y_q \leq \mathrm{Id}$ for every q. □

We following simple calculation will be useful.

Claim 131. *Let $Y_q \in \mathbb{C}^{d\times d}$, $0 \leq Y_q \leq \mathrm{Id}$, for $q \in Q^C$, and let $Y = \mathrm{E}_q\left[Y_q\right]$, $Y_{i,q_i} = \mathrm{E}_{q_{\neg i}}\left[Y_q\right]$ for $i \in [C]$. Then*

$$
\mathrm{E}_{(i,q_i)}\left[(Y - Y_{i,q_i})^2\right] \leq C^{-1}\mathrm{E}_q\left[Y_q^2\right].
$$

Proof. Write

$$\begin{aligned}0 &\leq \Big(Y_q - \sum_i (Y_{i,q_i} - Y)\Big)\Big(Y_q - \sum_i (Y_{i,q_i} - Y)\Big)\\ &= Y_q^2 - \sum_i \big(Y_q(Y_{i,q_i} - Y) + \big(Y_{i,q_i} - Y)Y_q\big) + \sum_{i,j} \big(Y_{i,q_i} - Y\big)\big(Y_{j,q_j} - Y\big)\end{aligned}$$

Taking the expectation over q, we obtain

$$\sum_i \mathrm{E}_{q_i}\left[(Y_{i,q_i} - Y)^2\right] \;\leq\; \mathrm{E}_q\big[Y_q^2\big]$$

Dividing by C on both sides proves the claim. □

Claim 132. *For every $q \in Q^C$ let $\{X_q^a\}_{a \in A^{C'}}$ be a POVM, and $\hat{X}_q^a := \sqrt{\pi(q)}\sqrt{X_q^a} \otimes \langle q, a|$ (as described in Section 7.2.1), and $\rho \geq 0$. Assume that $\hat{X}\hat{X}^\dagger = \sum_a \mathrm{E}_q\left[\hat{X}_q^a(\hat{X}_q^a)^\dagger\right] \leq \mathrm{Id}$. Then*

$$\sum_a \mathrm{E}_{(i,q_i)}\left[\Big|\mathrm{Tr}_\rho\big((\hat{X}^a)^\dagger \hat{X}^a (\hat{X}^a)^\dagger \hat{X}^a\big) - \mathrm{Tr}_\rho\big((\hat{X}_{q_i}^a)^\dagger \hat{X}_{q_i}^a (\hat{X}_{q_i}^a)^\dagger \hat{X}_{q_i}^a\big)\Big|\right] \;\leq\; 2\,C^{-1/2}\mathrm{Tr}(\rho).$$

Proof. Let $\tilde{X}_i^a = \big|\hat{X}^a(\hat{X}^a)^\dagger - \hat{X}_{q_i}^a(\hat{X}_{q_i}^a)^\dagger\big|$, and $\tilde{\rho}_i^a = \big|\hat{X}^a\rho(\hat{X}^a)^\dagger - \hat{X}_{q_i}^a\rho(\hat{X}_{q_i}^a)^\dagger\big|$, where the notation keeps the dependence on q_i implicit. Use the triangle inequality to write

$$\big|\mathrm{Tr}\big(\hat{X}^a(\hat{X}^a)^\dagger\hat{X}^a\rho(\hat{X}^a)^\dagger\big) - \mathrm{Tr}\big(\hat{X}_{q_i}^a(\hat{X}_{q_i}^a)^\dagger\hat{X}_{q_i}^a\rho(\hat{X}_{q_i}^a)^\dagger\big)\big| \leq \mathrm{Tr}\big(\tilde{X}_i^a\hat{X}^a\rho(\hat{X}^a)^\dagger\big) + \mathrm{Tr}\big(\hat{X}_{q_i}^a(\hat{X}_{q_i}^a)^\dagger\tilde{\rho}_i^a\big) \tag{A.24}$$

The expectation of the first term on the right-hand side of (A.24) can be bounded by Cauchy-Schwarz as

$$\begin{aligned}\mathrm{E}_{(i,q_i)}\Big[\mathrm{Tr}\big(\tilde{X}_i^a\hat{X}^a\rho(\hat{X}^a)^\dagger\big)\Big] &\leq \mathrm{E}_{(i,q_i)}\Big[\mathrm{Tr}_\rho((\hat{X}^a)^\dagger\hat{X}^a)^{1/2}\mathrm{Tr}\big((\tilde{X}_i^a)^2\hat{X}^a\rho(\hat{X}^a)^\dagger\big)^{1/2}\Big]\\ &\leq C^{-1/2}\mathrm{Tr}_\rho((\hat{X}^a)^\dagger\hat{X}^a)\end{aligned}$$

by Claim 129, applied to the (semi)-norm $\|A\|^2 := \mathrm{Tr}\big((A^\dagger A)\,(\hat{X}^a\rho(\hat{X}^a)^\dagger)\big)$ and the mapping $f : q \mapsto \hat{X}_q^a(\hat{X}_q^a)^\dagger$.

Regarding the second term on the right-hand side of (A.24), let A be the block-column matrix with blocks $\sqrt{\pi(q_i)}\tilde{\rho}_i^a$ for every (i, q_i) and a, and B with blocks $\sqrt{\pi(q_i)}\hat{X}_i^a(\hat{X}_i^a)^\dagger$. Then $B^\dagger B = \sum_a \mathrm{E}_{(i,q_i)}\Big[\big(\hat{X}_i^a(\hat{X}_i^a)^\dagger\big)^2\Big] \leq \mathrm{Id}$. Let $D = A^\dagger B = \sum_a \mathrm{E}_{(i,q_i)}\Big[\tilde{\rho}_i^a\hat{X}_i^a(\hat{X}_i^a)^\dagger\Big]$; the operator Cauchy-Schwarz inequality from Claim 115 gives

$$DD^\dagger \;\leq\; D(B^\dagger B)^{-1}D^\dagger \;\leq\; A^\dagger A = \sum_a \mathrm{E}_{(i,q_i)}\left[(\tilde{\rho}_i^a)^2\right]$$

Applying Claim 131 to $\hat{X}_q^a \rho (\hat{X}_q^a)^\dagger$ (for every a), we can then bound

$$DD^\dagger \leq C^{-1} \mathrm{E}_q \left[(\hat{X}_q \rho \hat{X}_q^\dagger)^2 \right] \leq C^{-1} \mathrm{E}_q \left[\hat{X}_q \rho^2 \hat{X}_q^\dagger \right] \tag{A.25}$$

where for the second inequality we used $\hat{X}_q^\dagger \hat{X}_q \leq \mathrm{Id}$. Since $\mathrm{Tr}(D) \leq \mathrm{Tr}\big(\sqrt{DD^\dagger}\big) = \|D\|_1$, taking the square root on both sides of (A.25) (the square root being operator monotone) and then the trace, we obtain

$$\sum_a \mathrm{E}_{(i,q_i)} \left[\mathrm{Tr}\big(\tilde{\rho}_i^a \hat{X}_i^a (\hat{X}_i^a)^\dagger \big) \right] \leq C^{-1/2} \mathrm{Tr} \sqrt{ \mathrm{E}_q \left[\hat{X}_q \rho^2 \hat{X}_q^\dagger \right] } = C^{-1/2} \big\| \hat{X} \rho \big\|_1$$

where $\hat{X}$ is the rectangular matrix with square blocks $\pi(q)^{-1/2} \hat{X}_q^a$ arranged in a column. By Holder's inequality $\big\|\hat{X}\rho\big\|_1 \leq \mathrm{Tr}(\rho) \|\hat{X}\|_\infty$, and $\|\hat{X}\|_\infty \leq 1$ since $\hat{X}^\dagger \hat{X} = \mathrm{E}_q \left[\hat{X}_q^\dagger \hat{X}_q \right] \leq \mathrm{Id}$. This finishes the proof of the claim. □

A.5 More on extractors

This appendix contains additional results related to Chapter 8, as well as some technical proofs omitted from that chapter. We first develop a bit more the general theory of extractors. In A.5.1 we define extractors for weakly random seeds, and in A.5.2 we show how to compose extractors to obtain more randomness from the same source. In A.5.3 we give technical lemmas: several min-entropy chain rules and the details of the reduction from Trevisan's construction to the underlying one-bit extractor. Section A.5.4 contains all previously known constructions for one-bit extractors and weak designs which we use in this work and plug into Trevisan's extractor. Finally, in A.5.5 we give a proof that list-decodable codes are one-bit extractors.

A.5.1 Weak random seed

In 8.2.1 we defined extractors as functions which take a uniformly random seed. This is the most common way of defining them, but not a necessary condition. Instead we can consider extractors which use a seed which is only weakly random, but with a bounded min-entropy. We extend Definition 76 this way.

Definition 133 (strong extractor with weak random seed). *A function* Ext : $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* (k, ε)-strong extractor with an s-bit seed, *if for all distributions* X *with* $H_{\min}(X) \geq k$ *and any seed* Y *independent from* X *with* $H_{\min}(Y) \geq s$, *we have*

$$\frac{1}{2} \|\rho_{\mathrm{Ext}(X,Y)Y} - \rho_{U_m} \otimes \rho_Y\|_{\mathrm{tr}} \leq \varepsilon,$$

where ρ_{U_m} *is the fully mixed state on a system of dimension* 2^m.

If quantum side information about the input is present in a system E, then as before, we require the seed and the output to be independent from that side-information.

Definition 134 (quantum-proof strong extractor with weak random seed)**.** *A function* Ext : $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* quantum-proof (k,ε)-strong extractor with an s-bit seed*, if for all states ρ_{XE} classical on X with $H_{\min}(X|E)_\rho \geq k$, and for any seed Y independent from XE with $H_{\min}(Y) \geq s$, we have*

$$\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YE} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} \leq \varepsilon,$$

where ρ_{U_m} is the fully mixed state on a system of dimension 2^m.

Lemma 78 says that any extractor will work with roughly the same parameters when classical side information about the input X is present. The same holds in the case of classical side information Z about the seed Y.

Lemma 135. *Let* Ext : $\{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *be a quantum-proof (k,ε)-strong extractor with an s-bit seed. Then for any classical X, Y and Z, and quantum E, such that XE and Y are independent, $Y \leftrightarrow Z \leftrightarrow E$ form a Markov chain,*[1] *$H_{\min}(Y|Z) \geq s + \log 1/\varepsilon$, and for all $z \in \mathcal{Z}$, $H_{\min}(X|EZ = z) \geq k$, we have*

$$\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YZE} - \rho_U \otimes \rho_{YZE}\|_{\mathrm{tr}} \leq 2\varepsilon.$$

Proof. For any two classical systems Y and Z, we have

$$2^{-H_{\min}(Y|Z)} = \mathbf{E}_{z\leftarrow Z}\left[2^{-H_{\min}(Y|Z=z)}\right],$$

so by Markov's inequality,

$$\Pr_{z\leftarrow Z}\left[H_{\min}(Y|Z=z) \leq H_{\min}(Y|Z) - \log 1/\varepsilon\right] \leq \varepsilon.$$

And since $Y \leftrightarrow Z \leftrightarrow E$ form a Markov chain, we have for all $z \in \mathcal{Z}$,

$$\rho_{YE|Z=z} = \rho_{Y|Z=z} \otimes \rho_{E|Z=z}.$$

Hence

$$\begin{aligned}
&\frac{1}{2}\|\rho_{\mathrm{Ext}(X,Y)YEZ} - \rho_U \otimes \rho_{YEZ}\|_{\mathrm{tr}} \\
&\qquad = \frac{1}{2}\sum_{z\in\mathcal{Z}} P_Z(z)\|\rho_{\mathrm{Ext}(X,Y)YE|Z=z} - \rho_U \otimes \rho_{YE|Z=z}\|_{\mathrm{tr}} \\
&\qquad = \frac{1}{2}\sum_{z\in\mathcal{Z}} P_Z(z)\|\rho_{\mathrm{Ext}(X,Y)YE|Z=z} - \rho_U \otimes \rho_{Y|Z=z} \otimes \rho_{E|Z=z}\|_{\mathrm{tr}} \leq 2\varepsilon.
\end{aligned}$$

The case of quantum side information correlated to both the input and the seed is out of the scope of this work.

[1] A ccq state ρ_{XYE} forms a Markov chain $X \leftrightarrow Y \leftrightarrow E$ if it can be expressed as $\rho_{XYE} = \sum_{x,y} P_{XY}(x,y)|x,y\rangle\langle x,y|x,y \otimes \rho_E^y$.

A.5.2 Composing extractors

If an extractor does not have optimal entropy loss, a useful approach to extract more entropy is to apply a second extractor to the original input, trying to extract the randomness that remains when the output of the first extractor is known. This was first proposed in the classical case by Wigderson and Zuckerman [128], and improved by Raz et al. [98]. König and Terhal [76] gave the first quantum version for composing m times quantum 1-bit extractors. We slightly generalize the result of König and Terhal [76] to the composition of arbitrary quantum extractors.

Lemma 136. *Let* $\mathrm{Ext}_1 : \{0,1\}^n \times \{0,1\}^{d_1} \to \{0,1\}^{m_1}$ *and* $\mathrm{Ext}_2 : \{0,1\}^n \times \{0,1\}^{d_2} \to \{0,1\}^{m_2}$ *be quantum-proof* (k, ε_1)*- and* $(k - m_1, \varepsilon_2)$*-strong extractors. Then the composition of the two, namely*

$$\begin{aligned}\mathrm{Ext}_3 :&\{0,1\}^n \times \{0,1\}^{d_1} \times \{0,1\}^{d_2} \to \{0,1\}^{m_1} \times \{0,1\}^{m_2}\\ &(x, y_1, y_2) \mapsto (\mathrm{Ext}_1(x, y_1), \mathrm{Ext}_2(x, y_2)),\end{aligned}$$

is a quantum-proof $(k, \varepsilon_1 + \varepsilon_2)$*-strong extractor.*

Proof. We need to show that for any state ρ_{XE} with $H_{\min}(X|E) \geq k$,

$$\frac{1}{2}\|\rho_{\mathrm{Ext}_1(X,Y_1)\,\mathrm{Ext}_2(X,Y_2)Y_1Y_2E} - \rho_{U_1} \otimes \rho_{U_2} \otimes \rho_{Y_1} \otimes \rho_{Y_2} \otimes \rho_E\|_{\mathrm{tr}} \leq \varepsilon_1 + \varepsilon_2. \tag{A.26}$$

The left-hand side of Eq. (A.26) can be upper-bounded by

$$\begin{aligned}\frac{1}{2}\|\rho_{\mathrm{Ext}_1(X,Y_1)Y_1E} \otimes \rho_{U_2} \otimes \rho_{Y_2} - \rho_{U_1} \otimes \rho_{Y_1} \otimes \rho_E \otimes \rho_{U_2} \otimes \rho_{Y_2}\|_{\mathrm{tr}}&\\ + \frac{1}{2}\|\rho_{\mathrm{Ext}_2(X,Y_2)Y_2\,\mathrm{Ext}_1(X,Y_1)Y_1E} - \rho_{U_2} \otimes \rho_{Y_2} \otimes \rho_{\mathrm{Ext}_1(X,Y_1)Y_1E}\|_{\mathrm{tr}}.&\end{aligned} \tag{A.27}$$

By the definition of Ext_1 the first term in Eq. (A.27) is upper-bounded by ε_1. For the second term we use Lemma 139 and get

$$\begin{aligned}H_{\min}(X|\,\mathrm{Ext}_1(X,Y_1)Y_1E) \geq H_{\min}(X|Y_1E) - H_0(\mathrm{Ext}_1(X,Y_1))&\\ = H_{\min}(X|E) - H_0(\mathrm{Ext}_1(X,Y_1)) \geq k - m_1.&\end{aligned}$$

By the definition of Ext_2 the second term in Eq. (A.27) can then be upper-bounded by ε_2. □

A.5.3 Technical lemmas

Min-entropy chain rules.

We use the following "chain-rule type" statement about the min-entropy. The proofs for the two first can be found in [100].

Lemma 137 ([100, Lemma 3.1.10]). *For any state ρ_{ABC},*

$$H_{\min}(A|BC) \geq H_{\min}(AC|B) - H_0(C),$$

where $H_0(C) = \log \operatorname{rank} \rho_C$.

Lemma 138 ([100, Lemma 3.1.9]). *For any state ρ_{ABZ} classical on Z,*

$$H_{\min}(AZ|B) \geq H_{\min}(A|B).$$

Lemma 139. *For any state ρ_{ABZ} classical on Z,*

$$H_{\min}(A|BZ) \geq H_{\min}(A|B) - H_0(Z),$$

where $H_0(Z) = \log \operatorname{rank} \rho_Z$.

Proof. Immediate by combining Lemma 137 and Lemma 138. □

Reduction step.

To show that a player who can distinguish the output of Ext_C (defined in Definition 83 on page 131) from uniform can also guess the output of the extractor C, we first show that such a player can guess one of the bits of the output of Ext_C given some extra classical information. This is a quantum version of a result by Yao [132].

Lemma 140. *Let ρ_{ZB} be a cq-state, where Z is a random variable on m-bit strings. If $\|\rho_{ZB} - \rho_{U_m} \otimes \rho_B\|_{\text{tr}} > \varepsilon$, then there exists an $i \in [m]$ such that*

$$\| \sum_{\substack{z \in \mathcal{Z} \\ z_i = 0}} p_z |z_{[i-1]}\rangle\langle z_{[i-1]}|z_{[i-1]} \otimes \rho_B^z - \sum_{\substack{z \in \mathcal{Z} \\ z_i = 1}} p_z |z_{[i-1]}\rangle\langle z_{[i-1]}|z_{[i-1]} \otimes \rho_B^z \|_{\text{tr}} > \frac{\varepsilon}{m}. \tag{A.28}$$

Using the fact that for any *binary* random variable X and quantum system Q with $\rho_{XQ} = \sum_{i=0,1} p_i |i\rangle\langle i|i \otimes \rho_Q^i$, the following equality holds: $\|\rho_{XQ} - \rho_{U_1} \otimes \rho_Q\|_{\text{tr}} = \|p_0 \rho_Q^0 - p_1 \rho_Q^1\|_{\text{tr}}$, Eq. (A.28) can be rewritten as $\|\rho_{Z_{i[i-1]}B} - \rho_{U_1} \otimes \rho_{Z_{[i-1]}B}\|_{\text{tr}} > \frac{\varepsilon}{m}$. Lemma 140 can thus be interpreted as saying that if a player holding B can distinguish Z from uniform with probability greater than ε, then there exists a bit $i \in [m]$ such that when given the previous $i-1$ bits of Z, he can distinguish the i^{th} bit of Z from uniform with probability greater than $\frac{\varepsilon}{m}$.

Proof. The proof uses a hybrid argument. Let

$$\sigma_i = \sum_{\substack{z \in \mathcal{Z} \\ r \in \{0,1\}^m}} \frac{p_z}{2^m} |z_{[i]}, r_{\{i+1,\ldots,m\}}\rangle\langle z_{[i]}, r_{\{i+1,\ldots,m\}}|z_{[i]}, r_{\{i+1,\ldots,m\}} \otimes \rho_B^z.$$

Then

$$\begin{aligned}\varepsilon &< \|\rho_{ZB} - \rho_{U_m} \otimes \rho_B\|_{\mathrm{tr}} \\ &= \|\sigma_m - \sigma_0\|_{\mathrm{tr}} \\ &\leq \sum_{i=1}^{m} \|\sigma_i - \sigma_{i-1}\|_{\mathrm{tr}} \\ &\leq m \max_i \|\sigma_i - \sigma_{i-1}\|_{\mathrm{tr}}.\end{aligned}$$

By rearranging $\|\sigma_i - \sigma_{i-1}\|_{\mathrm{tr}}$ we get the lhs of Eq. (A.28). □

We now need to bound the size of this extra information, the "previous $i-1$ bits", and show that when averaging over all the seeds of Ext_C, we average over all the seeds of C, which means that guessing a bit of the output of Ext_C corresponds to distinguishing the output of C from uniform. For the reader's convenience we now restate Proposition 1 and give its proof.

Proposition 3. *[Proposition 1] Let X be a classical random variable correlated to some quantum system E, let Y be a (not necessarily uniform) seed, independent from XE, and let*

$$\|\rho_{\mathrm{Ext}_C(X,Y)E} - \rho_{U_m} \otimes \rho_Y \otimes \rho_E\|_{\mathrm{tr}} > \varepsilon, \tag{A.29}$$

where Ext_C is the extractor from Definition 83. Then there exists a fixed partition of the seed Y in two substrings V and W, and a classical random variable G, such that G has size $H_0(G) \leq rm$, where r is one of the parameters of the weak design (Definition 82), $V \leftrightarrow W \leftrightarrow G$ form a Markov chain, and

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_{VWGE}\|_{\mathrm{tr}} > \frac{\varepsilon}{m}. \tag{A.30}$$

Proof. We apply Lemma 140 to Eq. (A.29) and get that there exists an $i \in [m]$ such that

$$\begin{aligned}&\Bigg\| \sum_{\substack{x,y \\ C(x,y_{S_i})=0}} p_x q_y |C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y\rangle\langle C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y| C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y \otimes \rho^x \\ &- \sum_{\substack{x,y \\ C(x,y_{S_i})=1}} p_x q_y |C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y\rangle\langle C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y| C(x,y_{S_1}) \cdots C(x,y_{S_{i-1}}), y \otimes \rho^x \Bigg\|_{\mathrm{tr}} \\ &> \frac{\varepsilon}{m},\end{aligned} \tag{A.31}$$

where $\{p_x\}_{x\in\mathcal{X}}$ and $\{q_y\}_{y\in\mathcal{Y}}$ are the probability distributions of X and Y respectively.

We split $y \in \{0,1\}^d$ in two strings of $t = |S_i|$ and $d - t$ bits, and write $v := y_{S_i}$ and $w := y_{[d]\setminus S_i}$. To simplify the notation, we set $g(w,x,j,v) := C(x, y_{S_j})$. Fix w, x and j, and consider the function $g(w,x,j,\cdot) : \{0,1\}^t \to \{0,1\}$. This function only depends on $|S_j \cap S_i|$ bits of v. So to describe this function we need a string of at most $2^{|S_j \cap S_i|}$ bits. And to describe $g^{w,x}(\cdot) := g(w,x,1,\cdot)\cdots g(w,x,i-1,\cdot)$, which is the concatenation of the bits of $g(w,x,j,\cdot)$ for $1 \leq j \leq i-1$, we need a string of length at most $\sum_{j=1}^{i-1} 2^{|S_j \cap S_i|}$. So a system G containing a description of $g^{w,x}$ has size at most $H_0(G) \leq \sum_{j=1}^{i-1} 2^{|S_j \cap S_i|}$. We now rewrite Eq. (A.31) as

$$\Bigg\| \sum_{\substack{x,v,w \\ C(x,v)=0}} p_x q_{v,w} |g^{w,x}(v), v, w\rangle\langle g^{w,x}(v), v, w|g^{w,x}(v), v, w \otimes \rho^x \\ - \sum_{\substack{x,v,w \\ C(x,v)=1}} p_x q_{v,w} |g^{w,x}(v), v, w\rangle\langle g^{w,x}(v), v, w|g^{w,x}(v), v, w \otimes \rho^x \Bigg\|_{\mathrm{tr}} > \frac{\varepsilon}{m}.$$

By providing a complete description of $g^{w,x}$ instead of its value at the point v, we can only increase the trace distance, hence

$$\Bigg\| \sum_{\substack{x,v,w \\ C(x,v)=0}} p_x q_{v,w} |g^{w,x}, v, w\rangle\langle g^{w,x}, v, w|g^{w,x}, v, w \otimes \rho^x \\ - \sum_{\substack{x,v,w \\ C(x,v)=1}} p_x q_{v,w} |g^{w,x}, v, w\rangle\langle g^{w,x}, v, w|g^{w,x}, v, w \otimes \rho^x \Bigg\|_{\mathrm{tr}} > \frac{\varepsilon}{m}.$$

By rearranging this a little more we finally get

$$\|\rho_{C(X,V)VWGE} - \rho_{U_1} \otimes \rho_{VWGE}\|_{\mathrm{tr}} > \frac{\varepsilon}{m},$$

where G is a classical system of size $H_0(G) \leq \sum_{j=1}^{i-1} 2^{|S_j \cap S_i|}$, and $V \leftrightarrow W \leftrightarrow G$ form a Markov chain. By the definition of weak designs, we have for all $i \in [m]$, $\sum_{j=1}^{i-1} 2^{|S_j \cap S_i|} \leq rm$ for some $r \geq 1$. So $H_0(G) \leq rm$. □

A.5.4 Known extractors and designs

In this section we list the known constructions for weak designs and 1-bit extractors, which we plug into Trevisan's extractor in 8.4.

Weak designs.

The following weak design allows nearly all the min-entropy of the source to be extracted, but requires a rather large seed (typically $O(\log^3 n)$ for an optimal 1-bit extractor).

Lemma 141 ([98, Lemma 17][2]). *For every $t, m \in \mathbb{N}$ there exists a weak $(t, 1)$-design $S_1, \ldots, S_m \subset [d]$ such that $d = t \left\lceil \frac{t}{\ln 2} \right\rceil \lceil \log 4m \rceil = O(t^2 \log m)$. Moreover, such a design can be found in time* $\operatorname{poly} m, d$ *and space* $\operatorname{poly} m$.

If we wish to minimize the length of the seed, we can use the following weak design with $\log r = \Theta(t)$. We then get a seed of length $O(\log n)$ (for an optimal 1-bit extractor), but only extract a sub-linear amount of min-entropy from the source.

Lemma 142 ([98, Lemma 15]). *For every $t, m \in \mathbb{N}$ and $r > 1$, there exists a weak (t, r)-design $S_1, \ldots, S_m \subset [d]$ such that $d = t \left\lceil \frac{t}{\ln r} \right\rceil = O\left(\frac{t^2}{\log r}\right)$. Moreover, such a design can be found in time* $\operatorname{poly} m, d$ *and space* $\operatorname{poly} m$.

The following weak design construction is much more efficient than the two previous ones, and ideal for a local extractor. It uses a seed of size $O(\log^2 n)$ and can extract a constant fraction of the min-entropy (for an optimal 1-bit extractor).

Lemma 143 ([51, Theorem 3]). *For every $m, t \in \mathbb{N}$, such that $m = \Omega(t^{\log t})$, and constant $r > 1$, there exists an explicit weak (t, r)-design $S_1, \ldots, S_m \subset [d]$, where $d = O(t^2)$. Such a design can be found in time* $\operatorname{polylog} m, t$ *and space* $\operatorname{polylog} m + \log t$.

Remark 144. *For the extractor from Lemma 146 and an error $\varepsilon = \operatorname{poly} 1/n$, this design requires $m = \Omega\left((\log n)^{\log \log n}\right)$. If we are interested in a smaller m, say $m = \operatorname{polylog} n$, then we can use the weak design from Lemma 142 with $r = n^\gamma$. This construction would require time and space $\operatorname{polylog} n = \operatorname{polylog} 1/\varepsilon$. The resulting seed would have length only $O(\log n)$ instead of $O(\log^2 n)$.*

One-bit extractors.

As 1-bit extractor, Raz et al. [98] (and Trevisan [119] too) used the bits of a list-decodable code. We give the parameters here as Proposition 4 and refer to A.5.5 for details on the construction and proof.

Proposition 4. *For any $\varepsilon > 0$ and $n \in \mathbb{N}$ there exists a (k, ε)-strong extractor with uniform seed $\mathrm{Ext}_{n,\varepsilon} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}$ with $d = O(\log(n/\varepsilon))$ and $k = 3 \log 1/\varepsilon$.*

[2] Hartman and Raz [51] give a more efficient construction of this lemma, namely in time $\operatorname{polylog} m, t$ and space $\operatorname{polylog} m + \log t$, with the extra minor restriction that $m > t^{\log t}$.

Local extractor. Local extractors are defined as follows.

Definition 145 (ℓ-local extractor). *An extractor* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is* locally computable *(or* ℓ-local*), if for every* $r \in \{0,1\}^d$*, the function* $x \mapsto \mathrm{Ext}(x,r)$ *depends on only* ℓ *bits of its input, where the bit locations are determined by* r*.*

Lu [79] modified Trevisan's scheme [119, 98] to use a local list-decodable code as 1-bit extractor. Vadhan [120] proposes another construction for local extractors, which is optimal up to constant factors. Both these constructions have similar parameters in the case of 1-bit extractors.[3] We state the parameters of Vadhan's construction here and Lu's constructions in A.5.5.

Lemma 146 ([120, Theorem 8.5]). *For any* $\varepsilon > \exp\left(-n/2^{O(\log^* n)}\right)$*,* $n \in \mathbb{N}$ *and constant* $0 < \gamma < 1$*, there exists an explicit* ℓ*-local* (k,ε)*-strong extractor with uniform seed* $\mathrm{Ext}_{n,\varepsilon,\gamma} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}$ *with* $d = O(\log(n/\varepsilon))$*,* $k = \gamma n$ *and* $\ell = O(\log 1/\varepsilon)$*.*

Weak random seed. Raz [97] shows how to transform any extractor which needs a uniform seed into one which can work with a weakly random seed.

Lemma 147 ([97, Theorem 4]). *For any* (k,ε)*-strong extractor* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}^m$ *with uniform seed, there exists a* $(k,2\varepsilon)$*-strong extractor* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^{t'} \to \{0,1\}^m$ *requiring only a seed with min-entropy* $H_{\min}(Y) \geq \left(\frac{1}{2}+\beta\right) t'$*, where* $t' = 8t/\beta$*.*

By applying this lemma to the 1-bit extractor given in Proposition 4, we obtain the following 1-bit extractor.

Corollary 148. *For any* $\varepsilon > 0$ *and* $n \in \mathbb{N}$ *there exists a* (k,ε)*-strong extractor* $\mathrm{Ext}_{n,\varepsilon} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}$ *requiring a seed with min-entropy* $\left(\frac{1}{2}+\beta\right) d$*, where* $d = O(\frac{1}{\beta}\log(n/\varepsilon))$ *and* $k = 3\log 1/\varepsilon + 3$*.*

A.5.5 List-decodable codes are one-bit extractors

Construction.

A standard error correcting code guarantees that if the error is small, any string can be uniquely decoded. A list-decodable code guarantees that for a larger (but bounded) error, any string can be decoded to a list of possible messages.

Definition 149 (list-decodable code). *A code* $C : \{0,1\}^n \to \{0,1\}^{\bar{n}}$ *is said to be* (ε, L)*-list-decodable if every Hamming ball of relative radius* $1/2 - \varepsilon$ *in* $\{0,1\}^{\bar{n}}$ *contains at most* L *codewords.*

[3] If the extractor is used to extract m-bits, then Vadhan's scheme reads less input bits and uses a shorter seed than Lu's.

Neither Trevisan [119] nor Raz et al. [98] state it explicitly, but both papers contain an implicit proof that if $C : \{0,1\}^n \to \{0,1\}^{\bar{n}}$ is a (ε, L)-list-decodable code, then

$$\begin{aligned}\mathrm{Ext} : \{0,1\}^n \times [\bar{n}] &\to \{0,1\}\\ (x,y) &\mapsto C(x)_y,\end{aligned}$$

is a $(\log L + \log 1/2\varepsilon, 2\varepsilon)$-strong extractor (according to Definition 76). We have rewritten their proof in A.5.6 for completeness.[4]

There exist list-decodable codes with following parameters.

Lemma 150. *For every $n \in \mathbb{N}$ and $\delta > 0$ there is a code $C_{n,\delta} : \{0,1\}^n \to \{0,1\}^{\bar{n}}$, which is $(\delta, 1/\delta^2)$-list-decodable, with $\bar{n} = \mathrm{poly} n, 1/\delta$. Furthermore, $C_{n,\delta}$ can be evaluated in time $\mathrm{poly} n, 1/\delta$ and $\bar{n}$ can be assumed to be a power of 2.*

For example, Guruswami et al. [49] combine a Reed-Solomon code with a Hadamard code, obtaining such a list-decodable code with $\bar{n} = O(n/\delta^4)$.

Such codes require all bits of the input x to be read to compute any single bit $C(x)_i$ of the output. If we are interested in so-called *local* codes, we can use a construction by Lu [79, Corollary 1].

A.5.6 List-decodable codes are strong extractors

Theorem 151. *Let $C : \{0,1\}^n \to \{0,1\}^{\bar{n}}$ be an (ε, L)-list-decodable code. Then the function*

$$\begin{aligned}C' : \{0,1\}^n \times [\bar{n}] &\to \{0,1\}\\ (x,y) &\mapsto C(x)_y,\end{aligned}$$

is a $(\log L + \log 1/2\varepsilon, 2\varepsilon)$-strong extractor.[5]

To prove this theorem we first show that a player who can distinguish the bit of $C'(X,Y)$ from uniform can construct a string α which is close to $C(X)$ on average (over X). Then using the error correcting properties of the code C, he can reconstruct X. Hence a player who can break the extractor must have low min-entropy about X.

Lemma 152. *Let X and Y be two independent random variables with alphabets $\{0,1\}^n$ and $[n]$ respectively. Let Y be uniformly distributed and X be distributed such that $\frac{1}{2}|X_Y \circ Y - U_1 \circ Y| > \delta$, where U_1 is uniformly distributed on $\{0,1\}$. Then there exists a string $\alpha \in \{0,1\}^n$ with*

$$\Pr\left[d(X,\alpha) \leq \frac{1}{2} - \frac{\delta}{2}\right] > \delta,$$

where $d(\cdot,\cdot)$ is the relative Hamming distance.

[4] A slightly more general proof, that *approximate* list-decodable codes are 1-bit extractors can be found in [33, Claim 3.7].

[5] This theorem still holds in the presence of classical side information with exactly the same parameters.

Proof. Define $\alpha \in \{0,1\}^n$ to be the concatenation of the most probable bits of X, i.e., $\alpha_y := \arg\max_b P_{X_y}(b)$, where $P_{X_y}(b) = \sum_{\substack{x\in\{0,1\}^n \\ x_y=b}} P_X(x)$.

The average relative Hamming distance between X and α is

$$\begin{aligned}\sum_{x\in\{0,1\}^n} P_X(x) d(x,\alpha) &= \frac{1}{n}\sum_{x\in\{0,1\}^n} P_X(x)\sum_{y=1}^{n}|x_y-\alpha_y| \\ &= \frac{1}{n}\sum_{\substack{x,y \\ x_y\neq\alpha_y}} P_X(x) = 1-\frac{1}{n}\sum_{y=1}^{n} P_X(\alpha_y).\end{aligned}$$

And since $\frac{1}{2}|X_Y \circ Y - U_1 \circ Y| > \delta$ is equivalent to $\frac{1}{n}\sum_{y=1}^{n}\max_{b\in\{0,1\}} P_{X_y}(b) > \frac{1}{2}+\delta$, we have

$$\sum_{x\in\{0,1\}^n} P_X(x) d(x,\alpha) < \frac{1}{2}-\delta. \tag{A.32}$$

We now wish to lower bound the probability that the relative Hamming distance is less than $\frac{1}{2}-\frac{\delta}{2}$. Let $B := \{x : d(x,\alpha) \leq \frac{1}{2}-\frac{\delta}{2}\}$ be the set of values $x \in \{0,1\}^n$ meeting this requirement. Then the weight of B, $w(B) := \sum_{x\in B} P_X(x)$, is the quantity we wish to lower bound. It is at its minimum if all $x \in B$ have Hamming distance $d(x,\alpha) = 0$. In which case the average Hamming distance is

$$\sum_{x\in\{0,1\}^n} P_X(x) d\left(x,\alpha\right) > (1-w(B))\left(\frac{1}{2}-\frac{\delta}{2}\right). \tag{A.33}$$

Combining Eqs. (A.32) and (A.33) we get

$$w(B) > \frac{\delta}{1-\delta} \geq \delta.$$

We are now ready to prove Theorem 151.

Proof of Theorem 151. We will show that if it is possible to distinguish $C'(X,Y)$ from uniform with probability at least 2ε, then X must have min-entropy $H_{\min}(X) < \log L + \log 1/2\varepsilon$.

If $\frac{1}{2}\left|C'(X,Y)\circ Y - U_1\circ Y\right| > 2\varepsilon$, then by Lemma 152 we know that there exists an $\alpha \in \{0,1\}^{\bar{n}}$ such that

$$\Pr\left[d\left(C(X),\alpha\right) \leq \frac{1}{2}-\varepsilon\right] > 2\varepsilon,$$

where $d(\cdot,\cdot)$ is the relative Hamming distance.

This means that with probability at least 2ε, X takes values x such that the relative Hamming distance $d(C(x),\alpha) \leq \frac{1}{2}-\varepsilon$. So for these values of X, if we choose one of the codewords in the Hamming ball of relative radius $\frac{1}{2}-\varepsilon$ around α uniformly at random as our guess for x, we will have chosen correctly with probability at least $1/L$, since the Hamming ball contains at most L code words. The total probability of guessing X is then at least $2\varepsilon/L$.

Hence by Eq. (3.1), $H_{\min}(X) < \log L + \log 1/2\varepsilon$.

A.6 Omitted proofs from Chapter 9

A.6.1 Identifying "good" blocks in Protocols A and B

In this section we prove Claim 97 and Claim 98, which play an analogous role for Theorem 93 and Theorem 94 respectively: that of identifying a special iteration of the protocol that will be useful to Alice and Bob in the guessing game.

Proof of Claim 97. Let BAD' be the set of strings $b \in \big(\{0,1\}^k\big)^m$ such that $\Pr(b|\mathrm{CHSH}) > 2^{-n}$. Assumption (i) together with Claim 8 show that $\Pr(\mathrm{BAD}'|\mathrm{CHSH}) \geq \varepsilon$, so using (ii) we get $\Pr(\mathrm{CHSH}|\mathrm{BAD}') \geq \varepsilon^2$. Define BAD to contain only those strings $b \in \mathrm{BAD}'$ such that $\Pr(\mathrm{CHSH}|b) \geq \varepsilon^2/2$; we have $\Pr(\mathrm{BAD}) \geq (\varepsilon^2/2)\Pr(\mathrm{BAD}') \geq \varepsilon^4/2$.

By definition of BAD, using Baye's rule we have that for every $b = (b_1, \ldots, b_m) \in \mathrm{BAD}$,

$$\Pr(B = b, \mathrm{CHSH}) \;=\; \prod_{i=1}^{m} \Pr(B_i = b_i, \mathrm{CHSH}_i|\mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) \;\geq\; 2^{-n}\varepsilon^2/2.$$

Taking logarithms on both sides,

$$\sum_{i=1}^{m} -\log \Pr(B_i = b_i \mathrm{CHSH}_i|\mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) \;\leq\; n + 3\log(1/\varepsilon) \;\leq\; 2n,$$

assuming as in the statement of the claim that ε is not too small. By an averaging argument at least $7/8$ of all $i \in [m]$ are such that a fraction at least $7/8$ of all $b \in \mathrm{BAD}$ are such that

$$\Pr(B_i = b_i|\mathrm{CHSH}_{<i}, B_{<i} = b_{<i}) \;\geq\; 2^{-128(n/m)} \;\geq\; 2^{-128/C}. \tag{A.34}$$

A similar argument, starting from $\Pr(\mathrm{CHSH}|b) \geq \varepsilon^2/2$ for all $b \in \mathrm{BAD}$, shows that

$$\sum_{i=1}^{m} -\log \Pr(\mathrm{CHSH}_i|\mathrm{CHSH}_{<i}, b) \;\leq\; 2\log(1/\varepsilon) + 1 \;\leq\; 3\log(1/\varepsilon)$$

for large enough n. By an averaging argument, a fraction at least $7/8$ of all $i \in [m]$ are such that a fraction $7/8$ of all all $b \in \mathrm{BAD}$ satisfy

$$-\log \Pr(\mathrm{CHSH}_i|\mathrm{CHSH}_{<i}, b) \leq \frac{192\log(1/\varepsilon)}{n}, \tag{A.35}$$

where here we used $m \geq n$. Let S be the set of $i \in [m]$ such that both (A.34) and (A.35) hold simultaneously for a fraction at least $3/4$ of $b \in \mathrm{BAD}$. By the union bound, we have $|S| \geq (3/4)m$.

We apply the same reasoning once more, focusing on the CHSH constraint being satisfied in a Bell block. Let N be a random variable equal to the number of Bell blocks that fall in S. Since S is a fixed set of indices, and each block is chosen to be a Bell block independently with probability $1/\ell$, N is concentrated around $\Delta(|S|/m) \geq \Delta/2$. By a Chernoff bound, the probability that N is less than $\Delta/4$ is at most $e^{-\Delta/16}$, which given our choice of Δ can be neglected in front of the other events we are considering. For the remainder of the proof we assume that $N \geq C/4$. Let T_j be a random variable denoting the index of the j-th Bell block, among those that fall in S. Starting from $\Pr(\mathrm{CHSH}|\mathrm{BAD}) \geq \varepsilon^2/2$ and using Baye's rule as before,

$$\sum_{j=1}^{N} -\log \Pr(\mathrm{CHSH}_{T_j}|\mathrm{CHSH}_{<T_j}, \mathrm{BAD}) \;\leq\; 2\log(1/\varepsilon) + 1 \;\leq\; 3\log(1/\varepsilon).$$

By an averaging argument and using our assumed lower bound on N this shows that a fraction at least $1/2$ of the Bell blocks in Protocol A are such that

$$\Pr(\mathrm{CHSH}_{T_j}|\mathrm{CHSH}_{<T_j}, \mathrm{BAD}) \geq \varepsilon^{24/C}. \tag{A.36}$$

Let $T_j \in T \cap S$ be a Bell block for which (A.36) holds. For a fraction at least $\varepsilon^{24/C}/2$ of $b \in \mathrm{BAD}$ it holds that

$$\Pr(\mathrm{CHSH}_{T_j}|\mathrm{CHSH}_{<T_j}, b) \geq \varepsilon^{24/C}/2. \tag{A.37}$$

By the union bound, at iteration T_j (A.37) will hold simultaneously with (A.34) and (A.35) for a subset G of BAD of size at least

$$\Pr(G) \;=\; \Pr(G|\mathrm{BAD})\Pr(\mathrm{BAD}) \;\geq\; \big(\varepsilon^{24/C}/2 - 1/4\big)\varepsilon^4/2 \;\geq\; \varepsilon^5$$

given our choice of parameters. By choosing C large enough, (A.34) implies item 1 in the claim, and (A.35) and (A.37) imply item 2, given the choice of Δ made in the claim. □

Proof of Claim 98. By definition, $\Pr\left(G^B\right) \geq p_s \Pr(\mathrm{CHSH}) \geq p_s\varepsilon$. Conditioned on G^B, by Markov's inequality it must be that $d_H(E^B, B) < 0.01$ on a fraction at least $1 - 100f_e$ of blocks in which the input to B was $(A, 0)$. Let $f'_e = 100f_e$. Let $\eta = 2^{-10^{-5}f'_e|T|/(2\cdot 100^2)}$, and assume C chosen large enough so that $\eta \leq p_s\varepsilon/6 = \Omega(n^{-8(1+\alpha)})$. This is possible since $|T|$ is sharply concentrated around $C\log^2\ell$ and $f'_e = \Omega(1/\log\ell)$.

Among the blocks in which Eve's prediction is correct, nothing distinguishes those Bell blocks in which $\mathcal{B}$'s input is $(A, 0)$: indeed, we may think of those only being designated as Bell blocks after Eve has made her prediction. By a Chernoff bound the probability that more than a fraction $2f'_e$ of such blocks fall into those for which G^B_j does not hold is upper-bounded by η. Hence the following holds

$$\Pr\left(\mathrm{E}_{j\in T:\,Y_j=(A,0)}\, G^B_j > 1-2f'_e | G^B\right) \,\geq\, 1-\eta. \tag{A.38}$$

Since V is a fixed subset of $[km]$ of size $|V| = O(m^\gamma \log^2 m)$, the probability that any of the randomly chosen $O(\log^2 \ell)$ Bell blocks intersects it is at most $O(m^{-1+\gamma}\log^4 m) = O(n^{2-1/\gamma}\log^4 n)$ for large enough n. We assume as in the statement of Theorem 94 that γ is chosen large enough so that this is much smaller than (our upper bound on) η, i.e. $\gamma < 1/(9+8\alpha)$. For the remainder of the proof we will neglect the chance of this happening.

Conditioning further on CHSH can only blow-up the error by a factor $1/\Pr(\mathrm{CHSH}|G^B) \leq 1/(p_s\varepsilon)$. In that case $G^A = G^B$ (Eve's prediction only depends on the advice bits she is given), so we obtain:

$$\frac{\Pr\left(\mathrm{E}_{j\in T:Y_j=0}\, G^B_j > 1-2f'_e, \mathrm{CHSH}|G^A\right)}{\Pr\left(\mathrm{CHSH}|G^A\right)} \,=\, \Pr\left(\mathrm{E}_{j\in T:\,Y_j=(A,0)}\, G^A_j > 1{-}2f'_e|G^A, \mathrm{CHSH}\right) \,\geq\, 1{-}\eta/(p_s\varepsilon). \tag{A.39}$$

Suppose Eve makes more than a fraction $5f'_e$ of errors in predicting $\mathcal{A}$'s output on those Bell blocks in which its input is $(A,0)$. Some of those will later be randomly chosen by Bob as Bell blocks, and by a Chernoff bound with probability at least $1-\eta$ the input to $\mathcal{B}$ will also be $(A,0)$ in at least 40% of those blocks. Whenever this happens, Eve's prediction will be wrong on a total fraction more than $2f'_e$ of $\mathcal{B}$'s $(A,0)$-input Bell blocks, contradicting (A.39). Indeed, whenever CHSH holds, if the input to both boxes is $(A,0)$ then Eve being correct in predicting $\mathcal{B}$'s output is equivalent to her being correct in predicting $\mathcal{A}$'s output. Hence the following holds:

$$\begin{aligned}\Pr\left(\mathrm{E}_{j\in T:\,X_j=(A,0)}\, G^A_j > 1-5f'_e, \mathrm{CHSH}|G^A\right) &\geq \Pr\left(\mathrm{E}_{j\in T:\,Y_j=(A,0)}\, G^A_j > 1-2f'_e, \mathrm{CHSH}|G^A\right) - \eta\\ &\geq (1-\eta/(p_s\varepsilon))\Pr\left(\mathrm{CHSH}|G^A\right) - \eta\\ &\geq (1-2\eta/(p_s\varepsilon))\Pr\left(\mathrm{CHSH}|G^A\right),\end{aligned} \tag{A.40}$$

where the last inequality uses $\Pr(\mathrm{CHSH}|G^A) \geq p_s\varepsilon$. As before, since $G^A \wedge \mathrm{CHSH} = G^B \wedge \mathrm{CHSH}$, (A.40) implies the following:

$$\Pr\left(\mathrm{E}_{j\in T:\,X_j=(A,0)}\, G^B_j > 1-5f'_e|G^B, \mathrm{CHSH}\right) \,\geq\, 1-2\eta/(p_s\varepsilon). \tag{A.41}$$

Next, suppose Eve makes a prediction that is wrong on a fraction at least $14f'_e$ of the Bell blocks, irrespective of Bob's inputs. Then again with high probability at least 40% of the inputs to $\mathcal{A}$ in those blocks will be $(A,0)$, implying that Eve is wrong on more than a fraction $5f'_e$ of $\mathcal{A}$'s $(A,0)$ inputs, and contradicting (A.41). Hence the following is proven just as (A.40) was:

$$\Pr\left(\mathrm{E}_{j\in T}\, G^B_j > 1-14f'_e|G^B, \mathrm{CHSH}\right) \,\geq\, 1-3\eta/(p_s\varepsilon). \tag{A.42}$$

Hence

$$\Pr\left(\mathrm{E}_{j\in T}\, G_j^A > 1 - 14f_e'|G^A, \mathrm{CHSH}\right) \geq 1 - 3\eta/(p_s\varepsilon),$$

which is greater than 1/2 given our choice of η. Removing all conditioning, whenever Eve is given advice bits by Alice, it holds that

$$\Pr\left(\mathrm{E}_{j\in T}\, G_j^A > 1 - 14f_e', \mathrm{CHSH}\right) \geq \Omega(p_s\varepsilon).$$

□

A.6.2 Proof of Lemma 95

In this section we give the proof of Lemma 95. The proof crucially uses properties of a specific extractor construction, first shown to be secure in the presence of quantum bounded-storage adversaries in [112], and in the more general setting of quantum bounded-information adversaries in [34]. We first describe the extractor.

The t-XOR extractor

The t-XOR extractor E_t, parametrized by an integer t, follows Trevisan's general extractor construction paradigm [119]. It is based on two main ingredients, the t-XOR code and a combinatorial design construction due to Hartman and Raz [51]. For us, only the details of the t-XOR code will be important.

The t-XOR code. Given integers m and $t \leq m$, let $C_t : \{0,1\}^m \to \{0,1\}^{\binom{m}{t}}$ map an m-bit string to the string of parities of all subsets of t out of its m bits. Two properties of this encoding will be relevant for us. The first is that it is locally computable: each bit of the code only depends on t bits of the input. The second is that it is approximately list-decodable (we summarize its parameters in Lemma 157 below).

Combinatorial designs. Given integers s, m, r and $\rho > 0$, a collection of subsets $S_1, \dots, S_r \subseteq [s]$ is called a (s, m, r, ρ) weak design if for all $i \in [r]$, $|S_r| = m$ and for all j, $\sum_{i<j} 2^{|S_i \cap S_j|} \leq \rho(r-1)$. For our purposes it will suffice to note that Hartman and Raz [51] proved the existence of a $(s, m, r, 1+\gamma)$ design for every m, $0 < \gamma < 1/2$, $s = O(m^2 \log 1/\gamma)$ and $r > s^{\Omega(\log s)}$.

The t-XOR extractor. We define the extractor that we will use in the proof of Lemma 95.

Definition 153. *Let m, r, t, s be given integers such that $t = O(\log m)$ and $s = O(\log^4 n)$. Then $E_t : \{0,1\}^m \times \{0,1\}^s \to \{0,1\}^r$ maps $(x, y) \in \{0,1\}^m \times \{0,1\}^s$ to $C_t(x)_{y_{S_1}}, \dots, C_t(x)_{y_{S_r}}$, where $(S_1, \dots, S_r)$ is a $(s, t\log m, r, 5/4)$ design and y_{S_i} designates the bits of y indexed by S_i, interpreted as a t-element subset of $[m]$.*

While, as shown in Corollary 5.11 in [34], E_t is a strong extractor with good parameters, we will not use this fact directly. Rather, we will use specific properties that arise from the "reconstruction paradigm"-based *proof* that it is an extractor secure against quantum adversaries, and one may argue that Lemma 95 is implicit in the proof of security of E_t given in [34]. Since it does not follow directly from the mere statement that E_t is an extractor, we give more details here. We will show the following lemma, which is more general than Lemma 95.

Lemma 154. *Let m, r, t be integers such that $t = O(\log^2 m)$ and $\varepsilon > 0$. Let ρ_{XE} be a cq-state such that X is a random variable distributed over m-bit strings. Let U_r be uniformly distributed over r-bit strings, and suppose that*

$$\left\|\rho_{Ext(X,Y)E} - \rho_{U_r} \otimes \rho_E\right\|_{tr} > \varepsilon, \tag{A.43}$$

i.e. an adversary Eve holding register E can distinguish the output of the extractor from a uniformly random r-bit string. Then there exists a fixed subset $V \subseteq [m]$ of size $|V| = O(tr)$ such that, given the string X_V as advice, with probability at least $\Omega(\varepsilon^2/r^2)$ over the choice of $x \sim p_X$ and her own randomness Eve can output a list of $\ell = O(r^4/\varepsilon^4)$ strings $\tilde{x}^1, \dots, \tilde{x}^\ell$ such that there is an $i \in [\ell]$, $d_H(\tilde{x}^i, x) \leq (2/t)\ln(4r/\varepsilon)$.

It is not hard to see why Lemma 154 implies Lemma 95. First note that if r is chosen in Lemma 154 so that $r > 2H_\infty^\varepsilon(X|E)$ then the assumption (A.43) is automatically satisfied.[6] The conclusion of Lemma 95 then follows from that of Lemma 154 by having Eve output a random string out of her ℓ predictions, and choosing $t = \Omega(\log^2 m)$ to ensure that $(2/t)\ln(4r/\varepsilon) \leq 1/\log m$.

In the remainder of this section we sketch the proof of Lemma 154. The first step, explained in Section A.6.2, consists in using a hybrid argument to show that, given (A.43), Eve can predict a random t-XOR of X's bits with reasonable success probability, given sufficiently many "advice bits" about X. In the second step, detailed in Section A.6.3, we show using an argument due to Koenig and Terhal [76] that this implies the adversary can in fact recover most t-XORs of X, simultaneously. Finally, in Section A.6.3 we use the list-decoding properties of the XOR code to show that as a consequence the adversary can with good probability produce a string that agree with X on a large fraction of coordinates.

The hybrid argument

Suppose that (A.43) holds. Proposition 4.4 from [34] shows that a standard hybrid argument, together with properties of Trevisan's extractor (specifically the use of the seed through combinatorial designs), can be used to show the following claim.

[6] The extra randomness coming from the seed of the extractor will be small, as its size can be taken to be $s = O(\log^4 m)$.

Claim 155. *There exists a subset $V \subseteq [m]$ of size $|V| = O(tr)$ such that, given the bits X_V, Eve can predict a random t-XOR of the bits of X with advantage ε/r. Formally,*

$$\left\|\rho_{C_t(X)_Y YVE} - \rho_{U_1} \otimes \rho_Y \otimes \rho_{VE}\right\|_{tr} > \frac{\varepsilon}{r}, \tag{A.44}$$

where Y is a random variable uniformly distributed over $\left[\binom{m}{t}\right]$ and V is a register containing the bits of X indexed by V.

A.6.3 Recovering all t-XORs.

The next step in the proof of Lemma 154 is to argue that Eq. (A.44) implies that an adversary given access to $E' = VE$ can predict not only a random XOR of X, but a string Z of length $\binom{m}{t}$ such that Z agrees with the string $C_t(X)$ of all t-XOR's of X in a significant fraction of positions. Classically this is trivial, as one can just repeat the single-bit prediction procedure guaranteed by (A.44) for all possible choices Y of the t bits whose parity one is trying to compute. In the quantum setting it is more tricky. We will follow an argument from [76] showing that (A.44) implies that there is a single measurement, independent of Y, that one can perform on E and using the (classical) result of which one can predict the bits $C_t(X)_Y$ with good success on average (over the measurement's outcome and the choice of Y).

Claim 156. *Suppose (A.44) holds. Then there exists a measurement $\mathcal{F}$, with outcomes in $\{0,1\}^m$, such that*

$$\Pr_{x\sim p_X,\, y\sim U_{t\log m}}\big(C_t(x)_Y \;=\; C_t(\mathcal{F}(VE))_y\big) \geq \frac{1}{2} + \frac{\varepsilon^2}{4r^2}\,, \tag{A.45}$$

where $\mathcal{F}(VE)$ denotes the outcome of $\mathcal{F}$ when performed on the cq-state ρ_{VE}.

Proof. Our argument closely follows the proof of Theorem III.1 from [76]. Given an arbitrary cq-state ρ_{ZQ}, define the non-uniformity of Z given Q as

$$d(Z \leftarrow Q) \;:=\; \left\|\rho_{ZQ} - \rho_{U_z} \otimes \rho_Q\right\|_{tr}.$$

Let ρ_x denote the state contained in registers VE, conditioned on $X = x$. For a fixed string y, define two states

$$\rho_0^y \;:=\; \sum_{x:\, C_t(x)_y = 0} p_X(x)\,\rho_x \qquad \text{and} \qquad \rho_1^y \;:=\; \sum_{x:\, C_t(x)_y = 1} p_X(x)\,\rho_x.$$

Then, by definition $d\big(C_t(X)_y \leftarrow VE\big) \;=\; \left\|\rho_0^y - \rho_1^y\right\|_{tr}$ is the adversary's maximum success probability in distinguishing those states ρ_x which correspond to an XOR of 0 from those which correspond to an XOR of 1. Let $\mathcal{E}_y \;=\; \left\{E_y^0, E_y^1\right\}$ be the pretty good measurement corresponding to the pair of states $\left\{\rho_0^y, \rho_1^y\right\}$:

$$E_y^0 \;=\; \rho_{VE}^{-1/2}\,\rho_0^y\,\rho_{VE}^{-1/2} \qquad \text{and} \qquad E_y^1 \;=\; \rho_{VE}^{-1/2}\,\rho_1^y\,\rho_{VE}^{-1/2},$$

where $\rho_{VE} = \sum_x P_X(x)\rho_x$. Lemma 2 from [76] (more precisely, Eq. (19)), shows that the following holds as a consequence of (A.44):

$$\sqrt{\mathrm{E}_y\big[\,2\,d\big(C_t(X)_y \leftarrow \mathcal{E}^y(VE)\big)\big]} + d(C_t(X)_Y \leftarrow Y) \;>\; \frac{\varepsilon}{r}, \tag{A.46}$$

where $\mathcal{E}^y(VE)$ is the result of the POVM $\mathcal{E}^y$ applied on ρ_{VE}, and $d(C_t(X)_Y \leftarrow Y)$ is the distance from uniform of the one-bit extractor's output, in the absence of the adversary. We may as well assume this term to be small: indeed, if it is more than $\varepsilon/(2r)$ then (A.45) is proved without even having to resort to the quantum system E. Hence (A.46) implies

$$\mathrm{E}_y\left[\,d\big(C_t(X)_y \leftarrow \mathcal{E}^y_{pgm}(VE)\big)\,\right] \;>\; \frac{\varepsilon^2}{2r^2}\,,$$

which can be equivalently re-written as

$$\mathrm{E}_y\left[\,\mathrm{Tr}\big(E^0_y\,\rho^0_y\big) + \mathrm{Tr}\big(E^1_y\,\rho^1_y\big)\,\right] \;>\; \frac{1}{2} + \frac{\varepsilon^2}{4r^2}\,. \tag{A.47}$$

Following the argument in [76], we define a new PGM $\mathcal{F}$ with outcomes in $\{0,1\}^m$ and POVM elements $F^x = P_X(x)\rho_{VE}^{-1/2}\rho_x\,\rho_{VE}^{-1/2}$. The important point to notice is that for $z \in \{0,1\}$ we have $E^z_y = \sum_{x:\,C_t(x)_y=z} F^x$, hence (A.47) can be re-written as

$$\mathrm{E}_y\Bigg[\sum_{b:\,C_t(b)_y=0} \mathrm{Tr}\big(F^x\,\rho^0_y\big) + \sum_{b:\,C_t(b)_y=1} \mathrm{Tr}\big(F^x\,\rho^1_y\big)\Bigg] \;>\; \frac{1}{2} + \frac{\varepsilon^2}{4r^2}\,,$$

which is exactly (A.45). □

List-decoding the XOR code.

The following lemma (for a reference, see [55], Lemma 42) states the list-decoding properties of the t-XOR code C_t that are important for us.

Lemma 157. *For every $\eta > 2t^2/2^m$ and $z \in (\{0,1\}^m)^t$, there is a list of $\ell \leq 4/\eta^2$ elements $x^1,\dots,x^\ell \in \{0,1\}^m$ such that the following holds: for every $z' \in \{0,1\}^m$ which satisfies*

$$\Pr_{\{y_1,\dots,y_t\}\in\binom{m}{t}}[z_{(y_1,\dots,y_t)} = \oplus_{i=1}^t z'_{y_i}] \geq \frac{1}{2} + \eta,$$

there is an $i \in [\ell]$ such that

$$\Pr_{y\sim\mathcal{U}_N}[x^i_y = z'_y] \geq 1-\delta,$$

with $\delta = (1/t)\ln(2/\eta)$.

Claim 156 implies that, with probability at least $\varepsilon^2/(8r^2)$ over the choice of x and over Eve's own randomness, when measuring her system with $\mathcal{F}$ she will obtain a string $\tilde{z}$ whose t-XORs agree with those of x with probability at least $1/2+\varepsilon^2/(8r^2)$. Lemma 157 shows that in that case she can recover a list of at most $2^8r^4/\varepsilon^4$ "candidate" strings $\tilde{z}^i$ such that there exists at least one of these strings which agrees with x at a (possibly adversarial) fraction $1-\delta$ of positions, where $\delta=(2/t)\ln(4r/\varepsilon)$ given our choice of parameters. Hence Lemma 154 is proved.