

Springer Theses

Recognizing Outstanding Ph.D. Research

Rotem Arnon-Friedman

Device-Independent Quantum Information Processing

A Simplified Analysis

 Springer

Springer Theses

Recognizing Outstanding Ph.D. Research

Aims and Scope

The series “Springer Theses” brings together a selection of the very best Ph.D. theses from around the world and across the physical sciences. Nominated and endorsed by two recognized specialists, each published volume has been selected for its scientific excellence and the high impact of its contents for the pertinent field of research. For greater accessibility to non-specialists, the published versions include an extended introduction, as well as a foreword by the student’s supervisor explaining the special relevance of the work for the field. As a whole, the series will provide a valuable resource both for newcomers to the research fields described, and for other scientists seeking detailed background information on special questions. Finally, it provides an accredited documentation of the valuable contributions made by today’s younger generation of scientists.

Theses are accepted into the series by invited nomination only and must fulfill all of the following criteria

- They must be written in good English.
- The topic should fall within the confines of Chemistry, Physics, Earth Sciences, Engineering and related interdisciplinary fields such as Materials, Nanoscience, Chemical Engineering, Complex Systems and Biophysics.
- The work reported in the thesis must represent a significant scientific advance.
- If the thesis includes previously published material, permission to reproduce this must be gained from the respective copyright holder.
- They must have been examined and passed during the 12 months prior to nomination.
- Each thesis should include a foreword by the supervisor outlining the significance of its content.
- The theses should have a clearly defined structure including an introduction accessible to scientists not expert in that particular field.

More information about this series at <http://www.springer.com/series/8790>

Rotem Arnon-Friedman

Device-Independent Quantum Information Processing

A Simplified Analysis

Doctoral Thesis accepted by
ETH Zurich, Zürich, Switzerland

 Springer

Author

Dr. Rotem Arnon-Friedman
Senior Scientist
Physics of Complex Systems
Weizmann Institute of Science
Rehovot, Israel

Supervisor

Prof. Renato Renner
Department of Physics
ETH Zurich
Zürich, Switzerland

ISSN 2190-5053

Springer Theses

ISBN 978-3-030-60230-7

<https://doi.org/10.1007/978-3-030-60231-4>

ISSN 2190-5061 (electronic)

ISBN 978-3-030-60231-4 (eBook)

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Neer, Avishai, and Omri

Supervisor's Foreword

Even the most sophisticated physics experiment suffers from the problem that our control and knowledge of the relevant parameters is limited. Take, for example, a single-particle interference experiment. To realise it, we would require a source that emits one single particle towards an interferometer and a detector that tells us where that particle arrived. Any real detector has, however, a finite efficiency and may thus only register a fraction, say $1/4$, of the incoming particles. Conversely, a realistic source may with a certain probability, say again $1/4$, accidentally emit two particles instead of only one at a time. Normally, we wouldn't care too much about such potential imperfections and, for example, simply ignore rounds of the experiment in which no particle was detected.

Nonetheless, doing so requires an assumption, which is known as “fair sampling”. Specifically, we would need to assume that the interference pattern seen in our actual data, which consists of the (roughly) one quarter of the rounds in which a particle was detected, would not significantly change if we also included the other three quarters of the rounds. But we obviously don't have data from them. The fair sampling assumption is thus inherently untestable!

However, without the fair sampling assumption the door is wide open for alternative explanations of the experimental outcome. For example, one may postulate the existence of a (hidden) mechanism that activates the detector exactly in those one quarter of rounds of the experiment in which the source accidentally emitted two particles instead of only one. The observed data would then solely reflect rounds of the experiment in which two particles were simultaneously in the interferometer. The interference pattern could thus be explained by a hypothesised interaction between these particles. To exclude this possibility, we really need the fair sampling assumption.

This puts us in a dilemma: To draw sensible conclusions from our experiment, we need extra assumptions, such as fair sampling. Conversely, these assumptions are not themselves experimentally testable. A way out could be to design experiments whose analysis does not require such untestable assumptions, or only weak ones. This is precisely the idea of “device independence”.

The loophole-free Bell tests that have been carried out recently by several experimental groups are excellent examples of this paradigm. The conclusions drawn from these experiments hold independently of any assumptions on the devices that have been used to generate the data, and in particular without the fair sampling assumption. This is crucial in this case, as the result is a very far-reaching one, namely that quantum mechanics cannot be turned into a deterministic theory without giving up locality.

While device independence may still appear to be a rather academic, if not paranoid, approach to physics, it has its natural place in quantum information processing and quantum cryptography. Here one would like to trust in the outcomes of computations or the security of encryption schemes even if the devices on which they are running are untrusted. Achieving such device-independent security is challenging, but recent progress, both on the theoretical and on the experimental side, is extremely promising. In fact, first proof-of-principle experiments are already underway.

This book by Dr. Rotem Arnon-Friedman provides an introduction into the theory of device-independent information processing. It explains general principles but also covers some of the latest research, to which the author contributed significantly. The book provides the toolbox that is necessary to understand and analyse current and future device-independent experiments, ranging from Bell tests to quantum cryptographic schemes.

Device independence may, on a first encounter, bear various surprises. Things that we normally take for granted have to be questioned—the fair sampling assumption described above is just one example. Rotem Arnon-Friedman elegantly guides the reader through this amazing subject. The path always returns to a set of “showcases” which explain how the abstract and general concepts are applied to concrete applications. It culminates in a full-fledged security proof of device-independent quantum key distribution, one of the most prominent applications in the field.

While this book focuses on information processing and cryptography, the device-independent approach ultimately also bears a message for physicists: It is worth questioning even the most naturally sounding and seemingly unavoidable assumptions. This message should of course not be new to them—the history of physics has told us that doing so is a key to deeper insights.

Zürich, Switzerland
September 2020

Renato Renner

Abstract

The field of *device-independent* quantum information processing concerns itself with devising and analysing protocols, such as quantum key distribution and quantum tomography, without referring to the quality of the physical devices utilised to execute the protocols. Instead, the analysis is based on the observed correlations that arise during repeated interactions with the devices and, in particular, their ability to violate the so-called Bell inequalities.

Since the analysis of device-independent protocols holds irrespectively of the underlying physical device, it implies that *any* device can be used to execute the protocols: If the apparatus is of poor quality, the users of the protocol will detect it and abort; otherwise, they will accomplish their goal. This strong statement comes at a price—the analysis of device-independent protocols is, a priori, extremely challenging. Having good techniques at hand is thus crucial.

The thesis presents an approach that can be taken to simplify the analysis of device-independent information processing protocols. The idea is the following: Instead of analysing the most general device leading to the observed correlations, one should first analyse a significantly simpler device that, in each interaction with the user, behaves in an identical way, independently of all other interactions. We call such a device an independently and identically distributed (IID) device. As the next step, special techniques are used to prove that, without loss of generality, the analysis of the IID device implies similar results for the most general device. Such techniques reduce the problem of analysing the general scenario to that of analysing an IID one and, hence, we term them *reductions to IID*.

We present two mathematical techniques that can be used as reductions to IID in the device-independent setting: *de Finetti reductions for correlations* and the *entropy accumulation theorem*. Each technique is accompanied by a showcase-application that exemplifies the reduction's usage and benefits. Specifically, we use our de Finetti reduction to prove a *non-signalling* (super-quantum) *parallel repetition theorem*, belonging to a family of theorems discussed in theoretical computer science. The entropy accumulation theorem is used to prove the security of *device-independent quantum cryptographic protocols*.

Performing the analysis via a reduction to IID instead of directly analysing the most general scenarios leads to simpler proofs and significant quantitative improvements, matching the tight results proven when analysing IID devices. In particular, our analysis of device-independent quantum key distribution protocols produces essentially optimal key rates and noise tolerance, crucial for all future experimental implementations of device-independent cryptography.

Acknowledgements

I am lucky to be surrounded by people who inspire me, believe in me, and allow me to grow. There is no better way of spending our most precious resource, time, and thus to them I am grateful.

I would like to thank Renato Renner, my supervisor, who offered me an opportunity of a lifetime and opened the door to the never-ending quest of solving intriguing and challenging questions. Renato was a great inspiration; I was constantly amazed by his contributions to science and his way of thinking, as well as how determined he is to come up with a proof even when it seems impossible (these days, whenever I get stuck, I just tell myself that Renato would never give up!). Most of all, I am thankful to Renato for believing in me, moving mountains to allow me to focus on my research, and supporting me when I needed to focus on other things.

I appreciate and thank the co-examiners, Andreas Winter, Nicolas Gisin, and Ran Raz, for taking the time to read my thesis, as well as Ernest Tan, Frédéric Dupuis, Jie Lin, Marco Tomamichel, and Thomas Vidick for their valuable comments on parts of the thesis.

I am thankful to my collaborators. Out of them, a special thanks goes to my office mate Christopher Portmann, who was always happy to discuss the details of the details, and to Thomas Vidick for being an inspiring and motivating collaborator, an invaluable mentor, and a great friend.

I was honored to be part of the QIT group at ETH, consisting of many talented people. Being part of this family allowed me to interact and learn from past and current members of the group to whom I thank.

My great appreciation to Marko Gebbers for enlightening me, helping me separate the wheat from the chaff, and challenging me to become a better version of myself. I am forever grateful.

Owing much more than that, I wish to thank Orna, Sheri, and Keren Arnon (mother, father, and sister, respectively), as well as my best friends, for always being there for me, even from far away, and for having an unbelievable amount of patience. I cannot thank enough my mother, Orna Arnon, and Tzvia Friedman who were always happy to fly to Zurich to help with the kids when I wanted to travel for

a conference and my father, Sheri Arnon, who illustrated an endless number of boxes for my papers, talks, and this thesis.

Finally, no words can describe my gratitude to Neer Friedman, who supported and believed in me from the first moment (pushing me into the building to meet Renato) to the last sentence of this thesis. I could have not done this without him.

Contents

1	Introduction	1
1.1	Motivation	1
1.1.1	Device-Independent Information Processing	1
1.1.2	Reductions to IID	4
1.2	Content of the Thesis	7
1.2.1	Reductions	8
1.2.2	Showcases	10
1.3	How to Read the Thesis	16
	References	17
2	Preliminaries: Basics and Notation	19
2.1	General Notation	19
2.2	Probability Distributions and Random Variables	20
2.2.1	Independent and Identical Random Variables	21
2.2.2	Concentration Inequalities	21
2.3	Quantum Formalism	22
2.3.1	Operators	22
2.3.2	Hilbert Spaces	23
2.3.3	Quantum States	24
2.3.4	Quantum Operations	27
2.4	Distance Measures	29
2.5	Entropies	30
2.5.1	Shannon and von Neumann Entropy	30
2.5.2	Min- and Max-Entropies	31
	References	32
3	Preliminaries: Device-Independent Concepts	33
3.1	Black Boxes	34
3.1.1	Non-signalling Boxes	34
3.1.2	Quantum Boxes	36

3.1.3	Classical Boxes	36
3.1.4	Correlations' Space	37
3.2	Bell Inequalities	38
3.2.1	Non-local Games	39
3.2.2	The CHSH Game	41
3.3	Untrusted Devices	42
	References	45
4	Introduction to the Showcases	47
4.1	Introduction to Non-signalling Parallel Repetition	47
4.1.1	Parallel Repeated Games	48
4.1.2	Threshold Theorems	49
4.2	Introduction to Device-Independent Quantum Cryptography	51
4.2.1	DIQKD Security Definitions	54
4.2.2	DIQKD Protocol	56
4.2.3	Main Task of a Security Proof	61
4.2.4	The Honest Implementation	61
4.2.5	Model of an Arbitrary Device	63
	References	63
5	Single-Round Box	67
5.1	The Model	69
5.1.1	Quantum Single-Round Boxes	70
5.2	Showcase: Device-Independent Quantum Cryptography	71
	References	73
6	Multi-round Box	75
6.1	Parallel Interaction	76
6.1.1	Non-signalling Parallel Boxes	77
6.2	Sequential Interaction	78
6.2.1	Without Communication Between the Rounds	79
6.2.2	With Communication Between the Rounds	82
	References	85
7	Working Under the IID Assumption	87
7.1	The IID Assumption	88
7.2	Asymptotic Equipartition Property	90
7.2.1	Classical Asymptotic Equipartition Property	90
7.2.2	Quantum Asymptotic Equipartition Property	93
7.3	Using the IID Assumption	94
7.3.1	Showcase: Non-signalling Parallel Repetition	94
7.3.2	Showcase: Device-Independent Quantum Cryptography	95
7.4	Beyond IID	99
	References	100

- 8 Reductions to IID: Parallel Interaction** 103
 - 8.1 Permutation Invariance 104
 - 8.2 de Finetti Reductions for Correlations 106
 - 8.3 Ways of Using the Reductions 109
 - 8.3.1 Post-selecting Permutation Invariant Boxes 110
 - 8.3.2 Failure Probability of a Test 111
 - 8.3.3 Diamond Norm 113
 - 8.4 Impossibility Results 116
 - 8.4.1 Restricted de Finetti Box 116
 - 8.4.2 Extension to an Adversary 118
 - 8.4.3 Other de Finetti-Type Theorems 120
 - References 121
- 9 Reductions to IID: Sequential Interaction** 123
 - 9.1 Sequential Quantum Processes 124
 - 9.2 Entropy Accumulation Theorem 125
 - 9.2.1 Conceptual Difficulties to Overcome 126
 - 9.2.2 Prerequisites of the Theorem 127
 - 9.2.3 Statement of the Theorem 131
 - References 135
- 10 Showcase: Non-signalling Parallel Repetition** 137
 - 10.1 Main Challenge and Goal 139
 - 10.2 Approximately Non-signalling Marginals 140
 - 10.2.1 Single-Round Boxes from Frequencies 141
 - 10.2.2 Signalling Test 142
 - 10.2.3 Guessing Game 146
 - 10.3 Threshold Theorem 151
 - 10.3.1 Winning Probability of Approximately Non-signalling Strategies 151
 - 10.3.2 Final Result 153
 - 10.4 Open Questions 156
 - References 157
- 11 Showcase: Device-Independent Quantum Cryptography** 159
 - 11.1 Main Challenge and Goal 160
 - 11.2 Device-Independent Entropy Accumulation 161
 - 11.2.1 EAT Channels 163
 - 11.2.2 Min-Tradeoff Function 165
 - 11.2.3 Smooth Min-Entropy Rate 169
 - 11.3 Device-Independent Quantum Key Distribution 172
 - 11.3.1 Completeness 173
 - 11.3.2 Soundness 174
 - 11.3.3 Key Rate Analysis 178

- 11.4 Open Questions 181
 - 11.4.1 Experimental Realisations 181
 - 11.4.2 Possible Extensions 182
 - 11.4.3 Bounding the von Neumann Entropy 184
- References 184
- 12 Outlook 187**
 - 12.1 Two-Party Device-Independent Quantum Cryptography 188
 - 12.2 Parallel Device-Independent Quantum Cryptography 189
 - 12.3 Device-Independent Tomography 189
- References 190

- Appendix A: Additional Proofs: de Finetti Reductions 191**
- Appendix B: Additional Proofs: Non-signalling Parallel Repetition 199**
- Appendix C: Additional Proofs: Device-Independent Quantum
Cryptography 207**

Chapter 1

Introduction



1.1 Motivation

1.1.1 Device-Independent Information Processing

The study of quantum information unveils new possibilities for remarkable forms of computation, communication, and cryptography by investigating different ways of manipulating quantum states. Crucially, the analysis of quantum information processing tasks must be based, in one way or another, on the actual physical processes used to implement the considered task; the physical processes must be inherently quantum as otherwise no advantage can be gained compared to classical information processing. In most applications, the starting point of the analysis is an explicit and exact characterisation of the quantum apparatus, or device, used to implement the task of interest.

As an example, consider the task of quantum key distribution (QKD). In a QKD protocol, the goal of the honest parties, called Alice and Bob, is to create a shared key, unknown to everybody else but them. The protocol is intrinsically quantum: To execute it Alice and Bob hold entangled quantum states in their laboratories and perform quantum operations, or measurements, on the quantum states. Informally, proving the security of a QKD protocol amounts to showing that no adversary can hold (significant) information about the produced key. To prove security one usually needs to have a complete description of the quantum devices, i.e., the quantum states and measurements, used by Alice and Bob. For example, the security proof of the celebrated BB84 protocol [1] builds on the assumptions that Alice and Bob hold two-qubit states and are able to measure them in a specific way. When these assumptions are dropped, the protocol is no longer secure [2]. Thus, if Alice and Bob wish to use their quantum devices in order to implement a QKD protocol they need to first make sure that the device is performing the exact operations described by the protocol.

Unfortunately, in practice we are unable to fully characterise the physical devices used in quantum information processing tasks. Even the most skilled experimentalist will recognise that a fully characterised, always stable, large-scale quantum device that implements a QKD protocol is extremely hard to build. If the honest users' device is different from the device analysed in the accompanying security proof, security is no longer guaranteed and imperfections can be exploited to attack the protocol.

Noise and imperfections cannot be completely avoided when implementing quantum information processing tasks. Furthermore, imperfections being imperfections, one also cannot expect to perfectly characterise them. That is, we cannot say for sure what exactly is about to go wrong in the quantum devices: Maybe the measurements are not well-calibrated, perhaps some noise introduces correlations between particles which are intended to be independent, or interaction with the environment may possibly lead to decoherence. Even the advent of fault-tolerant computation, if achievable one day, cannot resolve all types of errors if no promise is given regarding the number of errors and their, possibly adversarial, nature. Once we come to terms with the above, a natural question arises:

Can quantum information processing tasks be accomplished by utilising uncharacterised, perhaps even adversarial, physical devices?

An adversarial, or malicious, device is one implemented by a hostile party interested in, e.g., breaking the cryptographic protocol being executed. Clearly, this is an extreme scenario to consider. Note, however, that even if the manufacturer of the device is to be trusted, he may still be incompetent—the physical apparatus will be subject to uncharacterised imperfections even though the manufacturer is honest and has good intentions.

The field of device-independent information processing addresses the above question. In the device-independent framework we treat the physical devices, on which a minimal set of constraints is enforced,¹ as *black boxes*—Alice and Bob hold a box and can interact with it classically (as explained below) to execute the considered protocol, but they cannot open it to assess its internal workings.² They have no knowledge regarding the physical apparatus and do not trust that it works as alleged by the manufacturer of the device.

What *can* Alice and Bob do with the black box? They can interact with it by pushing buttons, each associated with some classical input (e.g., a bit) and record the classical outputs produced by the box in response to pressing its buttons. Thus, the

¹Clearly, one cannot perform any cryptographic task if the device includes a transmitter that just sends all the information to the adversary. Few minimal assumptions regarding the device will be needed; see Sect. 3.3. Depending on the considered task, some of the assumptions can be enforced in practice while others may require some minimal level of trust.

²Notice that even if Alice and Bob did have some information about the physical apparatus, the device-independent framework does not allow them to take advantage of this information in the analysis. For example, Alice and Bob may be able to distinguish a device that uses the polarisation of a photon to encode a qubit from one based on superconducting qubits (even the author is able to do that). Yet, this information is not to be used when treating the device as a black box.

only information available to Alice and Bob is the observed classical data created during their interaction with the black box. (Hence the name “device-independent”).

Since the device is not to be trusted, the classical information collected by Alice and Bob during the interaction with the box must allow them, somehow, to test the possibly faulty or malicious device and decide whether using it, e.g., to create their keys by executing a QKD protocol, poses any security risk. A protocol or task is said to be device-independent if it guarantees that by interacting with the device according to the specified steps the parties will either abort, if they detect a fault, or accomplish the desired task (with high probability).

The possibility of device-independent information processing is quite surprising. Indeed, restricting ourselves to classical physics and classical information, it is impossible to derive device-independent statements.³ The most important ingredients for device-independent protocols are the existence of Bell inequalities and quantum “non-local” correlations that violate them [3]. These two facts are far from trivial and play a fundamental role in quantum theory. In the context of device-independent information processing, a Bell inequality acts as a “test for quantumness” that allows the users of the device to verify that their device is “doing something quantum” and cannot be simulated by classical means. This “quantumness”, of a specific form discussed below, is what allows us to, e.g., prove security of a QKD protocol.

A Bell inequality can be thought of as a multi-player game, also called a non-local game, played by the parties using the device they share. A non-local game goes as follows. A referee asks each of the (cooperating) parties a question chosen according to a given probability distribution. The parties need to supply answers which fulfil a pre-determined requirement according to which the referee accepts or rejects the answers. In order to do so, they can agree on a strategy beforehand, but once the game begins communication between the parties is not allowed. If the referee accepts their answers the players win. The goal of the parties is, naturally, to maximise their winning probability in the game.

Different devices held by the parties implement different strategies for the game and may lead to different winning probabilities. In the device-independent setting we are interested in games that have a special “feature”—there exists a quantum device which achieves a winning probability in the game that is greater than all classical, local, devices.

Crucially, the winning probability in the game does not merely indicate that the device is doing something quantum but how non-classical it is. Relations are known between the probability of winning some non-local games and various other quantities. Some examples for quantities of interest are the entropy produced by the device, the amount of entanglement consumed to play the game, or the distance (under an appropriate distance measure) of the device from a specific fully characterised quantum device. Such relations lie at the heart of any analysis of device-independent information processing tasks.

³Consider for example the case of device-independent QKD. Classical devices can always be pre-programmed by the adversary to output a fixed key of her choice.

Although above we only mentioned device-independent QKD as an example for a device-independent task, the framework of device-independence does not only concern the more-than-average paranoid cryptographers. The framework fits any scenario in which, a priori, we do not want to assume anything about the utilised devices and their underlying physical nature. To reassure the reader, we give three additional examples.

Bell inequalities were originally introduced in the context of the foundations of quantum mechanics in order to resolve the EPR paradox [4]. When trying to test quantum theory against an alternative classical world that admits a “local hidden variable model” (or, in other words, falsify all classical explanations of a behaviour of a physical system), one cannot assume that quantum theory holds to begin with and must treat the device as a black box without assuming to know its internal workings.

A second example is that of blind tomography, also termed self-testing. Assume a quantum state is being produced in some experimental setting. Quantum tomography is the process of estimating which state is being created by performing measurements on copies of the state and collecting the statistics [5]. To get a meaningful estimation, a certain set of measurements needs to be used, depending on the dimension of the state. In other words, in order to estimate and characterise the quantum state, we must be able to first characterise the measurement devices. *Blind* quantum tomography refers to the process in which the measurements are also unknown. In such a case, nothing but the observed statistics can be used [6, 7].

Another interesting example is that of verification of computation—given a device claimed to be a quantum computer, how can human beings, who cannot perform quantum computations by themselves, verify that this is indeed the case? There are different ways of addressing this question, but in all cases we would like to make statements without presuming that the considered devices are performing any particular quantum operations (see, e.g., [8]).

The device-independent framework becomes relevant whenever one wishes to make concrete statements without referring to the underlying physical nature of the utilised devices and the types of imperfections or errors that may occur. The derived statements are extremely strong. Device-independent security, for example, is regarded as the gold standard for quantum cryptography, since attacks exploiting the mismatch between security proof and implementation are no longer an issue. Making such strong statements comes at a price. The analysis of device-independent tasks is, a priori, extremely challenging: We treat the devices as black boxes and thus the proofs need to account for an almost arbitrary, even adversarial, behaviour of the devices. Having good techniques for the analysis at hand is therefore crucial. This is further discussed in the following section.

1.1.2 Reductions to IID

In the device-independent setting one does not have a description of the specific device used in the considered task and, hence, must analyse the behaviour of arbitrary

devices. For example, when proving security of cryptographic protocols we clearly need to consider *any* possible device that the adversary may prepare. Unfortunately, analysing the behaviour of arbitrary devices can be wearying at best and infeasible at worst. Let us start by explaining why this is the case.

As mentioned above, the ability to achieve device-independent information processing tasks is based on the existence of non-local games and quantum strategies to play them that can beat any classical strategy. To perform complex tasks, such as device-independent cryptography, employing the device to play a *single* non-local game is clearly not enough; we cannot conclude any meaningful information regarding the device by asking it to produce outputs only for a single game. To put quantum information to work we must consider protocols in which the device is used to play *many* non-local games. This way, the parties executing the protocol can collect statistics and test their device. If the device does not pass the test the parties abort the protocol (see Protocol 1.1 below for an example).

The reason for the difficulty of the analysis lies in the fact that one needs to examine the overall behaviour of the device during the entire execution of the protocol, consisting of playing many games with the device, instead of its behaviour in a single game. As the device is uncharacterised its actions when playing one game may depend on other games.

In general, there are two families of devices able to play many games that one can consider—parallel and sequential devices. A parallel device is one which can be used to play *all* the games *at once*. That is, the parties executing the protocol are instructed to give all the inputs, for all the games, to the device and only then the device produces the outputs for all the games. In such a case, the actions of the device in one game may depend on *all* other games.

A sequential device, on the other hand, is used to play the games *one after the other*, i.e., the parties give the device the first inputs and wait for its outputs and only then proceed to play the next game. In between the games, some communication may be allowed between the parties and the different components of the device. In the case of a sequential device, the behaviour of the device in one game may depend on all *previous* games as well as communication taking place during the time between the games.⁴ In both cases, the input-output behaviour of the devices gets quite complicated.

One common assumption introduced to simplify the analysis of device-independent information processing tasks is the so called “independent and identically distributed” (IID) assumption. As the name suggests, a device is said to be an IID device if it plays each of the games independently of the others and utilises the same strategy for all games. An IID device is a special case of both parallel and sequential devices and, since it is highly structured, analysing its behaviour can be significantly simpler than analysing the more general devices; see Fig. 1.1.

The IID assumption heavily restricts the structure of the device. It is therefore not clear at all that analysing device-independent information processing tasks under the IID assumption is sufficient. Returning to the example of device-independent

⁴The formal definitions of parallel and sequential devices are given in Chap. 6.

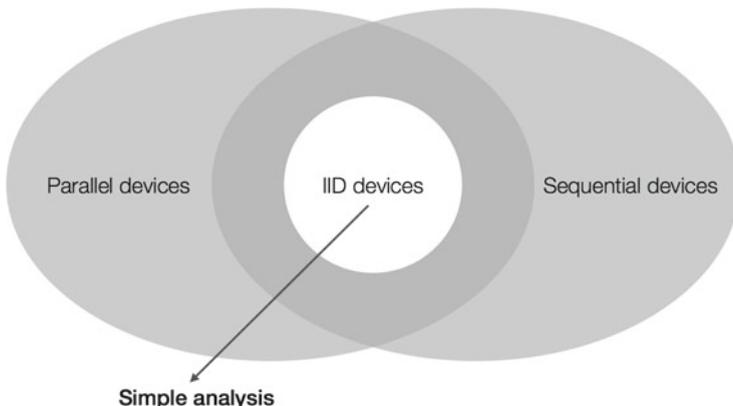


Fig. 1.1 The relation between the different sets of devices. The intersection of the sets of sequential and parallel devices includes the set of IID devices. The analysis of IID devices, i.e., that done under the IID assumption, is rather simple

cryptography, an adversary who can prepare arbitrary devices (let it be sequential or parallel) may be strictly stronger, i.e., can get more information about the outputs of the honest parties, than an adversary restricted to IID devices. Thus, simplifying the analysis by using the IID assumption comes at the cost of weakening the final statement.

The main question addressed in this thesis is the following:

Can the analysis of device-independent information processing tasks be reduced to that performed under the IID assumption?

The term *reduction* is widely used in theoretical computer science and is meant to describe the process of showing that one problem is as hard/easy as another. In our case, we ask whether analysing general devices is as easy as analysing IID devices or, in other words, does an analysis performed under the IID assumption imply results concerning general devices (i.e., statements which are not restricted to the IID case). A priori, it is not at all obvious that this is the case; clearly, not all devices are IID devices. A positive answer to the above question means that *even though* there exist devices that cannot be described as IID ones, it is sometimes possible to restrict the attention solely to IID devices and the rest will follow.

The idea of applying a reduction to IID as a proof technique was conceived⁵ in [9], following which a concrete reduction relevant for applications was developed in [10] and used to reduce the security proof of QKD protocols to that done under the IID assumption.⁶ As such, [10] acts as the first example for a proof using a reduction to IID.

⁵Perhaps surprisingly, as far as the author is aware the idea of a “reduction to IID” does not appear or used in classical information processing and cryptography.

⁶In the context of QKD, security under the IID assumption is called security against collective attacks.

Analysing information processing tasks via a reduction to IID has several significant advantages. Analysing IID devices is relatively easy and almost always intuitive. Thus, having tools that allow us to extend the analysis to the general case greatly simplifies proofs.⁷ The simplicity, in turn, allows for clear and modular statements as well as quantitatively strong results.⁸

The importance of quantitatively strong results is obvious, especially when discussing quantum information processing tasks: If we wish to benefit from the new possibilities brought by the study of quantum information, we must be able to implement the protocols in practice. Without strong quantitative bounds on, e.g., key rates and tolerable noise levels, we cannot take the device-independent field from theory to practice. Clarity and modularity should also not be dismissed. Science is not a “one-man’s job”; clarity and modularity are crucial when advancing science as a community. Indeed, complex and fine-tuned proofs are hard to verify, adapt to other cases of interest, and quantitatively improve.

Another advantage of reducing a general analysis to IID is that it allows us to separate the wheat from the chaff. The essence of the arguments used in proofs of information processing tasks almost always enter the game in the analysis of the IID case. Proofs that address the most general scenarios directly (i.e., not via a reduction to IID) are at risk of obscuring the “physics” by more technical mathematical steps. When using a reduction to IID this is (mostly) not the case—the essence, or the interesting part, lies in the analysis of IID devices while the technicalities are pushed into the reduction itself.

As we will show in the thesis, reductions to IID can also be developed and employed in device-independent quantum information processing. We present two techniques that can be used as reductions to IID, accompanied by two showcase-applications that illustrate how the reductions can be used and their benefits in terms of the derived theorems. The following section presents the content of the thesis in more detail.

1.2 Content of the Thesis

The goal of the thesis is to explain how reductions to IID can be performed in the context of device-independent information processing. To this end, after explaining the different mathematical objects that one needs to consider and their relevance, we discuss the IID assumption and its implications in the device-independent setting. We then present two techniques, or tools, that can be used as reductions to IID in

⁷The reductions themselves are not necessarily simple, but that is fine. They are technical tools that are only proved once and can then be used to simplify many other proofs. The researcher using the reduction does not need to reprove anything.

⁸This is in agreement with Occam’s razor; while there is no notion of the “right proof” out of several possible proofs (assuming they are all mathematically correct), the simplest proof usually turns out to be the most useful and insightful one.

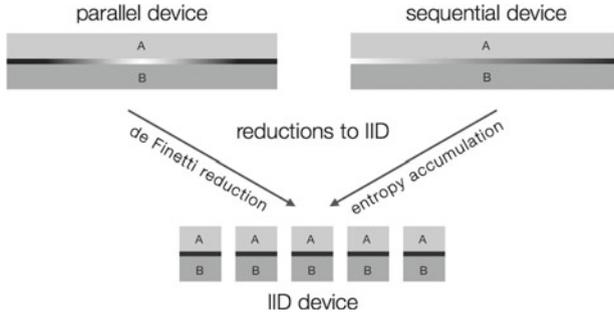


Fig. 1.2 Reductions to IID in device-independent information processing. de Finetti reductions can be used to reduce the study of parallel devices to IID device (see Chaps. 8 and 10), while the entropy accumulation theorem can be used when dealing with sequential devices (Chaps. 9 and 11)

the analysis of device-independent information processing tasks, one relevant for parallel devices and the other for sequential ones.

To better comprehend the topic and exemplify the usage of the two reductions, we consider two applications as showcases, namely, parallel repetition of non-local games and device-independent cryptography. These are studied in detail throughout the chapters of the thesis, while taking the perspective of reductions to IID.

1.2.1 Reductions

Two types of reductions are presented. The reductions are applicable in different scenarios and give statements of different forms; see Fig. 1.2.

1.2.1.1 de Finetti Reduction for Correlations

The first reduction, the topic of Chap. 8, is called “de Finetti reduction for correlations” and was developed in [11]. The de Finetti reduction is relevant for the analysis of *permutation invariant parallel devices*. Permutation invariance is an inherent symmetry in many information processing tasks, device-independent tasks among them. Thus, analysing permutation invariant devices is of special interest.

In short, in our context, a de Finetti reduction is a theorem that relates any permutation invariant parallel device to a special type of device, termed de Finetti device, which behaves as a convex combination of IID devices (see Chap. 8 for the formal definitions). The given relation acts as a reduction to IID when considering tasks admitting a permutation invariance symmetry and in which a parallel device needs to be analysed. Our showcase of parallel repetition of non-local games fits this description and thus can benefit from our de Finetti reduction.

Various quantum de Finetti theorems were known prior to our work and were successfully used to substantially simplify the analysis of many quantum information tasks. However, they cannot be applied in the device-independent setting, since they make many assumptions regarding the permutation invariant quantum states being analysed and therefore cannot accommodate uncharacterised devices. The unique property of the reduction presented in Chap. 8 is that, apart from permutation invariance, it makes no assumptions whatsoever regarding the systems of interest and is therefore applicable in the analysis of device-independent information processing.

For pedagogical reasons, we choose to present in the thesis a de Finetti reduction which is relevant to the case of bipartite devices, i.e., devices which are shared between two parties, Alice and Bob. The statements can be extended to any number of parties, as shown in [11]; the proofs of the general case do not include fundamental insights on top of those used in the bipartite case but require somewhat heavy notation. We therefore omit the more general theorems and proofs (while supplying the full analysis of the bipartite case in Chap. 8), with the hope of making the content more inviting for readers unfamiliar with the topic.

Apart from presenting the reduction and the possible ways of using it, Chap. 8 also includes a discussion of ways in which it may be possible to extend or modify the reduction (to be more specific, we mainly present impossibility results). This content does not appear in detail in other published papers and can be relevant for future studies of the topic.

1.2.1.2 Entropy Accumulation Theorem

The second reduction to IID that can be used in the device-independent setting is the entropy accumulation theorem (EAT) [12] and is the topic of Chap. 9. The EAT can be seen as an extension of the entropic formulation of the asymptotic equipartition property (AEP) [13, 14], applicable only under the IID assumption, to more general sequential processes.

The AEP, presented in Chap. 7, basically asserts that when considering IID random variables, the smooth min- and max-entropies of the random variables converge to their von Neumann (or Shannon, in the classical case) entropy, as the number of copies of the random variable increases. The AEP is of great importance when analysing, both classical and quantum, information processing tasks under the IID assumption: It explains why the von Neumann entropy is so important in information theory—the smooth entropies, which describe operational tasks, converge to the von Neumann entropy when considering a large number of independent repetitions of the relevant task.⁹

⁹A commonly used example is that of “data compression”. There, one would like to encode an n bit string using less bits. If we allow for some small error when decoding the data, the smooth max-entropy roughly describes the number of bits needed. However, for a large enough number of independent repetitions, less bits suffice and the exact amount is governed by the Shannon entropy.

Moving on from the IID setting, the EAT considers a certain class of quantum sequential processes. That is, in our context, it is relevant when studying *sequential devices*.¹⁰ Similarly to the AEP, when applicable, the EAT allows one to bound the total amount of the smooth min- and max-entropies using the same bound on the von Neumann entropy calculated for the IID analysis, i.e., the one used when applying the AEP. In this sense, the EAT can be seen as a reduction to IID—with the aid of the EAT the analysis done under the IID assumption using the AEP can be extended to the one relevant for sequential devices.

The proof of the EAT is not presented in the thesis (and should not be attributed to the author). We focus on motivating, presenting, and explaining the EAT in the form relevant for device-independent quantum information processing [15] (as well as quantum cryptography in general), so it can be later used in our showcase of device-independent cryptography. The pedagogical presentation of the EAT given in Chap. 9 does not appear in full in any other published material and we hope that it will make the theorem more broadly accessible.

Before presenting our showcases, let us remark that both of the reductions mentioned above are not “black box” reductions, in the sense that one cannot simply say that if a problem is solved under the IID assumption then it is solved in the general case. In particular, one should be familiar with the exact statements of the reductions (though not with their proofs) as well as the analysis of the considered task under the IID assumption in order to apply the reductions (or even just check whether they are applicable or not). When discussing the reductions in Chaps. 8 and 9, we explicitly explain in what sense the presented tools count as reductions to IID techniques.

1.2.2 Showcases

We use two showcases throughout the thesis in order to exemplify the approach of reductions to IID and the more technical usage of the presented reductions. The showcase of parallel repetition of non-local games uses the de Finetti reduction technique while the showcase of device-independent cryptography builds on the EAT. As mentioned in Sect. 1.1.2 above, we believe that analysing device-independent tasks using a reduction to IID has its benefits. The derived theorems are, arguably, more intuitive and insightful and, in addition, give strong quantitative results.

We shortly discuss below each of our showcases. We present informal theorems describing the results proven for the showcases. The informal theorems shed light on the fundamental nature and strength of the approach of reductions to IID.

¹⁰To be more precise, some requirements regarding the process, or protocol, in which the sequential device is to be used must hold. This is explained in details in Chap. 9.

1.2.2.1 Non-signalling Parallel Repetition

Our first showcase is that of non-signalling parallel repetition. Chapter 10 presents our formal statements and proofs, which previously appeared in [16]. As before, we focus in the thesis on the bipartite case for pedagogical reasons; [16] includes the general analysis, which is valid for any number of parties playing the game.

Non-local games, as mentioned in Sect. 1.1.1, are games played by several cooperating parties, also called players. A referee asks each of the players a question chosen according to a given probability distribution. The players need to supply answers which fulfil a pre-determined requirement according to which the referee accepts or rejects the answers. In order to do so, they can agree on a strategy beforehand, but once the game begins communication between the parties is no longer allowed. If the referee accepts their answers the players win.

In the language used so far, we can think of a device as implementing a strategy for the game. Depending on the field of interest, one can consider classical, quantum, or non-signalling devices, the latter referring to devices on which the only restriction is that they do not allow the players to communicate. We focus below on the case of non-signalling strategies, or devices.

One of the most interesting questions regarding non-local games is the question of parallel repetition. Given a non-local game with optimal winning probability $1 - \alpha$ using non-signalling strategies, we are interested in analysing the optimal winning probability of a non-signalling strategy in the repeated, or threshold, game. A threshold game is a game in which the referee asks the players to play $n \in \mathbb{N}$ instances of the non-local game, all at once, and the players' goal is to win more than $1 - \alpha + \beta$ fraction of the games, for $\beta > 0$ a parameter of the threshold game. The parallel repetition question concerns itself with upper-bounding the optimal winning probability in the threshold game, as the number of games n increases.¹¹

One trivial strategy that the players can use in the threshold game is a strategy employing a non-signalling IID device. That is, they simply answer each of the n questions independently using the optimal non-signalling device used to play a single game. Using an IID device, the fraction of successful answers is highly concentrated around $1 - \alpha$ and the probability to win more than a $1 - \alpha + \beta$ fraction of the games decreases exponentially fast with $n\beta^2$, as follows from the optimal formulation of the Chernoff bound.

However, since the players receive from the referee all the questions to the n instances of the non-local game at once, an IID device is not the most general device that they can use. Instead, they can use any non-signalling parallel device to implement their strategy. As parallel devices are strictly more general than IID ones, using parallel devices in fact allows them to win the threshold game with higher probability than in the IID case.¹² Still, one may ask how the winning probability behaves for a

¹¹This is actually a generalisation of the more commonly known parallel repetition question, in which one wishes to upper-bound the probability of winning *all* the n games.

¹²When first encountering the question of parallel repetition it may seem surprising that the players can do better using a parallel device, but this is indeed the case; see Sect. 4.1.2 a concrete example.

sufficiently large number of repetitions n and, especially, whether it decreases in a similar fashion as for IID strategies.

To answer the above question, we wish to reduce the study of strategies employing parallel devices to those using IID devices. A crucial observation that allows us to do so is that the threshold game itself admits a permutation invariance symmetry (i.e., the order of questions-answers tuples does not matter; see Chap. 10 for the details) and, therefore, we can assume without loss of generality that the optimal strategy is also permutation invariant. Now that we can restrict our attention to permutation invariant parallel devices, de Finetti reductions become handy and can be used as a tool for reduction to IID.

In Chap. 10 we consider the case of non-signalling strategies for complete-support games. A complete-support game is one in which all possible combinations of questions being sent to the players have some non-zero probability of being asked by the referee. We prove the following via a reduction to IID:

Theorem 1.1 (Informal) *Given a game with optimal non-signalling winning probability $1 - \alpha$, for any $\beta > 0$, the probability to win more than a fraction $1 - \alpha + \beta$ of n games played in parallel using a non-signalling strategy is exponentially small in $n\beta^2$, as in the IID case.*

Perhaps surprisingly, while the parallel repetition question is a well-investigated one, an exponential decrease that matches the IID case, as far as we are aware, was not known prior to our work (also not for classical or quantum strategies). In the context of reductions to IID, however, achieving the same behaviour as in the IID case is not unexpected.

To prove Theorem 1.1 we first prove another statement that has a “reduction to IID flavour” and is perhaps of more fundamental nature. To present it, however, we need to first set some notation.¹³

As mentioned above, we focus on two-player games, i.e., games played by Alice and Bob (and the referee). A parallel device used for the threshold game can be described using a conditional probability distribution $P_{AB|XY}$, where $\mathbf{A} = A_1, \dots, A_n$ is the random variable describing Alice’s answers in the threshold game (A_i being her answer in the i ’th game) and, similarly, $\mathbf{B} = B_1, \dots, B_n$ describes Bob’s answers, and $\mathbf{X} = X_1, \dots, X_n$ and $\mathbf{Y} = Y_1, \dots, Y_n$ are Alice’s and Bob’s questions, respectively.

When we say that a parallel device is non-signalling, we mean that it cannot be used as means of communication *between the parties*. The behaviour of the device in one

¹³We are jumping ahead now with the aim of being able to explain Theorem 1.2 to readers who are already somewhat familiar with device-independent information processing and non-signalling systems. For a reader unfamiliar with these topics, the mathematical statements may seem puzzling without further explanations. We will get back to the discussed theorem in Chap. 10, after giving all the preparatory information throughout the thesis. A reader unfamiliar with the used terminology can therefore skip the current discussion without the risk of missing out.

game, however, may depend on the other games.¹⁴ Mathematically, this means that, while the marginals $P_{A|X}$ and $P_{B|Y}$ are proper conditional probability distributions, objects such as $P_{A_i|X_i}$ are not well-defined.

During the threshold game, the device used by the players produces the observed data in the n games: $\mathbf{a} = a_1, \dots, a_n$, $\mathbf{b} = b_1, \dots, b_n$, $\mathbf{x} = x_1, \dots, x_n$, and $\mathbf{y} = y_1, \dots, y_n$. These are distributed according to $Q_{XY}^{\otimes n} P_{AB|XY}$, where Q_{XY} denotes the distribution used by the referee to choose the questions in a single non-local game. $Q_{XY}^{\otimes n}$ is then the IID distribution according to which the questions are chosen in the threshold game. The observed data $\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y}$ can be used to calculate frequencies and define a “frequencies’ conditional probability distribution”, which we denote by $O_{ABXY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$, as:

$$O_{ABXY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}(\tilde{a}\tilde{b}\tilde{x}\tilde{y}) = \frac{\left| \left\{ i : (a_i, b_i, x_i, y_i) = (\tilde{a}, \tilde{b}, \tilde{x}, \tilde{y}) \right\} \right|}{n}$$

and define

$$O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})} = \frac{O_{ABXY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}}{Q_{XY}}. \quad (1.1)$$

$O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ can be seen as a (not necessarily physical) device, or a strategy, for a single game. Starting with IID devices, which can be written in the form of¹⁵ $P_{AB|XY} = O_{AB|XY}^{\otimes n}$, it holds that if the device $O_{AB|XY}$ is non-signalling then $P_{AB|XY}$ is non-signalling and vice versa. This also implies that, for sufficiently large n , $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ is non-signalling with high probability.

For a non-IID, but non-signalling, device $P_{AB|XY}$, however, it is not clear at all that $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ should be non-signalling as well. Using a reduction to IID, the following theorem is proven:

Theorem 1.2 (Informal) *Let $P_{AB|XY}$ be a non-signalling permutation invariant parallel device and $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ as in Eq. (1.1). Then, for sufficiently large n , $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ is close to a non-signalling device with high probability. In particular, this means that the observed data produced by a non-signalling permutation invariant parallel device can be seen as if, with high probability, it was sampled using an IID device $O_{AB|XY}^{\otimes n}$ in which every single device $O_{AB|XY}$ is close to a non-signalling one.*

Theorem 1.1 follows directly from Theorem 1.2 by noting that the number of games won in a given use of the device can be directly read from $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ and that if $O_{AB|XY}^{\text{freq}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})}$ is close to a non-signalling device then its winning probability cannot be too far from the optimal non-signalling winning probability $1 - \alpha$.

¹⁴In other words, the local strategy of each player does require “communication between the games”: In order to (locally) answer the i ’th question received from the referee, the player needs to know his j ’th question (with $i \neq j$).

¹⁵An IID device is illustrated in the bottom of Fig. 1.2. We can then think of each copy $O_{AB|XY}$ as describing a single copy of the smaller boxes in the figure, while $P_{AB|XY} = O_{AB|XY}^{\otimes n}$ described the device including all the copies together.

Protocol 1.1 Device-independent quantum key distribution protocol (simplified example)

Given: A device for Alice and Bob that can play the chosen Bell game repeatedly

- 1: For every round $i \in [n]$ do Steps 2-3:
 - 2: Alice and Bob choose X_i, Y_i at random.
 - 3: They input X_i, Y_i to the device and record the outputs A_i, B_i .
 - 4: **Parameter estimation:** Alice and Bob estimate the average winning probability in the game from the observed data. If it is below the expected winning probability, ω_{exp} , they **abort**.
 - 5: **Classical post processing:** Alice and Bob apply an error correction protocol and a privacy amplification protocol (both classical) on their raw data \mathbf{A} and \mathbf{B} .
-

1.2.2.2 Device-Independent Quantum Cryptography

Chapter 11 is devoted to the analysis of our second showcase—device-independent cryptography. The chapter’s content previously appeared in [17]. The most challenging cryptographic task in which device-independent security has been considered is device-independent QKD (DIQKD); we will use this task as our main example. In DIQKD the goal of the honest parties, called Alice and Bob, is to create a shared key, unknown to everybody else but them. To execute the protocol they hold a device consisting of two parts: Each part belongs to one of the parties and is kept in their laboratories. Ideally, the device performs measurements on some entangled quantum states it contains.

The basic structure of a DIQKD protocol is presented as Protocol 1.1. The protocol consists of playing n non-local games, one after the other, with the given untrusted device and calculating the average winning probability from the observed data (i.e., Alice and Bob’s inputs and outputs). If the average winning probability is below the expected winning probability ω_{exp} defined by the protocol, Alice and Bob conclude that something is wrong and *abort* the protocol. Otherwise, they apply classical post-processing steps that allow them to create identical and uniformly distributed keys. (The full description of the considered DIQKD protocol is presented and discussed in the following chapters).

The central task when proving security of DIQKD consists in bounding the information that an adversary, called Eve, may obtain about Alice’s raw data $\mathbf{A} = A_1, \dots, A_n$ used to create the final key (see Protocol 1.1). More concretely, one needs to establishing a lower bound on the smooth conditional min-entropy $H_{\min}^{\varepsilon}(\mathbf{A}|E)$, where E is Eve’s quantum system, which can be initially correlated to the device used by Alice and Bob in the protocol and $\varepsilon > 0$ is one of the security parameters of the protocol (see Sect. 4.2). The quantity $H_{\min}^{\varepsilon}(\mathbf{A}|E)$ determines the maximal length of the secret key that can be created by the protocol. Hence, proving security amounts to lower-bounding $H_{\min}^{\varepsilon}(\mathbf{A}|E)$. Evaluating the smooth min-entropy $H_{\min}^{\varepsilon}(\mathbf{A}|E)$ of a large system is often difficult, especially in the device-independent setting where Alice and Bob are using an uncharacterised device, which may also be manufactured by Eve.

The IID assumption is commonly used in order to simplify the calculation of $H_{\min}^\varepsilon(A|E)$. In the IID case we can assume that Alice and Bob use an IID device to execute the protocol and, hence, each A_i is produced independently of all other outputs. Furthermore, one can assume that Eve’s quantum information also takes the IID form $E = E_1, \dots, E_n$, where each E_i holds information only regarding A_i . Then, the AEP, briefly mentioned above, can be used to calculate an upper-bound on $H_{\min}^\varepsilon(A|E)$ and, by this, prove security.

The most general adversarial device to consider is, clearly, not an IID one. Due to the sequential nature of the protocol, the relevant devices to consider are sequential devices. As sequential devices are more complex than IID ones, security proofs for DIQKD that proved security by addressing the most general device directly, e.g., [8, 18], had to use techniques which are far more complicated than the ones used for security proofs under the IID assumption, e.g., in [2]. Consequently, the derived security statements were of limited relevance for practical experimental implementations; they are applicable only in an unrealistic regime of parameters, e.g., small amount of tolerable noise and large number of signals.

We take the approach of reductions to IID in order to prove the security of our DIQKD protocol. In particular, we leverage the sequential nature of the protocol, as well as the specific way in which classical statistics are collected by Alice and Bob, to prove its security by reducing the analysis of sequential devices to that of IID devices using the EAT. The resulting theorem can be informally stated as follows:

Theorem 1.3 (Informal) *Security of DIQKD in the most general case follows from security under the IID assumption. Moreover, the dependence of the key rate on the number of rounds of the protocol, n , is the same as the one in the IID case, up to terms that scale like $1/\sqrt{n}$.*

On the fundamental level, the theorem establishes the a priori surprising fact that general quantum adversaries are no stronger than an adversary restricted to preparing IID devices. As mentioned in Sect. 1.1.2, this does not mean that the most general device that an adversary can prepare is an IID device. Instead, it means that the adversary (at least asymptotically) does not benefit from preparing more complex devices.

On the quantitative level, taking the path of a reduction to IID results in a proof with several advantages. In particular, it allows us to give simple and modular security proofs of DIQKD (as well as other device-independent protocols) and to extend tight results known for DIQKD under the IID assumption to the most general setting, thus deriving essentially optimal key rates and noise tolerance. This is crucial for experimental implementations of device-independent protocols. Our quantitative results have been applied to the analysis of the first experimental implementation of a protocol for randomness generation in the fully device-independent framework [19].

1.3 How to Read the Thesis

We review the structure of the thesis. Depending on the reader’s main interest and prior knowledge, different chapters of the thesis may or may not be relevant.

Chapters 2 and 3 give preliminary information. Chapter 2 presents general introductory information and notation. We remark that in most parts of the thesis, general intuition is sufficient and the exact mathematical definitions are not that important in order to understand the *essence*. Therefore, even a reader unfamiliar with, e.g., the quantum formalism or the mathematical definitions of the various entropies, may skip Chap. 2 in the first reading and get back to the relevant definitions appearing in it only when wishing to get a better understanding of the complete technical details.

Chapter 3 deals with basic information and terminology related to device-independent information processing. Readers who are unfamiliar with, e.g., non-locality, should first of all read this chapter. Readers already familiar with some device-independent tasks may skip the chapter and come back to it if needed.

Chapter 4 acts as an *introduction* to our showcases; no theorems or proofs are given there. Thus, readers who are familiar with the question of parallel repetition and the task of DIQKD may pass over this chapter.

Chapters 5 and 6 concern themselves with the mathematical objects that we consider in the thesis—the “black boxes” that model the different types of devices. Chapter 5 defines what we call a “single-round box”, which is, in a sense, a device that can be used to play only a single non-local game. The single-round box acts as an abstract object that allows us to study the fundamental aspects of non-locality, without needing to deal with complex protocols. As we will see, it captures the “physics” of the problem at hand. Hence, studying single-round boxes is the first step in any analysis of device-information processing task. In Chap. 6, we formally define parallel and sequential boxes, which give the mathematical model for parallel and sequential devices, and discuss the relations between them.

After setting the stage, we are ready to start discussing the method of reductions to IID. The first step in this direction is done in Chap. 7, where we discuss the IID assumption and see how it can be used to simplify the analysis of device-independent tasks and, in particular, our showcases. This chapter also presents the asymptotic equipartition property, which acts as a valuable mathematical tool when working under the IID assumption.

The tools used as reductions, i.e., the de Finetti reduction and the entropy accumulation theorem, are the topics of Chaps. 8 and 9, respectively. Chapters 10 and 11 are devoted to the analysis of our showcases via a reduction to IID.

Clearly, many open questions and directions for future works arise. We discuss open questions specific for our showcases within the relevant chapters. In addition, the thesis ends with an outlook in Chap. 12 including questions that, in order to answer, require further development of the toolkit of reductions to IID.

A reader interested in the topic of reductions to IID in general is recommended to read the thesis from the beginning to the end, following the order of the chapters. On the other hand, a reader who is mainly interested in one of the showcases may focus

Table 1.1 Reading suggestion according to the reader's main interest

Reader's interest	Recommended sections
Reductions to IID	All chapters
Parallel repetition	4.1, 5.1, 6.1, 7.1, 7.3.1, 8, 10
Device-independent cryptography	4.2, 5, 6.2, 7.1, 7.2, 7.3.2, 9, 11

only on the sections relevant for the showcase of interest. To assist such readers, we list in Table 1.1 the relevant sections (in the order in which they should be read) for each of the showcases.

References

1. Bennett, C. H. and Brassard, G. (1984). Proceedings of the IEEE international conference on computers, systems, and signal processing, Bangalore, India, 1984
2. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021
3. Bell JS (1964) On the Einstein-Podolsky-Rosen paradox. *Physics* 1(3):195–200
4. Einstein A, Podolsky B, Rosen N (1935) Can quantum-mechanical description of physical reality be considered complete? *Phys Rev* 47(10):777
5. Paris MG, Řeháček J (eds) (2004) Quantum state estimation. Springer, Berlin
6. Mayers D, Yao A (1998) Quantum cryptography with imperfect apparatus. In: Proceedings of 39th annual symposium on foundations of computer science, pp 503–509. IEEE
7. Bancal J-D, Navascués M, Scarani V, Vértesi T, Yang TH (2015) Physical characterization of quantum devices from nonlocal correlations. *Phys Rev A* 91(2):022115
8. Reichardt BW, Unger F, Vazirani U (2013) Classical command of quantum systems. *Nature* 496(7446):456–460
9. Christandl M, König R, Mitchison G, Renner R (2007) One-and-a-half quantum de Finetti theorems. *Commun Math Phys* 273(2):473–498
10. Renner R (2008) Security of quantum key distribution. *Int J Quantum Inf* 6(01):1–127
11. Arnon-Friedman R, Renner R (2015) de Finetti reductions for correlations. *J Math Phys* 56(5):052203
12. Dupuis F, Fawzi O, Renner R (2016) Entropy accumulation. [arXiv:1607.01796](https://arxiv.org/abs/1607.01796)
13. Holenstein T, Renner R (2011) On the randomness of independent experiments. *IEEE Trans Inf Theory* 57(4):1865–1871
14. Tomamichel M, Colbeck R, Renner R (2009) A fully quantum asymptotic equipartition property. *IEEE Trans Inform Theory* 55(12):5840–5847
15. Arnon-Friedman R, Dupuis F, Fawzi O, Renner R, Vidick T (2018) Practical device-independent quantum cryptography via entropy accumulation. *Nat Commun* 9(1):459
16. Arnon-Friedman R, Renner R, Vidick T (2016b) Non-signaling parallel repetition using de Finetti reductions. *IEEE Trans Inf Theory* 62(3):1440–1457
17. Arnon-Friedman R, Renner R, Vidick T (2019) Simple and tight device-independent security proofs. *SIAM J Comput* 48(1):181–225
18. Vazirani U, Vidick T (2014) Fully device-independent quantum key distribution. *Phys Rev Lett* 113(14):140501
19. Liu Y, Yuan X, Li M-H, Zhang W, Zhao Q, Zhong J, Cao Y, Li Y-H, Chen L-K, Li H, et al (2017) High speed self-testing quantum random number generation without detection loophole. *Frontiers in optics, pages FTh2E–1*. Optical Society of America

Chapter 2

Preliminaries: Basics and Notation



2.1 General Notation

The relevant notation for sets and vectors is summarised below.

- \mathbb{N} , \mathbb{R} , and \mathbb{C} are the sets of natural, real, and complex numbers, respectively.
- $[a, b]$ denotes the closed set of real numbers $a \leq x \leq b$.
- $[n]$ denotes the set $\{1, 2, \dots, n\}$.
- When an object x_i is defined for all $i \in [n]$, $\{x_i\}_{i \in [n]}$ denotes the set $\{x_1, x_2, \dots, x_n\}$.
- Other sets are mostly denoted by calligraphic letters, e.g., \mathcal{S} .
- $\mathcal{S} \subseteq \mathcal{P}$ means that \mathcal{S} is a subset of \mathcal{P} . $\mathcal{S} \subset \mathcal{P}$ means that \mathcal{S} is a proper subset of \mathcal{P} .
- $\mathcal{S} \setminus \mathcal{P} = \{s : s \in \mathcal{S} \wedge s \notin \mathcal{P}\}$ stands for the difference between the two sets.
- $\mathcal{S} \times \mathcal{P} = \{(s, p) : s \in \mathcal{S} \wedge p \in \mathcal{P}\}$ is the multiplication of the sets. Furthermore, $\mathcal{S} \times \mathcal{S}$ is denoted by \mathcal{S}^2 and \mathcal{S}^n is defined analogously for any n .
- For sets \mathcal{S}, \mathcal{P} we denote by $\text{Hom}(\mathcal{S}, \mathcal{P})$ the set of all homomorphisms from \mathcal{S} to \mathcal{P} . The set of all endomorphisms is denoted by $\text{End}(\mathcal{S})$, i.e., $\text{End}(\mathcal{S}) = \text{Hom}(\mathcal{S}, \mathcal{S})$.
- Vectors (of different objects) are marked in bold. For example, we use $\mathbf{x} = x_1, x_2, \dots, x_n$.
- Let $f : \mathcal{S} \rightarrow \mathbb{R}$ be a function over some set $\mathcal{S} \subset \mathbb{R}^n$. The infinity norm of the gradient of f is defined as

$$\|\nabla f\|_\infty = \sup \left\{ \frac{\partial}{\partial s_i} f(s) : s \in \mathcal{S}, i \in [n] \right\}.$$

We use the following general notation.

- $\wedge, \vee,$ and \neg denote the logical and, or, and negation, respectively.
- \oplus denotes the XOR operation.
- We denote by \log the logarithm in base 2.
- $\binom{n}{k_1, \dots, k_m}$ is the multinomial coefficient, i.e., $\binom{n}{k_1, \dots, k_m} = \frac{n!}{k_1! \dots k_m!}$, where $!$ is the factorial operation.

- A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is called negligible if for every positive polynomial $p(\cdot)$, there exists an n_0 such that for all $n > n_0$, $f(n) < \frac{1}{p(n)}$. In the thesis, in all cases where the term negligible is used $f(n)$ decreases exponentially fast with n .

2.2 Probability Distributions and Random Variables

We use both probability distributions and random variables (RV) and interchange the two when convenient. Specifically,

- Capital letters, e.g., X , denote RV. When implicit, a RV X takes values from the set denoted by the same letter, i.e., \mathcal{X} .
- P_X denotes the probability distribution corresponding to the RV X . To distinguish different probability distributions we sometimes replace P by other letters, such as O and Q .
- $P_X(x)$ is the probability that $X = x$.
- When a probability distribution is used without a need of referring to the event space etc., we simply use $\{p_i\}_{i \in \mathcal{I}}$ for some \mathcal{I} (usually clear from the context or irrelevant) while keeping in mind that $p_i \geq 0$ for all $i \in \mathcal{I}$ and $\sum_i p_i = 1$.
- When discussing more complex events $\Omega \subseteq \mathcal{X}$ over \mathcal{X} , we use $\Pr_{x \sim X}[\Omega]$ to denote the probability of the event Ω when sampling according to P_X . When it is clear from the context according to which probability distribution the sampling is done we may drop the subscript and write only $\Pr[\Omega]$. For example, when applying Chernoff-type bounds, we use standard notation such as $\Pr[\sum_i X_i > t]$ instead of $P_{X_1 \dots X_n}[\sum a_i > t]$.
- The expectation value $\mathbb{E}[X]$ of X is given by $\mathbb{E}[X] = \sum_{x \in \mathcal{X}} x P_X(x)$.

When considering two RVs X and Y , jointly distributed according to P_{XY} , the *marginal* P_X is defined via

$$P_X(x) = \sum_y P_{XY}(x, y) .$$

The *conditional distribution* of X given $Y = y$, given by

$$\forall x, \quad P_{X|Y=y}(x) = \frac{P_{XY}(x, y)}{P_Y(y)} . \quad (2.1)$$

We mostly use $P_{X|Y}(x|y)$ to denote $P_{X|Y=y}(x)$ and the shorthand notation $P_{X|Y} = P_{XY}/P_Y$ instead that of Eq. (2.1).

Throughout the thesis, we use the following operations on probability distributions:

- For any $q \in [0, 1]$, P_X , and R_X , the convex combination $S_X = qP_X + (1 - q)R_X$ is defined via

$$\forall x, \quad S_X(x) = qP_X(x) + (1 - q)R_X(x) .$$

- For any $n \in \mathbb{N}$ and $P_X, P_X^{\otimes n}$ denoted the probability distribution over \mathcal{X}^n defined via

$$\forall \mathbf{x}, \quad P_X^{\otimes n}(\mathbf{x}) = \prod_i P_X(x_i) ,$$

where $\mathbf{x} = x_1, x_2, \dots, x_n$.

2.2.1 Independent and Identical Random Variables

Consider two RV X and Y defined over \mathcal{X} and \mathcal{Y} respectively. We say that the two are independent if and only if for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $P_{XY}(x, y) = P_X(x) \cdot P_Y(y)$. For $\mathcal{X} = \mathcal{Y}$, we say that X and Y are identical if and only if $P_X = P_Y$.

A sequence of RVs $\mathbf{X} = X_1, X_2, \dots, X_n$, each over \mathcal{X} are said to be independently and identically distributed (IID) RVs if and only if they are all independent and identical to one another.

2.2.2 Concentration Inequalities

When considering IID RVs, concentration inequalities are of special importance. Roughly speaking, concentration inequalities give bounds on how fast the observed frequencies converge to the expected value when sampling IID RVs. The formal statements relevant for the thesis are given below.

Lemma 2.1 (Hoeffding's inequality) *Consider a RV X defined over $\mathcal{X} = \{0, 1\}$ and let X_1, X_2, \dots, X_n be a sequence of n identical and independent copies of X . Then,*

$$\Pr \left[\sum_i X_i - n\mathbb{E}[X] \geq tn \right] \leq \exp(-2nt^2)$$

and

$$\Pr \left[\left| \sum_i X_i - n\mathbb{E}[X] \right| \geq tn \right] \leq 2 \exp(-2nt^2) .$$

Sanov's inequality can be seen as a concentration inequality for conditional probability distributions, in the following sense. Let $O_{A|X}$ be a conditional probability distribution over \mathcal{A} and \mathcal{X} , Q_X be a probability distribution over \mathcal{X} and denote $O_{AX} = Q_X O_{A|X}$. Fix $n \in \mathbb{N}$. Consider a scenario in which we sample $\mathbf{a} = a_1, \dots, a_n$ and $\mathbf{x} = x_1, \dots, x_n$ using $O_{AX}^{\otimes n}$ and estimate $O_{A|X}$ from the sample by calculating $O_{AX}^{\text{freq}(\mathbf{a}, \mathbf{x})}$ defined by

$$\mathbb{O}_{AX}^{\text{freq}(a,x)}(\tilde{a}\tilde{x}) = \frac{|\{i : (a_i, x_i) = (\tilde{a}, \tilde{x})\}|}{n}$$

and define

$$\mathbb{O}_{A|X}^{\text{freq}(a,x)} = \frac{\mathbb{O}_{AX}^{\text{freq}(a,x)}}{\mathbb{Q}_X}. \quad (2.2)$$

Lemma 2.2 (Sanov’s inequality) *For every \mathbb{O}_{AX} and n ,*

$$\Pr_{\mathbf{ax} \sim \mathbb{O}_{AX}^{\otimes n}} \left[\left| \mathbb{O}_{A|X}^{\text{freq}(a,x)} - \mathbb{O}_{A|X} \right|_1 > \epsilon \right] \leq \delta(n, \epsilon)$$

where $\mathbb{O}_{A|X}^{\text{freq}(a,x)}$ is as in Eq. (2.2), $\delta(n, \epsilon) = (n + 1)^{|A| \cdot |\mathcal{X}| - 1} e^{-n\epsilon^2/2}$, and

$$\left| \mathbb{O}_{A|X}^{\text{freq}(a,x)} - \mathbb{O}_{A|X} \right|_1 = \sum_{\tilde{x}} \mathbb{Q}_X(\tilde{x}) \sum_{\tilde{a}} \left| \mathbb{O}_{A|X}^{\text{freq}(a,x)}(\tilde{a}|\tilde{x}) - \mathbb{O}_{A|X}(\tilde{a}|\tilde{x}) \right|.$$

2.3 Quantum Formalism

The basic notation used in the thesis related to the quantum formalism is listed below. We remark, however, that understanding what is meant by a “state” and “measurements” on the intuitive level will almost always suffice in order to understand the essence of the thesis. The exact definitions below are given for the sake of completeness. Clearly, they do not cover all concepts and definitions employed in quantum physics and quantum information theory. Readers who are not familiar with the topics and would like to get a more comprehensive understanding are directed to [1].

We use the Dirac notation: $|\psi\rangle$ denotes a column vector while $\langle\psi|$ is a row vector. $\langle\phi|\psi\rangle$ and $|\phi\rangle\langle\psi|$ denote inner and outer products of the two vectors, respectively.

2.3.1 Operators

We use the following standard notation and definitions.

- The identity matrix, or operator, of dimension d is denoted by \mathbb{I}_d . Alternatively, instead of indicating the dimension, we use, e.g., \mathbb{I}_X to denote the identity operator acting in a specific space associated to X (see below). When the space or dimension is clear from the context we simply write \mathbb{I} .
- A *Hermitian*, or self-adjoint, operator A is an operator satisfying $A = A^\dagger$.
- A *unitary* operator A is an operator satisfying $AA^\dagger = A^\dagger A = \mathbb{I}$.
- The trace of a square matrix A , i.e., the sum of the elements on the main diagonal of A , is denoted by $\text{Tr}(A)$.

- $A \succeq 0$, for A Hermitian, means that the eigenvalues of A are non-negative. $A \succeq B$ stands for $A - B \succeq 0$.
- The 1-norm is defined as $\|A\|_1 = \text{Tr}|A| = \text{Tr}\sqrt{A^\dagger A}$, where A^\dagger denotes the conjugate transpose of A .
- For a diagonal matrix A with eigenvalues $\{a_i\}_i$, $\log(A)$ is the diagonal matrix with eigenvalues $\{\log(a_i)\}_i$.

2.3.2 Hilbert Spaces

The postulates of quantum mechanics tell us that all quantum states “belong” to a complex vector space called a *Hilbert space*. All quantum states and operations will be defined with respect to the considered Hilbert spaces. We give the formal definitions below.

Definition 2.3 (*Hilbert space*) A Hilbert space \mathcal{H} is a complex vector space, i.e.,

$$|\psi\rangle, |\phi\rangle \in \mathcal{H} \text{ and } \lambda_1, \lambda_2 \in \mathbb{C} \rightarrow \lambda_1|\psi\rangle + \lambda_2|\phi\rangle \in \mathcal{H}$$

such that for all $|\psi\rangle, |\phi\rangle \in \mathcal{H}$, there exists $\langle\phi|\psi\rangle \in \mathbb{C}$ for which

1. it is linear in $|\psi\rangle$: $\langle\phi|\lambda_1\psi_1 + \lambda_2\psi_2\rangle = \lambda_1\langle\phi|\psi_1\rangle + \lambda_2\langle\phi|\psi_2\rangle$,
2. $\overline{\langle\phi|\psi\rangle} = \langle\psi|\phi\rangle$, where the bar denotes the complex conjugate,
3. for all $|\psi\rangle \in \mathcal{H}$, $\langle\psi|\psi\rangle \geq 0$ and $\langle\psi|\psi\rangle = 0 \leftrightarrow |\psi\rangle = 0$.

The norm of a vector $|\psi\rangle$ is defined as $\| |\psi\rangle \|_1 = \sqrt{\langle\psi|\psi\rangle}$.

Definition 2.4 (*Orthonormal basis*) An orthonormal basis of \mathcal{H} is a set of vectors $\{|\phi_i\rangle\}_{i \in I}$ such that

- $\langle\phi_i|\phi_j\rangle = \delta_{ij}$ for all $i, j \in I$ and
- $\langle\psi|\phi_i\rangle = 0$ for all $i \in I \rightarrow \psi = 0$.

We will usually consider Hilbert spaces of finite dimensions, meaning I is a set with a finite amount of elements.

Definition 2.5 (*Projector*) Let \mathcal{H} be a Hilbert space and \mathcal{H}' a subspace of \mathcal{H} with $\{|\phi_i\rangle\}_{i \in \mathcal{I}'}$ an orthonormal basis of \mathcal{H}' . The projector of \mathcal{H} onto \mathcal{H}' is the operator

$$P_{\mathcal{H}'} = \sum_{i \in \mathcal{I}'} |\phi_i\rangle\langle\phi_i|.$$

Given \mathcal{H}_A and \mathcal{H}_B , the tensor product Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ is defined such that for $|\psi\rangle \in \mathcal{H}_A$ and $|\phi\rangle \in \mathcal{H}_B$, it associates a vector $|\psi\rangle \otimes |\phi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ with the property that

1. $c \cdot (|\psi\rangle \otimes |\phi\rangle) = (c \cdot |\psi\rangle) \otimes |\phi\rangle = |\psi\rangle \otimes (c \cdot |\phi\rangle)$
 2. $(|\psi_1\rangle + |\psi_2\rangle) \otimes |\phi\rangle = |\psi_1\rangle \otimes |\phi\rangle + |\psi_2\rangle \otimes |\phi\rangle$
 3. $|\psi\rangle \otimes (|\phi_1\rangle + |\phi_2\rangle) = |\psi\rangle \otimes |\phi_1\rangle + |\psi\rangle \otimes |\phi_2\rangle$
- for all $c \in \mathbb{C}$, $|\psi_1\rangle, |\psi_2\rangle \in \mathcal{H}_A$ and $|\phi_1\rangle, |\phi_2\rangle \in \mathcal{H}_B$.

2.3.3 Quantum States

2.3.3.1 Pure and Mixed States

There are two “types” of quantum states one can consider—pure and mixed states.

A *pure quantum state* is associated with a vector belonging to an Hilbert space, $|\psi\rangle \in \mathcal{H}$, with normalisation $\| |\psi\rangle \|_1 = 1$.

Instead of working only with vectors, we can define quantum states as matrices, or operators.

Definition 2.6 (*Density operator*) A density operator, or simply a quantum state, $\rho \in \text{End}(\mathcal{H})$ is a Hermitian positive operator with trace 1. That is,

$$\rho = \rho^\dagger ; \quad \rho \succeq 0 ; \quad \text{Tr}(\rho) = 1 .$$

For a given Hilbert space \mathcal{H} , we denote by $\mathcal{S}(\mathcal{H})$ the set of all density operators defined over \mathcal{H} .

Any pure state $|\psi\rangle \in \mathcal{H}$ can be written as a density operator $\rho = |\psi\rangle\langle\psi| \in \mathcal{S}(\mathcal{H})$. Density operators can describe more general states, called *mixed quantum states*, which can be thought of as a convex combination of pure states:

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i| .$$

Note, however, that different convex combinations can result in the same mixed state ρ and, thus, ρ does not pin-down a specific decomposition to pure states.

A *qubit* is a quantum state belonging to $\mathcal{S}(\mathcal{H})$ for a two-dimensional Hilbert space \mathcal{H} . The basis states are denoted by $|0\rangle$ and $|1\rangle$.

2.3.3.2 Composite Systems

One can consider quantum states over tensor products of Hilbert spaces. Such states are called multipartite states. For example, a bipartite state is a quantum state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ for some Hilbert spaces \mathcal{H}_A and \mathcal{H}_B . The state ρ_{AB} can describe a state shared between two parties, Alice and Bob. The most important thing to notice in the context of the thesis is that given a bipartite state ρ_{AB} , its marginals are also quantum states; these are called the reduced density operators.

Definition 2.7 (*Reduced density operators*) Given $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, its reduced density operator over \mathcal{H}_A is given by

$$\rho_A = \text{Tr}_B(\rho_{AB}) = \sum_i (\mathbb{I}_A \otimes \langle \phi_i |) \rho_{AB} (\mathbb{I}_A \otimes |\phi_i\rangle)$$

where $\{|\phi_i\rangle\}_i$ is a basis of \mathcal{H}_B , and similarly for ρ_B .

Thinking of ρ_{AB} as shared between Alice and Bob, Alice's local state is then ρ_A while Bob's local state is ρ_B .

Given a state ρ_A we can consider its *purification*.

Definition 2.8 (*Purification*) The purification of a state $\rho_A \in \mathcal{S}(\mathcal{H}_A)$ is a *pure* bipartite state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ for which $\text{Tr}_B(\rho_{AB}) = \rho_A$.

Note that by applying a unitary on B the state on A is not being modified and the overall state remains pure. Thus, after the unitary operation, we are still holding a purification. In this sense, we usually say that all purifications are equivalent up to the application of a unitary on the purifying system B .

2.3.3.3 Classical Systems

A classical system, defined by a RV A with probability distribution P_A , can be represented by the density operator

$$\rho_A = \sum_{a \in \mathcal{A}} P_A(a) |a\rangle \langle a|,$$

where $\{|a\rangle\}_a$ is an orthonormal basis of a Hilbert space \mathcal{H}_A .

One example of a classical system that is of common use is the state associated with the uniform distribution U_m over $\{0, 1\}^m$. This distribution can be written as the state $\rho_{U_m} = \frac{1}{m} \mathbb{I}_m$, called the *completely mixed state* on m qubits.

A *classical-quantum state* is a bipartite state in which one register is classical and the other is quantum. Formally,

Definition 2.9 (*Classical-quantum state*) A classical-quantum state $\rho_{AE} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_E)$, classical on A , is a state of the form

$$\rho_{AE} = \sum_a P_A(a) |a\rangle \langle a| \otimes \rho_E^a,$$

where $\{|a\rangle\}_a$ is an orthonormal basis of the Hilbert space \mathcal{H}_A and, for all $a \in \mathcal{A}$, $\rho_E^a \in \mathcal{S}(\mathcal{H}_E)$.

Given a classical-quantum state ρ_{AE} as above, we can consider the quantum state arising from *conditioning* on an event defined over \mathcal{A} . For example, conditioning on the event $A = a$, the quantum state is ρ_E^a . Conditioning can also be done when considering more complicated events. For Ω some event over \mathcal{A} , the state conditioned on Ω is

$$\rho_{AE|\Omega} = \frac{1}{\Pr[\Omega]} \sum_{a \in \Omega} P_{A|\Omega}(a) \otimes \rho_E^a ,$$

where $\Pr[\Omega] = \sum_{a \in \Omega} P_A(a)$ is the probability of Ω according to ρ_{AE} and $P_{A|\Omega}(a) = \Pr[A = a \wedge \Omega] / \Pr[\Omega]$ is the probability of a given Ω .

2.3.3.4 Entanglement

Given a bipartite state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, shared between two parties, one can study the type of correlations that appear between the two parties. A state is said to be *separable* if it can be written as

$$\rho_{AB} = \sum_i p_i \rho_A^i \otimes \rho_B^i \tag{2.3}$$

for some probabilities p_i , $\rho_A^i \in \mathcal{S}(\mathcal{H}_A)$, and $\rho_B^i \in \mathcal{S}(\mathcal{H}_B)$. That is, a separable state is a convex combination of tensor product states. Using the above we notice that a pure state $|\psi\rangle_{AB}$ is separable if and only if it is a tensor product of two pure states $|\psi\rangle_{AB} = |\psi_A\rangle \otimes |\psi_B\rangle$.

Not all quantum states are separable. A bipartite state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is said to be *entangled* if it cannot be written in the form of Eq. (2.3). Such states exhibit correlations which cannot be explained by classical means.

Of specific interest to us are maximally entangled states of two qubits, also called Bell states, denoted by

$$\begin{aligned} |\Phi^+\rangle &= \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle) \quad , \quad |\Phi^-\rangle = \frac{1}{\sqrt{2}} (|00\rangle - |11\rangle) \quad , \\ |\Psi^+\rangle &= \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle) \quad , \quad |\Psi^-\rangle = \frac{1}{\sqrt{2}} (|01\rangle - |10\rangle) \quad . \end{aligned}$$

Here, $|00\rangle$ stands for $|0\rangle \otimes |0\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$, with \mathcal{H}_A and \mathcal{H}_B two-dimensional Hilbert spaces. $|01\rangle$, $|10\rangle$, and $|11\rangle$ are similarly defined.

2.3.4 Quantum Operations

2.3.4.1 Unitary Evolution

The evolution of a closed, or isolated, quantum system is described by unitary operations. By “a closed system” we mean that the transformation of the system of interest is independent of the “rest of the world”, or the environment. We have:

- For any unitary U , U evolves a pure state $|\psi^1\rangle \in \mathcal{H}$ to a pure state $|\psi^2\rangle \in \mathcal{H}$ according to $|\psi^2\rangle = U|\psi^1\rangle$.
- More generally, for mixed states, starting with $\rho^1 \in \mathcal{S}(\mathcal{H})$ we have $\rho^2 = U\rho^1 U^\dagger \in \mathcal{S}(\mathcal{H})$.
- For a bipartite state ρ_{AB}^1 , we can evolve each subsystem locally by $\rho_{AB}^2 = (U_A \otimes U_B)\rho_{AB}^1(U_A^\dagger \otimes U_B^\dagger)$.
- As unitary operations are reversible ($UU^\dagger = U^\dagger U = \mathbb{I}$), the evolution of closed systems is always reversible.

2.3.4.2 Quantum Measurements

To describe a quantum measurement one can use the so called *Kraus operators*.

Definition 2.10 (*Kraus operators*) A set of Kraus operators $\{K_i\}_{i \in \mathcal{I}}$ is a set of operators such that $\sum_{i \in \mathcal{I}} K_i^\dagger K_i = \mathbb{I}$.

Definition 2.11 (*Quantum measurement: Kraus representation*) Given a state ρ and a set of Kraus operator $\{K_i\}_{i \in \mathcal{I}}$ describing a measurement, the outcome of the measurement on ρ is a RV I , defined over the set \mathcal{I} , where each outcome $i \in \mathcal{I}$ is associated with the operator K_i . The probability of observing the outcome i when measuring ρ with $\{K_i\}_i$ is given by

$$\Pr(i) = \text{Tr}(K_i \rho K_i^\dagger) .$$

The post-measurement state is given by

$$\rho_i = \frac{K_i \rho K_i^\dagger}{\text{Tr}(K_i \rho K_i^\dagger)} .$$

We can further identify an operator $M_i = K_i^\dagger K_i$ and work with it, instead of the Kraus operators, to ease notation in some cases. These operators, called *positive operator valued measures* (POVMs), can then be used to describe the relevant measurements.

Definition 2.12 (*Positive operator valued measure*) A positive operator valued measure (POVM) is a set of positive Hermitian operators $\{M_i\}_{i \in \mathcal{I}}$ such that $\sum_{i \in \mathcal{I}} M_i = \mathbb{I}$.

Definition 2.13 (*Quantum measurement: POVM representation*) Given a state ρ and a POVM $\{M_i\}_{i \in \mathcal{I}}$ describing a measurement, the outcome of the measurement on ρ is a RV I , defined over the set \mathcal{I} , where each outcome $i \in \mathcal{I}$ is associated with the operator M_i . The probability of observing the outcome i when measuring ρ with $\{M_i\}_i$ is given by

$$\Pr(i) = \text{Tr}(M_i \rho) .$$

Given a POVM $\{M_i\}_{i \in \mathcal{I}}$ there are many different decomposition to Kraus operators. While the specific decomposition is not relevant for knowing the measurement statistics, they are needed in order to describe the post-measurement state.

In most of the scenarios considered in this thesis we will only be interested with the observed measurements statistics and therefore we will use POVMs to describe a measurement. When there will be a need to consider the post-measurement state we will switch to Kraus operators. Which form on quantum measurement is being used is usually clear from the context and hence we simply call all of them *measurement operators*.

The Pauli operators, denoted by σ_x , σ_y , and σ_z , are an example for measurement operators for qubits:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad ; \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad ; \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} . \quad (2.4)$$

2.3.4.3 Quantum Channels

Quantum channels, or maps, are functions describing the evolution of quantum states. In order for a map \mathcal{E} to describe a real physical process, transferring one quantum state $\rho \in \mathcal{S}(\mathcal{H}_A)$ to another $\mathcal{E}(\rho) \in \mathcal{S}(\mathcal{H}_B)$,¹ it must fulfil certain conditions. Specifically, it must be completely positive and trace preserving (CPTP).

Definition 2.14 (*Quantum channel*) A linear map $\mathcal{E} \in \text{Hom}(\mathcal{S}(\mathcal{H}_A), \mathcal{S}(\mathcal{H}_B))$ is a quantum channel if it is:

1. Completely positive (CP): for any $\rho_{AR} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_R)$ with $\rho_{AR} \succeq 0$,

$$(\mathcal{E} \otimes \mathbb{I}_R)(\rho_{AR}) \succeq 0 ,$$

where \mathcal{H}_R is any additional Hilbert space and \mathbb{I}_R is the identity map on that Hilbert space.

2. Trace preserving (TP): for any $\rho \in \mathcal{S}(\mathcal{H}_A)$, $\text{Tr}(\mathcal{E}(\rho)) = \text{Tr}(\rho)$.

¹Note that \mathcal{H}_B may be different than \mathcal{H}_A . For a unitary evolution, discussed before, this was not the case.

2.4 Distance Measures

The trace distance of two states is given by $\Delta(\rho, \tau) = \frac{1}{2}\|\rho - \tau\|_1$. Operationally, the trace distance quantifies the *distinguishing advantage* when trying to distinguish ρ from τ . Consider a situation in which either the state ρ or the state τ are chosen uniformly at random and given to someone who has no information as to which state was chosen and needs to output a guess. The probability of succeeding in this task depends on how far ρ and τ are from one another via

$$\Pr[\text{correct guess}] = \frac{1}{2} \left(1 + \frac{1}{2}\|\rho - \tau\|_1 \right) = \frac{1}{2} + \Delta(\rho, \tau).$$

We will be interested below in the so called purified distance. The purified distance involves sub-normalised states, i.e., states with $\text{Tr}(\rho) \leq 1$. For this, one first needs to extend the definition of the trace distance to describe also the distance between two sub-normalised states.

Definition 2.15 (*Generalised trace distance*) The trace distance between two sub-normalised states ρ and τ is given by

$$\Delta(\rho, \tau) = \frac{1}{2}\|\rho - \tau\|_1 + \frac{1}{2}|\text{Tr}(\rho - \tau)|.$$

Another important measure of distance (though not a metric) is the fidelity. The fidelity of two quantum states is given by $F(\rho, \tau) = (\text{Tr}|\sqrt{\rho}\sqrt{\tau}|)^2$. The fidelity is related to the trace distance by

$$1 - \sqrt{F(\rho, \tau)} \leq \Delta(\rho, \tau) \leq \sqrt{1 - F(\rho, \tau)}$$

Here, again, we can define the fidelity between two sub-normalised states.

Definition 2.16 (*Generalised fidelity*) The fidelity between two sub-normalised states ρ and τ is given by

$$F(\rho, \tau) = \left(\text{Tr}|\sqrt{\rho}\sqrt{\tau}| + \sqrt{(1 - \text{Tr}(\rho))(1 - \text{Tr}(\tau))} \right)^2.$$

The last distance measure that will be of importance for us is the purified distance [2]. This measure will be used to define the smooth entropies below and will always be considered with sub-normalised states.

Definition 2.17 (*Purified distance*) The purified distance between two sub-normalised states ρ and τ is given by

$$P(\rho, \tau) = \sqrt{1 - F(\rho, \tau)}.$$

2.5 Entropies

2.5.1 Shannon and von Neumann Entropy

Definition 2.18 (*Shannon entropy*) Given RVs A and B defined over \mathcal{A} and \mathcal{B} , respectively, the Shannon entropy of A is given by

$$H(A) = - \sum_{a \in \mathcal{A}} P_A(a) \log(P_A(a)) .$$

The conditional Shannon entropy of A given B is defined to be

$$H(A|B) = H(AB) - H(B) = \sum_{b \in \mathcal{B}} P_B(b) H(A|b) .$$

In the case of a RV defined over $\{0, 1\}$ with $P_A(0) = p$ the Shannon entropy is reduced to the so called “binary entropy” $h(p) = -p \log(p) - (1-p) \log(1-p)$.

The von Neumann entropy is the extension of the Shannon entropy to quantum states.

Definition 2.19 (*von Neumann entropy*) Given a quantum state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, the von Neumann entropy of A is given by

$$H(A)_\rho = -\text{Tr}(\rho \log \rho) .$$

The conditional von Neumann entropy of A given B is defined to

$$H(A|B)_\rho = H(AB)_\rho - H(B)_\rho .$$

When the state on which the entropy is evaluated is clear from the context we drop the subscript and write, e.g., $H(A|B)$.

Definition 2.20 (*Mutual information*) For a quantum state ρ_{ABC} , the conditional mutual information between A and B conditioned C is given by

$$\begin{aligned} I(A : B|C)_\rho &= H(A|C)_\rho + H(B|C)_\rho - H(AB|C)_\rho \\ &= H(A|C)_\rho - H(A|BC)_\rho . \end{aligned}$$

There are other equivalent ways of defining Markov chains for quantum states [3], but for our purposes this definition suffices.

The conditional mutual information fulfils the following properties:

1. Strong subadditivity: $I(A : B|C)_\rho \geq 0$ for any ρ .
2. Data processing: for any quantum channels $\mathcal{E} : \mathcal{S}(\mathcal{H}_A) \rightarrow \mathcal{S}(\mathcal{H}_{A'})$ and $\mathcal{F} : \mathcal{S}(\mathcal{H}_B) \rightarrow \mathcal{S}(\mathcal{H}_{B'})$,

$$I(A : B|C)_\rho \geq I(A' : B'|C)_{\rho'},$$

where $\rho'_{A'B'C} = \mathcal{E} \otimes \mathcal{F} \otimes \mathbb{I}_C(\rho_{ABC})$.

3. $I(A : B|C) = 0$ if and only if A and B are independent given C , i.e., $P_{AB|C} = P_{A|C} \cdot P_{B|C}$.

Definition 2.21 A tripartite quantum state ρ_{ABC} is said to fulfil the Markov chain condition $A \leftrightarrow C \leftrightarrow B$ if $I(A : B|C) = 0$.

2.5.2 Min- and Max-Entropies

We will work with the smooth min- and max-entropies, formally defined as follows.

Definition 2.22 (*Smooth conditional entropies*) For any $\varepsilon \in [0, 1]$ the ε -smooth conditional min- and max-entropy of a state ρ_{AB} are given by

$$H_{\min}^\varepsilon(A|B)_{\rho_{AB}} = \log \inf_{\sigma_{AB} \in \mathcal{B}^\varepsilon(\rho_{AB})} \inf_{\tau_B} \|\sigma_{AB}^{\frac{1}{2}} \tau_B^{-\frac{1}{2}}\|_\infty^2$$

$$H_{\max}^\varepsilon(A|B)_{\rho_{AB}} = \log \inf_{\sigma_{AB} \in \mathcal{B}^\varepsilon(\rho_{AB})} \sup_{\tau_B} \|\sigma_{AB}^{\frac{1}{2}} \tau_B^{-\frac{1}{2}}\|_1^2,$$

for $\mathcal{B}^\varepsilon(\rho_{AB})$ the set of sub-normalised states σ_{AB} with $P(\rho_{AB}, \sigma_{AB}) \leq \varepsilon$, where P is the purified distance as in Definition 2.17.

In practice, we will not need the fully general definitions above (which are stated for completeness). When considering the min-entropy, we will be interested in the case where the A system is classical. This leads to more intuitive definitions. When A is classical and B is trivial, one can simply write

$$H_{\min}(A) = -\log \left[\max_a P_A(a) \right].$$

For quantum B , the state can be written as $\rho_{AB} = \sum_a p_a |a\rangle\langle a| \otimes \rho_B^a$. Then, the conditional min-entropy is directly related to the guessing probability of A given B via

$$H_{\min}(A|B) = -\log p_{\text{guess}}(A|B),$$

where $p_{\text{guess}}(A|B)$ is the maximum probability of guessing A given the quantum system B :

$$p_{\text{guess}}(A|B) = \max_{\{M_B^a\}_a} \left| \sum_a p_a \text{Tr}(M_B^a \rho_B^a) \right|,$$

and the maximisation is performed over all POVMs $\{M_B^a\}_a$ on B . The smooth conditional min-entropy can be written by maximising the min-entropy over all close sub-normalised states, i.e.,

$$H_{\min}^{\varepsilon}(A|B)_{\rho_{AB}} = \max_{\sigma_{AB} \in \mathcal{B}^{\varepsilon}(\rho_{AB})} H_{\min}(A|B)_{\sigma_{AB}} .$$

Moving on to the max-entropy, we will mainly be interested in the case of classical registers. In the classical case, the following holds for the max-entropy

$$H_{\max}(A) \leq -\log \left[\min_{a | P_A(a) \neq 0} P_A(a) \right] .$$

Evaluating the smooth conditional max-entropy will be done by considering a closely related quantity, namely the classical smooth zero-entropy.

Definition 2.23 (*Classical zero-entropy*) For classical RVs A and B distributed according to P_{AB} ,

$$H_0(A|B) = \max_b \log |\text{Supp}(P_{A|B=b})| ,$$

where $\text{Supp}(P_{A|B=b}) = \{a : P_{A|B=b}(a) > 0\}$. The smooth version of the zero-entropy is given by

$$H_0^{\varepsilon}(A|B) = \min_{\Omega} \max_b \log |\text{Supp}(P_{A|\Omega, B=b})| ,$$

where the minimum ranges over all events Ω with probability at least $1 - \varepsilon$.

Finally, we remark that for any quantum state ρ_{AB} ,

$$H_{\max}(A|B) \geq H(A|B) \geq H_{\min}(A|B) .$$

The same ordering does not necessarily hold for the smooth entropies.

References

1. Nielsen MA, Chuang I (2002) Quantum computation and quantum information
2. Tomamichel M, Colbeck R, Renner R (2010a) Duality between smooth min- and max-entropies. IEEE Trans Inf Theory 56(9):4674–4681
3. Hayden P, Jozsa R, Petz D, Winter A (2004) Structure of states which satisfy strong subadditivity of quantum entropy with equality. Commun Math Phys 246(2):359–374

Chapter 3

Preliminaries: Device-Independent Concepts



The goal of this chapter is to present the basic information needed while reading the thesis. It is by no means a comprehensive review of the topic of device-independent information processing. A reader completely unfamiliar with the concepts of non-locality and device-independent protocols is encouraged to read the survey [1] and book [2].

As explained in the introduction, the device-independent framework allows one to examine certain properties of physical devices without referring to their internal workings. Instead of describing a device using its hardware and actions we think of it as a *box* with buttons, on which the user can press in order to give classical inputs to device, and a display, from which the user can read the classical outputs produced by the device. Then, the only information available to the user of the box is the observed data, i.e., the input-output behaviour of the box.

The input-output behaviour of the box can be described mathematically using a *conditional probability distribution* $P_{O|I}$, where I describes the possible inputs of the box and O the possible outputs. For example, if the box has three buttons we can think of I as being a random variable over $\{0, 1, 2\}$. If the box displays a bit as its output then O is a random variable over $\{0, 1\}$. $P_{O|I}$ then describes the, possibly probabilistic, actions of the box. For example, a box with $P_{O|I}(0|0) = 1/2$ and $P_{O|I}(1|0) = 1/2$ outputs 0 or 1, each with probability $1/2$, when the user presses the button associated with the input 0.

The following sections are devoted to explaining the types of boxes that one can consider and their properties. In Sect. 3.1 we define three important classes of boxes according to their input-output behaviour. In Sect. 3.2 we introduce the topic of Bell inequalities, which lies at the heart of all device-independent information processing tasks. In Sect. 3.3 we formally discuss the concept of untrusted devices and, in particular, how a possibly malicious box is modelled.

3.1 Black Boxes

In this thesis we mainly consider bipartite boxes. We think of a bipartite box as a box with two components, each belonging to a different party—one component for Alice and one for Bob. Crucially later on, the components are separated in space so Alice and Bob may locate their parts of the box in different places. Both of Alice’s and Bob’s component have buttons and a display. Alice has the possibility of supplying an input to *her component* and reading the output produced by *her component*. Bob has no access to Alice’s component. Similarly, Bob has the possibility of supplying an input to *his component* and reading the output produced by *his component*, while Alice has no access to Bob’s component.

Mathematically the bipartite nature of the box presents itself by considering conditional probability distributions $P_{AB|XY}$, where X and A denote Alice’s inputs and outputs, respectively, while Y and B denote Bob’s inputs and outputs, respectively. $P_{AB|XY}$ includes all the information about the input-output behaviour of the box and the correlations between Alice’s and Bob’s outputs.

A priori, there are no restrictions on $P_{AB|XY}$, i.e., it can be *any* conditional probability distribution. One may restrict the type of boxes being considered by imposing certain constraints on $P_{AB|XY}$ that depend on the physical theory being studied. Specifically, we are interested in boxes that describe classical, quantum, and non-signalling devices (as explained below). A quantum box, for example, may exhibit correlations between Alice and Bob that cannot be created by classical means. When considering the space of conditional probability distributions $P_{AB|XY}$ the constraints imposed on the box define sets to which the different type of boxes belong. The constraints defining the sets of interest are explained below.

3.1.1 Non-signalling Boxes

When considering general conditional probability distributions $P_{AB|XY}$ any dependence between A , B , X , and Y is allowed. In particular, even though we think of Alice and Bob as holding two separated parts of the box, Alice’s output A may depend on both inputs X and Y . In practice this means that in order for Alice’s box to produce an output, following Alice’s choice of input X , the box first needs to get Bob’s input Y as well. That is, until a signal including Bob’s input arrives to Alice’s component, no actions will be taken by Alice’s component of the box and Alice will need to wait.

In most cases, the above is not a desired behaviour; usually one expects the component of one user to produce an output as a response to pressing the button on *that component* alone. Mathematically this requirement is phrased using the so called “non-signalling conditions” that imply that the marginals $P_{A|X}$ and $P_{B|Y}$ are a well-defined conditional probability distribution. In other words, the behaviour of Alice’s part of the box is described by $P_{A|X}$, which is independent of Bob’s input Y . Thus, Alice’s box does not need to receive Y before producing A . A box fulfilling

the non-signalling conditions between Alice and Bob is called a non-signalling box and is defined as follows.

Definition 3.1 (*Non-signalling box*) A non-signalling box is a conditional probability distribution $P_{AB|XY}$ for which the non-signalling conditions

$$\sum_b P_{AB|XY}(a, b|x, y) = \sum_b P_{AB|XY}(a, b|x, y') \quad (3.1)$$

$$\sum_a P_{AB|XY}(a, b|x, y) = \sum_a P_{AB|XY}(a, b|x', y) \quad (3.2)$$

hold for all $a \in \mathcal{A}$, $b \in \mathcal{B}$, $x, x' \in \mathcal{X}$ and $y, y' \in \mathcal{Y}$.

Denote by $P_{A|X}^y(a|x, y) = \sum_b P_{AB|XY}(a, b|x, y)$ the behaviour of Alice's part of the box, which may a priori depend on Bob's choice of input y . Equation (3.1) states that $P_{A|X}^y = P_{A|X}^{y'}$ and, hence, the conditional probability distribution describing Alice's part of the box is independent of Bob's input, i.e., whether Bob inputs y or y' . We can therefore drop the superscript and simply consider $P_{A|X}$ —a well defined marginal. Similarly, Eq. (3.2) implies that $P_{B|Y}$ is independent of Alice's input and faithfully describes Bob's part of the box.

On the more fundamental level, the non-signalling conditions describe the assumption that the box cannot be used to send instantaneous signals between Alice and Bob. Alice and Bob may locate their components arbitrarily far away from one another. If we require the two components to produce outputs right away, then signals including information about the inputs used by the other party have no time to get from one part of the box to the other and influence its actions. In such a case, using the above notation, if $P_{A|X}^y \neq P_{A|X}^{y'}$, then Alice may conclude from her observed statistics whether Bob used y or y' , even though this information did not have enough time to travel from Bob to Alice. It follows that a non-signalling box is a device that cannot be used as means of communication between Alice and Bob.

A closely related definition that will be of use later is that of a non-signalling extension. Given Alice's component, one can consider an extension of it to an additional party Bob. Specifically, we will be interested in what we call a non-signalling extension of a box, defined below.

Definition 3.2 (*Non-signalling extension*) A non-signalling extension of a box $P_{A|X}$ is a non-signalling box $P_{AB|XY}$ such that for all $a \in \mathcal{A}$, $x \in \mathcal{X}$, and $y \in \mathcal{Y}$, $\sum_b P_{AB|XY}(a, b|x, y) = P_{A|X}(a|x)$.

In words, given $P_{A|X}$, $P_{AB|XY}$ is a non-signalling box with the “correct marginal” on Alice's side (while Bob's marginal $P_{B|Y}$ can be arbitrary).

Before moving on we point to the simplicity of the non-signalling conditions in Eqs. (3.1) and (3.2). The non-signalling conditions are linear. As a result, the set of non-signalling boxes is a polytope. The faces of the polytope are defined by the various non-signalling conditions as well as the positivity and normalisation constraints fulfilled by any conditional probability distribution.

3.1.2 Quantum Boxes

One can further restrict the modelled device by considering quantum boxes, i.e., boxes that exhibit quantum correlations. Such boxes are relevant when considering device-independent processing tasks in which all the resources are quantum.

Quantum correlations are correlations that can be explained within the formalism of quantum physics. To put it differently, we think of a box as a device that “holds” some bipartite quantum state $\rho_{Q_A Q_B}$, shared between Alice and Bob.¹ Alice’s component of the device performs some local quantum measurements on her marginal state ρ_{Q_A} and similarly for Bob. Formally:

Definition 3.3 (*Quantum box*) A quantum box is a conditional probability distribution $P_{AB|XY}$ such that there exist a bipartite state $\rho_{Q_A Q_B}$ and sets of POVMs for Alice and Bob $\{M_a^x\}_{a \in \mathcal{A}}$ for all $x \in \mathcal{X}$ and $\{M_b^y\}_{b \in \mathcal{B}}$ for all $y \in \mathcal{Y}$, respectively, for which

$$P_{AB|XY}(a, b|x, y) = \text{Tr} \left(M_a^x \otimes M_b^y \rho_{Q_A Q_B} \right) \quad \forall a, b, x, y. \quad (3.3)$$

We make the following remarks regarding Definition 3.3. First, while we assume that the bipartite box is quantum, we do not assume anything regarding its internal workings. In particular, we only assume here that the state $\rho_{Q_A Q_B}$ is defined over a bipartite Hilbert space² $\mathcal{H}_{Q_A} \otimes \mathcal{H}_{Q_B}$ (since we consider bipartite boxes) but we do not restrict the dimensions of \mathcal{H}_{Q_A} and \mathcal{H}_{Q_B} .

Second, the non-signalling assumption is “encoded” in the bipartite structure of Alice and Bob’s state $\rho_{Q_A Q_B}$ together with tensor product structure of their measurements as in Eq. (3.3). That is, the conditional probability distribution $P_{AB|XY}$ is by definition non-signalling. Hence, the set of quantum boxes is a subset of the set of non-signalling boxes.

3.1.3 Classical Boxes

A classical box is described by a conditional probability distribution that can be explained in terms of shared randomness alone. That is, we think of Alice’s and Bob’s component of the box as holding a shared random string (in contrast to a shared quantum state). Each component decides on its output depending on its input and the shared string. Formally:

¹We distinguish the quantum state from the correlations throughout the thesis: Q_A and Q_B denote quantum registers belonging to Alice and Bob while A and B denote their classical outputs.

²The definition of a quantum box over a bipartite Hilbert space $\mathcal{H}_{Q_A} \otimes \mathcal{H}_{Q_B}$ is the standard one in the context of non-relativistic quantum mechanics. When studying relativist quantum mechanics one considers a single Hilbert space \mathcal{H} and two commuting measurements acting on it (instead of tensor product measurements). The two definitions coincide when restricting the attention to finite dimensional Hilbert spaces but otherwise different in general [3].

Definition 3.4 (*Classical box*) A classical box is a conditional probability distribution $P_{AB|XY}$ that can be written in the form

$$P_{AB|XY}(a, b|x, y) = \int_{\Lambda} d\lambda \Pr[\Lambda = \lambda] P_{A|X\Lambda}(a|x\lambda) \cdot P_{B|Y\Lambda}(b|y\lambda), \quad (3.4)$$

where Λ is the random variable describing the randomness shared by the two components of the box.

One can assume without loss of generality that $P_{A|X\Lambda}$ and $P_{B|Y\Lambda}$ are deterministic. That is, for all λ , x , and a , either $P_{A|X\Lambda}(a|x, \lambda) = 0$ or $P_{A|X\Lambda}(a|x, \lambda) = 1$, and similarly for Bob. This holds since we can always “push” the non-deterministic behaviour of the components to the shared randomness λ itself. As the number of *deterministic* assignments of a to each x is finite (assuming \mathcal{A} and \mathcal{X} are finite), it follows that one can also express all classical boxes as

$$P_{AB|XY}(a, b|x, y) = \sum_{\lambda} \Pr[\Lambda = \lambda] P_{A|X\Lambda}(a|x\lambda) \cdot P_{B|Y\Lambda}(b|y\lambda),$$

for λ belonging to a finite set and deterministic $P_{A|X\Lambda}$ and $P_{B|Y\Lambda}$.

In the context of Bell inequalities, discussed below, Λ is called the “hidden variable” that explains the correlations between Alice’s and Bob’s parts of the box. Conditioned on the value of Λ , the two components, $P_{A|X\Lambda}$ and $P_{B|Y\Lambda}$, are independent of one another, as seen in Eq. (3.4). Classical boxes, or correlations, are also termed “local correlations”.³

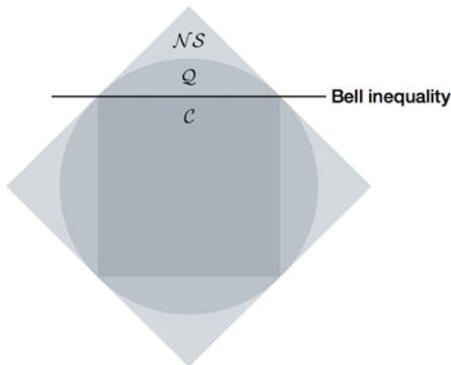
It is easy to see that when considering scenarios with a *single* party, i.e., boxes $P_{A|X}$, all conditional probability distributions can be written in the form of Eq. (3.4). Thus, all single-party boxes are classical boxes. This is not an interesting scenario and, in particular, no device-independent information processing task can be performed in such a case. Thus, boxes of two parties or more are always considered.

3.1.4 Correlations’ Space

Let \mathcal{C} , \mathcal{Q} , and \mathcal{NS} denote the sets of classical, quantum, and non-signalling boxes, respectively. It is easy to see that all of these sets are convex: given two classical boxes $P_{AB|XY}^1$ and $P_{AB|XY}^2$, the box $P_{AB|XY} = pP_{AB|XY}^1 + (1 - p)P_{AB|XY}^2$ is also classical, and similarly for quantum and non-signalling boxes. The convex sets of classical and non-signalling boxes can be described as the convex combination of a finite number of extremal point and hence \mathcal{C} and \mathcal{NS} are polytopes. This is not the case for the quantum set \mathcal{Q} . See Fig. 3.1 for an illustration.

³Though common, this is a rather confusing and unjustified terminology. As clear from Eq. (3.3), quantum correlations are also local, in the sense that each component performs a local operation on its part of the state.

Fig. 3.1 Illustration of the sets of boxes. \mathcal{C} , \mathcal{Q} , and \mathcal{NS} denote the sets of classical, quantum, and non-signalling boxes, respectively. All sets are convex, \mathcal{C} and \mathcal{NS} being polytopes, and the relation $\mathcal{C} \subsetneq \mathcal{Q} \subsetneq \mathcal{NS}$ holds. Bell inequalities can be used to separate classical boxes from quantum ones



As clear from the definition of the various types of boxes, any classical box is also a quantum box and any quantum box is also a non-signalling box. Furthermore, there are examples for quantum boxes that *cannot* be written in the form of Eq. (3.4) and for non-signalling boxes that *cannot* be written in the form of Eq. (3.3). It follows that the sets fulfil the relation

$$\mathcal{C} \subsetneq \mathcal{Q} \subsetneq \mathcal{NS},$$

as in Fig. 3.1.

Bell inequalities, discussed in the next section, give us a way of separating classical boxes from quantum ones⁴—an essential tool in any device-independent information processing task.

3.2 Bell Inequalities

From now on our discussion is restricted to boxes which fulfil, at the least, the non-signalling conditions. That is, we are considering only boxes which cannot be used as means of communication between two separated parties, Alice and Bob.

The polytope of classical boxes \mathcal{C} is a strict subset of the set of quantum and non-signalling boxes. As such, some of the affine hyperplanes defining \mathcal{C} separate \mathcal{C} from \mathcal{Q} .⁵ Informally, when we say that a hyperplane separates \mathcal{C} from \mathcal{Q} we mean that all classical boxes are on one side of the hyperplane, while the other side can only include quantum and non-signalling boxes.

The condition of being on “one side of the hyperplane” is written in the form of an inequality

⁴Separating quantum boxes from non-signalling ones is a far more complicated task; see, e.g., [4].

⁵Other hyperplanes represent the trivial conditions of positivity of normalisation of the conditional probability distributions which are relevant for all sets.

$$\forall P_{AB|XY} \in \mathcal{C}, \quad \sum_{a,b,x,y} s(a,b,x,y)P_{AB|XY}(a,b|x,y) \leq S, \quad (3.5)$$

for some given constants S and $s(a,b,x,y)$, for all a,b,x,y .

Given a box $P_{AB|XY}$, the simple form of Eq. (3.5) allows us to test, by calculating $\sum_{a,b,x,y} s(a,b,x,y)P_{AB|XY}(a,b|x,y)$, if the box *cannot* be a classical one. In other words, if the inequality is violated, i.e.,

$$\sum_{a,b,x,y} s(a,b,x,y)P_{AB|XY}(a,b|x,y) > S,$$

then $P_{AB|XY}$ cannot be written in the form of Eq. (3.4).

As first noticed by [5], some quantum boxes, arising from measurements performed on entangled states, are capable of violating inequalities as in Eq. (3.5).⁶ Bell suggested to use such states in an experiment, proposed to test the EPR paradox [6], that will allow us to check whether there is some classical piece of information, that we are just unaware of or cannot observe, that can explain the apparent “non-local” correlations exhibit by certain quantum states. Such experiments, called today “loophole-free Bell tests” [7–9], have verified the violation of Bell inequalities and by this refuted the possibility of classical explanations of the behaviour of some entangled quantum states.

The inequalities which are fulfilled by any classical box while being violated by *some* quantum boxes are called *Bell inequalities*; see Fig. 3.1. All of the above implies that a Bell inequality acts as a test for “quantumness” or, more precisely, “non-classicalness” and its violation acts as a certificate for passing the test. As such, it is crucial for any device-independent information processing task in which we need to rule out the possibility of executing the considered task with a classical device.

3.2.1 Non-local Games

Bell inequalities, as in Eq. (3.5), can also be phrased as special types of games, called non-local, or Bell, games. In a game, a referee asks Alice and Bob, the players of the game, a question each, chosen according to a given probability distribution; each player only sees her/his question. The players then need to supply answers which fulfil a pre-determined requirement according to which the referee accepts or rejects the answers. To win the game the players can agree on a strategy beforehand but,

⁶Notice that the statement that some quantum states violate Bell inequalities is independent from the statement that classical boxes cannot violate the inequality; it could have been the case that no box is able to violated such inequalities. This would have implied that all quantum correlations can be written in the form of Eq. (3.4) and, hence, can be described as arising from some shared randomness, or an “hidden variable”, λ .

once the game begins, communication between the players is not allowed. If the referee accepts their answers the players win.

Formally, a game $G = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, Q_{XY}, w)$ is defined by sets of possible questions \mathcal{X} and answers \mathcal{A} for Alice, sets of possible questions \mathcal{Y} and answers \mathcal{B} for Bob, a probability distribution Q_{XY} over the questions, according to which the referee chooses the questions, and a winning condition $w : \mathcal{A} \times \mathcal{B} \times \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$, where $w(a, b, x, y) = 0$ means that the referee rejects (a, b, x, y) , i.e., the players lose, and $w(a, b, x, y) = 1$ means that the players win with (a, b, x, y) .

A strategy for the game is naturally described by a box $P_{AB|XY}$ held by the players—the referee chooses questions $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and the players need to reply with answers $(a, b) \in \mathcal{A} \times \mathcal{B}$. The fact that the players are not allowed to communicate during the game means that, at the least, the box $P_{AB|XY}$ is constrained by non-signalling conditions between Alice and Bob, i.e., $P_{AB|XY}$ belongs to the non-signalling polytope \mathcal{NS} .⁷ The winning probability of a box $P_{AB|XY}$ in the game G is given by

$$\omega(P_{AB|XY}) = \sum_{a,b,x,y} Q_{XY}(x, y) P_{AB|XY}(a, b|x, y) w(a, b, x, y) .$$

When the discussed box $P_{AB|XY}$ is clear from the context we simply write ω to denote its winning probability.⁸

In the context of device-independent information processing we interpret a Bell inequality as a special type of game. What makes the game special is that it is designed so that any classical box used by the players leads to a winning probability of at most $\omega_c < 1$, while there exists a quantum box that can be used by the players to achieve a greater winning probability, $\omega_q > \omega_c$. Instead of a Bell inequality as in Eq. (3.5) we have

$$\forall P_{AB|XY} \in \mathcal{C}, \quad \omega(P_{AB|XY}) \leq \omega_c . \quad (3.6)$$

Violating a Bell inequality then translates to violating Eq. (3.6) by winning the respective game with probability greater than ω_c . In both cases, the conclusion is the same; if $P_{AB|XY}$ violates Eq. (3.6) then $P_{AB|XY} \notin \mathcal{C}$.

Before discussing an explicit example of a non-local game, one remark is in order. Above, we thought of Alice and Bob as the ones preparing the box, according to their strategy in the game, and the referee was asking them questions to test their winning probability. Alternatively, we can think of Alice and Bob as holding an uncharacterised box and they are the ones testing the box, by choosing the questions themselves. In that case, Alice and Bob basically take the role of the referee (while the box takes the role of Alice and Bob). (In the showcase of non-signalling parallel

⁷Depending on the context, one can further restrict the allowed strategies by considering classical or quantum boxes.

⁸Notice the notation: w denotes a winning condition (function) while ω is the winning probability (a number). W will be used to denote the random variable describing whether a game is won or lost. In any case, the difference between these three objects is always clear from the text.

repetition, in Chap. 10, we use the first terminology, while the showcase of device-independent quantum cryptography, in Chap. 11, the second is terminology is the more appropriate one).

3.2.2 The CHSH Game

We now present an explicit non-local game that will be of use in the thesis. The Clauser–Horne–Shimony–Holt (CHSH) game [10] is probably the most famous non-local game. In the game, Alice’s and Bob’s inputs and outputs are bits, $a, b, x, y \in \{0, 1\}$ and the inputs are distributed uniformly at random, i.e., $Q_{XY}(x, y) = 1/4$ for all x and y . The winning conditions is given by:

$$w_{\text{CHSH}} = \begin{cases} 1 & \text{and } a \oplus b = x \cdot y \\ 0 & \text{otherwise.} \end{cases} \quad (3.7)$$

The optimal classical box, or strategy, achieves a winning probability of 0.75. An example for such a strategy is one in which the outputs are always $(a, b) = (0, 0)$.

An optimal quantum strategy consists in measuring the maximally entangled state $|\Phi^+\rangle = (|00\rangle + |11\rangle) / \sqrt{2}$ with the following measurements: Alice’s measurements $x = 0$ and $x = 1$ correspond to the Pauli operators σ_z and σ_x respectively and Bob’s measurements $y = 0$ and $y = 1$ to $(\sigma_z + \sigma_x) / \sqrt{2}$ and $(\sigma_z - \sigma_x) / \sqrt{2}$ respectively. A box implementing the above achieves winning probability $\omega = \frac{2+\sqrt{2}}{4} \approx 0.85$. Perhaps surprisingly, any box that achieves the optimal quantum winning probability (or close to it) must be implementing a strategy identical to the above up to local isometries (or close to such a strategy) [11–13].

The CHSH game can also phrased in the form of a Bell *inequality*. The most common way of writing the CHSH inequality is as follows. Given $P_{AB|XY}$, for any pair of inputs (x, y) , let

$$E_{xy} = P_{AB|XY}(0, 0|x, y) + P_{AB|XY}(1, 1|x, y) \\ - P_{AB|XY}(0, 1|x, y) - P_{AB|XY}(1, 0|x, y)$$

and denote the CHSH *value* by

$$\beta(P_{AB|XY}) = E_{00} + E_{01} + E_{10} - E_{11} .$$

The CHSH inequality reads

$$\forall P_{AB|XY} \in \mathcal{C}, \quad \beta(P_{AB|XY}) \leq 2 .$$

The interesting regime is $\beta \in [2, 2\sqrt{2}]$, where $\beta = 2$ is the optimal classical violation while $\beta = 2\sqrt{2}$ is the quantum one. The relation between the winning probability in

the CHSH game and the CHSH value is given by $\omega = 1/2 + \beta/8$ and we have $\omega \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4} \right]$.

When discussing device-independent quantum cryptography in Chap. 11, we use a variant of the CHSH game previously used in [14, 15]. In this game Alice has two inputs $\mathcal{X} = \{0, 1\}$ while Bob has three possible inputs $\mathcal{Y} = \{0, 1, 2\}$. The output sets are $\mathcal{A} = \mathcal{B} = \{0, 1\}$. The winning condition is the following⁹:

$$w_{\text{CHSH}} = \begin{cases} 1 & x, y \in \{0, 1\} \text{ and } a \oplus b = x \cdot y \\ 1 & (x, y) = (0, 2) \text{ and } a = b \\ 1 & (x, y) = (1, 2) \\ 0 & \text{otherwise.} \end{cases} \quad (3.8)$$

The optimal quantum strategy for this game is the same as in the standard CHSH game, except that if Bob's input is a $y = 2$ he applies the same measurement as Alice's measurement for input 0. Since the underlying state is maximally entangled this ensures that their outputs always match when $(x, y) = (0, 2)$.

Conditioned on Bob's input not being 2, the game played is the CHSH game and the optimal classical and quantum winning probabilities are as above.

3.3 Untrusted Devices

Formally defining an *untrusted device*, or untrusted box, is essential when analysing device-independent information processing tasks. The current section is devoted to explaining what is meant by this term and what are the assumptions regarding an untrusted device. To understand the definition of an untrusted device, it is perhaps best to consider a cryptographic scenario in which the device may be manufactured by the malicious party, the adversary, and hence is not to be trusted. The same definition of a device is used also when we do not have an explicit adversary in mind; the device itself is still uncharacterised (in the sense explained below) but we are free to ignore the additional subsystem given to the adversary in what follows. We therefore employ below the terminology used in the cryptographic setting.

As before, we consider the case of two honest parties. A device D is modelled by a bipartite box $P_{AB|XY}$, shared between the honest parties, Alice and Bob, who try to accomplish a certain task. We think of the box as being prepared by the adversary Eve and hence we call it untrusted. Since Eve is the one manufacturing the device it allows her, in particular, to keep an extension of Alice and Bob's device. Formally, we consider a non-signalling extension of $P_{AB|XY}$ to a tripartite box $P_{ABC|XYZ}$ (recall Definition 3.2): we have

⁹For the inputs $(x, y) = (1, 2)$ one can set either $w_{\text{CHSH}} = 1$ or 0 (it is not relevant later on); for completeness we choose $w_{\text{CHSH}} = 1$ in this case, following previous works.

$$P_{AB|XY}(a, b|x, y) = \sum_c P_{ABC|XYZ}(abc|xyz) \quad \forall a, b, x, y, z$$

and Eve “holds” the marginal $P_{C|Z}$. Eve can use her component $P_{C|Z}$ as she wishes. For example, in a cryptographic protocol, Eve can eavesdrop on all the classical communication between the honest parties during the run of the protocol and only later choose to use her box with input z that depends on all other information available to her.

When considering non-signalling (super-quantum) boxes, the only constraint on the extension $P_{ABC|XYZ}$ is that it is non-signalling between the three parties and that the marginal of Alice and Bob is equal to the box $P_{AB|XY}$. In the quantum case, $P_{AB|XY}$ describes both the state shared between Alice and Bob $\rho_{Q_A Q_B}$ and measurements devices used to measure $\rho_{Q_A Q_B}$. Eve then holds a purification¹⁰ of Alice and Bob’s quantum state in a quantum register in her possession. The tripartite box $P_{ABC|XYZ}$ describes the pure state $\rho_{Q_A Q_B E}$ together with the measurements of Alice and Bob as well as the measurements that can be used by Eve to measure her marginal ρ_E .¹¹

Although the device is untrusted, we always assume that the following requirements hold.

Alice and Bob can interact with the device as expected. In any considered scenario, the type of interaction with the device D is defined explicitly. In particular, every protocol clearly states how the users should interact with the device utilised to run the protocol; for example, a protocol may require the users to play n games with the device (by pressing buttons and recording the outputs) *one after the other*. The different types of interactions and the resulting conditions on the untrusted device are discussed in Chap. 6. Note that this requirement can be verified—if the honest parties try to use the device in the specified way and the device does not react as expected (e.g., it does not produce outputs or produces outputs from a different alphabet) then it is clear that something is wrong. In this sense, the requirement that it is possible to interact with the device as expected is not really an assumption regarding the device, but rather a formality that allows us to be explicit when talking about untrusted devices.

Communication (signalling) between the components of the device. The communication between Alice, Bob, and Eve’s components of the device is restricted in the following way:

1. Alice and Bob’s components of D cannot signal to Eve’s component.
2. Alice and Bob can decide when to allow communication (if any) between their components.

¹⁰A purification $\rho_{Q_A Q_B E}$ is the most general extension of a quantum state $\rho_{Q_A Q_B}$, in the sense that it gives Eve the maximal amount of information regarding Alice and Bob’s marginal state. Hence, in the cryptographic setting we always say that Eve holds the purifying system E , without loss of generality—any adversary holding a system E' which is not the purifying system E can only be weaker than that holding E .

¹¹We emphasise again that Eve is not required to measure her quantum state at any particular point.

3. Alice and Bob can decide when to receive communication (if any) from Eve’s component.

The requirement given in Item 1 is necessary for device-independent *cryptogra-phy*; without it the device could directly send to Eve all the raw data it generated.

Item 2 implies that Alice and Bob’s component must be (at least) bipartite. This is necessary to assure that the violation of the considered Bell inequality is meaningful. In the quantum case, this requirement is identified with the “assumption” that we can write Alice and Bob’s quantum state as a bipartite state $\rho_{Q_A Q_B}$ and that the measurements made in Alice’s and Bob’s components of the device are in tensor product with one another.

Items 2 and 3 give Alice, Bob, and Eve’s components the *possibility* to communicate in certain stages of the protocol (see Sect. 4.2.5 for an explicit example). This is not a restrictive nor necessary assumption. This possibility to communicate is added since it is advantageous to actual implementations of certain protocols. For instance, allowing the different components of the device to communicate in certain stages of some protocols opens the possibility of distributing resources, such as entanglement, “on the fly” for each round of the protocol, instead of maintaining large quantum memories.

Other assumptions. Apart from the above description of the untrusted device, the following list includes the standard assumptions used in device-independent information processing tasks (in particular, device-independent cryptography):

1. The honest parties have a trusted random number generator (that can be used to choose the inputs for playing the games, for example).
2. The honest parties have a trusted classical post-processing units to make the necessary (classical) calculations during the considered task.
3. There is a public, but authenticated, classical channel connecting the honest parties (if the considered task requires that the parties communicate classically with one another).
4. In cryptographic scenarios—the honest parties’ physical locations are secure and can be isolated if needed (unwanted information cannot leak outside to Eve or between their devices).
5. Depending on the considered scenario—the actions of the device can be described within the non-signalling or quantum formalism.

In contrast to an *untrusted device*, we sometimes use the terminology *honest device* or *honest implementation*. A device is said to be honest if it implements the considered protocol by using a certain pre-specified strategy. In that case, the actions of the device are known and fixed. See Sect. 4.2.4 for an example.

References

1. Brunner N, Cavalcanti D, Pironio S, Scarani V, Wehner S (2014) Bell nonlocality. *Rev Mod Phys* 86(2):419
2. Scarani V (2019) Bell nonlocality. Oxford University Press, Oxford
3. Slofstra W (2017) The set of quantum correlations is not closed. [arXiv:1703.08618](https://arxiv.org/abs/1703.08618)
4. Navascués M, Pironio S, Acín A (2008) A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J Phys* 10(7):073013
5. Bell JS (1964) On the Einstein-Podolsky-Rosen paradox. *Physics* 1(3):195–200
6. Einstein A, Podolsky B, Rosen N (1935) Can quantum-mechanical description of physical reality be considered complete? *Phys Rev* 47(10):777
7. Hensen B, Bernien H, Dréau A, Reiserer A, Kalb N, Blok M, Ruitenberg J, Vermeulen R, Schouten R, Abellán C, et al (2015) Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* 526(7575):682–686
8. Shalm LK, Meyer-Scott E, Christensen BG, Bierhorst P, Wayne MA, Stevens MJ, Gerrits T, Glancy S, Hamel DR, Allman MS et al (2015) Strong loophole-free test of local realism. *Phys Rev Lett* 115(25):250402
9. Giustina M, Versteegh MA, Wengerowsky S, Handsteiner J, Hochrainer A, Phelan K, Steinlechner F, Kofler J, Larsson J-Å, Abellán C et al (2015) Significant-loophole-free test of Bell's theorem with entangled photons. *Phys Rev Lett* 115(25):250401
10. Clauser JF, Horne MA, Shimony A, Holt RA (1969) Proposed experiment to test local hidden-variable theories. *Phys Rev Lett* 23(15):880
11. Popescu S, Rohrlich D (1992) Which states violate bell's inequality maximally? *Phys Lett A* 169(6):411–414
12. Mayers D, Yao A (2003) Self testing quantum apparatus. [arXiv:quant-ph/0307205](https://arxiv.org/abs/quant-ph/0307205)
13. McKague M, Yang TH, Scarani V (2012) Robust self-testing of the singlet. *J Phys A: Math Theory* 45(45):455304
14. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021
15. Vazirani U, Vidick T (2014) Fully device-independent quantum key distribution. *Phys Rev Lett* 113(14):140501

Chapter 4

Introduction to the Showcases



4.1 Introduction to Non-signalling Parallel Repetition

Non-local games, introduced in Sect. 3.2.1, are relevant in many areas of both theoretical physics and theoretical computer science. In the context of parallel repetition, we think of a game G as follows. A referee asks each of the cooperating parties, also called players, a question chosen according to a given probability distribution. The players then need to supply answers which fulfil a pre-determined requirement according to which the referee accepts or rejects the answers. In order to do so, they can agree on a strategy beforehand, but once the game begins communication between the players is not allowed. If the referee accepts their answers the players win. The goal of the players is, naturally, to maximise their winning probability in the game.

According to the field of interest, one can analyse any non-local game under different restrictions on the players (in addition to not being allowed to communicate). In classical computer science the players are usually assumed to have only classical resources, or strategies. That is, they can use only local operations and shared randomness. In contrast, one can also consider quantum strategies: before the game starts the players create a multipartite quantum state that can be shared among them. When the game begins each player locally measures their own part of the state and bases the answer on their measurement result. Another, more general, type of strategies are those where the players can use any type of correlations that do not allow them to communicate, i.e., non-signalling strategies.

One of the most interesting questions regarding non-local games is the question of parallel repetition. Given a game G with optimal winning probability $1 - \alpha$ (using either classical, quantum, or non-signalling strategies), we are interested in analysing the winning probability in the repeated game, denoted by G^n . In G^n the referee gives the players n independent tuples of questions at once, to which the players should reply. The players win G^n if they win all of the n games. Another, more general and natural, winning criterion is that the players answer a certain fraction $1 - \alpha + \beta$ of the n game instances correctly. One can then ask what is the probability that the

players succeed in the repeated game, as the number of repetitions n increases and whether, in particular, this probability decreases exponentially fast with n , similarly to what happens when playing each of the games independently. While the question of parallel repetition is easy to phrase, its answer is far from trivial (and, in fact, up to date there is no general answer that holds for all games).

The device-independent framework fits perfectly to the study of the parallel repetition question: We can think of a box, i.e., a conditional probability distribution, as describing a strategy of the players. The requirement that the players are not allowed to communicate easily translates to non-signalling conditions between the parties holding the box. Furthermore, the claims that we wish to make regarding the probability of winning the repeated game are oblivious of the exact description of the strategy and hence treating the strategy as a *black box* makes sense. In particular,

g the behaviour of the strategy without having an explicit description of it is necessary in order to be able to use parallel repetition results to, e.g., analyse experiments that aim at ruling-out local realism while performing several Bell violation experiments in parallel or for hardness amplification [1] in complexity theory and cryptography.

We define and explain the question of parallel repetition below. Our showcase, presented in Chap. 10, focuses on the case of non-signalling parallel repetition. Note, however, that all statements made in the following two sections are general and applicable to any type of strategies. (One only needs to interchange the words non-signalling and quantum or classical).

4.1.1 Parallel Repeated Games

For simplicity and as in the rest of the thesis, we consider two player non-local games. All of the statements below can be extended to an arbitrary number of players.

We define a two-player game, with the players named Alice and Bob, similarly to a non-local game.¹

Definition 4.1 (*Two-player game*) A two-player game $G = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, Q_{XY}, R)$ is defined by:

1. A set of possible questions for each player: \mathcal{X} for Alice and \mathcal{Y} for Bob.
2. A probability distribution Q_{XY} over the questions, according to which the referee choses the questions.
3. A set of possible answers for each player: \mathcal{A} for Alice and \mathcal{B} for Bob.
4. A winning condition $R : \mathcal{A} \times \mathcal{B} \times \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$.

In the repeated game, denoted by G^n , the referee asks Alice and Bob n questions, all at once. The questions are chosen independently for each game and the answers are

¹Multi-player games and non-local games are one and the same; we define a two-player game here just to set the terminology used when discussing the showcase of parallel repetition.

checked independently. In most works dealing with parallel repetition, the winning condition of the repeated game is defined such that Alice and Bob win G^n if and only if they win *all* n repetitions of G (hence the name). We will use a more general winning condition in which only a certain fraction of the games needs to be won. We call such games *threshold games*.

Definition 4.2 (*Threshold game*) Any two-player game $G = (\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}, Q_{XY}, R)$ induces a two-player threshold game $G_{1-\gamma}^n = (\mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n, Q_{XY}^{\otimes n}, R_{1-\gamma}^n)$, for $0 \leq \gamma \leq 1$, where the winning criterion $R_{1-\gamma}^n$ is defined by:

$$R_{1-\gamma}^n(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y}) = 1 \Leftrightarrow \frac{1}{n} \sum_{i \in [n]} R(a_i, b_i, x_i, y_i) \geq 1 - \gamma .$$

A strategy for a game is simply a box, i.e., a conditional probability distribution, defining the input-output behaviour of the players. Throughout the thesis, a strategy for a single game G is denoted by $O_{AB|XY}$. The winning probability of a strategy $O_{AB|XY}$ in game G is given by

$$w(O_{AB|XY}) = \sum_{a,b,x,y} Q_{XY}(x, y) O_{AB|XY}(a, b|x, y) R(a, b, x, y) \quad (4.1)$$

When we say that the optimal non-signalling winning probability in a game G is $1 - \alpha$ we mean that

$$\max_{O_{AB|XY}} w(O_{AB|XY}) = 1 - \alpha ,$$

where the maximisation is over all non-signalling strategies $O_{AB|XY}$.

A strategy for the threshold game $G_{1-\gamma}^n$ is denoted by $P_{AB|XY}$. $P_{AB|XY}$'s winning probability in the threshold game is given by

$$w(P_{AB|XY}) = \sum_{a,b,x,y} Q_{XY}^{\otimes n}(x, y) P_{AB|XY}(a, b|x, y) R_{1-\gamma}^n(a, b, x, y) .$$

4.1.2 Threshold Theorems

The parallel repetition question is the following. Let G be a non-local game whose optimal non-signalling winning probability is $1 - \alpha$ and let $G_{1-\gamma}^n$ be the threshold game defined by G (Definition 4.2). We then ask—*what is the probability that a non-signalling strategy $P_{AB|XY}$ wins $G_{1-\gamma}^n$* ? The more “standard” parallel repetition question is retrieved by setting $\gamma = 0$. The interesting scenario to consider is the one in which $1 - \gamma = 1 - \alpha + \beta$ for $\beta > 0$.

The players can always use the trivial independent and identically distributed (IID) strategy for $G_{1-\alpha+\beta}^n$: they simply answer each of the n questions independently

according to the optimal non-signalling strategy for G . In this case, the fraction of successful answers is highly concentrated around $1 - \alpha$ and the probability to win all games simultaneously is $(1 - \alpha)^n$. Thus, for any $\beta > 0$, the winning probability in $G_{1-\alpha+\beta}^n$ decreases exponentially fast with n .

Can the players do better when using a correlated, i.e., a *non-IID*, strategy? There are many examples showing that the answer is yes for certain games. One of the most interesting examples to a person studying non-locality is that of two repetitions of the CHSH game. As mentioned in Sect. 3.2.2, the optimal *classical* strategy in the CHSH game achieves a winning probability of $3/4$. If the players play two CHSH games in parallel and use the optimal classical strategy of a single game twice, the probability that they win both games is $9/16$. However, there exists a better classical strategy [2]:

$$\begin{aligned} \text{Alice's actions:} & \quad \begin{cases} (a_1, a_2) = (1, 1) & (x_1, x_2) \in \{(0, 0), (1, 0), (1, 1)\} \\ (a_1, a_2) = (1, 0) & (x_1, x_2) = (0, 1) \end{cases} \\ \text{Bob's actions:} & \quad \begin{cases} (b_1, b_2) = (1, 1) & (y_1, y_2) \in \{(0, 0), (0, 1), (1, 1)\} \\ (b_1, b_2) = (0, 1) & (y_1, y_2) = (1, 0) \end{cases} \end{aligned}$$

One can easily check that this strategy wins both games with probability $10/16$ and, hence, is better than playing the two games independently. More examples are known for classical games, e.g., [3, 4], as well as for quantum and non-signalling games [5].

Still, one may ask whether the players can achieve a *significantly* higher winning probability compared to the IID strategy as n increases. In the IID case the probability of winning more than a fraction $1 - \alpha + \beta$ of the games decreases exponentially fast with n and β^2 ; see Sect. 7.3.1 for the simple analysis. Does this type of decrease also hold when considering strategies $P_{AB|XY}$ that may correlate the different rounds? If correlated strategies for $G_{1-\alpha+\beta}^n$ are not substantially better than independent ones, even in an asymptotic manner, we learn that “one cannot fight independence with correlations”. As long as the questions are asked, and the answers are verified, in an independent way, creating correlations between the different answers using a correlated strategy cannot help much.

The first exponential parallel repetition theorem was derived for classical two-player games and appeared in [6]: it was shown that if the classical optimal winning probability in a game G is smaller than 1, then the probability to win all the games in the repeated game, using a classical strategy, decreases exponentially with the number of repetitions n . This was improved and adapted to the non-signalling case in [7]. Another improvement was made in [8], where a threshold theorem for the classical two-player case was proven: the probability to win more than a fraction $1 - \alpha + \beta$ of the games for any $\beta > 0$ is exponentially small in n .

Following the same proof technique as [6–9] gave a threshold theorem for multi-player non-signalling complete-support games. Their threshold theorem was the first result where more than two players were considered. In [10, 11] a completely different proof technique, based on de Finetti reductions, was used to derive similar

(improved in some respects) results as [7, 9].² Holmgren and Yang [12] gave a counter example to a general non-signalling parallel repetition—they show that for a certain three-player game without complete-support the probability of winning n instances of the game played in parallel remains *constant*. This implies that the results of [9–11] cannot be extended to games *without complete-support* (as they hold for any number of parties).

The question of parallel repetition in the quantum case is less well understood than its classical and non-signalling versions. The only results applicable to all two-player games is that of [13], which states that the probability of winning all n instances of the game decreases inverse polynomially with n . Exponential decrease is known for different classes of two-player games [14–17] or modifications thereof [18–20].

4.2 Introduction to Device-Independent Quantum Cryptography

Classical cryptography relies on computational assumptions, such as the hardness of factoring, to deliver a wide range of functionalities. The advent of quantum information brought forward a completely different possibility: security based only on the fundamental laws of physics. For example, the quantum key distribution (QKD) protocols by Bennett and Brassard [21] and Ekert [22] allow mutually trustful users connected only by an authenticated classical channel, and an arbitrary quantum channel, to establish a private key whose security is guaranteed by the laws of quantum mechanics. With their private key, the users can communicate with perfect security using, e.g., a one-time pad.

The security of cryptographic protocols such as QKD relies on certain assumptions regarding the physical implementation, such as the quantum states and measurements used in the apparatus implementing the protocol. In real life, however, the manufacturer of the device can have limited technological abilities (and hence cannot guarantee that the device’s actions are exact and non-faulty) or even be malicious. Furthermore, the quantum device may be far too complex for the honest parties running the protocol to open and assess whether it works as alleged. In the cryptographic setting, imperfections in the physical apparatus are of a real concern, even when the manufacturer himself is honest and has good intentions. Indeed, when trying to implement quantum devices we find that creating perfect states and measurements is practically impossible. In the presence of an adversary, imperfections and noise in the implementation can and are being exploited to gain information on the outputs of the cryptographic protocols [23–26]. This means that if one does not trust that the quantum devices are exactly as supposed to be, due to a potentially incompetent or malicious manufacture, then the security of the protocols no longer holds.

²The de Finetti reduction used in these proofs is the topic of Chap. 8; the non-signalling threshold theorem of [10] acts as one of our showcases and is discussed in Chap. 10.

To solve this issue the quantum cryptography community took one step further. In contrast to “standard” quantum cryptographic protocols that are proven to be secure only for specific implementations of the used devices, device-independent quantum cryptographic protocols achieve an unprecedented level of security with guarantees that hold (almost) irrespective of the quality, or trustworthiness, of the physical devices used to implement them and hence count as the “gold standard” of quantum cryptography [27]. In device-independent cryptography we let the adversary, called Eve, prepare the quantum devices used in the protocol. The honest parties, Alice and Bob, must therefore treat the possibly faulty or malicious device as an untrusted device (as defined in Sect. 3.3) with which they can only interact according to the protocol. The protocol must allow them to test the untrusted device and decide whether using it to run the considered cryptographic protocol poses any security risk. The protocol guarantees that by interacting with the device according to the specified steps the honest parties will either abort, if they detect a fault, or produce secure outputs (with high probability). Clearly, the security proof cannot rely on the inner-workings of the device as it may be malicious. Hence, if we are able to prove that the produced outcomes are secure to use, then the statement is inherently independent of the implementation of the physical device (hence the name “device-independent security”).

At first sight, it seems impossible to prove that the outputs of a cryptographic protocol are secure to use when the adversary is the one to manufacture the device. As known for quite some time now, the solution is to base device-independent protocols on the violation of Bell inequalities [22, 28, 29]. As explained in Sect. 3.2.1, a Bell inequality can be thought of as a game played by the honest parties using the device they hold. Different devices lead to different winning probabilities when playing the game. The game has a special property—there exists a quantum device which achieves a winning probability ω greater than all classical, local, devices. Hence, if the honest parties observe that their device wins the game with probability ω they conclude it must be non-local.³ A non-local game therefore acts as a “test for quantumness”. The idea of basing the security of cryptographic protocols (QKD especially) on the violation of Bell inequalities originates in the celebrated work of Ekert [22]. Later, Mayers and Yao [28] recognised that devices that maximally violate a certain Bell inequality could be fully characterised, up to local degrees of freedom, and thus need not be trusted a priori.

Device-independent security relies on the following deep but well-established facts. High winning probability in a non-local game not only implies that the measured system is non-local but, more importantly, that the kind of non-local correlations it exhibits are “private”—the higher the winning probability, the less information any adversary can have about the devices’ outcomes. The amount of entropy, or secrecy, generated in a single round of the protocol can therefore be calculated

³A recent sequence of breakthrough experiments have verified the quantum advantage in non-local games in a loophole-free way [30–32]. In the context of device-independent cryptography, the fact that the experiments are “loophole-free” means that the experiments were executed without making assumptions that could otherwise be exploited by Eve to compromise the security of a cryptographic protocol.

from the winning probability in the game. Let us gain some intuition regarding the relation between the winning probability in a non-local game and the knowledge of the adversary about, e.g., Alice’s output in the game by considering two extreme cases—the optimal classical and quantum strategies for the CHSH game.

If the device is classical its local strategy used to win the game can be written as (recall Definition 3.4):

$$P_{AB|XY}(ab|xy) = \int_{\Lambda} d\lambda \Pr[\Lambda = \lambda] P_{A|X\Lambda}(a|x\lambda) \cdot P_{B|Y\Lambda}(b|y\lambda) ,$$

where λ describes the “hidden variable” (or shared randomness). $\Pr[\Lambda = \lambda]$ as well as $P_{A|X\Lambda}$ and $P_{B|Y\Lambda}$ are chosen by the adversary and, in particular, $P_{A|X\Lambda}$ and $P_{B|Y\Lambda}$ may be deterministic. It is easy to see that in such a case Eve may simply keep a copy of λ for herself by extending $P_{AB|XY}$ to include her system E in the following way:

$$P_{ABE|XY}(abe|xy) = \int_{\Lambda} d\lambda \Pr[\Lambda = \lambda] P_{A|X\Lambda}(a|x\lambda) \cdot P_{B|Y\Lambda}(b|y\lambda) \cdot P_{E|\Lambda}(e|\lambda) ,$$

where $P_{E|\Lambda}(e|\lambda) = 1$ if $e = \lambda$ and $P_{E|\Lambda}(e|\lambda) = 0$ otherwise. Since $P_{A|X\Lambda}$ is deterministic, λ (and x , which is considered to be known to the adversary in most cryptographic protocols) reveals all the information about Alice’s outcome a . Hence, Eve has full information about Alice’s outcome.

However, if the device is implementing the *optimal* quantum strategy then the underlying quantum state and measurements are fully characterised (recall Sect. 3.2.2). In particular, the state shared between Alice and Bob must be the maximally entangled state. As such, any quantum state held by Eve must be completely uncorrelated with Alice and Bob’s state, i.e., it is of the form

$$\rho_{Q_A Q_B E} = |\Phi^+\rangle\langle\Phi^+|_{Q_A Q_B} \otimes \rho_E ,$$

and hence is uncorrelated with Alice’s measurement outcome. Furthermore, measuring the maximally entangled state using the optimal measurements employed by the device results in a uniformly distributed bit on Alice’s side. In total we get that Alice’s output is completely random from Eve’s perspective. In Sect. 5.2 we will see a quantitative relation between the knowledge of any adversary and the winning probability in the CHSH game which goes beyond the above two extreme cases.

We use the task of device-independent QKD (DIQKD) as one of the showcases considered in this thesis. In DIQKD the goal of the honest parties, called Alice and Bob, is to create a shared key, unknown to everybody else but them. To execute the protocol they hold a device consisting of two parts: each part belongs to one of the parties and is kept in their laboratories. Ideally, the device performs measurements on some entangled quantum states it contains. The basic structure of a DIQKD protocol was presented as Protocol 1.1. The protocol consists of playing n non-local games with the given untrusted device and calculating the average winning probability from

the observed data (i.e., Alice and Bob’s inputs and outputs). If the average winning probability is below the expected winning probability ω_{exp} defined by the protocol, Alice and Bob conclude that something is wrong and *abort* the protocol. Otherwise, they apply classical post-processing steps that allow them to create identical and uniformly distributed keys. (The full description of the DIQKD protocol considered in the analysis performed in the following chapters is given in Sect. 4.2.2 below).

Barrett et al. [29] were the first to derive a “proof of concept”⁴ of the security of DIQKD. Following that an extended line of research has explored the application of the device-independence paradigm to multiple cryptographic tasks. A partial list includes QKD [29, 33–36], randomness expansion [36–40] and amplification [41–45], verified quantum computation [35, 46–48], bit commitment [49] and weak string erasure [50].

The following sections present the preliminary knowledge needed when considering our showcase of device-independent quantum cryptography in the upcoming chapters. Specifically, in Sect. 4.2.1 we explain what is meant when talking about the *security* of DIQKD and present the formal security definitions. Sections 4.2.2 and 4.2.3 describe our DIQKD protocol and explain what is the main challenge in any security proof. Section 4.2.4 includes a possible implementation of the protocol in the honest (i.e., non-adversarial) case while Sect. 4.2.5 describes the assumptions made regarding a potentially malicious device. The security analysis itself is presented as a showcase in later chapters. In particular, the full security proof, which previously appeared in [51], is given in Chap. 11.

4.2.1 DIQKD Security Definitions

A DIQKD protocol consists of an interaction between two trusted parties, Alice and Bob, and an untrusted device as defined in Sect. 3.3. At the end of the protocol each party outputs a key of length ℓ , \tilde{K}_A for Alice and \tilde{K}_B for Bob. The goal of the adversary, Eve, is to gain as much information as possible about Alice and Bob’s keys without being detected (i.e., in the case where the protocol is not being aborted).

Correctness, secrecy, and overall security of a DIQKD protocol are defined as follows (see also [52, 53]):

Definition 4.3 (*Correctness*) A DIQKD protocol is said to be $\varepsilon_{\text{corr}}$ -correct, when implemented using a device D , if Alice and Bob’s keys, \tilde{K}_A and \tilde{K}_B respectively, are identical with probability at least $1 - \varepsilon_{\text{corr}}$. That is, $\Pr(\tilde{K}_A \neq \tilde{K}_B) \leq \varepsilon_{\text{corr}}$.⁵

⁴The protocol of [29] could not tolerate any amount of noise and produced just one secret bit when using the device many times (i.e. the key rate is zero); we therefore consider it to be a “proof of concept” showing that device-independent security is possible to achieve.

⁵We use the convention that when the protocol aborts, $\tilde{K}_A = \tilde{K}_B = \perp$.

Definition 4.4 (*Secrecy*) A DIQKD protocol is said to be ε_{sec} -secret, when implemented using a device D , if for a key of length ℓ ,⁶

$$(1 - \text{Pr}[\text{abort}]) \|\rho_{\tilde{K}_{AE}} - \rho_{U_\ell} \otimes \rho_E\|_1 \leq \varepsilon_{sec} ,$$

where $\text{Pr}[\text{abort}]$ is the probability that the protocol aborts when running using device D and $\rho_{\tilde{K}_{AE}}$ is Alice and Eve's quantum state in the end of the protocol, conditioned on not aborting, with E a quantum register holding Eve's state that may initially be correlated with D .

ε_{sec} in the above definition can be understood as the probability that some non-trivial information leaks to the adversary [52].

If a protocol is ε_{corr} -correct and ε_{sec} -secret (for a given D), then it is $\varepsilon_{\text{QKD}}^s$ -correct-and-secret for any $\varepsilon_{\text{QKD}}^s \geq \varepsilon_{corr} + \varepsilon_{sec}$.

Definition 4.5 (*Security*) A DIQKD protocol is said to be $(\varepsilon_{\text{QKD}}^s, \varepsilon_{\text{QKD}}^c)$ -secure if:

1. (Completeness) There exists an honest implementation of the device D such that the protocol does not abort with probability greater than $1 - \varepsilon_{\text{QKD}}^c$.
2. (Soundness) For *any* implementation of the device D , the protocol is $\varepsilon_{\text{QKD}}^s$ -correct-and-secret.

The protocols that we consider below take into account possible noise in the honest implementation. That is, even when there is no adversary at all, the actual implementation of the devices might not be perfect. This should be taken into account when proving the *completeness* of the protocol—completeness must be proven for noisy but honest devices (as otherwise the protocol is of no real use). By doing so we get that the completeness of the protocol implies its *robustness* to the desired amount of noise.

Lastly, a remark regarding the *composability* of this security definition is in order. A security definition is said to be composable [52, 54, 55] if it implies that the protocol can be used arbitrarily and composed with other protocols (proven secure by themselves), without compromising security. Obviously, if Alice and Bob wish to use the keys they produced in a DIQKD protocol in some other cryptographic protocol (i.e., they compose the two protocols), it is necessary for them to use protocols which were proven to have composable security.

For the case of (device-dependent) QKD, Definition 4.5 was rigorously proven to be composable [52]. This suggests that the same security definition should also be the relevant one in the device-independent context and, indeed, as far as we are aware, it is the sole definition used in works on DI cryptography. Nevertheless, the claim that Definition 4.5 is composable for device-independent protocols as well has never been rigorously proven. Even worse, there is some evidence indicating that the definition is *not* composable when the same devices are being reused in the composition. Let us briefly explain that.

⁶ ℓ can be thought of as a parameter of the protocol. In what follows, we set ℓ in terms of the other parameters of the protocol, such that secrecy holds for the protocol.

Barrett et al. [56] highlighted a simple fact: A malicious device may store the raw data used to create the key in a first execution of the DIQKD protocol and then, when reusing the device to execute the protocol *for the second time* (or any other protocol for that matter), leak the raw data from the first run.⁷ Our security definition, Definition 4.5, deals only with a single execution of the protocol and, hence, does not address this type of attack. In other words, even when proving that the considered protocol is secure according to Definition 4.5, the above attack can still be performed by a malicious device when composing two protocols that utilise the same device. This implies that, as is, the security definition is not composable. Note that the same issue does *not* arise when considering device-*dependent* protocols; there, by assumption, the devices do not keep any information in their memory after the end of the execution of the protocol.

Even given the above, Definition 4.5 seems like the most promising security definition to date. We therefore stick to it here. This implies that, as in all other works, *after the end of the protocol* the device cannot be used again in an arbitrary way.

4.2.2 DIQKD Protocol

Our protocol for DIQKD is described as Protocol 4.1. An honest implementation of a device that can be used to run the protocol is described in Sect. 4.2.4.

In the first part of the protocol Alice and Bob use their devices to produce the raw data by playing n CHSH games one after the other. Specifically, in each round Alice and Bob randomly choose whether the round is going to be a test round or a generation round ($T_i = 1$ or $T_i = 0$, respectively, in Protocol 4.1). This can be done using classical communication or shared public randomness. In both cases, this information becomes available to Eve during the execution of the protocol. (Crucially, she does not know in advance, i.e., before supplying the devices to Alice and Bob, which rounds are going to be test rounds). The inputs used by Alice and Bob in each round depend on whether it is a test or generation round; see Protocol 4.1.

In the second part of the protocol Alice and Bob apply classical post-processing steps to produce their final keys. We choose classical post-processing steps that optimise the key rate but may not be optimal in other aspects, e.g., computation time. The protocol and the analysis presented in Chap. 11 can easily be adapted for other choices of classical post-processing.

We now describe the three post-processing steps, error correction, parameter estimation, and privacy amplification in detail.⁸

⁷This should not be confused with “reusing” the device in a given execution of the protocol, i.e., playing many non-local games with the same physical device.

⁸In many QKD protocols there is an additional step called “sifting”; in the sifting step Alice and Bob announce their choice of measurements in the different rounds so that they can ignore the rounds that do not contribute to parameter estimation or the generation of the key (for example,

Protocol 4.1 CHSH-based DIQKD protocol

Arguments:

- D – untrusted device of two components that can play CHSH repeatedly
- $n \in \mathbb{N}_+$ – number of rounds
- $\gamma \in (0, 1]$ – expected fraction of test rounds
- ω_{exp} – expected winning probability in an honest implementation
- $\delta_{\text{est}} \in (0, 1)$ – width of the confidence interval for parameter estimation
- EC – error correction protocol
- PA – privacy amplification protocol

- 1: For every round $i \in [n]$ do Steps 2-4:
 - 2: Alice and Bob choose a random $T_i \in \{0, 1\}$ such that $\Pr(T_i = 1) = \gamma$.
 - 3: If $T_i = 0$, Alice and Bob choose $(X_i, Y_i) = (0, 2)$ and otherwise $X_i, Y_i \in \{0, 1\}$ uniformly at random.
 - 4: Alice and Bob use D with X_i, Y_i and record their outputs as A_i and \tilde{B}_i respectively.
- 5: **Error correction:** Alice and Bob apply the error correction protocol EC. If EC aborts they abort the protocol. Otherwise, they obtain raw keys denoted by K_A and K_B .
- 6: **Parameter estimation:** Using \tilde{B} and K_B , Bob sets $W_i = w_{\text{CHSH}}(K_{B_i}, \tilde{B}_i, X_i, Y_i)$ for the test rounds and $W_i = \perp$ otherwise. He aborts if $\sum_{j:T_j=1} W_j < (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n$.
- 7: **Privacy amplification:** Alice and Bob apply the privacy amplification protocol PA on K_A and K_B to create their final keys \tilde{K}_A and \tilde{K}_B of length ℓ .
-

4.2.2.1 Error Correction

An essential property of any QKD protocol is its correctness—Alice and Bob should hold identical keys in the end of the protocol (see Definition 4.3). Since the raw data of the two parties may differ in parts, Alice and Bob need to run an error correction protocol (also termed an “information reconciliation protocol” in the literature). An error correction protocol⁹ starts by the exchange of classical information between Alice and Bob that should help the parties agree on the final key. When the communication is only from one party to the other, the protocol is said to be a “one-way error correction protocol”. By sending classical information about the raw data over a public classical channel the uncertainty of the adversary regarding the key decreases. A good error correction protocol therefore needs to minimise the amount of communication, or leakage, while still allowing to correct the errors with high probability.

in protocols like BB84 [21] Alice and Bob ignore the rounds in which they chose non-identical measurements). Sifting is not necessary in our case since in Step 2 of Protocol 4.1 Alice and Bob choose T_i together (or exchange its value between them) in every round of the protocol and choose their inputs accordingly. This is in contrast to choosing Alice and Bob’s inputs from a product distribution and then adding a sifting step. It follows from our proof technique that making T_i public as we do does not compromise the security of the protocol.

⁹Note that we are discussing *classical* error correction protocols, not to be confused with the task of quantum error correction [57].

In the considered DIQKD protocol, Alice and Bob use an error correction protocol EC to obtain identical raw keys K_A and K_B from their raw data A, \tilde{B} .¹⁰ We use a one-way error correction protocol, based on universal hashing, which minimises the amount of leakage to the adversary [58, 59] (see also [53, Sect. 3.3.2] for more details). To implement EC Alice chooses an hash function and sends the chosen function and the hashed value of her bits to Bob. We denote this classical communication by O and the number of bits of O by leak_{EC} . Bob uses O , together with all his prior knowledge $\tilde{B}XYT$, to compute a guess \hat{A} for Alice's bits A .¹¹ If EC fails to produce a good guess the protocol aborts; in an *honest* implementation this happens with probability at most $\varepsilon_{\text{EC}}^c$. The probability of Alice and Bob not aborting and while holding non-identical keys is at most ε_{EC} .

The following guarantee holds for the described protocol [59, 60]:

$$\text{leak}_{\text{EC}} \leq H_0^{\varepsilon_{\text{EC}}'}(A|\tilde{B}XYT)_{\rho^{\text{honest}}} + \log\left(\frac{1}{\varepsilon_{\text{EC}}}\right), \quad (4.2)$$

for any $\varepsilon_{\text{EC}}^c, \varepsilon_{\text{EC}}', \varepsilon_{\text{EC}} \in [0, 1]$ such that $\varepsilon_{\text{EC}}' = \varepsilon_{\text{EC}}^c - \varepsilon_{\text{EC}}$ and where $H_0^{\varepsilon_{\text{EC}}'}(A|\tilde{B}XYT)_{\rho^{\text{honest}}}$ is the smooth zero-entropy (Definition 2.23) evaluated on the state ρ^{honest} used in an *honest* implementation of the protocol.¹² Equation (4.2) presents the tradeoff between the probability of having non-identical keys after the end of the protocol (ε_{EC}), the probability of the protocol not succeeding in the honest case ($\varepsilon_{\text{EC}}^c$), and the number of bits leaked to the adversary in the process (leak_{EC}). The amount of communication during the error correction protocol is chosen, before running the DIQKD protocol, such that Eq. (4.2) holds. If more errors than expected in the honest implementation occur when running the DIQKD protocol (due to the use of adversarial or too noisy devices), then Bob may not have a sufficient amount of information to obtain a good guess of Alice's bits and hence will not be able to correct the errors. If so, this will be detected with probability at least $1 - \varepsilon_{\text{EC}}$ and the protocol will abort.

4.2.2.2 Parameter Estimation

The goal of the parameter estimation step is to check whether the device D , used to run the protocol, is sufficiently good in order to produce a *secret* key. In the case of device-independent protocols the quantity to be considered is the number of games

¹⁰It will become clear in Sect. 11.3 why we use here \tilde{B}_i rather than B_i . Although it is not relevant at the moment, we keep it like this for the sake of consistency.

¹¹The idea is basically the following—given the output of the hash function, there is a small set of possible strings (from the domain of the function) compatible with it; Bob then chooses the one which is most compatible to his prior knowledge about Alice's key [58, Sect. 4].

¹²For quantum channels with an IID noise model $H_0^{\varepsilon_{\text{EC}}'}(A|\tilde{B}XYT)_{\rho^{\text{honest}}}$ can be bounded by above using the asymptotic equipartition property, discussed in Sect. 7.2.2. The explicit calculation is done in Sect. 11.3.3.

won during the run of the protocol. If the number of games won is not large enough, the honest parties conclude that the device cannot be used to produce a secure key (an adversary may be present). Specifically, we require that the number of games won, $\sum_{j:T_j=1} W_j$, fulfils

$$\sum_{j:T_j=1} W_j \geq (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n, \quad (4.3)$$

where ω_{exp} , γ , and δ_{est} are parameters of the protocol. γ is the probability of a test round while ω_{exp} is the expected winning probability (of an honest device). Thus, the multiplication $\omega_{\text{exp}}\gamma$ gives the expected fraction of games won out of *all* rounds of the protocol. δ_{est} describes the desired confidence interval (which cannot be zero since we consider a finite number of rounds n).

After the error correction step described above, Bob has all of the relevant information to perform parameter estimation from his data alone, without any further communication with Alice.¹³ Using his raw data $\tilde{\mathbf{B}}$ and his guess of Alice's key K_B , Bob sets

$$W_i = \begin{cases} \perp & T_i = 0 \\ w_{\text{CHSH}}(\hat{A}_i, \tilde{B}_i, X_i, Y_i) = w_{\text{CHSH}}(K_{B_i}, \tilde{B}_i, X_i, Y_i) & T_i = 1, \end{cases}$$

where w_{CHSH} is the CHSH winning condition given in Eq.(3.8). Bob aborts if the fraction of successful game rounds is too low, that is, if Eq. (4.3) is *not* fulfilled.

As Bob does the estimation using his guess of Alice's bits, the probability of aborting in this step in an honest implementation, $\varepsilon_{\text{PE}}^c$, is bounded by

$$\begin{aligned} \varepsilon_{\text{PE}}^c \leq & \Pr\left(\sum_{j:T_j=1} W_j < (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n \mid K_A = K_B\right) \\ & + \Pr(K_A \neq K_B \text{ and EC does not abort}). \end{aligned} \quad (4.4)$$

4.2.2.3 Privacy Amplification

The final classical post-processing step is that of privacy amplification. The goal of privacy amplification is to take Alice's raw key¹⁴ \mathbf{A} , on which the adversary may have partial information, and transform it to a secret final key, as required by the secrecy

¹³In many QKD protocols error correction is performed *after* the parameter estimation step. In such cases, Alice and Bob reveal the data collected in the test rounds and use it for parameter estimation. Further information is then communicated during the error correction step.

¹⁴Note that Alice's and Bob's raw keys, \mathbf{A} and $\hat{\mathbf{A}}$ respectively, are identical with high probability, due to the error correction step. As we now explain, in the privacy amplification step Alice and Bob can perform the exact same actions so that they end with identical final keys (assuming that the error correction step was successful). Thus, we describe here only Alice's actions, while keeping in mind that Bob is going to perform the same steps on his raw key.

definition of the protocol (Definition 4.4). To this end, Alice applies a quantum-proof randomness extractor, defined as follows.

Definition 4.6 (*quantum-proof strong extractor*) A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ that takes as an input a string $A \in \{0, 1\}^n$ together with a seed $S \in \{0, 1\}^d$ and outputs a string $\tilde{K}_A \in \{0, 1\}^\ell$ (for $\ell \leq n$) is called a *quantum-proof* $(m, \varepsilon_{\text{PA}})$ -strong extractor if for any ρ_{AE} with $H_{\min}(A|E)_\rho \geq m$ and uniformly distributed seed S we have

$$\|\rho_{\text{Ext}(A,S)SE} - \rho_{U_\ell} \otimes \rho_{SE}\|_1 \leq \varepsilon_{\text{PA}}. \quad (4.5)$$

Several constructions of extractors have been shown to fulfil the above definition, among them [61–64]. Different constructions are used in different scenarios; for example, some constructions minimise the length d of the seed while others maximise the output length ℓ or the computation time needed to apply the extractor.

Before continuing, an important (though somewhat technical) remark is in order. An extractor, as above, is defined with respect to the min-entropy. However, it is the *smooth* min-entropy $H_{\min}^{\varepsilon_s}(A|E)_\rho$, rather than the min-entropy, that is known to give a tight bound on the maximum amount of uniform randomness that can be extracted from A while being independent from E [65]. If one is interested in using an extractor when starting with a lower bound on the smooth min-entropy, then some parameters should be adapted. In particular, ε_{PA} appearing in Eq. (4.5) is the error probability of the extractor when it is applied on a normalised state satisfying the relevant min-entropy condition. For universal hashing [61] for example, when only a bound on the smooth min-entropy is supplied the smoothing parameter ε_s should be added to the error ε_{PA} (as done below). When working with other extractors one should adapt the parameters accordingly; see [66, Sect. 4.3].

For simplicity we use universal hashing [61, 67] as our privacy amplification protocol PA.¹⁵ The secrecy of the final key $\tilde{K}_A = \text{Ext}(A, S)$ depends only on the privacy amplification protocol used and the value of $H_{\min}^{\varepsilon_s}(A|XYTOE)$, evaluated on the state at the end of the protocol, conditioned on not aborting. For universal hashing, for any $\varepsilon_{\text{PA}}, \varepsilon_s \in (0, 1)$ a secure key of maximal length [67]

$$\ell = H_{\min}^{\varepsilon_s}(A|XYTOE) - 2 \log \frac{1}{\varepsilon_{\text{PA}}} \quad (4.6)$$

is produced with probability at least $1 - \varepsilon_{\text{PA}} - \varepsilon_s$.

¹⁵Any other quantum-proof strong extractor, e.g., Trevisan’s extractor [64], can be used for this task and the analysis done in Chap. 11 can be easily adapted.

4.2.3 Main Task of a Security Proof

After presenting the DIQKD protocol and the relevant security definition we are equipped with the necessary information needed to explain what the main task is when proving security of a DIQKD protocol. First, note that in order to prove security one needs to prove both the correctness (Definition 4.3) and the secrecy (Definition 4.4) of the protocol. Correctness follows almost directly from the error correction step performed in the protocol. We therefore focus below on the secrecy of the protocol.

Returning to the secrecy requirement of a DIQKD protocol given in Definition 4.4 and the definition of a quantum-proof extractor as in Definition 4.6, we see that by applying the extractor we assure that the output of the extractor $\tilde{K}_A = \text{Ext}(A, S)$ is ε_{PA} -close to an ideal key, i.e., a uniform key of m bits that is completely independent of the overall side-information SE ¹⁶ and hence the protocol is secret.

For the extractor to work, the raw data A must exhibit a sufficient amount of min-entropy (by definition). Relations for specific extractor, such as the one given in Eq. (4.6), determine the length of the key that can be extracted for a given amount of (smooth) min-entropy. Therefore, the main task of any security proof of a protocol applying an extractor boils down to computing a lower bound on the (smooth) min-entropy. Indeed, the security proofs presented in Sect. 7.3.2 and Chap. 11 are focused on deriving such bounds.

4.2.4 The Honest Implementation

The honest implementation of the device D describes the way the device acts when an adversary is not present. In other words, this is the device Alice and Bob expect to share when the manufacture of the device is not malicious and “everything goes according to the plan”. In the analysis of DIQKD the description of the honest implementation is used in two places. Firstly, the completeness of the protocol (recall Definition 4.5) is proven with respect to the chosen honest implementation. Secondly, it is used to set the amount of communication between Alice and Bob during the error correction step, according to the relation presented in Eq. (4.2). We remark that these are the only two places in the proof where the choice of honest implementation is taken into account and both are used solely for *choosing* the parameters of the protocol. Critically, the soundness proof does not depend in any way on the choice of honest implementation.

¹⁶We include the seed S as part of the side-information and ask that the output of the extractor is close to uniform even conditioned on the seed S . Extractors that fulfil this requirement are called “strong extractors” (while those that fulfil the weaker condition $\|\rho_{\text{Ext}(A,S)E} - \rho_{U_\ell} \otimes \rho_E\|_1 \leq \varepsilon_{\text{PA}}$ are termed “weak extractors”). When considering QKD protocols, one needs to use a strong extractor since the seed S is to be communicated between Alice and Bob and hence should be considered as information which leaks to the adversary.

The chosen honest implementation may also be noisy. In fact, in an experiment, the mathematical description of the honest device, or honest boxes, should be chosen to fit the behaviour of the physical systems as accurately as possible. An accurate description allows us carefully choose the parameters of the DIQKD protocol (e.g., ω_{exp}) such that the produced key rate is maximised while keeping the probability of the protocol aborting, when utilising the honest device, small. That is, an accurate description allows us to construct a protocol which is useful in practice.

Most commonly, one chooses the honest implementation to be an IID one. That is, that device D acts in an IID manner: in *every round* $i \in [n]$ of the protocol D performs the measurements $\mathcal{M}_{x_i}^{a_i} \otimes \mathcal{M}_{y_i}^{b_i}$ on Alice and Bob's state $\sigma_{Q_A Q_B}$. That is, the device is initialised with an IID bipartite state, $\sigma_{Q_A Q_B}^{\otimes n}$, on which the device makes IID measurements. The state $\sigma_{Q_A Q_B}$ and measurements are such that the winning probability achieved in the CHSH game in a single round is ω_{exp} .¹⁷

As a concrete example, one possible realisation of such an implementation is the following. Alice and Bob share the two-qubit Werner state

$$\sigma_{Q_A Q_B} = (1 - \nu)|\phi^+\rangle\langle\phi^+| + \nu\frac{\mathbb{I}}{4}$$

for $|\Phi^+\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$ and $\nu \in [0, 1]$. The state $\sigma_{Q_A Q_B}$ arises, e.g., from the state $|\Phi^+\rangle$ after going through a depolarisation channel. We can therefore think of the over all state $\sigma_{Q_A Q_B}^{\otimes n}$ as resulting from the transmission of $|\Phi^+\rangle^{\otimes n}$ using an IID noisy channel. For every $i \in [n]$, Alice's measurements $X_i = 0$ and $X_i = 1$ correspond to the Pauli operators¹⁸ σ_z and σ_x respectively and Bob's measurements $Y_i = 0$, $Y_i = 1$, and $Y_i = 2$ to the Pauli operators $\frac{\sigma_z + \sigma_x}{\sqrt{2}}$, $\frac{\sigma_z - \sigma_x}{\sqrt{2}}$ and σ_z respectively. The winning probability in the CHSH game (restricted to $X_i, Y_i \in \{0, 1\}$) using these measurements on $\sigma_{Q_A Q_B}$ is

$$\omega_{\text{exp}} = \frac{2 + \sqrt{2}(1 - \nu)}{4}$$

and the quantum bit error rate is given by

$$Q = \Pr[A_i \neq B_i | (X_i, Y_i) = (0, 2)] = \frac{\nu}{2}.$$

¹⁷Note that in our notation, the noise that affects the winning probability in the CHSH game is already included in ω_{exp} .

¹⁸Even though both are denoted by σ , do not confuse our bipartite state $\sigma_{Q_A Q_B}$ describing the honest state with the Pauli operators σ_x and σ_z defined in Eq.(2.4).

4.2.5 *Model of an Arbitrary Device*

As previously mentioned, Alice and Bob's device is considered to be an untrusted device, as defined in Sect. 3.3.

On top of the general statements made in Sect. 3.3, we can further describe the untrusted device in the case of DIQKD as follows. Alice and Bob interact with D according to Protocol 4.1. Alice and Bob's components of D implement the protocol by making sequential measurements on quantum states. In each round of the protocol, we say that the device is implementing some strategy for the CHSH game. The device may have memory, and thus apply a different strategy each time the game is played, depending on the previous rounds. Therefore, the measurement operators may change in each round, and the state on which the measurements are performed may be the post-measurement state from the previous round, a new state, or any combination of these two.

To be specific, we consider the following scenario. *In-between* different rounds of the protocol, Alice and Bob's components of the device are allowed to communicate freely. During the execution of a single round, however, no communication is allowed. In particular, when the game is being played, there is no communication between the components once the honest parties' inputs are chosen and until the outputs are supplied by the device. That is, communication is allowed in every round i right after Step 4 is done, and until the beginning of round $i + 1$, i.e., before T_{i+1} is chosen in Step 2. Furthermore, in-between rounds Eve may send information to the device, but not receive any from it. In actual implementations this implies that entanglement can be distributed "on the fly" for each round of the protocol, instead of maintaining large quantum memories.

Section 3.3 includes a list of standard assumptions made when working with device-independent protocols. The following list includes the assumptions that are made when proving the security of DIQKD:

1. Alice and Bob have a trusted random number generator.
2. Alice and Bob have trusted classical post-processing units.
3. There is a public, but authenticated, classical channel connecting the honest parties.
4. Alice's and Bob's physical locations are secure (unwanted information cannot leak outside to Eve).
5. Quantum physics is correct.

References

1. Håstad J (2001) Some optimal inapproximability results. J ACM (JACM) 48(4):798–859
2. Barrett J, Collins D, Hardy L, Kent A, Popescu S (2002) Quantum nonlocality, bell inequalities, and the memory loophole. Phys Rev A 66(4):042111

3. Feige, U (1991) On the success probability of the two provers in one-round proof systems. In: 1991 proceedings of the sixth annual structure in complexity theory conference. IEEE, pp 116–123
4. Raz R (2011) A counterexample to strong parallel repetition. *SIAM J Comput* 40(3):771–777
5. Kempe, J, Regev, O (2010) No strong parallel repetition with entangled and non-signaling provers. In: 2010 IEEE 25th annual conference on computational complexity (CCC). IEEE, pp 7–15
6. Raz R (1998) A parallel repetition theorem. *SIAM J Comput* 27(3):763–803
7. Holenstein T (2007) Parallel repetition: simplifications and the no-signaling case. In: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing. ACM, pp 411–419
8. Rao A (2011) Parallel repetition in projection games and a concentration bound. *SIAM J Comput* 40(6):1871–1891
9. Buhman H, Fehr S, Schaffner C (2013) On the parallel repetition of multi-player games: the no-signaling case. arXiv preprint [arXiv:1312.7455](https://arxiv.org/abs/1312.7455)
10. Arnon-Friedman R, Renner R, Vidick T (2016) Non-signaling parallel repetition using de finetti reductions. *IEEE Trans Inf Theory* 62(3):1440–1457
11. Lancien C, Winter A (2016) Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de finetti reduction. *Chicago J Theor Comput Sci* (11)
12. Holmgren J, Yang L (2017) (A counterexample to) parallel repetition for non-signaling multi-player games. In: Electronic colloquium on computational complexity (ECCC), vol 24, p 178
13. Yuen H (2016) A parallel repetition theorem for all entangled games. *Int Colloq Autom, Lang, Program*
14. Cleve R, Slofstra W, Unger F, Upadhyay S (2008) Perfect parallel repetition theorem for quantum xor proof systems. *Comput Complex* 17(2):282–299
15. Kempe J, Regev O, Toner B (2010) Unique games with entangled provers are easy. *SIAM J Comput* 39(7):3207–3229
16. Dinur I, Vidick Steurer DT (2015) A parallel repetition theorem for entangled projection games. *Comput Complex* 24(2):201–254
17. Chung K-M, Wu X, Yuen H (2015) Parallel repetition for entangled k-player games via fast quantum search. In: Proceedings of the 30th conference on computational complexity. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp 512–536
18. Kempe J, Vidick T (2011) Parallel repetition of entangled games. In: Proceedings of the forty-third annual ACM symposium on theory of computing. ACM, pp 353–362
19. Bavarian M, Vidick T, Yuen H (2017) Anchoring games for parallel repetition. *Symposium on the theory of computing*
20. Bavarian M, Vidick T, Yuen H (2017) Parallel repetition via fortification: analytic view and the quantum case. *Innovat Theor Comput Sci*
21. Bennett CH, Brassard G (1984) Proceedings of the IEEE international conference on computers, systems, and signal processing, bangalore, india, 1984
22. Ekert AK (1991) Quantum cryptography based on Bell’s theorem. *Phys Rev Lett* 67(6):661
23. Fung C-HF, Qi B, Tamaki K, Lo H-K (2007) Phase-remapping attack in practical quantum-key-distribution systems. *Phys. Rev. A* 75(3):032314
24. Lydersen L, Wiechers C, Wittmann C, Elser D, Skaar J, Makarov V (2010) Hacking commercial quantum cryptography systems by tailored bright illumination. *Nat. Photon.* 4(10):686–689
25. Weier H, Krauss H, Rau M, Fürst M, Nauerth S, Weinfurter H (2011) Quantum eavesdropping without interception: an attack exploiting the dead time of single-photon detectors. *New J. Phys.* 13(7):073024
26. Gerhardt I, Liu Q, Lamas-Linares A, Skaar J, Kurtsiefer C, Makarov V (2011) Full-field implementation of a perfect eavesdropper on a quantum cryptography system. *Nat Commun* 2:349
27. Ekert A, Renner R (2014) The ultimate physical limits of privacy. *Nature* 507(7493):443–447
28. Mayers D, Yao A (1998) Quantum cryptography with imperfect apparatus. In: 39th annual symposium on foundations of computer science, 1998. Proceedings. IEEE, pp 503–509

29. Barrett J, Hardy L, Kent A (2005) No signaling and quantum key distribution. *Phys Rev Lett* 95(1):010503
30. Hensen B, Bernien H, Dréau A, Reiserer A, Kalb N, Blok M, Ruitenberg J, Vermeulen R, Schouten R, Abellán C, et al (2015) Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* 526(7575):682–686
31. Shalm LK, Meyer-Scott E, Christensen BG, Bierhorst P, Wayne MA, Stevens MJ, Gerrits T, Glancy S, Hamel DR, Allman MS et al (2015) Strong loophole-free test of local realism. *Phys Rev Lett* 115(25):250402
32. Giustina M, Versteegh MA, Wengerowsky S, Handsteiner J, Hochrainer A, Phelan K, Steinlechner F, Kofler J, Larsson J-Å, Abellán C et al (2015) Significant-loophole-free test of Bell’s theorem with entangled photons. *Phys Rev Lett* 115(25):250401
33. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021
34. Vazirani U, Vidick T (2014) Fully device-independent quantum key distribution. *Phys Rev Lett* 113(14):140501
35. Reichardt BW, Unger F, Vazirani U (2013) Classical command of quantum systems. *Nature* 496(7446):456–460
36. Miller CA, Shi Y (2014) Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. In: *Proceedings of the 46th annual ACM symposium on theory of computing*. ACM, pp 417–426
37. Colbeck, R (2006) Quantum and relativistic protocols for secure multi-party computation. PhD thesis, Trinity College, University of Cambridge
38. Pironio S, Acín A, Massar S, de La Giroday AB, Matsukevich DN, Maunz P, Olmschenk S, Hayes D, Luo L, Manning TA et al (2010) Random numbers certified by Bell’s theorem. *Nature* 464(7291):1021–1024
39. Vazirani U, Vidick T (2012) Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In: *Proceedings of the forty-fourth annual ACM symposium on theory of computing*. ACM, pp 61–76
40. Coudron M, Yuen H (2013) Infinite randomness expansion and amplification with a constant number of devices. arXiv preprint [arXiv:1310.6755](https://arxiv.org/abs/1310.6755)
41. Colbeck R, Renner R (2012) Free randomness can be amplified. *Nat Phys* 8(6):450–453
42. Gallego R, Masanes L, De La Torre G, Dhara C, Aolita L, Acín A (2013a) Full randomness from arbitrarily deterministic events. *Nat Commun* 4:2654
43. Chung K-M, Shi Y, Wu X (2014) Physical randomness extractors: Generating random numbers with minimal assumptions. arXiv preprint [arXiv:1402.4797](https://arxiv.org/abs/1402.4797)
44. Brandão FG, Ramanathan R, Grudka A, Horodecki K, Horodecki M, Horodecki P, Szarek T, Wojewódka H (2016) Realistic noise-tolerant randomness amplification using finite number of devices. *Nat Commun* 7:11345
45. Kessler M, Arnon-Friedman R (2017) Device-independent randomness amplification and privatization. arXiv preprint [arXiv:1705.04148](https://arxiv.org/abs/1705.04148)
46. Gheorghiu A, Kashefi E, Wallden P (2015) Robustness and device independence of verifiable blind quantum computing. *New J Phys* 17(8):083040
47. Hajdušek M, Pérez-Delgado CA, Fitzsimons JF (2015) Device-independent verifiable blind quantum computation. arXiv preprint [arXiv:1502.02563](https://arxiv.org/abs/1502.02563)
48. Coladangelo A, Grilo A, Jeffery S, Vidick T (2017) Verifier-on-a-leash: new schemes for verifiable delegated quantum computation, with quasilinear resources. arXiv preprint [arXiv:1708.07359](https://arxiv.org/abs/1708.07359)
49. Aharon N, Massar S, Pironio S, Silman J (2015) Device-independent bit commitment based on the CHSH inequality. arXiv preprint [arXiv:1511.06283](https://arxiv.org/abs/1511.06283)
50. Kaniewski J, Wehner S (2016) Device-independent two-party cryptography secure against sequential attacks. *New J Phys* 18(5):055004
51. Arnon-Friedman R, Renner R, Vidick T (2019) Simple and tight device-independent security proofs. *SIAM J Comput* 48(1):181–225

52. Portmann C, Renner R (2014) Cryptographic security of quantum key distribution. arXiv preprint [arXiv:1409.3525](https://arxiv.org/abs/1409.3525)
53. Beaudry NJ (2015) Assumptions in quantum cryptography. arXiv preprint [arXiv:1505.02792](https://arxiv.org/abs/1505.02792)
54. Canetti R (2001) Universally composable security: a new paradigm for cryptographic protocols 2005. Revision 3 of ECCC report
55. Ben-Or M, Mayers D (2004) General security definition and composability for quantum and classical protocols. arXiv preprint [arXiv:quant-ph/0409062](https://arxiv.org/abs/quant-ph/0409062)
56. Barrett J, Colbeck R, Kent A (2013) Memory attacks on device-independent quantum cryptography. *Phys Rev Lett* 110(1):010503
57. Gottesman D (2010) An introduction to quantum error correction and fault-tolerant quantum computation. *Quantum information science and its contributions to mathematics. Proc Symp Appl Math* 68:13–58
58. Brassard G, Salvail L (1993) Secret-key reconciliation by public discussion. In: *Advances in cryptology EUROCRYPT 93*. Springer, pp 410–423
59. Renner R, Wolf S (2005) Simple and tight bounds for information reconciliation and privacy amplification. In: *Advances in cryptology-ASIACRYPT 2005*. Springer, pp 199–216
60. Renner R (2008) Security of quantum key distribution. *Int J Quantum Inf* 6(01):1–127
61. Renner R, König R (2005) Universally composable privacy amplification against quantum adversaries. In: *Theory of cryptography*. Springer, pp 407–425
62. König RT, Terhal BM (2008) The bounded-storage model in the presence of a quantum adversary. *IEEE Trans Inf Theory* 54(2):749–762
63. Fehr S, Schaffner C (2008) Randomness extraction via δ -biased masking in the presence of a quantum attacker. In: *Theory of cryptography conference*. Springer, pp 465–481
64. De A, Portmann C, Vidick T, Renner R (2012) Trevisan’s extractor in the presence of quantum side information. *SIAM J Comput* 41(4):915–940
65. König R, Renner R, Schaffner C (2009) The operational meaning of min-and max-entropy. *IEEE Trans Inf Theory* 55(9):4337–4347
66. Arnon-Friedman R, Portmann C, Scholz VB (2016) Quantum-proof multi-source randomness extractors in the markov model. In: *11th conference on the theory of quantum computation, communication and cryptography*, p 1
67. Tomamichel M, Renner R, Schaffner C, Smith A (2010) Leftover hashing against quantum side information. In: *2010 IEEE international symposium on information theory proceedings (ISIT)*. IEEE, pp 2703–2707

Chapter 5

Single-Round Box



In the device-independent framework we use “boxes” to describe the physical devices, or resources, of interest. A box, formally modelled as a conditional probability distribution (recall Sect. 3.1), is always defined with respect to a *specific* task or protocol. More specifically, note the following:

1. To define a box $P_{AB|XY}$ we need to fix the sets of the inputs \mathcal{X} , \mathcal{Y} and the outputs \mathcal{A} , \mathcal{B} of the box. These sets are chosen according to the task in which the box is being used. For example, if a box is used to play a single CHSH game then the sets are all chosen to be $\{0, 1\}$. The box’s action is undefined when it is used with, e.g., the input $x = 2$.
2. The location of the used devices in space (or space-time) also sets the conditions that the box describing the devices must fulfil. For example, if a protocol demands two devices, separated in space, that cannot communicate during the execution of the protocol then the defined box should fulfil certain non-signalling conditions.¹
3. When considering boxes that are used to execute a complex protocol, in which many games are being played with the box (as done in the succeeding chapters), we also need to take into account the type of interaction when defining the box. For example, some protocols require boxes with which we can interact sequentially—in each round of the protocol we give one input to the box, wait for the output, and only then give the next input. Other protocols involve boxes which accepts all the

¹Interestingly, if one considers protocols with more than two parties in which the devices can only be used in specific space-time coordinates and merely assumes that the box modelling the devices respects relativistic causality (in the sense that it cannot lead to casual loops) then the conditions defining the box are different than the non-signalling ones [1]. This acts as another example for how the specific use of the devices effects the mathematical model of the box.

inputs and only then produces all the outputs. If we only give one input to such a box we do not expect it to output anything and its action is undefined. Thus, these differences in the behaviour of the boxes depend on the way we intend to use it in the task of interest and effect the mathematical model of the considered boxes.

To grasp the dependence of the box on the considered task, as described above, one can contrast it with the standard formalism used to define quantum states and measurements. For example, the definition of a quantum state in terms of a density operator is completely independent of the way we might want to measure it. Consider, for example, a quantum state used to play the CHSH game with the measurements σ_x and σ_z for one of the parties. Even though we only intend to perform these measurements, the formalism also tells us what will happen if we choose to measure σ_y instead. This stands in contrast to Item 1 above.²

The current chapter as well as Chap. 6 are devoted to the way one models the different boxes used in device-independent information processing, depending on the considered setting and interaction with the boxes. In Chap. 6 we will be interested in boxes, or devices, which can be used to implement certain protocols. Before we explain how such boxes can be described let us focus on a simpler object—the “single-round box”.

We think of a single-round box as illustrated in Fig. 5.1, as a small device that can be used to play a *single* round of a Bell game. That is, in the case of the CHSH game, for example, Alice and Bob can input their bits $x, y \in \{0, 1\}$ to the box and receive the outcomes $a, b \in \{0, 1\}$. After that the box can no longer be used (i.e., Alice and Bob cannot play another game with it). Mathematically, such a box can be described by a non-signalling conditional probability distribution $P_{AB|XY}$ as explained in Sect. 3.1. Physically, an example of a single-round box is a single EPR pair together with a set of possible measurements for each party.

A single-round box is *not* a useful resource in the *operational sense*. Since our starting point in the device-independent setting is that we do not know how the device operates, we must interact with it to test it. However, since a single-round box allows us to play just a single game we can hardly conclude anything regarding its inner-working. One can imagine Alice and Bob playing the CHSH game with their box and observing $(a, b, x, y) = (0, 0, 0, 0)$. Then what? It can always be the case that they are sharing a classical device that always outputs $(a, b) = (0, 0)$ for the inputs $(x, y) = (0, 0)$. Thus, Alice and Bob cannot learn anything regarding, e.g., the randomness of their outputs, from this single game. As the information collected in a single game is not sufficient to test the box we start, instead, with an *assumption* regarding the box, e.g., that it can be used to win the CHSH game with winning probability ω . As will be shown below, various fundamental properties can be concluded by starting with such an assumption.

²One can rightfully say that this property of boxes, among several other properties, renders them an “unphysical description” of real systems and resources. With this respect, the formalism of the so called “generalised probabilistic theories” [2, 3] is a more appropriate mathematical setting to discuss physical theories which extend, or abstract, quantum physics. In contrast, boxes are merely a simplified mathematical model sufficient for certain analyses.

Fig. 5.1 A single-round box. We think of a single-round box as a small device, shared between Alice and Bob, which can be used to play a single round of a Bell game, such as the CHSH game. It is described by a conditional probability distribution $P_{AB|XY}$



Although a single-round box is not a valuable resource in practice, it is useful as a simple abstract object that allows us to study the fundamental implications of violating a Bell inequality (while putting aside many technical details that arise when considering the complex devices used in protocols). Furthermore, it is the goal of this thesis to explain how “single-round box statements” can be lifted to operational statements regarding more complex scenarios such as the analysis of device-independent protocols.

5.1 The Model

Mathematically, we model a single-round black box by a non-signalling conditional probability distribution $P_{AB|XY}$ that can be used to play a single Bell game G defined over the sets of inputs \mathcal{X}, \mathcal{Y} and outputs \mathcal{A}, \mathcal{B} for Alice and Bob (see Sect. 3.2.1 for complete definitions). $P_{AB|XY}$ is also sometimes referred to as a strategy for G .

As mentioned above, when considering single-round boxes one usually assumes that the box $P_{AB|XY}$ can be used to win the game with a certain winning probability ω . That is, $P_{AB|XY}$ is such that

$$\mathbb{E}_{x,y} \sum_{\substack{a,b \\ w(a,b,x,y)=1}} P_{AB|XY}(ab|xy) = \omega, \quad (5.1)$$

where the expectation $\mathbb{E}_{x,y}$ is defined with respect to the input distribution of the considered game and $w : \mathcal{A} \times \mathcal{B} \times \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ is the winning function of the game.

Depending on the context, one can consider quantum single-round boxes or non-signalling ones.

5.1.1 Quantum Single-Round Boxes

When we say that a single-round box is quantum we mean that its inner-working can be described within the quantum formalism. Specifically:

Definition 5.1 (*Quantum single-round box*) Given a Bell game G , a quantum single-round box is a quantum box $P_{AB|XY}$, as in Definition 3.3, defined for the inputs and outputs of the game $G - \mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}$. That is, there exist a bipartite state $\rho_{Q_A Q_B}$ and measurements $\{M_a^x\}$ and $\{M_b^y\}$ such that

$$P_{AB|XY}(ab|xy) = \text{Tr} \left(M_a^x \otimes M_b^y \rho_{Q_A Q_B} \right) \quad \forall a, b, x, y. \quad (5.2)$$

The quantum single-round box is said to win G with winning probability ω when the state and measurements are such that Eq. (5.1) holds.

Note that mathematically a quantum single-round box is merely a quantum box (Definition 3.3). What makes it *single-round* is that $P_{AB|XY}$ is defined for the inputs and outputs of a single game G .

When considering cryptographic applications where a quantum adversary is present we extend the box to the adversary. That is, we let $\rho_{Q_A Q_B E}$ be the purification of $\rho_{Q_A Q_B}$ where E is a quantum register belonging to the adversary and $\rho_{Q_A Q_B} = \text{Tr}_E (\rho_{Q_A Q_B E})$ is Alice and Bob's marginal satisfying Eqs. (5.1) and (5.2).

5.1.1.1 Non-signalling Single-Round Boxes

Instead of restricting our attention to quantum boxes we can also consider non-signalling single-round boxes. These are defined in a similar way to their quantum counterparts.

Definition 5.2 (*Non-signalling single-round box*) Given a Bell game G , a non-signalling single-round box is a non-signalling box $P_{AB|XY}$, as in Definition 3.1, defined for the inputs and outputs of the game $G - \mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}$. That is, for all $a \in \mathcal{A}$, $b \in \mathcal{B}$, $x, x' \in \mathcal{X}$ and $y, y' \in \mathcal{Y}$,

$$\begin{aligned} \sum_b P_{AB|XY}(a, b|x, y) &= \sum_b P_{AB|XY}(a, b|x, y') \\ \sum_a P_{AB|XY}(a, b|x, y) &= \sum_a P_{AB|XY}(a, b|x', y). \end{aligned}$$

The non-signalling single-round box is said to win G with winning probability ω when $P_{AB|XY}$ is such that Eq. (5.1) holds.

Here as well one can consider an extension of the single-round box to an additional party describing a non-signalling (super-quantum) adversary. This will not be needed in this thesis so we do not explain how this is done. The interested reader is referred to [4, Sect. 3.2].

5.2 Showcase: Device-Independent Quantum Cryptography

As mentioned above, a single-round box is useful as a simple abstract object that allows us to study the fundamental implications of violating a Bell inequality. More specifically, certain properties of the box can be concluded if we assume to know the probability of winning a Bell game using a single-round box described by $P_{AB|XY}$. We consider our showcase of device-independent cryptography as an example.

The most crucial observation when considering device-independent cryptographic protocols is the fact that high winning probability in a Bell game not only implies that the measured system is non-local, but more importantly that the kind of non-locality it exhibits cannot be shared: the higher the winning probability, the less information any eavesdropper can have about the outcomes produced by the box.

There are different ways of making such a statement quantitative. One possible way (that will also be of relevance later on) is to consider the conditional von Neumann entropy $H(A|XYE)$ where A is the random variable describing Alice's outcome bit, X and Y are the random variables describing the inputs of Alice and Bob and E is a quantum register holding the quantum side information belonging to the adversary. If the adversary is completely oblivious to the value of a bit A even given X , Y and E then takes its maximal value $H(A|XYE) = 1$.

A tight trade-off between the winning probability of a single-round box ω and the entropy $H(A|XYE)$ generated by the box was derived in [5, 6] and is stated in the following lemma.

Lemma 5.3 ([5, 6]³) *For any quantum single-round box $P_{AB|XY}$ with winning probability $\omega \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right]$ in the CHSH game,*

$$H(A|XYE) \geq 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\omega(\omega-1)+3}\right), \quad (5.3)$$

where E denotes the quantum side-information belonging to the adversary and $h(\cdot)$ is the binary entropy function.

The relation stated in Eq. (5.3) is plotted in Fig. 5.2. One can see that the entropy increases as the winning probability ω increases. That is, the amount of secret randomness in Alice's outcome is directly related to the winning probability of the single-round box. In particular, we observe that $H(A|XYE) = 0$ (i.e., the adversary knows the value of A) for the optimal classical winning probability and $H(A|XYE) = 1$ (i.e., A looks completely random to the adversary) for the optimal quantum winning probability.⁴ Note that there can be many different boxes $P_{AB|XY}$ (and hence extensions to the adversary) with the same winning probability ω . That is, the assumption

³Lemma 5.3 is stated in the form appearing in [7]. To see how the original results of [5] can be used to derive the lemma as we state it, follow the proof given in Appendix C.1.

⁴These two extreme cases are easy to understand. When the box employs a classical strategy the adversary can simply hold a copy of A . When the box employs the optimal quantum strategy the

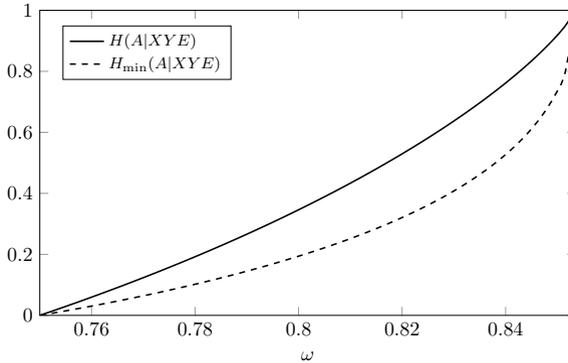


Fig. 5.2 Secrecy versus winning probability ω in the CHSH game for a *single-round box*. Two lower-bounds are shown: one for the conditional von Neumann entropy $H(A|XYE)$ [5] and the other for the conditional min-entropy $H_{\min}(A|XYE)$ [8]; both bounds are tight. As soon as the winning probability is above the classical threshold of 75% some secret randomness is produced

regarding the winning probability of the box does not pin down the full probability distribution. The bound given in Eq. (5.3) is thus very strong—it says that for *any* single-round box with winning probability ω and *any* purification to the adversary the stated lower bound holds.

Instead of considering the von Neumann entropy as above, one can also study lower-bounds on the conditional min-entropy $H_{\min}(A|XYE)$ as a function of the winning probability of a single-round box—as was done in [8]. We plot the resulting bound in Fig. 5.2. As can be seen in the figure, for non-optimal Bell violation the min-entropy can be significantly lower than the von Neumann entropy. Indeed, the min-entropy is always upper-bounded by the von Neumann entropy (hence the name). Still, in some cases a bound on the min-entropy, rather than the von Neumann entropy, is needed or, at the least, is easier to derive. In particular, lower-bounds on the min-entropy for *single-round boxes* can be found using general techniques based on the semidefinite programming hierarchies of [9] while, up to date, there is no general technique to derive (or even estimate) such bounds on the von Neumann entropy.

Similar bounds were derived also for other Bell inequalities. For example, lower-bounds on the min-entropy produced by a single-round box were found as a function of the violation of the Mermin inequality [10, Eq. (6)] and the tilted-CHSH inequality [11, Lemma 2]. Another result in the same spirit is that of [12, Sect. 5], where a bound on the min-entropy is derived as a function of several Bell inequalities all at once.⁵ Lower-bounds on the von Neumann entropy were derived as a function of

used state is the maximally entangled state. Then, due to monogamy of entanglement, the adversary is completely decoupled from the Alice and Bob’s state. For more details see Sect. 4.2.

⁵That is, instead of assuming that we know just the winning probability of the single-round box in a specific game, we assume we know its winning probabilities in several different games. In the context of single-round boxes this is a stronger assumption regarding the device. However, in actual application this is not an issue, as will be mentioned later on.

the violation of the MDL inequalities [13, Sect. 3] and the MABK inequality [14, Lemma S5].

Before continuing to the next chapter, we emphasise once again that single-round statements as mentioned above should not be understood as operational statements. If we are given a single-round box but we do not assume to know its winning probability ω then we cannot conclude anything about its properties (e.g., the entropy of the outputs). When considering, for example, device-independent cryptographic protocols one must test the device in order to estimate whether it can violate a Bell inequality or not. This is done by playing several games with the device and collecting statistic regarding its input-output behaviour. For this purpose we need to consider *multi-rounds* boxes, as done in the following sections.

References

1. Horodecki P, Ramanathan R (2016) Relativistic causality versus no-signaling as the limiting paradigm for correlations in physical theories. arXiv preprint [arXiv:1611.06781](https://arxiv.org/abs/1611.06781)
2. Barrett J (2007) Information processing in generalized probabilistic theories. *Phys Rev A* 75(3):032304
3. Chiribella G, D'Ariano GM, Perinotti P (2010) Probabilistic theories with purification. *Phys Rev A* 81(6):062348
4. Hänggi E (2010) Device-independent quantum key distribution. PhD thesis
5. Pironio S, Acín A, Massar S, de La Giroday AB, Matsukevich DN, Maunz P, Olmschenk S, Hayes D, Luo L, Manning TA et al (2010) Random numbers certified by Bell's theorem. *Nature* 464(7291):1021–1024
6. Acín A, Massar S, Pironio S (2012) Randomness versus nonlocality and entanglement. *Phys Rev Lett* 108(10):100402
7. Arnon-Friedman R, Renner R, Vidick T (2019) Simple and tight device-independent security proofs. *SIAM J Comput* 48(1):181–225
8. Masanes L, Pironio S, Acín A (2011) Secure device-independent quantum key distribution with causally independent measurement devices. *Nat Commun* 2:238
9. Navascués M, Pironio S, Acín A (2008) A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J Phys* 10(7):073013
10. Gallego R, Masanes L, De La Torre G, Dhara C, Aolita L, Acín A (2013) Full randomness from arbitrarily deterministic events. *Nat Commun* 4
11. Bamps C, Massar S, Pironio S (2017) Device-independent randomness generation with sub-linear shared quantum resources. arXiv preprint [arXiv:1704.02130](https://arxiv.org/abs/1704.02130)
12. Nieto-Silleras O, Bamps C, Silman J, Pironio S (2016) Device-independent randomness generation from several bell estimators. arXiv preprint [arXiv:1611.00352](https://arxiv.org/abs/1611.00352)
13. Kessler M, Arnon-Friedman R (2017) Device-independent randomness amplification and privatization. arXiv preprint [arXiv:1705.04148](https://arxiv.org/abs/1705.04148)
14. Ribeiro J, Murta G, Wehner S (2017) Fully device independent conference key agreement. arXiv preprint [arXiv:1708.00798](https://arxiv.org/abs/1708.00798)

Chapter 6

Multi-round Box



In the previous chapter we discussed the *single-round box*, which can be seen as a simple abstract object that allows us to study the fundamental aspects of non-locality. When studying actual device-independent information processing tasks, however, one must consider more complex objects that describe the behaviour of the devices while performing the task of interest. More concretely, in actual applications we usually interact with a device by playing *many* games. Even in the simplest setting where one would like to merely verify the violation of a Bell inequality, as in experiments performing loophole-free Bell tests, a Bell game is played many times so that sufficient amount of data can be collected to estimate the violation in a satisfactory statistical manner. Playing just a single game is clearly not enough. Another example is device-independent protocols, such as quantum key distribution. All protocols include a phase in which the users (or honest parties) are playing many games with their device in order to decide whether it can be used for the considered task. Hence, considering boxes that can be used to play just a single game is not enough. Instead, we need to work with *multi-round boxes*.

Multi-round boxes can be described using a conditional probability distribution $P_{AB|XY}$ over the inputs and outputs of many rounds of a game. That is, for n the number of games which one would like to play with the box (e.g., the number of rounds of a protocol), $\mathbf{A} = A_1 A_2 \dots A_n$ is a random variable over \mathcal{A}^n and \mathbf{B} , \mathbf{X} , and \mathbf{Y} are similarly defined.

As explained in the beginning of Chap. 5, the way we model a box, and in particular a multi-round box, depends on the type of interaction that we would like to perform with it. We consider two different forms of interactions: parallel and sequential interactions. Different tasks require different types of boxes. Parallel boxes are used, for example, in self-testing [1], parallel quantum key distribution [2], and certification of entanglement [3]. Some examples for settings in which sequential boxes

are considered are delegated computation [4] and randomness amplification [5]. In the scope of this thesis, Chaps. 8 and 10 deal with parallel boxes while Chaps. 9 and 11 focus on sequential boxes.

6.1 Parallel Interaction

The simplest to describe form of interaction is the “parallel interaction”. In such an interaction the box is “expecting” to get the n inputs of all the rounds, \mathbf{x} and \mathbf{y} , at the same time and is expected to give all the outputs, \mathbf{a} and \mathbf{b} , together; see Fig. 6.1. If the box is only given inputs of a single game, e.g., x_1, y_1 , it is not expected to return any output. This behaviour of the box will present itself in the mathematical model of the box, as we explain below.

For a given a game G , a parallel multi-round box is a device with which Alice and Bob can play n instances of G in parallel (i.e., at the same time). Mathematically this translates to a conditional probability distribution $P_{AB|XY}$, non-signalling between Alice and Bob, defined over the inputs and outputs of n games. For example, when considering the CHSH game, $\mathbf{A}, \mathbf{B}, \mathbf{X}$, and \mathbf{Y} are all random variables over $\{0, 1\}^n$.

As explained in Sect. 3.1.1, the non-signalling conditions between Alice and Bob imply that Alice and Bob’s marginals, $P_{A|X}$ and $P_{B|Y}$ respectively, are well-defined. The fact that we are talking about a *parallel* multi-round box means that no further structure can be assumed. In particular, other marginals, e.g., $P_{A_1|X_1}$ or $P_{A_2B_2|X_2B_2}$, are not necessarily well-defined. Intuitively this stands for the fact that the box is expecting to get all the inputs together and only then it produces the outputs; the output for A_1 can therefore depend, for example, on the value of X_5 and not on

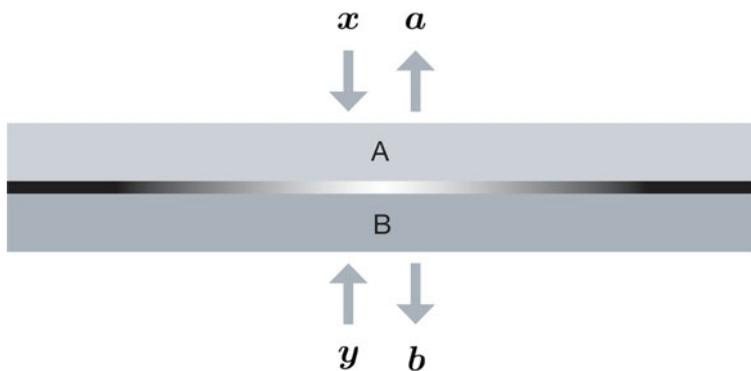


Fig. 6.1 Parallel multi-round box. We think of a parallel multi-round box as a large device, shared between Alice and Bob, which can be used to play many rounds of a Bell game, all at once. Such a box is expecting to get the inputs for all rounds, \mathbf{x} and \mathbf{y} , at the same time, and it will then produce all the outputs, \mathbf{a} and \mathbf{b} for Alice and Bob

just that of X_1 . Hence the conditional probability distribution $P_{A_1|X_1}$ is not properly defined.

6.1.1 Non-signalling Parallel Boxes

One can consider a parallel multi-round box which is only restricted by the non-signalling conditions. We then get the following definition.

Definition 6.1 (*Non-signalling parallel multi-round box*) Given a Bell game G , a non-signalling parallel multi-round box is a non-signalling box $P_{AB|XY}$, as in Definition 3.1, defined for the inputs and outputs of n rounds of the game $G - \mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$. That is, for all $\mathbf{a} \in \mathcal{A}^n, \mathbf{b} \in \mathcal{B}^n, \mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$ and $\mathbf{y}, \mathbf{y}' \in \mathcal{Y}^n$,

$$\begin{aligned} \sum_b P_{AB|XY}(\mathbf{ab}|\mathbf{xy}) &= \sum_b P_{AB|XY}(\mathbf{ab}|\mathbf{xy}') \\ \sum_a P_{AB|XY}(\mathbf{ab}|\mathbf{xy}) &= \sum_a P_{AB|XY}(\mathbf{ab}|\mathbf{x}'\mathbf{y}) . \end{aligned} \tag{6.1}$$

As mentioned above, the only non-signalling conditions restricting the parallel box, are those between Alice and Bob appearing in Definition 6.1; we do not set any other assumptions regarding the box apart from that.

6.1.1.1 Quantum Parallel Boxes

Similarly to a quantum single-round box, as in Definition 5.1, a quantum parallel multi-round box is just a quantum box (Definition 3.3) defined for the inputs and outputs of n rounds of G .

Definition 6.2 (*Quantum parallel multi-round box*) Given a Bell game G , a quantum parallel multi-round box is a quantum box $P_{AB|XY}$, as in Definition 3.3, defined for the inputs and outputs of n rounds of the game $G - \mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$. That is, there exist a bipartite state $\rho_{Q_A Q_B}$ and measurements $\{M_a^x\}$ and $\{M_b^y\}$ such that

$$P_{AB|XY}(\mathbf{ab}|\mathbf{xy}) = \text{Tr} (M_a^x \otimes M_b^y \rho_{Q_A Q_B}) \quad \forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} . \tag{6.2}$$

The non-signalling conditions in Eq. (6.1) are automatically fulfilled by quantum parallel boxes defined above. We remark again that there are no further assumptions regarding the structure of the state and measurements apart from what appears in Eq. (6.2). Specifically, ρ_{Q_A} and ρ_{Q_B} are not assumed to have some further subsystem structure and the measurements need not have a tensor product form such as $M_{a_1}^{x_1} \otimes \dots \otimes M_{a_n}^{x_n}$.

6.2 Sequential Interaction

In the previous section we discussed *parallel* multi-rounds boxes. These are boxes that allow (and “expect”) to be interacted with in a parallel way, i.e., by giving all the inputs to the box at the same time. As the parallel multi-round box receives all the inputs at once, the output for, e.g., the first game, A_1 , can depend on the inputs for all games X_1, X_2, \dots, X_n .

In this section we consider a different type of multi-round boxes – *sequential multi-round boxes*. Such boxes are, in some sense, more structured than parallel multi-round boxes and accurately model the devices used in many device-independent scenarios. As such, sequential multi-round boxes are of relevance for applications. Furthermore, the additional structure of sequential multi-round boxes will allow us to derive stronger results than those derived for their parallel counterparts.

As mentioned above, the way we model a multi-round box depends on how we would like to interact with it. Most device-independent protocols proceed in rounds which are performed one after the other: Alice and Bob use their box in the first round of the protocol and only once they receive the outputs from the box they proceed to the second round, and so on; See Protocol 1.1 for an example. We call such an interaction with the box “sequential interaction”. This is illustrated in Fig. 6.2 (the reader may compare Fig. 6.2 to the single-round box in Fig. 5.1 and the parallel multi-round box in Fig. 6.1).

The chronological order which is implied by the sequential interaction enforces certain constraints on the behaviour of the box. In particular, while past events can influence future ones, the future cannot change the past. For example, the first output A_1 can depend on the first input X_1 but not on the inputs of the next rounds X_2, \dots, X_n . The second output A_2 can depend both on X_2 and past events, such as the values assigned to A_1 and X_1 , but not on the following inputs X_3, \dots, X_n .

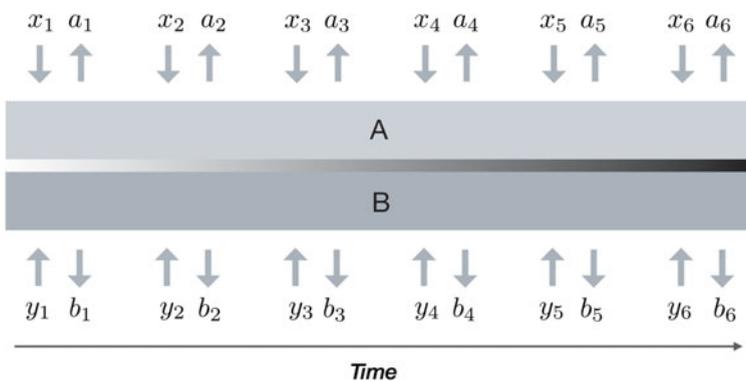


Fig. 6.2 Sequential interaction with a multi-round box. Alice and Bob start by playing the first game with the box and only once they receive the outputs from the box they proceed to the second game, and so on

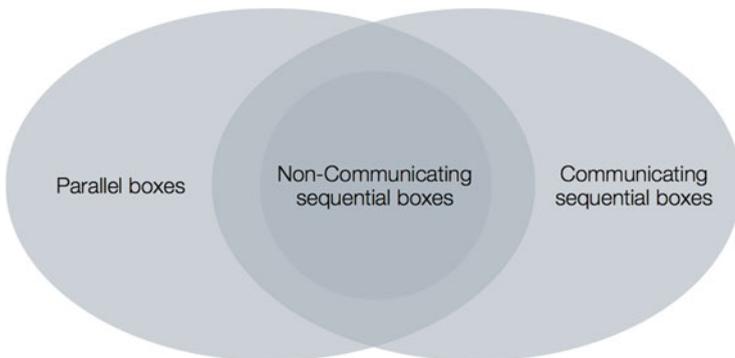


Fig. 6.3 The relation between the different multi-round boxes

We define two different types of sequential boxes – one which allows for communication between the rounds of interactions and one which does not. A box that allows for communication between the rounds is a box in which Alice and Bob’s devices can exchange classical or quantum information after finishing playing a game and before starting the next one. Such boxes should be considered when entanglement is to be distributed “on the fly”, e.g., in protocols where Alice is expected to send half of an entangled state to Bob in each round, or when the devices are located far enough so they cannot communicate during a *single* game but too close to make sure signals from one round cannot arrive to the other device until the end of *all* games. A box that does not allow for communication can be considered, e.g., in cryptographic settings in which any communication between the devices implies that *all* information can leak to the adversary. We remark that parallel boxes and sequential boxes that allow for communications are incomparable to one another, while both are more general than sequential boxes without communication; see Fig. 6.3. This is explained in more detail after formally defining the two types of sequential boxes.

6.2.1 Without Communication Between the Rounds

As in the case of a parallel multi-round box, a sequential multi-round box is described by a conditional probability distribution $P_{AB|XY}$ defined over the inputs and outputs of n rounds of the game $G - \mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$. The special thing about a sequential box is that the marginals describing the individual rounds of the game are well-defined and non-signalling between Alice and Bob. That is, they are boxes by themselves.

In this section we consider a model of sequential boxes in which Alice’s and Bob’s components are not allowed to communicate between the rounds of the game. For short, we call such boxes *non-communicating sequential boxes*. Formally, to define a non-communicating sequential box we consider the marginals of $P_{AB|XY}$ describing a round $i \in [n]$. The relevant marginals are

$$\mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}} \quad (6.3)$$

where $H^{i, \text{Alice}} = X_{1, \dots, i-1} \mathbf{A}_{1, \dots, i-1}$ and $H^{i, \text{Bob}} = Y_{1, \dots, i-1} \mathbf{B}_{1, \dots, i-1}$ denote the “histories” of Alice and Bob’s boxes in round i . These histories basically describe all the information that can be kept by the boxes from the previous rounds (we can think of such boxes as devices which record past events in their memory). The history may include more information¹ than past inputs and outputs; for simplicity we stick to the above choice.

A first requirement on a sequential box is that the marginals (6.3) are well-defined. This can be mathematically described by a set of non-signalling conditions. Explicitly, for every $i \in [n]$, we denote:

1. $\mathcal{P} = [i - 1]$, $\mathbf{a}_{\mathcal{P}} = a_1, \dots, a_{i-1}$, and similarly for $\mathbf{b}_{\mathcal{P}}$, $\mathbf{x}_{\mathcal{P}}$, and $\mathbf{y}_{\mathcal{P}}$.
2. $\mathcal{F} = \{i + 1, \dots, n\}$, $\mathbf{a}_{\mathcal{F}} = a_{i+1}, \dots, a_n$, and similarly for $\mathbf{b}_{\mathcal{F}}$, $\mathbf{x}_{\mathcal{F}}$, and $\mathbf{y}_{\mathcal{F}}$.
3. For any $\mathbf{x}_{\mathcal{P}}$, $\mathbf{y}_{\mathcal{P}}$, x_i , y_i , $\mathbf{x}_{\mathcal{F}}$, $\mathbf{y}_{\mathcal{F}}$, $\mathbf{x}'_{\mathcal{F}}$, and $\mathbf{y}'_{\mathcal{F}}$,
 - (a) $\mathbf{x} = \mathbf{x}_{\mathcal{P}}, x_i, \mathbf{x}_{\mathcal{F}}$
 - (b) $\mathbf{x}' = \mathbf{x}_{\mathcal{P}}, x_i, \mathbf{x}'_{\mathcal{F}}$

and similarly for \mathbf{y} and \mathbf{y}' .

Then, we require that the following non-signalling conditions hold for all $\mathbf{a}_{\mathcal{P}}$, $\mathbf{b}_{\mathcal{P}}$, $\mathbf{x}_{\mathcal{P}}$, $\mathbf{y}_{\mathcal{P}}$, a_i , b_i , x_i , y_i , $\mathbf{x}_{\mathcal{F}}$, $\mathbf{x}'_{\mathcal{F}}$, $\mathbf{y}_{\mathcal{F}}$, and $\mathbf{y}'_{\mathcal{F}}$,

$$\begin{aligned} \sum_{\mathbf{a}_{\mathcal{F}}, \mathbf{b}_{\mathcal{F}}} \mathbb{P}_{A_i B_i A_{\mathcal{F}} B_{\mathcal{F}} | A_{\mathcal{P}} B_{\mathcal{P}} X Y} (a_i, b_i, \mathbf{a}_{\mathcal{F}}, \mathbf{b}_{\mathcal{F}} | \mathbf{a}_{\mathcal{P}}, \mathbf{b}_{\mathcal{P}}, \mathbf{x}, \mathbf{y}) = \\ \sum_{\mathbf{a}_{\mathcal{F}}, \mathbf{b}_{\mathcal{F}}} \mathbb{P}_{A_i B_i A_{\mathcal{F}} B_{\mathcal{F}} | A_{\mathcal{P}} B_{\mathcal{P}} X Y} (a_i, b_i, \mathbf{a}_{\mathcal{F}}, \mathbf{b}_{\mathcal{F}} | \mathbf{a}_{\mathcal{P}}, \mathbf{b}_{\mathcal{P}}, \mathbf{x}', \mathbf{y}') . \end{aligned} \quad (6.4)$$

Now that the marginals $\mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}}$ are well-defined for all $i \in [n]$, we further ask that they are non-signalling between Alice and Bob, when each party holds only its own history. That is, $\mathbb{P}_{A_i | X_i H^{i, \text{Alice}}}$ and $\mathbb{P}_{B_i | Y_i H^{i, \text{Bob}}}$ need to be well-defined as well. Explicitly, for each round $i \in [n]$, for all $a \in \mathcal{A}$, $b \in \mathcal{B}$, $x, x' \in \mathcal{X}$, $y, y' \in \mathcal{Y}$ and histories $h^{i, \text{Alice}}, h^{i, \text{Alice}'} \in \mathcal{X}^{i-1} \times \mathcal{A}^{i-1}$ and $h^{i, \text{Bob}}, h^{i, \text{Bob}'} \in \mathcal{Y}^{i-1} \times \mathcal{B}^{i-1}$,

$$\begin{aligned} \sum_b \mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}} (a, b | x, y, h^{i, \text{Alice}}, h^{i, \text{Bob}}) = \\ \sum_b \mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}} (a, b | x, y', h^{i, \text{Alice}}, h^{i, \text{Bob}'}) \\ \sum_a \mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}} (a, b | x, y, h^{i, \text{Alice}}, h^{i, \text{Bob}}) = \\ \sum_a \mathbb{P}_{A_i B_i | X_i Y_i H^{i, \text{Alice}} H^{i, \text{Bob}}} (a, b | x', y, h^{i, \text{Alice}'}, h^{i, \text{Bob}}) . \end{aligned} \quad (6.5)$$

¹For example, in device-independent quantum key distribution protocols the parties randomly choose in each round whether the round is used for testing the device or for generating key bits. This information can also be included in the history H^i .

The fact that the boxes cannot communicate between the rounds presents itself by having two different histories, one for Alice and one for Bob. The above equations then imply that the actions of Alice's box in round i depend only on Alice's history, i.e., on what happened in the previous rounds on Alice's side (while she is oblivious to Bob's history), and similarly for Bob.²

Note that we only ask the marginals $P_{A_i|X_i H^{i,Alice}}$ and $P_{B_i|Y_i H^{i,Bob}}$ to be well-defined. $P_{A_i|X_i}$, on the other hand, are not necessarily valid boxes.

6.2.1.1 Non-signalling Non-communicating Sequential Boxes

A non-signalling non-communicating sequential multi-round box is simply a box $P_{AB|XY}$ fulfilling the above non-signalling constraints; there are no further requirements.

Definition 6.3 (*Non-signalling non-communicating sequential multi-round box*) Given a Bell game G , a non-signalling non-communicating sequential multi-round box is a conditional probability distribution $P_{AB|XY}$ defined for the inputs and outputs of n rounds of the game $G - \mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$ fulfilling the non-signalling conditions given in Eqs. (6.4) and (6.5).

6.2.1.2 Quantum Non-communicating Sequential Boxes

The simplest way of defining a quantum non-communicating sequential box is to consider the initial state shared by Alice and Bob and the sequence of measurements that they perform.

More specifically, in each round Alice and Bob's boxes can perform a measurement on the post-measurement state of the previous round. We denote the state in the beginning of round $i \in [n]$ (i.e., before performing the measurements of the i 'th round) by $\rho_{Q_A Q_B}^{i, h^{i,Alice}, h^{i,Bob}}$. As clear from the notation, this state depends on the histories $h^{i,Alice}, h^{i,Bob}$. We identify $\rho_{Q_A Q_B}^1 = \rho_{Q_A Q_B}$ as the initial state of the box.

Furthermore, we denote the (Kraus) measurements performed in each round by $\{K_a^x\}$ and $\{K_b^y\}$.³ One can think of the measurements $\{K_a^x\}$ as depending on the history $h^{i,Alice}$ and similarly for Bob. Alternatively, we can imagine that the history is already kept in some classical registers within the quantum state $\rho_{Q_A Q_B}^{i, h^{i,Alice}, h^{i,Bob}}$, i.e., ρ_{Q_A} includes also the information $h^{i,Alice}$ and similarly for Bob. The measurements can thus be defined as first reading the history and then applying the relevant measurement depending on the history. This allows us to use the shorter notation in which the operators do not depend on the histories explicitly.

²This should be compared to the next section, where we will have just a single history H^i for Alice and Bob together.

³Note that in contrast to the previous definitions, the measurement operators K are now written as Kraus operators and not POVMs, since we are interested in the post-measurement state. See Sect. 2.3 for more details.

Using the above notation, the relation between the state in round i to that of round $i - 1$ is simply (up to normalisation of the state)

$$\rho_{Q_A Q_B}^{i, h^{i, \text{Alice}}, h^{i, \text{Bob}}} \propto \left(K_{a_{i-1}}^{x_{i-1}} \otimes K_{b_{i-1}}^{y_{i-1}} \right) \rho_{Q_A Q_B}^{i-1, h^{i-1, \text{Alice}}, h^{i-1, \text{Bob}}} \left(\left(K_{a_{i-1}}^{x_{i-1}} \right)^\dagger \otimes \left(K_{b_{i-1}}^{y_{i-1}} \right)^\dagger \right), \quad (6.6)$$

where $h^{i, \text{Alice}}$ and $h^{i, \text{Bob}}$ uniquely determine x_{i-1} , a_{i-1} , $h^{i-1, \text{Alice}}$ and y_{i-1} , b_{i-1} , $h^{i-1, \text{Bob}}$, respectively (i.e., the values on the righthand-side of Eq. (6.6) should be consistent with the histories on the lefthand-side). The conditions stated in Eq. (6.5) follow directly.

Definition 6.4 (*Quantum non-communicating sequential multi-round box*) Given a Bell game G , a quantum non-communicating sequential multi-round box is a conditional probability distribution $P_{AB|XY}$ defined for the inputs and outputs of n rounds of the game G , \mathcal{X}^n , \mathcal{Y}^n , \mathcal{A}^n , \mathcal{B}^n , such that there exist a bipartite state $\rho_{Q_A Q_B}$ and measurements $\{K_a^x\}$ and $\{K_b^y\}$ defining a sequence of bipartite states for $i \in [n]$ as in Eq. (6.6).

As mentioned before, a non-communicating sequential box is also a parallel one. Indeed, it is easy to see that a parallel box can always simulate the behaviour of a non-communicating sequential box.

6.2.2 With Communication Between the Rounds

In the previous section we considered sequential boxes in which Alice’s and Bob’s components are not allowed to communicate between the rounds. This implies that Alice’s and Bob’s components evolve separately in time and each of them has their own “history”: $h^{i, \text{Alice}}$ for Alice and $h^{i, \text{Bob}}$ for Bob. Now, we consider a scenario in which Alice’s and Bob’s components are allowed to communicate between the different games, i.e., after the outputs of round $i - 1$ were supplied by the box and before the i ’th inputs are given.⁴ Considering boxes that are allowed to communicate is, in particular, relevant when considering realistic application of, e.g., device-independent cryptography. There, one would like to allow the experimentalists to distribute entanglement “on the fly” during the protocol. To send a new quantum state in each round the communication channels need to be open and an adversarial box may use this opportunity to communicate.

Mathematically this setting can be formalised by allowing Alice and Bob to keep a common history register that includes the classical information of all past events *on both sides*. More specifically, the marginal describing the i ’th round of the game,

⁴In Protocol 1.1, for example, “between the different games” refers to the time *after* Step 3 of round $i - 1$ and *before* Step 2 of round i , for all $i \in [n]$.

for $i \in [n]$, is given by $P_{A_i B_i | X_i Y_i H^i}$, where H^i denotes the history defined by the previous rounds. H^i includes $X_{1,\dots,i-1} Y_{1,\dots,i-1} A_{1,\dots,i-1} B_{1,\dots,i-1}$ as well as any other information available to Alice's and Bob's component. For simplicity we assume that $H^i = X_{1,\dots,i-1} Y_{1,\dots,i-1} A_{1,\dots,i-1} B_{1,\dots,i-1}$ similarly to what was done before. The only non-trivial communication to consider is one which depends on the history, since any other information could have been included as part of the box to begin with. Therefore, we can assume without loss of generality that the communicated information is simply the entire history.

As before, we first require that $P_{A_i B_i | X_i Y_i H^i}$ are well-defined, i.e., Eq. (6.4) is fulfilled. In addition, $P_{A_i B_i | X_i Y_i H^i}$ needs to be non-signalling between Alice and Bob, when they both hold their common history. That is, $P_{A_i | X_i H^i}$ and $P_{B_i | Y_i H^i}$ are well-defined. Formally: for each round $i \in [n]$, for all $a \in \mathcal{A}$, $b \in \mathcal{B}$, $x, x' \in \mathcal{X}$, $y, y' \in \mathcal{Y}$ and $h^i \in \mathcal{A}^{i-1} \times \mathcal{B}^{i-1} \times \mathcal{X}^{i-1} \times \mathcal{Y}^{i-1}$,

$$\begin{aligned} \sum_b P_{A_i B_i | X_i Y_i H^i}(a, b | x, y, h^i) &= \sum_b P_{A_i B_i | X_i Y_i H^i}(a, b | x, y', h^i) \\ \sum_a P_{A_i B_i | X_i Y_i H^i}(a, b | x, y, h^i) &= \sum_a P_{A_i B_i | X_i Y_i H^i}(a, b | x', y, h^i). \end{aligned} \quad (6.7)$$

In contrast to Eq. (6.5), in the above equations the behaviour of Alice's component in the i 'th round may depend also on past events on Bob's side, as H^i includes also $Y_{1,\dots,i-1} B_{1,\dots,i-1}$, and similarly for Bob's part of the box.

6.2.2.1 Non-signalling Communicating Sequential Boxes

A non-signalling communicating sequential multi-round box is a box $P_{AB|XY}$ fulfilling the above non-signalling constraints.

Definition 6.5 (*Non-signalling communicating sequential multi-round box*) Given a Bell game G , a non-signalling communicating sequential multi-round box is a conditional probability distribution $P_{AB|XY}$ defined for the inputs and outputs of n rounds of the game $G - \mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$ fulfilling the non-signalling conditions given in Eqs. (6.4) and (6.7).

It is perhaps instructive to note that $P_{AB|XY}$ itself is *not* a non-signalling box.; communication (i.e., signalling) between the rounds may be *necessary* in order to implement the box. We give a trivial example in the end of the section.

6.2.2.2 Quantum Communicating Sequential Boxes

When we say that a communicating sequential multi-round box is quantum we mean that in each round the behaviour of the box can be described within the formalism of quantum physics.

Definition 6.6 (*Quantum communicating sequential multi-round box*) Given a Bell game G , a quantum sequential multi-round box is a conditional probability distribution $P_{AB|XY}$ defined for the inputs and outputs of n rounds of the game G , $\mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$, such that for all $i \in [n]$ the marginal $P_{A_i B_i | X_i Y_i H^i}$, for $H^i = X_{1, \dots, i-1} Y_{1, \dots, i-1} A_{1, \dots, i-1} B_{1, \dots, i-1}$, is a quantum box as in Definition 3.3. That is, there exist a bipartite state $\rho_{Q_A Q_B}^{h^i}$ and measurements $\{M_a^{h^i, x}\}$ and $\{M_b^{h^i, y}\}$ such that

$$P_{A_i B_i | X_i Y_i H^i}(ab|xyh^i) = \text{Tr} \left(M_a^{h^i, x} \otimes M_b^{h^i, y} \rho_{Q_A Q_B}^{h^i} \right) \quad \forall a, b, x, y, h^i. \quad (6.8)$$

The box in Eq. (6.8) is written as $P_{A_i B_i | X_i Y_i H^i}$ so it is mathematically clear which marginals of $P_{AB|XY}$ are being discussed. On the level of the state and measurements one thinks of $\rho_{Q_A Q_B}^{h^i}$, $\{M_a^{h^i, x}\}$, and $\{M_b^{h^i, y}\}$ as depending on the history h^i , which allows the actions in each round to depend on the past. As in Sect. 6.2.1, we may also consider a state $\rho_{Q_A Q_B}^{h^i}$ that keeps h^i in one of its registers and measurements that first read the history and then apply the relevant operations; in such a case we may think of $\{M_a^x\}$, and $\{M_b^y\}$ independent of the history.

It may seem from Definition 6.6 that only the individual rounds are considered. The sequential nature of the box is concealed in the relations between the different rounds. It becomes apparent when noting that all the marginals describing the individual rounds should be consistent with the same overall box $P_{AB|XY}$. Alternatively, one can consider an equivalent definition of a quantum communicating sequential multi-round box that is perhaps more intuitive (but mathematically more complex): Similarly to the evolution described in Eq. (6.6), we start with some initial quantum state and make sequential measurements. In contrast to Eq. (6.6), however, we allow for an additional general operation, which may depend on the history, to be performed on the post-measurement state of each round. The general operation between the rounds is what models the communication between the two parts of the box.

Before concluding this section, let us mention the relations between the different types of multi-round boxes. The relations are shown in Fig. 6.3. It is obvious to see that communicating sequential boxes are more general than non-communicating sequential boxes. In contrast to non-communicating sequential boxes, parallel boxes cannot simulate a general communicating sequential box. A trivial example is a communicating sequential box that always outputs $b_2 = x_1$. Clearly, since a parallel box must, in particular, fulfill Eq. (6.1), it cannot simulate such a box. On the other hand, communicating sequential boxes cannot simulate a general parallel box. For example, a communicating sequential box cannot simulate a parallel box for which $a_1 = x_2$. Thus, the two types of boxes are incomparable.

References

1. Natarajan A, Vidick T (2017) A quantum linearity test for robustly verifying entanglement. In: Proceedings of the 49th annual ACM SIGACT symposium on theory of computing. ACM, pp. 1003–1015
2. Jain R, Miller CA, Shi Y (2017) Parallel device-independent quantum key distribution. arXiv preprint [arXiv:1703.05426](https://arxiv.org/abs/1703.05426)
3. Arnon-Friedman R, Yuen H (2018) Noise-tolerant testing of high entanglement of formation. Int Coll Autom, Lang, Program
4. Reichardt BW, Unger F, Vazirani U (2013) Classical command of quantum systems. Nature 496(7446):456–460
5. Kessler M, Arnon-Friedman R (2017) Device-independent randomness amplification and privatization. arXiv preprint [arXiv:1705.04148](https://arxiv.org/abs/1705.04148)

Chapter 7

Working Under the IID Assumption



In this thesis, we are interested in analysing the behaviour of multi-round boxes when such boxes are used to play many non-local games, e.g., while running a cryptographic protocol. In the previous chapter we discussed the different models of multi-round boxes (the parallel and sequential ones). As we saw, their behaviour can be quite complex. As a consequence, the analysis of protocols which use such boxes is (a priori) tedious in the good case and infeasible in the worst.

In this chapter we discuss an assumption that can make the analysis of the scenarios of interest much simpler—the so called “independent and identically distributed” (IID) assumption. The assumption states that the boxes behave independently and identically when playing the n games. The IID assumption is commonly made in the literature as it significantly simplifies the behaviour of the considered boxes and allows us to gain better intuition and understanding of the problem at hand. As we explain below, there is no reason to believe that the IID assumption can be enforced in the device-independent setting; we use it just as a first stage before moving on to the general analysis. In Chaps. 8–9 we will see that, in certain scenarios, some techniques can be used to reduce the general analysis to the one made under the IID assumption.

We start by explaining the assumption itself. Following that, we present a mathematical tool, namely the “quantum asymptotic equipartition property”, which is of great use when considering IID random variables and quantum systems. Finally, we discuss the analysis of our showcases under the IID assumption.

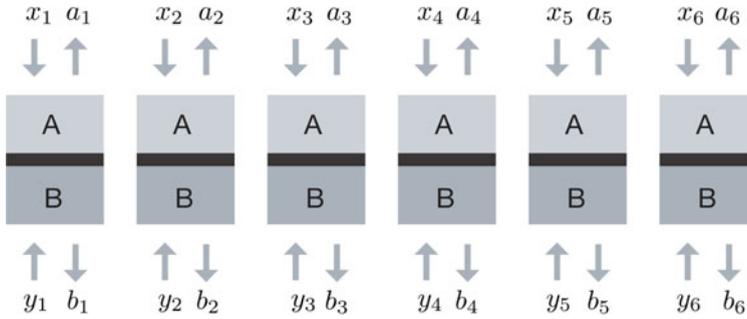


Fig. 7.1 IID box. The box $P_{AB|XY}$ can be described as n identical and independent copies of a single-round box. Each game is played with a different copy of the box. We can see each copy as a well defined subsystem

7.1 The IID Assumption

As in the previous chapter, we consider multi-round boxes. An IID box, as the name suggests, is a multi-round box which behaves identically and independently in each game played with it. Pictorially, we can think of an IID box as n identical and independent copies of a single-round box, as shown in Fig. 7.1. Comparing this to Figs. 6.1 and 6.2, one sees that an IID box has more structure than the other, more general, multi-round boxes. In particular, in the case of an IID box we can talk about a subsystem structure of the box. In the quantum case, for example, if σ denotes the state of a *single-round box* in Fig. 7.1 then the overall state of the IID box is $\sigma^{\otimes n}$. A similar tensor product structure also holds for the measurements describing the box. Mathematically, an IID box is defined as follows.

Definition 7.1 (*IID box*) Given a non-local game G , an IID box is a conditional probability distribution $P_{AB|XY}$ defined for the inputs and outputs of n rounds of the game G , $\mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$, such that

$$P_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = \prod_{i \in [n]} P_{AB|XY}(a_i, b_i|x_i, y_i) \quad (7.1)$$

for some *single-round box* $P_{AB|XY}$. An IID box is said to be quantum or non-signalling if the single round box $P_{AB|XY}$ is quantum (Definition 5.1) or non-signalling (Definition 5.1), respectively.

Note that the single-round box $P_{AB|XY}$ in Eq. (7.1) is the same for every round $i \in [n]$. This means that the behaviour of the box is identical in each round and independent of all other rounds.¹ Hence, the behaviour of an IID box $P_{AB|XY}$ is solely characterised by the single-round box $P_{AB|XY}$ and, thus, the substance of any

¹As always, a box is a *conditional* probability distribution and its definition is therefore independent of the distribution of the inputs, \mathbf{x} and \mathbf{y} , which can be arbitrary (depending on how the box is being

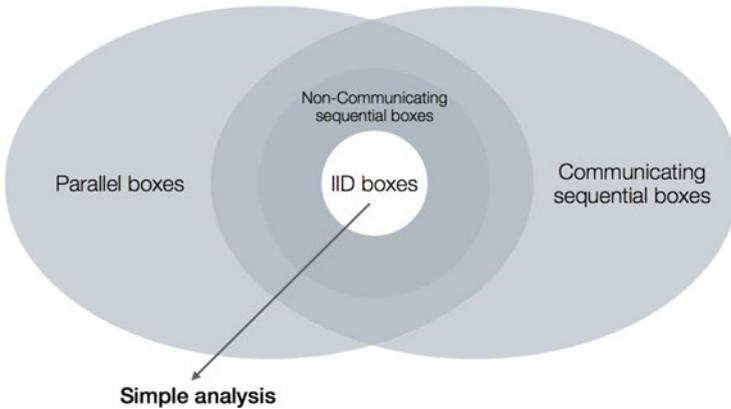


Fig. 7.2 The relation between the sets of multi-round boxes. The intersection of the sets of sequential and parallel boxes includes the set of IID boxes. The analysis of IID boxes is rather simple

analysis done for the IID box is the study of the single-round box. This also implies that the box behaves exactly the same whether we give it all the inputs at once (parallel interaction) or one after the other (sequential interaction). An IID box is therefore both a parallel multi-round box and a sequential multi-round box; see Fig. 7.2. Given all of the above, it indeed makes sense that any analysis done solely for IID boxes can be much simpler than the general analysis in which one needs to deal with parallel or sequential multi-round boxes.

When considering device-independent protocols one usually encounters IID boxes in two different contexts—the so called “completeness” and “soundness” of the protocols (recall, e.g., Sect. 4.2). When proving the completeness of a protocol one shows that an “honest implementation” of the box does not cause the protocol to abort (with high probability). The honest implementation is the implementation of the box that one would like to have if the manufacture of the device is to be trusted and, most commonly, it is described as an IID box. Thus, investigating the behaviour of the protocol when an IID box is being used allows us to see what happens in the honest scenario when “everything goes according to the plan”.

The second context to discuss IID boxes is that of the soundness proof. There, one ought to show that the protocol acts as required for *any* box, i.e., even for adversarial ones.² Clearly, not all boxes are IID boxes and hence analysing the situation only for IID boxes is not sufficient. That is, by *assuming* that all boxes behave in an

used). It is perhaps helpful to note that the idea here is that, while the inputs of the different rounds may be correlated in general (i.e., not IID), the box itself does not “create” further correlations between the rounds (in contrast to parallel and sequential boxes). In any case, in most scenarios the inputs are usually taken to be IID random variables as well.

²Recall that in the device-independent setting we assume that the adversary is the one constructing the box. Device-independent protocols are expected to abort, with high probability, when an adversarial device is detected.

IID manner we weaken the final statement. Still, working under the IID assumption allows us to gain better understanding of the full question at hand.

It is important to remark that, even though quite convenient for the soundness analysis, the IID assumption cannot be justified a priori. Assuming that the box behaves in an IID way goes against the spirit of device-independence by imposing severe restrictions on the implementation of the box. In particular, the assumption implies that the multi-round box does not include any, classical or quantum, internal memory (i.e., its actions when playing one game cannot depend on the other games) and cannot display time-dependent behaviour. We therefore emphasise that working under the IID assumption is only a first step in the process of proving full soundness (as will be shown in the proceeding chapters).

7.2 Asymptotic Equipartition Property

When analysing IID processes a useful mathematical tool is the so called “asymptotic equipartition property” (AEP). The entropic formulation of the AEP used in this thesis basically asserts that when considering IID RV $\mathbf{A} = A_1, A_2, \dots, A_n$, all identical copies of the RV A , the smooth min- and max-entropies rates, $H_{\min}^\varepsilon(\mathbf{A})/n$ and $H_{\max}^\varepsilon(\mathbf{A})/n$, converge to $H(A)$ [1]. Similarly, the quantum version of the AEP asserts that the same is true for IID quantum states $(\sigma_A)^{\otimes n}$ and, even more, it holds also when considering conditional entropies [2–4].

In many information theoretic tasks one needs to bound the smooth min- and max-entropies, as they describe operational quantities. In particular, this will be the case in one of the showcases investigated in the thesis. When considering IID processes, as done in this chapter, the AEP allows us to reduce the analysis of the smooth entropies for *IID boxes* to the analysis of the von Neumann entropy for a *single-round box*.³ This explains why the AEP is a useful tool when working under the IID assumption.

To comprehend the statement of the AEP and its significance we start by presenting and explaining the classical AEP. The quantum variant is then presented as an extension of the classical one.

7.2.1 Classical Asymptotic Equipartition Property

The (classical) AEP can be seen as the “information theoretic version” of the law of large numbers. Given IID RV A_1, A_2, \dots, A_n the law of large numbers states that for large enough number of samples n , the average is close to the expected value in probability. Formally this can be written as

³An example of the analysis of the von Neumann entropy for single-round boxes was presented in Sect. 5.2. The AEP motivates the analysis done in that section when working under the IID assumption.

$$\forall \mu > 0 \quad \lim_{n \rightarrow \infty} \Pr \left[\left| \frac{1}{n} \sum_i A_i - \mathbb{E}[A] \right| > \mu \right] = 0, \quad (7.2)$$

for $A = A_1 = \dots = A_n$ a single copy of the RV. Similarly, the AEP, which is a direct consequence of the law of large numbers,⁴ states that for IID RV

$$\forall \mu > 0 \quad \lim_{n \rightarrow \infty} \Pr \left[\left| -\frac{1}{n} \log(\mathbb{P}_A[\mathbf{a}]) - H(A) \right| > \mu \right] = 0, \quad (7.3)$$

where we denoted $\mathbf{A} = A_1 A_2 \dots A_n$.

Assume that we sample a sequence \mathbf{a} . What can we say about its probability $\mathbb{P}_A[\mathbf{a}]$? We learn from Eq.(7.3) that, for large enough n ,

$$2^{-n(H(A)+\mu)} < \mathbb{P}_A[\mathbf{a}] < 2^{-n(H(A)-\mu)} \quad (7.4)$$

with high probability. This allows us to talk about “typical sequences” and “typical sets”. A typical sequence is a sequence \mathbf{A} for which Eq.(7.4) holds and the typical set includes all typical sequences. Denote by $1 - \varepsilon$ the probability that Eq.(7.4) holds or, in other words, the probability of the typical set. In the limit $n \rightarrow \infty$, $\varepsilon \rightarrow 0$, the typical set has probability approximately 1, all elements of it appear with approximately $2^{-nH(A)}$ probability, and, hence, it includes approximately $2^{nH(A)}$ elements. (For formal proofs see [5, Chap. 3]). Thus, the AEP implies that when analysing probabilistic statements regarding a sequence of IID RV, one can focus on the typical events (and ignore the non-typical ones) without introducing much of an error.

Equation(7.4) can be used to state the AEP in terms of the smooth min- and max-entropies; this form of the AEP is the one used in this thesis.

Theorem 7.2 (AEP⁵ (direct part)) *Let $\mathbf{A} = A_1 A_2 \dots A_n$ be a sequence of IID RV. Then, for any $\varepsilon \in (0, 1)$ and n large enough,*

$$\begin{aligned} \frac{1}{n} H_{\min}^{\varepsilon}(\mathbf{A}) &\geq H(A) - \frac{\delta}{\sqrt{n}} \\ \frac{1}{n} H_{\max}^{\varepsilon}(\mathbf{A}) &\leq H(A) + \frac{\delta}{\sqrt{n}}, \end{aligned}$$

where δ depends on ε and A .⁶

⁴To see this, one can define a new RV, \tilde{A}_i , which, for all $a \in \mathcal{A}$ takes the value $\log(\Pr[a])$ with probability $\Pr[a]$. Applying Eq.(7.2) for the new IID RV $\tilde{A}_1, \dots, \tilde{A}_n$, Eq.(7.3) follows.

⁵Note that this theorem is actually a *non*-asymptotic version of the AEP, as it describes also the convergence rate for finite n (i.e., it includes also the second order term). The limit, stated as Eq. (7.5) below, follows trivially from the presented theorem.

⁶For the time being we are not interested in the explicit form of δ ; this will be discussed when relevant.

To gain some intuition of how the smooth entropies enter to the above theorem we sketch the main arguments here in a somewhat hand-waving way. For the more accurate analysis we refer the interested reader to [1, 2, 4]. Recall that

$$H_{\min}(\mathbf{A}) = \min_{\mathbf{a}} -\log [\mathbb{P}_{\mathbf{A}}[\mathbf{a}]]$$

$$H_{\max}(\mathbf{A}) \leq \max_{\mathbf{a} | \mathbb{P}_{\mathbf{A}}[\mathbf{a}] \neq 0} -\log [\mathbb{P}_{\mathbf{A}}[\mathbf{a}]] .$$

Thus, when considering only typical events, it follows from Eq. (7.4) that

$$\frac{1}{n} H_{\min}(\mathbf{A}) > H(A) - \mu$$

$$\frac{1}{n} H_{\max}(\mathbf{A}) < H(A) - \mu .$$

To account for non-typical events we need to incorporate their probability ε . We do so by switching to the smooth versions of the entropies while using ε as the smoothing parameter.⁷

Theorem 7.2, in combination with a converse bound,⁸ implies that when n goes to infinity both smooth entropies converge to the Shannon entropy:

$$\lim_{n \rightarrow \infty} \frac{1}{n} H_{\min}^{\varepsilon}(\mathbf{A}) = \lim_{n \rightarrow \infty} \frac{1}{n} H_{\max}^{\varepsilon}(\mathbf{A}) = H(A) . \quad (7.5)$$

This explains why the Shannon entropy is so important in information theory—the smooth entropies, which describe operational tasks (recall, for example, Sect. 4.2), converge to the Shannon entropy when considering a large number of independent repetitions of the relevant task. A commonly used example is that of “data compression”. There, one would like to encode an n bit string using less bits. If we allow for some small error when decoding the data, roughly $H_{\max}^{\varepsilon}(\mathbf{A})$ bits are needed [7]. For a large enough IID sequence A_1, A_2, \dots, A_n , however, $nH(A)$ bits suffices [8].

A final important comment about the entropic formulation of the AEP is with regards to the so called “chain rules”. The Shannon entropy respects the chain rule $H(\mathbf{A}) = \sum_i H(A_i | \mathbf{A}_{<i})$, where $\mathbf{A}_{<i}$ denotes the sequence of all RV A_j with $j < i$. In the case of IID RV this is reduced to $H(\mathbf{A}) = nH(A)$. That is, *the total amount of entropy of \mathbf{A} is n times the entropy of a single copy of A* . Thus, in order to calculate $H(\mathbf{A})$ we only need to know $H(A)$. In contrast to the Shannon entropy, the smooth min- and max-entropies do not fulfil a similar chain rule. Theorem 7.2 tells us that, to first order in n , $H_{\min}^{\varepsilon}(\mathbf{A}) = H_{\max}^{\varepsilon}(\mathbf{A}) = nH(A)$. Therefore, for sufficiently large n , *the total amount of the smooth min- and max- entropies of \mathbf{A} are n times*

⁷The above only (roughly) explains why the *smooth* entropies are considered, without addressing the second order term of the AEP. The second order term does not come from the law of large numbers but its refinement—the central limit theorem.

⁸The converse bound roughly follows from the monotonicity of the so called α -entropies. For details see [6, Sect. 6.4].

the Shannon entropy of a single copy of A and, here as well, we only need to know $H(A)$ to calculate $H_{\min}^\varepsilon(A)$ and $H_{\max}^\varepsilon(A)$.

7.2.2 Quantum Asymptotic Equipartition Property

As the name suggests, the quantum AEP is an extended version of the AEP that applies to IID quantum states $\rho = (\sigma_{AB})^{\otimes n}$ (the classical variant is then a special case of the quantum one). The following theorem, developed in [4, Result 5] (see also [3, Theorem 9]), acts as the generalisation of Theorem 7.2 above; it extends the theorem to quantum states and, at the same time, incorporates conditioning on quantum systems.⁹

Theorem 7.3 (Quantum AEP (direct part) [4]) *Let $\rho = (\sigma_{AB})^{\otimes n}$ be an IID quantum state. Then, for any $\varepsilon \in (0, 1)$ and n large enough,*

$$\frac{1}{n} H_{\min}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho \geq H(A|B)_\sigma - \frac{\delta(\varepsilon, \nu)}{\sqrt{n}} \quad (7.6)$$

$$\frac{1}{n} H_{\max}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho \leq H(A|B)_\sigma + \frac{\delta(\varepsilon, \nu)}{\sqrt{n}}, \quad (7.7)$$

where $\delta(\varepsilon, \nu) = 4 \log \nu \sqrt{\log(2/\varepsilon^2)}$ for $\nu = 2\sqrt{2^{H_{\max}(A|B)}} + 1$.

In combination with a converse bound [6, Corollary 6.3], we get the asymptotic equality of the conditional entropies:

$$\lim_{n \rightarrow \infty} \frac{1}{n} H_{\min}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho = \lim_{n \rightarrow \infty} \frac{1}{n} H_{\max}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho = H(A|B)_\sigma.$$

For the proofs of the quantum AEP the reader is directed to [6, Sect. 6.4].

The quantum AEP reveals the same important facts as its classical counterpart when considering IID quantum states—it justifies the use of the von Neumann entropy in quantum information processing and tell us that, for sufficiently large n , the total amount of the conditional smooth entropies, $H_{\min}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho$ and $H_{\max}^\varepsilon(\mathbf{A}|\mathbf{B})_\rho$, are n times the von Neumann entropy $H(A|B)_\sigma$ of a single copy of σ . That is, instead of calculating the entropies of the full state ρ one only needs to analyse the von Neumann entropy for a single copy of σ .

When considering applications in which the analysis should be done for a finite number of repetitions n , it is not sufficient to know that the smooth entropies converge

⁹The classical AEP, given as Theorem 7.2, can be easily written also in terms of conditional entropies if the conditioning is done on *classical* systems (then one can directly define the probability distribution of A as the conditional one). This is not the case when conditioning on quantum systems. That is to say that the statement of the theorem which includes conditional entropies does not follow directly from a “non-conditional” variant.

to the von Neumann entropy; we also need to know how fast they converge. The second order terms appearing in Eqs. (7.6) and (7.7), i.e., the terms that scale with $1/\sqrt{n}$, account for the “finite-size effects”. While the $1/\sqrt{n}$ dependency is optimal, the constant δ is not tight.¹⁰

7.3 Using the IID Assumption

In this section we discuss the analysis of our showcases when working under the IID assumption. The analysis of the parallel repetition question, presented in Sect. 7.3.1, is somewhat trivial. Our showcase of device-independent cryptography, considered in Sect. 7.3.2, demonstrates the use of the quantum AEP in device-independent information processing tasks.

7.3.1 Showcase: Non-signalling Parallel Repetition

In our terminology, parallel repetition results aim to upper-bound the probability that a parallel multi-round box can simultaneously win all the n games played with it; recall Sect. 4.1. The discussion below holds for classical, quantum, and non-signalling strategies. The word “optimal” then refers to the considered type of players.

One simple strategy for the parallel repeated game is the IID strategy. This strategy takes the form of an IID box, which plays each of the n games independently and identically, as in Definition 7.1. That is, the box does not take advantage of the fact that it gets all the inputs at the same time. For an *optimal IID strategy*, i.e., the strategy which achieves the maximal probability of winning all games out of all IID strategies, the single-round box $P_{AB|XY}$ appearing in Eq. (7.1) is the *optimal single-game strategy*, that is, the one achieving winning probability of $1 - \alpha$.

It is easy to see that the probability that an IID box wins all the n games simultaneously decreases exponentially fast with n . Specifically, consider an IID box (or strategy) $P_{AB|XY}$ with

$$P_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = \prod_{i \in [n]} P_{AB|XY}(a_i, b_i|x_i, y_i)$$

for some single-round box $P_{AB|XY}$. Let W_i denote the RV describing whether the i 'th game is won ($W_i = 1$) or not ($W_i = 0$) and denote by $1 - \alpha = \Pr[W_i = 1]$ the winning probability of the *single-round box* $P_{AB|XY}(a_i b_i|x_i y_i)$ in a single game (as in Eq. (5.1)). Due to the IID assumption, all the RV W_i are independent and identically distributed. Thus, the probability that all the n games are won is given by

¹⁰To see that the $1/\sqrt{n}$ dependency is optimal follow, e.g., the proof of [2, Theorem 3.3.3]. Second order terms with constants better than δ can be derived from [9].

$$\Pr \left[\sum_i W_i = n \right] = \prod_i \Pr [W_i = 1] = (1 - \alpha)^n ,$$

Clearly, for any $1 - \alpha < 1$, $(1 - \alpha)^n$ decreases exponentially fast in n .

It is easy to show that also a concentration bound holds for any IID box: for any $0 \leq \beta \leq \alpha$, Hoeffding's inequality tells us that¹¹

$$\Pr \left[\sum_i W_i \geq (1 - \alpha + \beta)n \right] \leq \exp(-2n\beta^2) ,$$

which decreases exponentially fast in n as well. The answer to the parallel repetition question, *under the IID assumption*, is therefore almost trivial.

7.3.2 Showcase: Device-Independent Quantum Cryptography

Following the first proof of concept of the security of device-independent quantum key distribution derived in [10], a long line of works [11–20] considered the security of device-independent quantum and non-signalling cryptography under the IID assumption. In this section we explain how the IID assumption is used when analysing device-independent quantum cryptographic protocols. Specifically, we consider here the task of device-independent randomness certification in the presence of a quantum adversary, which acts as the main building block of many device-independent cryptographic protocols, e.g., device-independent quantum key distribution. We focus only on the parts of the security proof in which the IID assumption plays a crucial role and present them in a slightly simplified form. In particular, we consider large enough number of rounds n and neglect finite-size effects for the moment. In Chap. 11 we give full security proofs (which do not rely on the IID assumption) and contrast the relevant parts with the analysis done here.

When dealing with device-independent cryptography we first need to model the box used by the honest parties, Alice and Bob, and the adversary's knowledge about it. Under the IID assumption, the state of Alice and Bob has an IID structure $\rho_{Q_A Q_B} = (\sigma_{Q_A Q_B})^{\otimes n}$ where each copy of $\sigma_{Q_A Q_B}$ is a bipartite state shared between Alice and Bob. Moreover, we assume that the measurements performed in each round of the protocol are all identical and independent of one another, i.e., for all \mathbf{a}, \mathbf{x} , $M_{\mathbf{a}}^{\mathbf{x}} = (M_{\mathbf{a}}^{\mathbf{x}})^{\otimes n}$ and similarly for Bob's measurements. The most general quantum adversary holds a purification of Alice and Bob's state. As all purifications are equivalent up to

¹¹Hoeffding's inequality tells us even more; it says that when using the optimal IID strategy the probability of winning *less* than $1 - \alpha - \beta$ fraction of the games is also decreasing exponentially fast.

local unitaries on Eve's state, we can assume without loss of generality that the overall state of Alice, Bob, and Eve takes the IID form¹²

$$\rho_{Q_A Q_B E} = (\sigma_{Q_A Q_B E})^{\otimes n} . \quad (7.8)$$

We remark that while we assume that ρ has the above IID structure, the state $\sigma_{Q_A Q_B E}$ is unknown.

Equation (7.8), together with the IID form of the quantum measurements describing the device, indeed leads to an IID box $P_{AB|XY}$ as in Eq. (7.1). In particular, this implies that A_1, A_2, \dots, A_n are IID RV. Furthermore, it follows from Eq. (7.8) that, for all $i \in [n]$, the quantum system E_i holds information *only* regarding the output A_i of the same round (that is, A_1 and E_2 , for example, are independent of one another).

Recall from Sect. 4.2.3 that the central task when proving security of quantum cryptographic protocols is to bound the amount of information that Eve may obtain about certain values generated by the protocol, which are supposed to be unknown to her. In the case of randomness certification the main technical step of all soundness proofs is to lower-bound the smooth min-entropy of Alice's outputs $\mathbf{A} = A_1, A_2, \dots, A_n$ (see, e.g., Protocol 1.1). Our goal is therefore to lower-bound $H_{\min}^\varepsilon(\mathbf{A}|\mathbf{E})$, where $\mathbf{E} = E_1, E_2, \dots, E_n$ are Eve's IID quantum systems appearing in Eq. (7.8).

The rough idea behind a security proof under the IID assumption is illustrated in Fig. 7.3 and is rather simple. The first step is the estimation of the winning probability ω of the single-round box defining the IID box, i.e., the unknown state $\sigma_{Q_A Q_B}$. Alice and Bob play the n games with each of their independent quantum boxes and collect the statistics. Denoting by W_i the RV describing whether the i 'th game is won or not, the IID assumption implies that W_1, W_2, \dots, W_n are, as well, IID RV. Thus, it follows from Chernoff's bound that the average $\frac{1}{n} \sum_i W_i$ is close to the expected winning probability $\mathbb{E}[W]$, which is no other than the winning probability ω of a single copy of the state (see Eq. (5.1)). That is,

$$\omega \approx \frac{1}{n} \sum_i W_i .$$

The second step is to lower-bound the conditional smooth min-entropy $H_{\min}^\varepsilon(\mathbf{A}|\mathbf{E})$ as a function of ω . Due to the IID assumption, we can do so using the quantum AEP presented as Theorem 7.3 above. Specifically, for large enough n we have

¹²It is the equivalence of all purifications that allows us to go from an IID assumption regarding Alice and Bob's state $\rho_{Q_A Q_B}$ to an IID assumption regarding the state $\rho_{Q_A Q_B E}$, which also includes Eve. Interestingly, the same thing cannot be done when considering non-signalling boxes and adversaries. It follows from [21] that the extension of a non-signalling IID box to the adversary does not necessarily have an IID structure as well. (See also [20], where the box itself is assumed to have a subsystem structure similar to that of an IID box while the structure of the adversary's system is unrestricted).

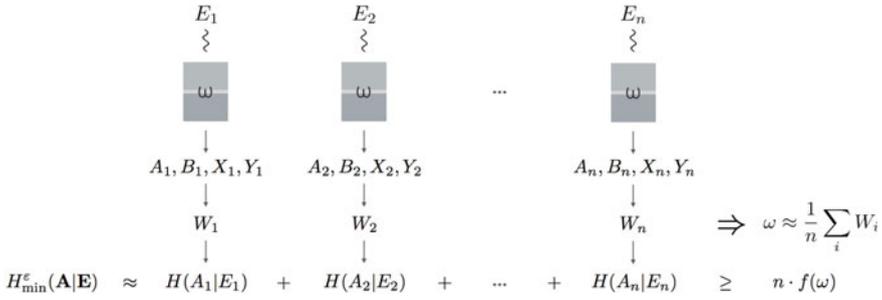


Fig. 7.3 Sketch of a security proof under the IID assumption and for large enough n . The honest parties hold an IID box. Each quantum system E_i , belonging to the adversary, can be entangled only to the i 'th box. The non-local game is being played with each of the independent and identical boxes. The statistics are then collected and used to estimate the winning probability ω of the single-round boxes. According to the quantum AEP, for large enough n , the total amount of smooth min-entropy is the sum of the von Neumann entropy of each round, which can be bounded as a function of the estimated winning probability ω

$$\begin{aligned}
 H_{\min}^\epsilon(\mathbf{A}|\mathbf{E}) &\approx nH(A|E) \\
 &\geq nf(\omega),
 \end{aligned} \tag{7.9}$$

where $f(\omega)$ is some function of ω that lower-bounds the conditional von Neumann entropy $H(A|E)_\sigma$ for any state σ with the estimated winning probability ω .¹³ For the CHSH game, such a function $f(\omega)$ was given in Lemma 5.3 as part of the discussion of single-round boxes.

In certain protocols one would like to use different copies of the boxes in different ways. For example, in device-independent quantum key distribution the protocol includes “test rounds” and “generation rounds”. Alice’s usage of the box in a test round may be different than her usage in a generation round. The winning probability ω is estimated from the statistics collected in the test rounds, as discussed above. Using the IID assumption we can conclude that the other boxes, utilised in the generation rounds, could have also been used to win the game with probability ω , even though Alice and Bob do not test these boxes.

Clearly, the IID assumption plays a crucial role in the above proof sketch; it allows us to talk about a *single-round box*, estimate its winning probability in a meaningful way, and, furthermore, to bound the total amount of smooth min-entropy of the outputs as the number of games played times the von Neumann entropy of the output of a single game. In total, the IID assumption allows us to reduce the analysis of the multi-round box to that of a single-round box—the “physics” enters the analysis

¹³We previously wrote σ as the tripartite state $\sigma_{Q_A Q_B E}$ while here we are referring also to the classical register A . What is meant by this notation is that σ is a state which can lead to winning probability ω when measured with some given measurements $\{M_a^x\}$ and $\{M_b^y\}$. The result of measuring Q_A with $\{M_a^x\}$ defines the RV A .

only in the single-round statement (e.g., Lemma 5.3) while the rest is done using standard mathematical tools such as Chernoff’s bound and the AEP.

7.3.2.1 Quantum Key Distribution Key Rates

The main fundamental difference between device-independent randomness certification and device-independent quantum key distribution is that in the latter Alice and Bob should share identical secret keys in the end of the protocol. To this end, they need to apply an additional classical post-processing step, namely, error correction. The goal of the error correction step is to reconcile the differences between Alice’s and Bob’s keys so they share the same final key with high probability.

In classical error correction protocols utilising one-way communication, Alice sends some classical information about her key to Bob. This information, together with all of Bob’s prior information, helps Bob conclude which key Alice is most likely to hold. If the information sent is not sufficient in order for Bob to derive a conclusion, the parties abort the protocol. Since Alice sends the additional information to Bob over a public (but authenticated) classical channel, this information also leaks to the adversary and hence increases her knowledge about Alice’s key. In other words, the leakage reduces the conditional smooth min-entropy—we now need to consider $H_{\min}^{\varepsilon}(A|EO)$, where O denotes the leaked information, instead of $H_{\min}^{\varepsilon}(A|E)$ appearing in Eq. (7.9).

Notice the resulting tradeoff. To get good key rates, we wish to leak as little information as possible to the adversary (so we do not reduce the min-entropy by too much). On the other hand, we want the error correction step to succeed when Alice and Bob use the *honest* box¹⁴ and, thus, Alice needs to send a sufficient amount of information that will allow Bob to correct the errors. We therefore wish to minimise the amount of leakage needed for successful error correction. As explained in Sect. 4.2, this turns out to be quantified by the conditional smooth zero-entropy $H_0^{\varepsilon}(A|B)$ [22], which is closely related to the conditional smooth max-entropy $H_{\max}^{\varepsilon}(A|B)$.

In many cases the honest box, which also incorporates the considered honest noise model, is chosen to be an IID box. (For example, a common choice is a box describing n independent pairs of maximally entangled states which are being distributed over an IID noisy quantum channel.) Hence, one can use the AEP to get

$$H_{\max}^{\varepsilon}(A|B) \approx nH(A|B) \tag{7.10}$$

and by this upper-bound the amount of leakage due to error correction. All and all, under the IID assumption and for sufficiently large n , the key rate is governed by

¹⁴If the box is malicious or simply noisier than we wished for, we anyhow expect the protocol to abort. Thus, we only ask that the error correction protocol does not abort with high probability when the honest implementation of the devices is used since, otherwise, it will affect the *completeness* of the protocol.

$$r \gtrsim \frac{1}{n} (H_{\min}^\varepsilon(A|E) - H_{\max}^\varepsilon(A|B)) \quad (7.11)$$

$$\gtrsim H(A|E) - H(A|B) . \quad (7.12)$$

Equation (7.12) is usually referred to as the DW-formula since it first appeared in [23, Theorem 2.1]. In [22], the smooth entropies were used to describe the optimal key rates without employing the IID assumption and, by the use of the AEP, the results of [22] imply Eq. (7.12) (as we sketched above). Interestingly, [23, Theorem 2.8] states that, up to some possible classical post-processing, Eq. (7.12) is tight for any protocol utilising error correction with one-way communication.

Two last remarks are in order. Firstly, we would like to emphasise that Eq. (7.10) does not rely on the IID *assumption* that we are making in order to simplify the soundness analysis. Here we are allowed to use the AEP since we *choose* to consider an IID box as our *honest* box. Other choices can also be made (if one, for example, wishes to analyse the protocol under a different honest noise model) and then the AEP might no longer be relevant. In Chap. 11 we will drop the IID assumption used for the soundness analysis but will still choose an IID honest implementation for the completeness analysis.

Secondly, since an adversary limited to preparing IID boxes is weaker than one that can make general multi-round boxes, tight key rates achieved under the IID assumption act as upper-bounds on the achievable key rates in the general setting. Thus, by calculating key rates using Eq. (7.12) we usually already get a feeling of what is the best we can hope for when performing the general analysis. Indeed, [16] used Eq. (7.12) to derive tight key rates for device-independent quantum key distribution under the IID assumption for $n \rightarrow \infty$. These will act as an upper-bound when considering the most powerful quantum adversary in Chap. 11.

7.4 Beyond IID

In this chapter we studied the behaviour of IID boxes and saw how the main ingredient in an analysis of IID boxes is the analysis of a single-round box. In a way, one can say that once we understand the behaviour of a single-round box, we understand the “physics”, or the essence, of the problem at hand. Unfortunately, IID boxes are far from being the most general ones and so we are enforced to go beyond the IID analysis and consider more complicated objects, namely, the different types of multi-round boxes that we encountered in Chap. 6.

As explained in Chap. 1, it is the goal of this thesis to show that the analysis of IID boxes can be almost directly extended, at least in some cases, to the analysis of multi-round boxes via a *reduction to IID*. In the following chapters we present two techniques that can be used to reduce the analysis of parallel and sequential boxes to that of IID boxes; see Fig. 7.4. Specifically, Chap. 8 deals with a technique called “de Finetti reduction” that relates permutation invariant parallel boxes to IID

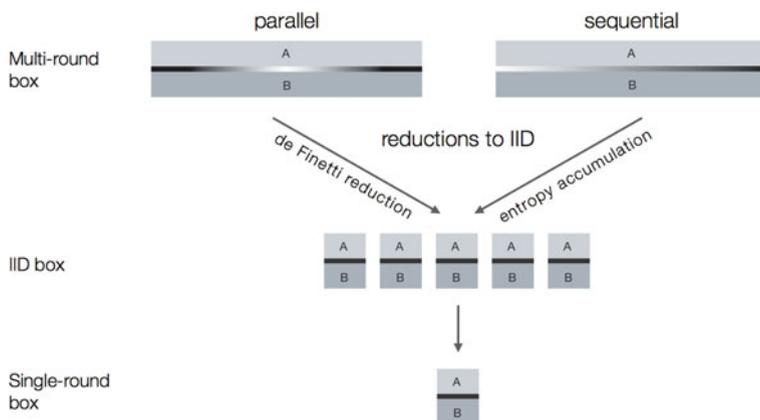


Fig. 7.4 The big picture. The single-round box, at the bottom of the figure, is the simplest object to consider (Chap. 5). The IID box consists of many copies of the single-round box and, thus, can be easily analysed once we understand the behaviour of the single-round box (Chap. 7). The multi-round boxes are the most complex objects (Chap. 6). “Reductions to IID” techniques can be used to simplify the analysis of multi-round boxes by *reducing* it to that of IID boxes. The “de Finetti reduction” technique (Chap. 8) is used when dealing with parallel boxes while the “entropy accumulation theorem” (Chap. 9) is relevant for sequential boxes. In total, with the help of the different reductions, the main thing to study when considering device-independent information processing tasks is the behaviour of single-round boxes

boxes [24]. Chapter 9 presents the so called “entropy accumulation theorem” that relates sequential boxes, fulfilling certain Markov-chain conditions, to IID ones [25].

With the help of those techniques one can show that, in certain scenarios, the analysis of IID boxes is sufficient without loss of generality. This is not to say that all multi-round boxes are IID boxes; clearly this is not the case. Instead, we claim that *even though* there exist multi-round boxes that can not be described as IID boxes, one can sometimes restrict the attention solely to IID boxes and the rest will follow. This will be clarified with the aid of our showcases in Chaps. 10 and 11.

References

1. Holenstein T, Renner R (2011) On the randomness of independent experiments. *IEEE Trans Inf Theory* 57(4):1865–1871
2. Renner R (2008) Security of quantum key distribution. *Int J Quantum Inf* 6(01):1–127
3. Tomamichel M, Colbeck R, Renner R (2009) A fully quantum asymptotic equipartition property. *IEEE Trans Inf Theory* 55(12):5840–5847
4. Tomamichel M (2012) A framework for non-asymptotic quantum information theory. [arXiv:1203.2142](https://arxiv.org/abs/1203.2142)
5. Cover TM, Thomas JA (2012) *Elements of information theory*. Wiley, New York
6. Tomamichel M (2015) *Quantum information processing with finite resources: mathematical foundations*, vol 5. Springer, Berlin

7. Renner R, Wolf S (2004) Smooth rényi entropy and applications. In: International symposium on information theory, 2004. ISIT 2004. Proceedings, p. 233. IEEE
8. Shannon CE (1948) A mathematical theory of communication. *Bell Syst Techn J* 27
9. Tomamichel M, Hayashi M (2013) A hierarchy of information quantities for finite block length analysis of quantum tasks. *IEEE Trans Inf Theory* 59(11):7693–7710
10. Barrett J, Hardy L, Kent A (2005) No signaling and quantum key distribution. *Phys Rev Lett* 95(1):010503
11. Acín A, Gisin N, Masanes L (2006) From Bell’s theorem to secure quantum key distribution. *Phys Rev Lett* 97(12):120405
12. Acín A, Massar S, Pironio S (2006) Efficient quantum key distribution secure against no-signalling eavesdroppers. *New J Phys* 8(8):126
13. Scarani V, Gisin N, Brunner N, Masanes L, Pino S, Acín A (2006) Secrecy extraction from no-signaling correlations. *Phys Rev A* 74(4):042339
14. Acín A, Brunner N, Gisin N, Massar S, Pironio S, Scarani V (2007) Device-independent security of quantum cryptography against collective attacks. *Phys Rev Lett* 98(23):230501
15. Masanes L (2009) Universally composable privacy amplification from causality constraints. *Phys Rev Lett* 102(14):140501
16. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021
17. Hänggi E, Renner R, Wolf S (2010) Efficient device-independent quantum key distribution. *Advances in Cryptology-EUROCRYPT 2010*. Springer, Berlin, pp 216–234
18. Hänggi E, Renner R (2010) Device-independent quantum key distribution with commuting measurements. [arXiv:1009.1833](https://arxiv.org/abs/1009.1833)
19. Masanes L, Pironio S, Acín A (2011) Secure device-independent quantum key distribution with causally independent measurement devices. *Nature Commun* 2:238
20. Masanes L, Renner R, Christandl M, Winter A, Barrett J (2014) Full security of quantum key distribution from no-signaling constraints. *IEEE Trans Inf Theory* 60(8):4973–4986
21. Arnon-Friedman R, Ta-Shma A (2012) Limits of privacy amplification against nonsignaling memory attacks. *Phys Rev A* 86(6):062333
22. Renner R, Wolf S (2005) Simple and tight bounds for information reconciliation and privacy amplification. *Advances in cryptology-ASIACRYPT 2005*. Springer, Berlin, pp 199–216
23. Devetak, I. and Winter, A. (2005). Distillation of secret key and entanglement from quantum states. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 461, pages 207–235. The Royal Society
24. Arnon-Friedman R, Renner R (2015) de Finetti reductions for correlations. *J Math Phys* 56(5):052203
25. Dupuis F, Fawzi O, Renner R (2016) Entropy accumulation. [arXiv:1607.01796](https://arxiv.org/abs/1607.01796)

Chapter 8

Reductions to IID: Parallel Interaction



Multi-round parallel boxes, discussed in Sect. 6.1, can display an almost arbitrary behaviour and hence are complicated to analyse. However, some additional structure of the boxes can be assumed when certain types of symmetries are present in the considered information processing task. In this chapter we focus on the analysis of parallel boxes that are *permutation invariant*. Permutation invariance is an inherent symmetry in many information processing tasks, device-independent tasks among them. Thus, analysing permutation invariant boxes (as defined below) is of special interest.

A well known family of tools used to study permutation invariant systems¹ is the family of “de Finetti-type theorems”. A de Finetti-type theorem is any theorem that relates (in one way or another) *permutation invariant systems* to a more structured system, having the form of a *convex combination of IID systems*, called a de Finetti system (or state). The relation given by the theorem can be used, in certain cases, to argue that instead of analysing permutation invariant systems one can restrict the attention to the simpler to analyse (convex combination of) IID systems. A de Finetti theorem therefore acts as a reduction to IID.

The first de Finetti theorem [1] established that the collection of infinitely exchangeable sequences, i.e., distributions on infinite strings that are invariant under all permutations, exactly coincides with the collection of all convex combinations of IID distributions. Subsequent results gave quantitative bounds of different forms [2–9]. de Finetti-type theorems had proven to be useful in various proofs. The quantum de Finetti theorems, for example, enable a substantially simplified analysis of many quantum information tasks such as quantum cryptography [7, 10], tomography [11], channel capacities [12] and complexity [9].

¹Depending on the context, the term system may refer to a probability distribution, a quantum state, or a box.

The de Finetti theorems listed above cannot be used in the device-independent setting for various reasons.² In this chapter we present a de Finetti-type theorem, which was introduced in [13], that is applicable when working with parallel boxes. Our de Finetti theorem, termed “de Finetti reduction for correlations”, is then used in the analysis of one of our showcases, namely, non-signalling parallel repetition, in Chap. 10.

The chapter is arranged as follows. We start by explaining the notion of permutation invariance in the device-independent context in Sect. 8.1. The de Finetti reduction is presented and proven in Sect. 8.2. Section 8.3 exemplifies how the reductions can be used in two different general ways (while Chap. 10 deals with a specific application). The theorems proven in Sect. 8.3 clarify in what sense we think of a de Finetti reduction as a reduction to IID in the device-independent setting.

In accordance with the rest of the thesis, the chapter focuses only on the case of two parties. All the statements can be extended to any number of parties, as can be seen in [13].

8.1 Permutation Invariance

As mentioned above, we are interested in considering permutation invariant parallel multi-round boxes. Let n be the number of games that can be played with the parallel box of interest $P_{AB|XY}$. A permutation π is a bijective function $\pi : [n] \rightarrow [n]$. We denote $\pi(\mathbf{x}) = x_{\pi^{-1}(1)}, x_{\pi^{-1}(2)}, \dots, x_{\pi^{-1}(n)}$ and similarly for $\pi(\mathbf{y})$, $\pi(\mathbf{a})$, and $\pi(\mathbf{b})$. A permutation invariant box³ is defined as follows.

Definition 8.1 (*Permutation invariant box*) Given a parallel multi-round box $P_{AB|XY}$, defined over \mathcal{X}^n , \mathcal{Y}^n , \mathcal{A}^n , \mathcal{B}^n , and a permutation $\pi : [n] \rightarrow [n]$ we denote by $P_{AB|XY} \circ \pi$ the box defined by

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad (P_{AB|XY} \circ \pi)(\mathbf{a}, \mathbf{b} | \mathbf{x}, \mathbf{y}) = P_{AB|XY}(\pi(\mathbf{a}), \pi(\mathbf{b}) | \pi(\mathbf{x}), \pi(\mathbf{y})) . \quad (8.1)$$

A parallel multi-round box $P_{AB|XY}$ is said to be permutation invariant if and only if

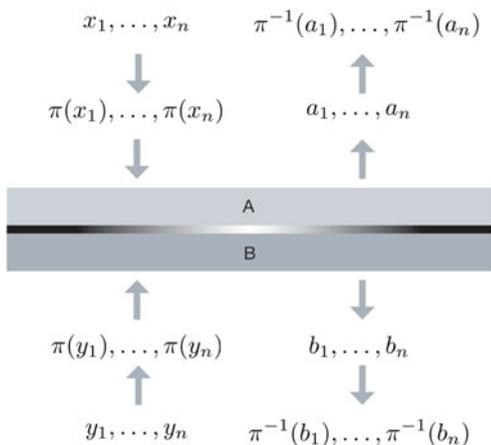
$$\forall \pi \quad P_{AB|XY} = P_{AB|XY} \circ \pi .$$

Figure 8.1 illustrate the action of permuting a parallel box. The action of the permuted box can be understood as follows: First, the box applies the permutation π on the inputs. Second, it uses the initial box $P_{AB|XY}$ to produce the intermediate outputs. Lastly, it applies the inverse permutation π^{-1} on the intermediate outputs

²The mentioned theorems rely on some initial subsystem structure and/or a bound on the dimension of the subsystems. In the device-independent setting one cannot start with such assumptions regarding the considered boxes in general.

³The definition and the derived theorem are independent of the nature of the box, i.e., if it is classical, quantum, non-signalling, or even signalling. This will be addressed in Sect. 8.2.

Fig. 8.1 Permutation of a box $P_{AB|XY}$. The permuted box, $P_{AB|XY} \circ \pi$ acts by first applying the permutation π on the inputs, then producing the outputs using the initial box $P_{AB|XY}$, and lastly applying the inverse permutation on the outputs. The input output distribution of the box is then defined according to Eq. (8.1). A box is said to be permutation invariant if for all π , $P_{AB|XY} = P_{AB|XY} \circ \pi$



and returns these final strings as the ultimate outputs. Note that only the inputs and the outputs of the box are being permuted, all using the same permutation π . In particular, we do not permute the parties, that is, Alice and Bob do not swap their inputs and outputs with one another.

As we are merely permuting the classical inputs and outputs, the box itself need not to have a subsystem structure. That is, we do not require, e.g., $P_{A_1|X_1}$ to be a valid system (i.e., a conditional probability distribution). This is in contrast to, e.g., quantum de Finetti-type theorems such as [5, 7], where the permutation is applied on the quantum states themselves.⁴ This distinction is relevant when wishing to discuss general parallel boxes (recall Sect. 6.1).

In some applications (e.g., the showcase considered in Chap. 10) one can easily show that it is sufficient to consider permutation invariant boxes without loss of generality. If this is not the case, it is also possible to *enforce* permutation invariance. A protocol, for example, can be modified to enforce the symmetry by adding a step in which a random permutation is applied⁵ on the box and by this make it permutation invariant. Precisely: given *any* parallel box $P_{AB|XY}$, let

$$\tilde{P}_{AB|XY} = \frac{1}{n!} \sum_{\pi} P_{AB|XY} \circ \pi$$

be the result of applying a permutation π , chosen uniformly at random out of all permutations, on the original box. It can be easily verified that $\tilde{P}_{AB|XY}$ is indeed a permutation invariant box.

⁴In a quantum de Finetti statement a permutation takes a state $|\phi_1\rangle \otimes \dots \otimes |\phi_n\rangle$ to $|\phi_{\pi^{-1}(1)}\rangle \otimes \dots \otimes |\phi_{\pi^{-1}(n)}\rangle$. That is, the quantum states themselves are being permuted.

⁵Depending on the considered scenario, the application of the permutation may be a purely theoretical step or needs to be done in practice.

8.2 de Finetti Reductions for Correlations

A de Finetti-type theorem is any theorem that relates a permutation invariant system to a much more structured system called a de Finetti system. In our context, we consider permutation invariant and de Finetti boxes. A *de Finetti box* is defined as follows.

Definition 8.2 (*de Finetti box*⁶) A de Finetti box is any box of the form of a convex combination of IID boxes. That is, it is a box $\tau_{AB|XY}$, defined over $\mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$, such that

$$\tau_{AB|XY} = \int \mathcal{O}_{AB|XY}^{\otimes n} d\mathcal{O}_{AB|XY},$$

where $d\mathcal{O}_{AB|XY}$ is some measure on the space of bipartite boxes over $\mathcal{A}, \mathcal{B}, \mathcal{X}$, and \mathcal{Y} and $\mathcal{O}_{AB|XY}^{\otimes n}$ is the IID box defined by $\mathcal{O}_{AB|XY}$, i.e.,

$$\mathcal{O}_{AB|XY}^{\otimes n}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = \prod_{i \in [n]} \mathcal{O}_{AB|XY}(a_i, b_i|x_i, y_i).$$

As seen from the above definition, by choosing different measures $d\mathcal{O}_{AB|XY}$ we define different de Finetti boxes. Depending on the measure, $\tau_{AB|XY}$ may be classical, quantum, non-signalling, or even signalling between the two parties. If the measure $d\mathcal{O}_{AB|XY}$ assigns weight only to, e.g., non-signalling boxes $\mathcal{O}_{AB|XY}$, then the de Finetti box $\tau_{AB|XY}$ is non-signalling as well. The other direction does not necessarily hold—there are convex combinations of signalling boxes that result in over-all non-signalling boxes.

A de Finetti reduction is a de Finetti-type theorem of a specific form: it sets an *inequality relation* between any permutation invariant box to a certain de Finetti box. Specifically, the following theorem is a de Finetti reduction for any permutation invariant conditional probability distribution [13].⁷

Theorem 8.3 (de Finetti reduction for conditional probability distributions) *For any $\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}$, and n there exists a de Finetti box $\tau_{AB|XY}$, defined over $\mathcal{X}^n, \mathcal{Y}^n, \mathcal{A}^n, \mathcal{B}^n$, such that for every permutation invariant box $P_{AB|XY}$*

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad P_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \leq (n+1)^{|\mathcal{X}||\mathcal{Y}|(|\mathcal{A}||\mathcal{B}|-1)} \tau_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}). \quad (8.2)$$

To see why Theorem 8.3 is not trivial and what needs to be done to prove it, let us first consider a “bad choice” of a de Finetti box, $\tau_{AB|XY}^{\text{bad}}$. Imagine that we choose our de Finetti box to be the uniform distribution over $\mathcal{A}^n \times \mathcal{B}^n$ for all \mathbf{x} and \mathbf{y} . With

⁶As previously mentioned, we focus on the case of two parties. The definition extends to any number of parties trivially.

⁷In [13], a more general version of Theorem 8.3 was proven, in which further symmetries of $P_{AB|XY}$ (on top of permutation invariance) can be exploited to construct more structured de Finetti boxes and prove de Finetti reductions with improved parameters. Theorem 8.3 was then derived as a corollary. To keep things (relatively) concise, we present in this thesis a direct proof of Theorem 8.3.

this choice, $\tau_{AB|XY}^{\text{bad}}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = (|\mathcal{A}||\mathcal{B}|)^{-n}$ for all $\mathbf{a}, \mathbf{b}, \mathbf{x}$, and \mathbf{y} . Then, the only inequality relation that holds is

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad P_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \leq (|\mathcal{A}||\mathcal{B}|)^n \tau_{AB|XY}^{\text{bad}}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}),$$

i.e., a relation with a pre-factor exponential in n . By choosing a “good” de Finetti box, we are able to get a pre-factor polynomial in n instead; this is crucial for applications of de Finetti reductions. In Sect. 8.3 we show how Theorem 8.3 can be utilised as a reduction to IID in certain scenarios.⁸

The proof of the theorem proceeds in two steps. In the first, an explicit de Finetti box $\tau_{AB|XY}$ is constructed and a lower-bound on its entries is calculated. In the second step the permutation invariance of $P_{AB|XY}$ is used to upper-bound its entries. The theorem follows by combining the two bounds.

In the proofs below we use the following notation.

1. $|\mathcal{X}||\mathcal{Y}| = l$ and we identify each pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with a label $j \in [l]$ by writing $(x, y) = j$.
2. $|\mathcal{A}||\mathcal{B}| = m$ and we identify each pair $(a, b) \in \mathcal{A} \times \mathcal{B}$ with a label $k \in [m]$ by writing $(a, b) = k$.
3. For all $j \in [l]$ and $k \in [m]$, $p_k^j \in [0, 1]$ such that $\sum_k p_k^j = 1$.
4. For all $j \in [l]$ and $k \in [m]$, $c_k^j = 1 - \sum_{t < k} p_t^j$.
5. For all \mathbf{x}, \mathbf{y} , and $j \in [l]$, $n^j = |\{i : (x_i, y_i) = j\}|$, i.e., n^j denotes the number of indices of (\mathbf{x}, \mathbf{y}) in which the type of inputs is $(x, y) = j$.
6. For all $\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}$, $j \in [l]$, and $k \in [m]$, $n_k^j = |\{i : (x_i, y_i) = j \wedge (a_i, b_i) = k\}|$, i.e., n_k^j denotes the number of indices of $(\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b})$ in which the type of inputs is $(x, y) = j$ and the type of outputs is $(a, b) = k$.

Note that by definition:

1. For all $j \in [l]$ and $k \in [m - 1]$, $p_k^j \in [0, c_k^j]$ and $p_m^j = c_m^j$.
2. For all $j \in [l]$, $n_m^j = n^j - \sum_{k=1}^{m-1} n_k^j$.

According to Definition 8.2, a de Finetti box is defined via the choice of measure $dO_{AB|XY}$. We think of a bipartite box $O_{AB|XY}$ as a set of probabilities p_k^j , with the identification $O_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = p_k^j$ for $(x, y) = j$ and $(a, b) = k$. Thus, we can define a measure over $O_{AB|XY}$ by a measure over the probabilities p_k^j . Our chosen measure is

$$dO_{AB|XY} = \prod_{j=1}^l \frac{dp_1^j}{c_1^j} \cdots \frac{dp_{m-1}^j}{c_{m-1}^j},$$

where dp_k^j is the uniform measure over $[0, c_k^j]$ for c_k^j defined above. The resulting de Finetti box is given by

⁸A curious reader may already take a glimpse of Theorems 8.11 and 8.15.

$$\begin{aligned}
\tau_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) &= \int \mathbf{O}_{AB|XY}^{\otimes n} d\mathbf{O}_{AB|XY} \\
&= \prod_{j=1}^l \left[\int_0^{c_1^j} \frac{dp_1^j}{c_1^j} (p_1^j)^{n_1^j} \right] \cdots \left[\int_0^{c_{m-1}^j} \frac{dp_{m-1}^j}{c_{m-1}^j} (p_{m-1}^j)^{n_{m-1}^j} \right] \\
&\quad \cdot (p_m^j)^{n^j - \sum_{k=1}^{m-1} n_k^j}.
\end{aligned} \tag{8.3}$$

The measure $d\mathbf{O}_{AB|XY}$ assigns some weight to *all* conditional probability distributions $\mathbf{O}_{AB|XY}$. As a result, the de Finetti box in Eq.(8.3) is *signalling*. This is discussed in Sect. 8.4 below.

The following lower-bound on the entries of the above de Finetti box is proven in Appendix A.1:

Lemma 8.4 *For all $\mathbf{a}, \mathbf{b}, \mathbf{x}$, and \mathbf{y} ,*

$$\tau_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \geq \prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1} \frac{1}{(n^j + 1)^{m-1}},$$

where $\tau_{AB|XY}$ is as in Eq.(8.3).

Next, we exploit the permutation invariance of $\mathbf{P}_{AB|XY}$ to prove the following upper-bound on it:

Lemma 8.5 *For every permutation invariant box $\mathbf{P}_{AB|XY}$, as in Definition 8.1, and for all $\mathbf{a}, \mathbf{b}, \mathbf{x}$, and \mathbf{y} ,*

$$\mathbf{P}_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \leq \prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1}.$$

Proof To prove the lemma we bound the value of a specific entry $\mathbf{P}_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y})$ by counting how many entries $\mathbf{P}_{AB|XY}(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}|\mathbf{x}, \mathbf{y})$ must have the same value as $\mathbf{P}_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y})$ due to permutation invariance. The normalisation of $\mathbf{P}_{AB|XY}$ then implies a bound on the value of the entries.

Denote

$$\mathcal{N}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y}) = \left| \left\{ (\tilde{\mathbf{a}}, \tilde{\mathbf{b}}) \in \mathcal{A} \times \mathcal{B} : \mathbf{P}_{AB|XY}(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}|\mathbf{x}, \mathbf{y}) = \mathbf{P}_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \right\} \right|.$$

The permutation invariance of $\mathbf{P}_{AB|XY}$ implies that $\mathcal{N}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})$ is lower-bounded by the number permutations π for which $\pi(\mathbf{x}) = \mathbf{x}$, $\pi(\mathbf{y}) = \mathbf{y}$. To keep $\pi(\mathbf{x}) = \mathbf{x}$ and $\pi(\mathbf{y}) = \mathbf{y}$, the relevant permutations π are only allowed to permute indices with the same input type (x, y) . The number of such permutations is exactly $\prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}$. Thus,

$$\mathcal{N}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y}) \geq \prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}$$

and

$$P_{AB|XY}(\mathbf{ab}|\mathbf{xy}) \leq \frac{1}{\mathcal{N}(\mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y})} \leq \prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1}. \quad \square$$

Proof of Theorem 8.3. Using Lemmas 8.4 and 8.5 one can easily prove Theorem 8.3. For all $\mathbf{a}, \mathbf{b}, \mathbf{x}$, and \mathbf{y} ,

$$\begin{aligned} \frac{P_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y})}{\tau_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y})} &\leq \frac{\prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1}}{\prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1} (n^j + 1)^{-(m-1)}} \\ &\leq \prod_{j=1}^l (n^j + 1)^{m-1} \\ &\leq (n + 1)^{l(m-1)}. \quad \square \end{aligned}$$

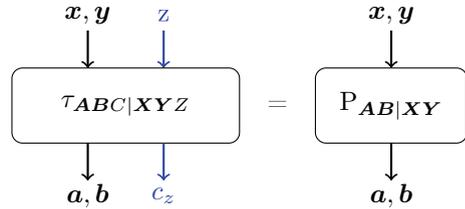
To end this section let us give a last remark regarding Theorem 8.3. Notice the order of the quantifiers; *there exists one* de Finetti box for which Eq. (8.2) holds for *all* permutation invariant box. For the purpose of applications, one could also imagine a different statement in which for each permutation invariant box a de Finetti box is constructed (i.e., different permutation invariant boxes may be related to different de Finetti boxes). Such a statement has the potential of improving the obtained parameters and simplifying the use of the reduction in applications (see also [13] for examples of such statements).

8.3 Ways of Using the Reductions

The main motivation for considering de Finetti reductions as in Theorem 8.3 is to allow us to simplify the analysis of device-independent information processing tasks. However, it is a priori not clear how one can bring an inequality as that in Eq. (8.2) into work. The aim of this section is to exemplify the usage of the inequality in a mathematical way by considering two types of abstract applications. Chapter 10 discusses a more concrete application of the reduction to prove a non-signalling parallel repetition theorem.

To derive the results presented in this section we use an alternative, but equivalent, version of the de Finetti reduction; this is the topic of Sect. 8.3.1 below. Sections 8.3.2

Fig. 8.2 post-selecting a box $P_{AB|XY}$ from an extension of $\tau_{AB|XY}$. Conditioned on the output c_z , the resulting box is $P_{AB|XY}$. After [13]



and 8.3.3 present two ways of using the de Finetti reduction via the alternative formulation.

8.3.1 Post-selecting Permutation Invariant Boxes

Lemma 8.6 *There exists a de Finetti box $\tau_{AB|XY}$ and a non-signalling extension⁹ of it (Definition 3.2) to a larger box $\tau_{ABC|XYZ}$ such that for every permutation invariant box $P_{AB|XY}$ there exists an input z and an output of this input c_z for which*

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad \tau_{ABC|XYZ}(\mathbf{a}, \mathbf{b}, c_z | \mathbf{x}, \mathbf{y}, z) = \frac{1}{(n+1)^{l(m-1)}} P_{AB|XY}(\mathbf{a}, \mathbf{b} | \mathbf{x}, \mathbf{y}),$$

where $l = |\mathcal{X}||\mathcal{Y}|$ and $m = |\mathcal{A}||\mathcal{B}|$.

This lemma states that there exists a de Finetti box $\tau_{AB|XY}$ and a non-signalling extension of it $\tau_{ABC|XYZ}$ such that any permutation invariant box $P_{AB|XY}$ can be *post-selected* from it with probability greater or equal to $\frac{1}{(n+1)^{l(m-1)}}$. When we say that $P_{AB|XY}$ can be post-selected we mean that there exists an input z to $\tau_{ABC|XYZ}$ and an output c_z of this input such that with probability $\tau_{C|Z}(c_z|z) \geq \frac{1}{(n+1)^{l(m-1)}}$ the resulting box (the “post-measurement box”, using terminology borrowed from quantum physics) is $P_{AB|XY}$ (see Fig. 8.2). Note that we consider a single extension $\tau_{ABC|XYZ}$ of the box $\tau_{AB|XY}$, and by choosing different inputs z we can post-select different boxes $P_{AB|XY}$.

It is easy to see how to derive Lemma 8.6 from Theorem 8.3 by using the formalism introduced in [14, 15] of partitions of a conditional probability distribution. We repeat here the relevant statements.

Definition 8.7 A partition of a box $Q_{AB|XY}$ is a family of pairs $\left\{ \left(q_c, Q_{AB|XY}^c \right) \right\}_c$ where $q_c \geq 0$, $\sum_c q_c = 1$, and the boxes $Q_{AB|XY}^c$ are such that

$$Q_{AB|XY} = \sum_c q_c \cdot Q_{AB|XY}^c.$$

⁹Note that $\tau_{AB|XY}$ may be signalling, as in our previous statements. The fact that we are considering non-signalling extensions only means that the marginals $\tau_{AB|XY}$ and $\tau_{C|Z}$ of $\tau_{ABC|XYZ}$ are well defined.

Lemma 8.8 (Lemma 9 in [14]) *Given a box $Q_{AB|XY}$, there exists a partition with element $(q_c, Q_{AB|XY}^c)$ if and only if*

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad q_c \cdot Q_{AB|XY}^c(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) \leq Q_{AB|XY}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) .$$

Lemma 8.9 (Lemma 3.2 in [15]) *Given a box $Q_{AB|XY}$, let Z be the set of all partitions $\left\{ (q_{c_z}, Q_{AB|XY}^{c_z}) \right\}_{c_z}$ of $Q_{AB|XY}$. Then, there exist a non-signalling extension $Q_{ABC|XYZ}$ of $Q_{AB|XY}$, an input z , and an output c_z such that*

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad Q_{ABC|XYZ}(\mathbf{a}, \mathbf{b}, c_z|\mathbf{x}, \mathbf{y}, z) = q_{c_z} \cdot Q_{AB|XY}^{c_z}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) .$$

Using the lemmas above and Theorem 8.3 we can now prove Lemma 8.6.

Proof of Lemma 8.6. The above lemmas together with Theorem 8.3 imply that for any permutation invariant box $P_{AB|XY}$, $\left(\frac{1}{(n+1)^{l(m-1)}}, P_{AB|XY} \right)$ is an element of a partition of $\tau_{AB|XY}$. Moreover, there exists a box $\tau_{ABC|XYZ}$ and an input z such that with probability $\frac{1}{(n+1)^{l(m-1)}}$ the resulting box is $P_{AB|XY}$:

$$\forall \mathbf{a}, \mathbf{b}, \mathbf{x}, \mathbf{y} \quad \tau_{ABC|XYZ}(\mathbf{a}, \mathbf{b}, c_z|\mathbf{x}, \mathbf{y}, z) = \frac{1}{(n+1)^{l(m-1)}} P_{AB|XY} . \quad \square$$

Lemma 8.6 is used in the following sections to illustrate two ways in which de Finetti reductions can be used in applications.

8.3.2 Failure Probability of a Test

We start by considering the following abstract application. Let \mathcal{T} be a test which interacts with a box $P_{AB|XY}$ and outputs “success” or “fail” with some probabilities. One can think about this test, which can be chosen according to the application being considered, as a way to quantify the success probability of a protocol when the box $P_{AB|XY}$ is given as input. For example, if one considers an estimation, or a tomography, protocol a test can be chosen to output “success” when the estimated box is close to the actual box [7]. Another type of test will be considered explicitly in Sect. 10.2.

A test \mathcal{T} interacts with $P_{AB|XY}$ by supplying it with inputs \mathbf{x}, \mathbf{y} , according to some probability distribution $\Pr_{\mathcal{T}}(\mathbf{x}, \mathbf{y})$ over $\mathcal{X}^n \times \mathcal{Y}^n$, and collecting its outputs \mathbf{a}, \mathbf{b} . This is illustrated in Fig. 8.3. The test then decides whether to output 0 or 1 depending on $\mathbf{x}, \mathbf{y}, \mathbf{a}$, and \mathbf{b} . Given a test \mathcal{T} , we denote by $\Pr_{\text{fail}}(P_{AB|XY})$ the failure probability of the test, i.e., the probability that \mathcal{T} outputs 0 after interacting with $P_{AB|XY}$:

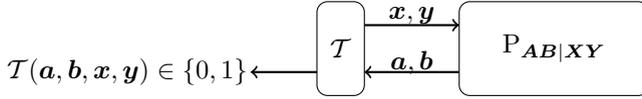


Fig. 8.3 The test \mathcal{T} interacts with $P_{AB|XY}$ by supplying it with inputs x, y and collecting its outputs a, b . The test then decides whether to output 0 or 1 depending on $x, y, a,$ and b . If the output is 0 then we say that the test failed. After [13]

$$\Pr_{\text{fail}}(P_{AB|XY}) = \sum_{x,y} \Pr_{\mathcal{T}}(x, y) \sum_{a,b:\mathcal{T}(a,b,x,y)=0} P_{AB|XY}(a, b|x, y) .$$

The event of failing the test can therefore be defined as an event over $\mathcal{X}^n \times \mathcal{Y}^n \times \mathcal{A}^n \times \mathcal{B}^n$.

We consider permutation invariant tests, defined as follows.

Definition 8.10 A test \mathcal{T} is permutation invariant if and only if for all boxes $P_{AB|XY}$ and all permutations π we have

$$\Pr_{\text{fail}}(P_{AB|XY}) = \Pr_{\text{fail}}(P_{AB|XY} \circ \pi) .$$

Using the de Finetti reduction in Theorem 8.3 we can prove upper bounds of the following type:

Theorem 8.11 *Let \mathcal{T} be a permutation invariant test. Then for every box $P_{AB|XY}$*

$$\Pr_{\text{fail}}(P_{AB|XY}) \leq (n + 1)^{l(m-1)} \Pr_{\text{fail}}(\tau_{AB|XY}) ,$$

where $\tau_{AB|XY}$ is the de Finetti box given in Eq. (8.3).

The importance of de Finetti reductions is already obvious from Theorem 8.11— if one wishes to prove an upper bound on the failure probability of the test \mathcal{T} , then instead of proving it for all boxes $P_{AB|XY}$, it is sufficient to prove it for the de Finetti box $\tau_{AB|XY}$ and “pay” for it with the additional polynomial pre-factor of $(n + 1)^{l(m-1)}$. Since the de Finetti box can be written as a convex combination of IID boxes, this can highly simplify the calculations of the bound. In this sense the de Finetti reduction acts as a *reduction to IID*.

In many cases one finds that the bound on $\Pr_{\text{fail}}(\tau_{AB|XY})$ is exponentially small in n . For an estimation protocol, the failure probability of the test, when interacting with an IID box, can be shown to be exponentially small in the number of boxes n used for the estimation, using Chernoff bounds. This is also the case when dealing with security proofs—the failure probability of a protocol, when a de Finetti box is given as input, is usually exponentially small in the number of boxes n used in the protocol. If this is indeed the case then the polynomial pre-factor of $(n + 1)^{l(m-1)}$ becomes irrelevant in the asymptotic limit of large n . In other words, an exponentially small bound on $\Pr_{\text{fail}}(\tau_{AB|XY})$ implies an exponentially small bound on $\Pr_{\text{fail}}(P_{AB|XY})$.

Let us prove Theorem 8.11 using the de Finetti reduction given as Theorem 8.3.

Proof of Theorem 8.11. We follow here a similar proof given in [16] for the quantum post-selection theorem [7]. First, since the test \mathcal{T} is permutation invariant it is sufficient to consider only permutation invariant boxes. To see this recall that for any box $P_{AB|XY}$ and permutation π we have $\Pr_{\text{fail}}(P_{AB|XY}) = \Pr_{\text{fail}}(P_{AB|XY} \circ \pi)$ according to Definition 8.10. Therefore we also have by linearity¹⁰

$$\Pr_{\text{fail}}(P_{AB|XY}) = \frac{1}{n!} \sum_{\pi} \Pr_{\text{fail}}(P_{AB|XY} \circ \pi) = \Pr_{\text{fail}}\left(\frac{1}{n!} \sum_{\pi} P_{AB|XY} \circ \pi\right).$$

The box $\frac{1}{n!} \sum_{\pi} P_{AB|XY} \circ \pi$ is permutation invariant and therefore we can consider only permutation invariant boxes without loss of generality.

Next we define the following probabilities. Let $\Pr_{\text{fail} \wedge c_z}(\tau_{ABC|XYZ})$ be the probability that the second part of the box, $\tau_{C|Z}$, is used with the input z and the output is c_z and that the first part of the box, $\tau_{AB|XY}$, fails the test \mathcal{T} at the same time. That is,

$$\Pr_{\text{fail} \wedge c_z}(\tau_{ABC|XYZ}) = \Pr_{\text{fail}}(\tau_{AB|XY}) \cdot \tau_{C|Z}(c_z|z).$$

In a similar way we define $\Pr_{\text{fail}|c_z}(\tau_{ABC|XYZ})$ to be the probability that $\tau_{AB|XY}$ fails the test \mathcal{T} given that c_z is the output of $\tau_{C|Z}$ when used with the input z . We have

$$\Pr_{\text{fail}|c_z}(\tau_{ABC|XYZ}) = \frac{\Pr_{\text{fail} \wedge c_z}(\tau_{ABC|XYZ})}{\tau_{C|Z}(c_z|z)} \leq \frac{\Pr_{\text{fail}}(\tau_{AB|XY})}{\tau_{C|Z}(c_z|z)}$$

since $\Pr_{\text{fail} \wedge c_z}(\tau_{ABC|XYZ}) \leq \Pr_{\text{fail}}(\tau_{AB|XY})$ always holds.

Lemma 8.6 implies that $\tau_{C|Z}(c_z|z) \geq \frac{1}{(n+1)^{l(m-1)}}$ and that $\Pr_{\text{fail}|c_z}(\tau_{ABC|XYZ}) = \Pr_{\text{fail}}(P_{AB|XY})$ (given that the output was c_z , the resulting box is $P_{AB|XY}$). All together we get $\Pr_{\text{fail}}(P_{AB|XY}) \leq (n+1)^{l(m-1)} \Pr_{\text{fail}}(\tau_{AB|XY})$ as required. \square

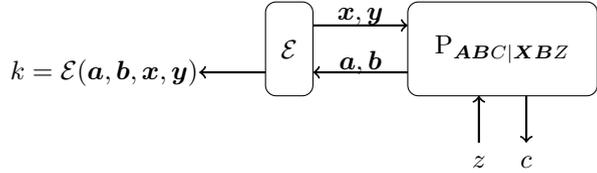
8.3.3 Diamond Norm

Theorem 8.3 allows for a simple treatment of cases that can be analysed using the notation of a test. In some information processing tasks this is not possible and different ways of utilising the reductions are needed. In this section we consider the task of distinguishing two channels acting on boxes. The channels can describe, for example, a cryptographic protocol.¹¹

¹⁰Linearity refers here to the linearity of the test in the box $P_{AB|XY}$, which follows from the fact that the test interacts only once with $P_{AB|XY}$ (or, in other words, the test gets only a single copy of the box).

¹¹Let us briefly explain why the notation of a test considered in Sect. 8.3.1 is not appropriate in the cryptographic setting. When considering tests, we were interested in events defined over $\mathcal{X}^n \times \mathcal{Y}^n \times \mathcal{A}^n \times \mathcal{B}^n$. Whether an output of a protocol (a key, for example) is secure to use cannot

Fig. 8.4 The channel $\mathcal{E} \otimes \mathbb{I}$ acts on an extension $P_{ABC|XBZ}$ of $P_{AB|XY}$ and outputs a classical string $k \in \{0, 1\}^t$ according to the probability $E_K(k)$. After [13]



When considering quantum protocols the distinguishing advantage is given by the diamond norm [17]. The distance between two channels \mathcal{E} and \mathcal{F} which act on quantum states ρ_{AB} is given by $\|\mathcal{E} - \mathcal{F}\|_\diamond = \max_{\rho_{ABC}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I} \rho_{ABC}\|_1$ where ρ_{ABC} is a purification of ρ_{AB} and $\|\cdot\|_1$ is the trace distance. Informally, the idea is that in order to distinguish two channels we are not only allowed to choose the input state to the channels, ρ_{AB} , but also keep to ourselves a purifying state ρ_C .

Although the definition of the diamond norm includes a maximisation over all states ρ_{ABC} it was proven, using the quantum post-selection theorem [7], that when considering permutation invariant channels it is sufficient to calculate the distance for a specific quantum de Finetti state. Motivated by this, we give a similar bound on a distance analogous to the diamond norm for channels which act on boxes (instead of quantum states).

In the following, we denote by \mathcal{P} the set of all boxes $P_{AB|XY}$ and by \mathcal{K} the set of all probability distributions P_K over $\{0, 1\}^t$ for some $t \in \mathbb{N}$. We consider channels of the form $\mathcal{E} : \mathcal{P} \rightarrow \mathcal{K}$ which interact with boxes $P_{AB|XY}$ and output a classical bit string $k \in \{0, 1\}^t$ of some length $t \geq 0$ with some probability $P_K(k)$. The connection between the channel and the box is illustrated in Fig. 8.4.¹²

The probability distribution of the output depends on the channel \mathcal{E} itself and is given by the following definition.

Definition 8.12 The probability that a channel \mathcal{E} outputs a string $k \in \{0, 1\}^t$ when interacting with $P_{AB|XY}$ is

$$E_K(k) = \sum_{x,y} \Pr_{\mathcal{E}}(x, y) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} P_{AB|XY}(a, b|x, y)$$

where $\Pr_{\mathcal{E}}(x)$ is the probability that \mathcal{E} inputs x, y to $P_{AB|XY}$ and $\mathcal{E}(a, b, x, y)$ is the function according to which the output of the channel is determined. Analogously,

$$E_{K|C}(k|c) = \sum_{x,y} \Pr_{\mathcal{E}}(x, y) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} P_{AB|XYC}(a, b|x, y, c).$$

be defined as an event. Security depends on the *process* of producing the key rather on the specific *data* that was produced during the run of the protocol.

¹²Figure 8.4 is almost identical to Fig. 8.3, describing a test. The difference between the two scenarios lies in the quantity that we wish to bound; see the previous footnote.

Definition 8.13 The distance between two channels $\mathcal{E}, \mathcal{F} : \mathcal{P} \rightarrow \mathcal{K}$ according to the diamond norm is

$$\|\mathcal{E} - \mathcal{F}\|_{\diamond} = \max_{\mathbb{P}_{ABC|XYZ}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1,$$

where the maximisation is over all boxes $\mathbb{P}_{AB|XY}$ and all possible extensions of them and

$$\begin{aligned} \mathcal{E} \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ}) &= \mathcal{E} \otimes \mathbb{I}(\mathbb{P}_{AB|XY} \cdot \mathbb{P}_{C|Z}) \\ &= \mathbb{E}_{K|C} \cdot \mathbb{P}_{C|Z}. \end{aligned}$$

$\mathcal{F} \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})$ is defined in a similar way.

Similarly to the concept of a permutation invariant test presented in Definition 8.10, we define a permutation invariant channel:

Definition 8.14 A channel \mathcal{E} is permutation invariant if for all boxes $\mathbb{P}_{AB|XY}$ and all permutations π we have

$$\mathcal{E}(\mathbb{P}_{AB|XY}) = \mathcal{E}(\mathbb{P}_{AB|XY} \circ \pi).$$

Using the de Finetti reduction, Theorem 8.3, we prove the following theorem.

Theorem 8.15 For any two permutation invariant channels $\mathcal{E}, \mathcal{F} : \mathcal{P} \rightarrow \mathcal{K}$

$$\|\mathcal{E} - \mathcal{F}\|_{\diamond} \leq (n+1)^{l(m-1)} \max_{\tau_{ABC|XYZ}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ})\|_1 \quad (8.4)$$

where $\tau_{ABC|XYZ}$ is a non-signalling extension of the de Finetti box $\tau_{AB|XY}$ where given in Eq. (8.3).

Theorem 8.15 tells us that when looking to bound the diamond norm for permutation invariant channels, one does not need to optimise over all possible boxes (as in Definition 8.13) but can consider only extensions of de Finetti boxes¹³ without loss of generality. This gives us another example as to why a de Finetti reduction is a *reduction to IID* technique. As in the case of Theorem 8.11 if one is able to find an exponentially small upper bound on $\|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ})\|_1$, an exponentially small upper bound on $\|\mathcal{E} - \mathcal{F}\|_{\diamond}$ follows. That is, the polynomial pre-factor $(n+1)^{l(m-1)}$ does not affect the asymptotic behaviour.

The proof of Theorem 8.15 builds on the following lemma.

Lemma 8.16 For every two permutation invariant channels $\mathcal{E}, \mathcal{F} : \mathcal{P} \rightarrow \mathcal{K}$ where \mathbb{P}_K is a probability distribution over $k \in \{0, 1\}^t$ for some $t > 0$, and all $\mathbb{P}_{ABC|XYZ}$,

$$\|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1 \leq (n+1)^{l(m-1)} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ}^{\mathbb{P}_{ABC|XYZ}})\|_1$$

¹³Note, however, that the extension $\tau_{ABC|XYZ}$ itself cannot be written as a convex combination of IID boxes, only its marginal $\tau_{AB|XY}$ is a de Finetti box. Furthermore, $\tau_{AB|XY}$ may be signalling in general, as before.

where $\tau_{ABC|XYZ}^{\mathbb{P}_{ABC|XYZ}}$ is a non-signalling extension of $\tau_{AB|XY}$ which depends on the specific box $\mathbb{P}_{ABC|XYZ}$.

The proof of the lemma follows by using Lemma 8.6 in order to construct a specific convex decomposition of $\tau_{AB|XY}$ from a convex decomposition of $\mathbb{P}_{AB|XY}$. A detailed proof is given in Appendix A.2.

Theorem 8.15 now easily follows from Lemma 8.16:

Proof of Theorem 8.15 Using Lemma 8.16,

$$\begin{aligned} \|\mathcal{E} - \mathcal{F}\|_{\diamond} &= \max_{\mathbb{P}_{ABC|XYZ}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1 \\ &\leq (n+1)^{l(m-1)} \max_{\substack{\mathbb{P}_{ABC|XYZ} \\ \tau_{ABC|XYZ}}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ}^{\mathbb{P}_{ABC|XYZ}})\|_1 \\ &\leq (n+1)^{l(m-1)} \max_{\tau_{ABC|XYZ}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ})\|_1 \end{aligned}$$

where $\tau_{ABC|XYZ}$ is a non-signalling extension of $\tau_{AB|XY}$. □

8.4 Impossibility Results

Before concluding this chapter, let us discuss the directions in which one could hope to further develop the technique of device-independent de Finetti reductions. We do so by presenting several impossibility results with regards to different variants of Theorem 8.3.

8.4.1 Restricted de Finetti Box

First, as explained in the above sections, our de Finetti box, given in Eq. (8.3), is a *signalling* box. Clearly, this raises some difficulties when coming to use the different theorems presented in this chapter.¹⁴ Ideally, we would have wished to have a de Finetti reduction in which the de Finetti box $\tau_{AB|XY}$ can be quantum or non-signalling when starting with a quantum or non-signalling box $\mathbb{P}_{AB|XY}$. That is, we wish to find reductions of the form (with some c polynomial¹⁵ in n):

$$\mathbb{P}_{AB|XY}^{\text{quant}} \leq c \cdot \tau_{AB|XY}^{\text{quant}} \quad ; \quad \mathbb{P}_{AB|XY}^{\text{ns}} \leq c \cdot \tau_{AB|XY}^{\text{ns}}, \quad (8.5)$$

where $\mathbb{P}_{AB|XY}^{\text{quant}}$ and $\tau_{AB|XY}^{\text{quant}}$ are quantum boxes and, similarly, $\mathbb{P}_{AB|XY}^{\text{ns}}$ and $\tau_{AB|XY}^{\text{ns}}$ are non-signalling boxes.

¹⁴Though this does not make them useless; see Chap. 10.

¹⁵Weaker statements, e.g., with a pre-factor sub-exponential in n , may also be of interest in certain applications.

Sadly, such reductions cannot be true when considering general permutation invariant boxes $P_{AB|XY}^{\text{quant}}$ and $P_{AB|XY}^{\text{ns}}$. One way to see that this is the case is by considering the task of parallel repetition of games (which acts as one of our showcases; see Sect. 4.1). Reductions as those in Eq. (8.5) will imply very strong parallel repetition results. Indeed, if, e.g., $P_{AB|XY}^{\text{quant}} \leq c \cdot \tau_{AB|XY}^{\text{quant}}$ holds for any permutation invariant quantum box $P_{AB|XY}^{\text{quant}}$, then it follows that, for *any* game,

$$w \left(P_{AB|XY}^{\text{quant}} \right) \leq c \cdot w \left(\tau_{AB|XY}^{\text{quant}} \right) = \text{poly}(n) \cdot \omega^n, \quad (8.6)$$

where $w(\circ)$ is the winning probability of the considered box in the repeated game, ω is the winning probability of the optimal quantum strategy in a single game, and $\text{poly}(n)$ is some polynomial of n , possibly depending on the alphabet of the RVs A , B , X , and Y . However, there are examples of games (in the classical, quantum, and non-signalling case) for which a strong decrease in the winning probability with the number of games played n , as in Eq. (8.6), does not hold; recall Sect. 4.1. Thus, reductions as in Eq. (8.5) cannot be true.

Knowing that Eq. (8.5) is not more than a wishful thinking, one could hope for the next best thing, i.e., an approximate version of the reduction. Concretely, we are interested in reductions of the form

$$P_{AB|XY}^{\text{quant}} \leq c \cdot \tau_{AB|XY}^{\text{approx-quant}} \quad ; \quad P_{AB|XY}^{\text{ns}} \leq c \cdot \tau_{AB|XY}^{\text{approx-ns}}, \quad (8.7)$$

where $\tau_{AB|XY}^{\text{approx-quant}}$ is an *approximately*-quantum de Finetti box and $\tau_{AB|XY}^{\text{approx-ns}}$ is an *approximately*-non-signalling one. By approximately-quantum (and analogously for the non-signalling case) we mean that the de Finetti box can be written as

$$\tau_{AB|XY}^{\text{approx-quant}} = \int \left(\mathcal{O}_{AB|XY}^{\text{quant}} \right)^{\otimes n} d\mathcal{O}_{AB|XY}^{\text{quant}} + \int \left(\mathcal{O}_{AB|XY}^{\text{non-quant}} \right)^{\otimes n} d\mathcal{O}_{AB|XY}^{\text{non-quant}},$$

where $d\mathcal{O}_{AB|XY}^{\text{quant}}$ and $d\mathcal{O}_{AB|XY}^{\text{non-quant}}$ are measures over quantum and non-quantum single-round boxes, respectively, and $\int d\mathcal{O}_{AB|XY}^{\text{non-quant}}$ is, say, exponentially small in n and/or assigns weight only to boxes $\mathcal{O}_{AB|XY}^{\text{non-quant}}$ which are close to quantum boxes, under some distance measure.¹⁶

Parallel repetition results can, again, be used to show that such reductions cannot hold in general, at least in the non-signalling case. Here the reason lies in the observation that the reductions in Eq. (8.7) are independent of the choice of *distribution* over the inputs \mathcal{X}^n and \mathcal{Y}^n (while they may depend on the *alphabet* of the inputs). Thus, they would imply general parallel repetition results which hold for any distribution over the inputs to the parallel boxes. As there are games for which such non-signalling parallel repetition results do not hold [18], at best $P_{AB|XY}^{\text{ns}} \leq c \cdot \tau_{AB|XY}^{\text{approx-ns}}$ cannot be true in general.

¹⁶The hope here is that by adding the additional weight on non-quantum or signalling boxes one could account for the ‘‘gap’’ between Eq. (8.6) and the known parallel repetition results.

By this we learn that we ought to consider reductions that also include the input distribution P_{XY} :

$$P_{XY} P_{AB|XY}^{\text{quant}} \leq c \cdot P_{XY} \tau_{AB|XY}^{\text{approx-quant}}, \quad (8.8)$$

$$P_{XY} P_{AB|XY}^{\text{ns}} \leq c \cdot P_{XY} \tau_{AB|XY}^{\text{approx-ns}}. \quad (8.9)$$

The case of $P_{XY} = Q_{XY}^{\otimes n}$ is of special interest. For such distributions, two results are known. In Sect. 10.2 we prove a result in the *flavour* of Eq. (8.9) using the de Finetti reduction given as Theorem 8.3. The result, which originally appeared as part of [19], is stated informally as Theorem 10.2. Roughly speaking, it says that observed data that is sampled using a permutation invariant non-signalling parallel box looks *as if* it was sampled using an approximately non-signalling IID box.

In [20] a reduction similar to that of Eq. (8.9) was proven by combining the de Finetti reduction in Theorem 8.3 together with another de Finetti-type theorem, presented in [21]. Their theorem can be written as follows¹⁷:

Theorem 8.17 (Theorem 4.3 in [20]) *For any non-signalling permutation invariant parallel box $P_{AB|XY}$ and distribution Q_{XY}*

$$Q_{XY}^{\otimes n} P_{AB|XY} \leq \int \tilde{F}(O_{ABXY})^{2n} O_{ABXY}^{\otimes n} dO_{ABXY}, \quad (8.10)$$

where

$$\tilde{F}(O_{ABXY}) = \min \left\{ \max_{R_{A|X}} F(Q_{XY} R_{A|X}, O_{AXY}), \max_{R_{B|Y}} F(Q_{XY} R_{B|Y}, O_{BXY}) \right\}$$

for F the fidelity.

To see that Eq. (8.10) is in the spirit of Eq. (8.9) note that $\tilde{F}(O_{ABXY})$ is some measure of how far O_{ABXY} is from $Q_{XY} \tilde{O}_{AB|XY}$ for a non-signalling box $\tilde{O}_{AB|XY}$. Recall that the fidelity is small when the distributions are far from one another; thus, $\tilde{F}(O_{ABXY})^{2n}$ assures that only negligible weight is assigned to distributions O_{ABXY} originating from highly signalling boxes (or with marginals O_{XY} far from Q_{XY}).

We conjecture that reductions similar to Eq. (8.8), relevant for quantum boxes, should also hold. Yet, up to date there are no proofs in this direction (the difficulty in deriving such a statement is discussed in Chap. 10).

8.4.2 Extension to an Adversary

Another direction in which one may wish to extend our de Finetti reductions is relevant for device-independent cryptographic protocols. To explain what we aim

¹⁷We present only the bipartite case; [20, Theorem 4.3] is stated for an arbitrary number of parties.

for, let us first discuss the quantum variant of Theorem 8.15, also called the post-selection technique, developed in [7].¹⁸ The post-selection theorem implies that for any two permutation invariant quantum channels, \mathcal{E} and \mathcal{F} , acting on quantum states $\rho_{Q_A Q_B} \in \mathcal{S}(\mathcal{H}_{Q_A Q_B}^{\otimes n})$ for some bipartite Hilbert space $\mathcal{H}_{Q_A Q_B}$ of dimension d ,

$$\|\mathcal{E} - \mathcal{F}\|_{\diamond} \leq (n+1)^{d^2-1} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{Q_A Q_B E})\|_1 \quad (8.11)$$

where $\tau_{Q_A Q_B E}$ is a purification of a given de Finetti state. Equation (8.11) should be compared to Eq. (8.4); while Eq. (8.4) includes a maximisation over all possible non-signalling extensions of the de Finetti box, in Eq. (8.11) we consider only a single purification. The reason is simple—in the quantum case all purifications of a state are equivalent up to local unitaries. Furthermore (and crucially for applications), there exists a purification of a de Finetti state that has a very special form. To purify

$$\tau_{Q_A Q_B} = \int (\sigma_{Q_A Q_B})^{\otimes n} d\sigma_{Q_A Q_B}$$

we can first purify the states $\sigma_{Q_A Q_B}$ to get

$$\tau_{Q_A Q_B E'} = \int (\sigma_{Q_A Q_B E'})^{\otimes n} d\sigma_{Q_A Q_B E'}$$

and then purify the state $\tau_{Q_A Q_B E'}$ using an additional system E'' to account for the convex combination of the pure states $(\sigma_{Q_A Q_B E'})^{\otimes n}$. This defines us the pure state $\tau_{Q_A Q_B E' E''}$. Denoting $E = E' E''$ we get our pure $\tau_{Q_A Q_B E}$.

In the cryptographic setting the quantum register E is considered to belong to the adversary. Hence, any information about the structure of the system kept in it could be useful when analysing security. Equation (8.11) in combination with the observation regarding the structure of the purification, $\tau_{Q_A Q_B E' E''}$, we learn that the main task when proving security is to analyse the IID case, as in Chap. 7 (see [7, 16] for the detailed explanation). That is, the quantum de Finetti reduction can be used as a reduction to IID in quantum cryptography.

In contrast, in general, it is impossible to prove a modified version of Theorem 8.15 in which the extension $\tau_{ABC|XYZ}$ of our de Finetti box $\tau_{AB|XY}$ will be as structured as the quantum state $\tau_{Q_A Q_B E}$. In particular, even if we can start with a de Finetti reduction where both $P_{AB|XY}$ and $\tau_{AB|XY}$ are non-signalling,¹⁹ it is impossible to derive a theorem which would imply that the analysis of device-independent cryptography in the presence of a non-signalling adversary can be reduced to the analysis under the

¹⁸Reference [7] presented the first de Finetti reduction, i.e., an inequality relation between permutation invariant systems and de Finetti systems (all previous de Finetti-type theorems gave other types of relations between the two systems). The term “de Finetti reduction” was not used at that time and the authors chose the name “post-selection technique” as they first proved the quantum analogue of Lemma 8.6.

¹⁹In the presence of certain types of symmetries (in addition to permutation invariance) one can derive such de Finetti reductions; see [13].

IID assumption. This is due to the impossibility result of [22], which asserts that, while exponential privacy amplification in the presence of a non-signalling adversary is possible under the IID assumption [23], it is impossible when the IID assumption is dropped.

8.4.3 Other de Finetti-Type Theorems

A final remark is with regards to the more common type of de Finetti theorem, in which one bounds the trace distance between an n -exchangeable system and a de Finetti one. More specifically, let us first consider the classical case, i.e., a system is a probability distribution. P_{A_1, \dots, A_k} is permutation invariant if it is invariant under any permutation of A_1, \dots, A_k (as before). We say that P_{A_1, \dots, A_k} is n -exchangeable, for $n \geq k$, if it is a marginal of some permutation invariant P_{A_1, \dots, A_n} . In [24] a bound on the *distance* between an n -exchangeable probability distribution and a de Finetti distribution was proven.²⁰ Results of this type were also proven for quantum states [6, 25] and non-signalling boxes [8].

Let us focus on the non-signalling case [8]. There, a conditional probability distribution $P_{A_1, \dots, A_n | X_1, \dots, X_n}$ is said to be non-signalling if the box cannot be used to signal from any subset of parties $I \subset [n]$ to the rest of the parties $[n] \setminus I$. Permutation invariance is defined with respect to permutations $\pi : [n] \rightarrow [n]$. Similarly to the classical case described above, $P_{A_1, \dots, A_k | X_1, \dots, X_k}$ is n -exchangeable, for $n \geq k$, when it is the marginal of a permutation invariant non-signalling box $P_{A_1, \dots, A_n | X_1, \dots, X_n}$. We then have the following bound [8, Theorem 3] (using the above notation):

Theorem 8.18 ([8]) *For any permutation invariant non-signalling box $P_{A_1, \dots, A_n | X_1, \dots, X_n}$ and any $k < n$ there exists a de Finetti box $\tau_{A_1, \dots, A_k | X_1, \dots, X_k}$ such that*

$$\left| P_{A_1, \dots, A_k | X_1, \dots, X_k} - \tau_{A_1, \dots, A_k | X_1, \dots, X_k} \right| \leq \min \left\{ \frac{2k|\mathcal{X}| |\mathcal{A}|^{|\mathcal{X}|}}{n}, \frac{k(k-1)|\mathcal{X}|}{n} \right\}.$$

The crucial thing to note here is that the boxes $P_{A_1, \dots, A_k | X_1, \dots, X_k}$ and $P_{A_1, \dots, A_n | X_1, \dots, X_n}$ are a very special type of parallel boxes: the non-signalling conditions must hold for any division of the indices in $[n]$. This implies that for any $i, j \in [n]$, A_i is independent of the inputs X_j for $j \neq i$. Theorems such as Theorem 8.18 cannot be proven for general parallel boxes since they study exchangeable boxes, which inherently require the ability to consider the marginals of the boxes.

²⁰In this language, the original result of de Finetti [1] stated that all infinitely-exchangeable distributions (i.e., distributions that are n -exchangeable for any $n \geq k$) are equal to distributions of the form of a convex combination of IID distributions.

References

1. de Finetti B (1969) Sulla proseguibilità di processi aleatori scambiabili. *Rend Matem Trieste* 53–67
2. Diaconis P, Freedman D (1980) Finite exchangeable sequences. *Annal Probab* 745–764
3. Raggio G, Werner R (1989) Quantum statistical mechanics of general mean field systems. *Helv Phys Acta* 62(8):980–1003
4. Caves CM, Fuchs CA, Schack R (2002) Unknown quantum states: the quantum de-Finetti representation. *J Math Phys* 43:4537
5. Renner R (2007) Symmetry of large physical systems implies independence of subsystems. *Nat Phys* 3(9):645–649
6. Christandl M, König R, Mitchison G, Renner R (2007) One-and-a-half quantum de Finetti theorems. *Commun Math Phys* 273(2):473–498
7. Christandl M, König R, Renner R (2009) Postselection technique for quantum channels with applications to quantum cryptography. *Phys Rev Lett* 102(2):020504
8. Christandl M, Toner B (2009) Finite de Finetti theorem for conditional probability distributions describing physical theories. *J Math Phys* 50:042104
9. Brandao FG, Harrow AW (2013) Quantum de Finetti theorems under local measurements with applications. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 861–870. ACM
10. Leverrier A (2014) Composable security proof for continuous-variable quantum key distribution with coherent states. [arXiv:1408.5689](https://arxiv.org/abs/1408.5689)
11. Christandl M, Renner R (2012) Reliable quantum state tomography. *Phys Rev Lett* 109(12):120403
12. Berta M, Christandl M, Renner R (2011) The quantum reverse Shannon Theorem based on one-shot information theory. *Commun Math Phys* 306(3):579–615
13. Arnon-Friedman R, Renner R (2015) de Finetti reductions for correlations. *J Math Phys* 56(5):052203
14. Hänggi E, Renner R, Wolf S (2010) Efficient device-independent quantum key distribution. In: *Advances in cryptology—EUROCRYPT 2010*, pp 216–234. Springer
15. Hänggi E, Renner R (2010) Device-independent quantum key distribution with commuting measurements. [arXiv:1009.1833](https://arxiv.org/abs/1009.1833)
16. Renner R (2010) Simplifying information-theoretic arguments by post-selection. In: *NATO advanced research workshop quantum cryptography and computing: theory and implementation*, vol 26, pp 66–75. IOS Press
17. Kitaev AY (1997) Quantum computations: algorithms and error correction. *Russ Math Surv* 52(6):1191–1249
18. Holmgren J, Yang L (2017) (a counterexample to) parallel repetition for non-signaling multi-player games. In: *Electronic colloquium on computational complexity (ECCC)*, vol 24, p 178
19. Arnon-Friedman R, Renner R, Vidick T (2016) Non-signaling parallel repetition using de finetti reductions. *IEEE Trans Inf Theory* 62(3):1440–1457
20. Lancien C, Winter A (2016) Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de finetti reduction. *Chic J Theor Comput Sci* (11)
21. Lancien C, Winter A (2017) Flexible constrained de finetti reductions and applications. *J Math Phys* 58(9):092203
22. Arnon-Friedman R, Ta-Shma A (2012) Limits of privacy amplification against nonsignaling memory attacks. *Phys Rev A* 86(6):062333
23. Hänggi E, Renner R, Wolf S (2009) Quantum cryptography based solely on bell’s theorem. [arXiv:0911.4171](https://arxiv.org/abs/0911.4171)
24. Diaconis P, Freedman D (1980) Finite exchangeable sequences. *Ann Probab* 745–764
25. König R, Renner R (2005) A de finetti representation for finite symmetric quantum states. *J Math Phys* 46(12):122108

Chapter 9

Reductions to IID: Sequential Interaction



Many device-independent protocols proceed in rounds and, hence, require devices with which the honest parties can interact sequentially, i.e., one round after the other. A particular example for such protocols is our showcase dealing with device-independent quantum cryptography. When analysing the showcase under the IID assumption (Sect. 7.3.2), we observed that the quantum AEP, given as Theorem 7.3, plays a crucial role in the proof of soundness. Specifically, the quantum AEP allowed us to bound the total amount of the relevant smooth entropy using a bound on the von Neumann entropy calculated for a single round of the protocol.

The focus of this chapter is the so called “entropy accumulation theorem” (EAT) [1, 2]. The EAT is a generalisation of the AEP to a scenario in which, instead of the raw data being produced by an IID process, it is produced by certain sequential quantum processes of interest.¹ In particular, similarly to the AEP, the EAT allows one to bound the total amount of the considered smooth entropy using the same bound on the von Neumann entropy calculated for the IID analysis. In this sense, the EAT can be seen as a *reduction to IID*—with the aid of the EAT the analysis done under the IID assumption using the AEP is directly extended to the one relevant for *multi-round sequential boxes*.

Below, we motivate, present, and explain the EAT in the form relevant for device-independent quantum information processing [3]. The EAT is later used in the analysis of our showcase in Chap. 11. For the most general statement of the EAT and its proof the interested reader is referred to [1, 2].

¹We remark that the EAT, as the quantum AEP, is only relevant when assuming that everything can be described within the quantum formalism. In particular, it cannot be used when talking about, e.g., cryptographic protocols in the presence of a non-signalling (super-quantum) adversary.

9.1 Sequential Quantum Processes

We are interested in multi-round quantum sequential boxes (with communication between the rounds; see Definition 6.6) fulfilling certain conditions. The simplest way of describing the relevant conditions is by looking at the sequential quantum process defining the boxes, i.e., the process that results in the input-output distribution of the boxes.

Consider a sequential process as illustrated in Fig. 9.1. We start with some initial state $\rho_{R_0 E}^{\text{in}}$; the distinction between R_0 and E is such that R_0 is the part of the state which may change during the considered process while E denotes the “environment” register, i.e., the part of the state not being modified. The marginal $\rho_{R_0}^{\text{in}}$ undergoes a sequence of operations in which a sequence of (non-IID) registers $\mathbf{O} = O_1, \dots, O_n$ and $\mathbf{S} = S_1, \dots, S_n$ are being created. We treat the registers \mathbf{O} as the “output registers” while \mathbf{S} are the “side-information registers”. Our ultimate objective is to bound the conditional smooth entropies $H_{\min}^\varepsilon(\mathbf{O}|\mathbf{S}E)$ and $H_{\max}^\varepsilon(\mathbf{O}|\mathbf{S}E)$.

To be able to bound the above entropies, some statistical data must be collected during the protocol. Specifically, we consider additional classical registers $\mathbf{C} = C_1, \dots, C_n$ holding the information relevant for the estimation phase performed in the protocol. For every round $i \in [n]$, C_i is derived by performing some action on the registers O_i and S_i . For example, the value of C_i can be the result of applying a function on some classical information included in O_i and S_i .

To gain a bit of intuition regarding all the different registers, let us quickly consider the cryptographic setting (a more precise discussion can be found in Chap. 11). When analysing cryptographic protocols one may make the following choices: E acts as the register belonging to the adversary, \mathbf{O} as the raw data which is supposed to be secret, \mathbf{S} as the side-information leaked during the protocol (e.g., all classical information which is communicated between Alice and Bob), and \mathbf{C} —the indicators of whether the test rounds were successful or not (e.g., $C_i = 1$ when the i 'th game was won).

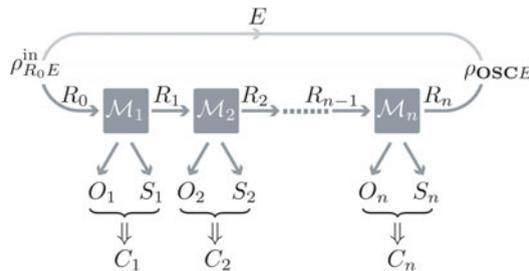


Fig. 9.1 Sequential quantum process. The initial state $\rho_{R_0 E}^{\text{in}}$ is transformed to the final one ρ_{OSCE} by applying a sequence of maps on the marginal $\rho_{R_0}^{\text{in}}$. Each map \mathcal{M}_i in the sequence outputs the registers O_i and S_i , from which C_i is created. The “memory system” R_i is being passed to the next map as input

The goal is then to lower bound $H_{\min}^\varepsilon(\mathbf{O}|\mathbf{S}E)$ and this should be done by using the statistics kept in \mathbf{C} .

The sequential process itself is formally defined by a sequence of quantum channels, namely, CPTP maps (Definition 2.14),

$$\mathcal{M}_i : R_{i-1} \rightarrow R_i O_i S_i C_i , \quad (9.1)$$

for all $i \in [n]$. As seen from Eq. (9.1), when we say that a process is *sequential* we not only mean that the maps act one after the other, but also that, in each round $i \in [n]$, the output O_i , the side information S_i , and the estimation data C_i are being created by the map \mathcal{M}_i applied in *that* round. That is, \mathbf{O} denotes a sequence of registers created one after the other and similarly for \mathbf{S} and \mathbf{C} .² The state of interest in the end of the process is given by³

$$\rho_{OSCE} = (\text{Tr}_{R_n} \circ \mathcal{M}_n \circ \cdots \circ \mathcal{M}_1) \otimes \mathbb{I}_E \rho_{R_0E}^{\text{in}} . \quad (9.2)$$

We remark that in the device-independent setting the initial state $\rho_{R_0E}^{\text{in}}$, the maps $\{\mathcal{M}_i\}_{i \in [n]}$, and the final state ρ_{OSCE} are unknown; we merely require that some quantum states and maps $\{\mathcal{M}_i\}_{i \in [n]}$, describing the overall process, exist. In particular, this implies that we do not restrict the content of the registers $\{R_i\}_{i \in [n]}$ and, thus, they may include information which is being passed from previous rounds to the next ones. That is, we can think of these registers as holding some quantum memory. This is in stark contrast to what happens when working under the IID assumption.

9.2 Entropy Accumulation Theorem

As mentioned above, the EAT [1] acts as a generalisation of the quantum AEP to scenarios in which certain sequential processes are considered, rather than IID ones. The EAT, as the name suggests, quantifies the amount of entropy accumulated during the considered quantum processes. Moreover, as in the case of the AEP, the total amount of smooth entropies can be bounded by calculating certain von Neumann entropies (the precise statements are given below). By this, the EAT justifies the use of the von Neumann entropy in quantum information processing even outside of the IID regime. While we discuss the EAT in the context of device-independent quantum information processing, we remark that the EAT is a general information-theoretic tool, which can also be applied in other contexts.

²The reader may be concerned by the distinction between, e.g., O_1 and O_2 —clearly, we can also consider a situation in which \mathcal{M}_1 does not output O_1 but transfers it to \mathcal{M}_2 that later outputs $O_1 O_2$. We will soon assume that the different registers fulfil certain conditions and then the distinction will become clear.

³We will be interested below in bounding the smooth entropies evaluated on a state closely related to the final state ρ_{OSCE} , namely, the final state conditioned on the event of not aborting the considered protocol; see Sect. 9.2.3.

9.2.1 Conceptual Difficulties to Overcome

Before stating the theorem itself, let us explain the conceptual difficulties that arise when seeking for an “AEP-style” theorem in non-IID scenarios. Specifically, we would like to understand what is the form of the theorem we are aiming for and what is non-trivial about it. To keep this section concise we focus on the smooth min-entropy; the same statements are relevant for the smooth max-entropy as well.

Our goal is to have a theorem resembling the quantum AEP appearing as Theorem 7.3. That is, we look for a statement of the form

$$H_{\min}^{\varepsilon}(\mathbf{O}|\mathbf{S}E) \geq nt - \mu\sqrt{n} . \quad (9.3)$$

for some t and μ (independent of n but otherwise unrestricted).

The EAT aims at providing a lower-bound on $H_{\min}^{\varepsilon}(\mathbf{O}|\mathbf{S}E)$ which scales linearly with the number of rounds n (to first order in n , i.e., up to finite statistic effects, as in the AEP). As $\mathbf{O} = O_1, \dots, O_n$ and $\mathbf{S} = S_1, \dots, S_n$ are being created in a sequential manner, we intuitively wish to say that in each round $i \in [n]$ we accumulate an additional constant amount of entropy due to the production of O_i (while taking into account S_i and E) until, in the end of the process, the total amount of entropy is linear in n . Consider, however, a sequential process in which S_1, \dots, S_{n-1} are all empty (i.e., do not reveal any information about the outputs) while the side-information S_n produced by the last map \mathcal{M}_n includes all of the outputs O_1, \dots, O_n . Clearly, even though we may have accumulated entropy in the rounds $i \in [n-1]$, all of it is lost after S_n is leaked in the last round n . This implies that we can only hope to prove a statement like the one given in Eq. (9.3) under some restrictions on the considered sequential processes.

A more fundamental difficulty to overcome is the following. In the case of the AEP, i.e., when considering IID processes, or boxes, it is clear what t , appearing in Eq. (9.3), is—it is a quantity describing the *single-round box* defining the IID box. (And, as it turns out, this quantity is the relevant conditional von Neumann entropy evaluated on a single-round box; recall Sect. 7.2.2).

Moving to the sequential processes, or multi-round boxes, it is not obvious at all which quantity t should describe. We would like to find a quantity related to some “single-round property”, but how can we even define such a thing in a meaningful way? Since the behaviour of the box in each round may depend on everything that happened in previous rounds (see Definition 6.6), we cannot directly define a multi-round box in terms of single-round ones. To put it differently, when holding a multi-round device, there is no physical system that we can point to and treat as an isolated subsystem. Thus, t cannot refer to such a system as in the IID case.

Keeping in mind the conceptual difficulties that one needs to overcome when phrasing the theorem, we are now ready to discuss the EAT on a more concrete level. In particular, the following section unveils the resolutions of the issues presented above.

9.2.2 Prerequisites of the Theorem

Before presenting the EAT, we need to define two objects to which the theorem refers—EAT channels and tradeoff functions. The “correct” definition of these objects is what allows us to overcome the conceptual difficulties discussed above. Furthermore, when coming to use the EAT, the choice, or construction, of these objects is what allows one to derive a strong bound on the considered entropy. Thus, understanding the prerequisites of the theorem is of great importance.

9.2.2.1 EAT Channels

As mentioned in the previous section, entropy does not accumulate in any general sequential process. Therefore, we must restrict our attention to processes which fulfil certain conditions. Specifically, we work with processes defined via the following type of maps, called “EAT channels”:

Definition 9.1 (*EAT channels*) Quantum channels $\{\mathcal{M}_i : R_{i-1} \rightarrow R_i O_i S_i C_i\}_{i \in [n]}$ are said to be EAT channels if the following requirements hold:

1. $\{O_i\}_{i \in [n]}$ are finite dimensional quantum systems of dimension d_O and $\{C_i\}_{i \in [n]}$ are finite-dimensional classical systems (RV). $\{S_i\}_{i \in [n]}$ and $\{R_i\}_{i \in [n]}$ are arbitrary quantum systems.
2. For any $i \in [n]$ and any input state $\sigma_{R_{i-1}}$, the output state $\sigma_{R_i O_i S_i} = \mathcal{M}_i(\sigma_{R_{i-1}})$ has the property that the classical value C_i can be measured from the marginal $\sigma_{O_i S_i}$ without changing the state. That is, for the map $\mathcal{T}_i : O_i S_i \rightarrow O_i S_i C_i$ describing the process of deriving C_i from O_i and S_i , it holds that $\text{Tr}_{C_i} \circ \mathcal{T}_i(\sigma_{O_i S_i}) = \sigma_{O_i S_i}$.
3. For any initial state $\rho_{R_0 E}^{\text{in}}$, the final state $\rho_{O S C E} = (\text{Tr}_{R_n} \circ \mathcal{M}_n \circ \dots \circ \mathcal{M}_1) \otimes \mathbb{I}_E \rho_{R_0 E}^{\text{in}}$ fulfils the Markov-chain conditions (Definition 2.21)

$$O_1, \dots, O_{i-1} \leftrightarrow S_1, \dots, S_{i-1}, E \leftrightarrow S_i \quad (9.4)$$

for all $i \in [n]$.

In words, Eq. (9.4) states that in each round, the previous outcomes O_1, \dots, O_{i-1} are independent of the future side-information S_i given all the past side information S_1, \dots, S_{i-1}, E . That is, the side-information of any given round does not reveal new information about previous outcomes. When coming to use the EAT, one is free to choose the different systems as one wishes. By choosing O_i and S_i properly the required Markov chain condition can be satisfied by sequential protocols such as device-independent quantum key distribution, as will be shown in Chap. 11.⁴

⁴In some cases the obvious choices for O_i and S_i are such that Eq. (9.4) does not hold. Still, sometimes, one can overcome the problem by considering related protocols in which the Markov-chain conditions are “enforced”. This is done, for example, in [4].

Equation (9.4) acts as the additional constraint on the sequential process which allows us to avoid processes in which entropy does not accumulate.⁵ We remark that the above requirements, and Eq. (9.4) in particular, give sufficient, but perhaps not necessary, conditions for the entropy to accumulate. That is, there might be sets of weaker or incomparable conditions that can also be used to show that entropy accumulates.

9.2.2.2 Tradeoff Functions

As explained in Sect. 9.2.1, since we cannot directly define a multi-round box in terms of single-round ones, it is not clear what quantity should replace t in Eq. (9.3). The tradeoff functions, defined below, give an adequate way to quantify the amount of entropy which is accumulated in a single step of the process, i.e., in an application of just one channel, and by this allow us to define t in a meaningful way. We first present the formal definition of the functions and then explain.

Definition 9.2 (*Tradeoff functions*) Let $\{\mathcal{M}_i\}_{i \in [n]}$ be a family of EAT channels and \mathcal{C} denote the common alphabet of C_1, \dots, C_n . A differentiable and convex function f_{\min} from the set of probability distributions p over \mathcal{C} to the real numbers is called a *min-tradeoff function* for $\{\mathcal{M}_i\}_{i \in [n]}$ if it satisfies⁶

$$f_{\min}(p) \leq \inf_{\sigma_{R_{i-1}R'}: \mathcal{M}_i(\sigma)_{C_i} = p} H(O_i | S_i R')_{\mathcal{M}_i(\sigma)} \quad (9.5)$$

for all $i \in [n]$, where the infimum is taken over all purifications of input states of \mathcal{M}_i for which the marginal on C_i of the output state is the probability distribution p .

Similarly, a differentiable and concave function f_{\max} from the set of probability distributions p over \mathcal{C} to the real numbers is called a *max-tradeoff function* for $\{\mathcal{M}_i\}$ if it satisfies

$$f_{\max}(p) \geq \sup_{\sigma_{R_{i-1}R'}: \mathcal{M}_i(\sigma)_{C_i} = p} H(O_i | S_i R')_{\mathcal{M}_i(\sigma)} \quad (9.6)$$

for all $i \in [n]$, where the supremum is taken over all purifications of input states of \mathcal{M}_i for which the marginal on C_i of the output state is the probability distribution p .

Figure 9.2 illustrates the considered scenario—a single step in the sequential process (compare to Fig. 9.1). For any possible input state $\sigma_{R_{i-1}} \in \mathcal{S}(\mathcal{H}_{R_{i-1}})$ of the map \mathcal{M}_i , we denote by $\sigma_{R_{i-1}R'} \in \mathcal{S}(\mathcal{H}_{R_{i-1}} \otimes \mathcal{H}_{R'})$ its purification. Note that the register R' is not being affected by the map (similarly to E in Fig. 9.1). O_i and S_i denote the output and side-information registers of the output state $\mathcal{M}_i(\sigma)$. C_i can be inferred from O_i and S_i as before.

⁵One can easily verify that the problematic process described in Sect. 9.2.1 does not fulfil the Markov-chain conditions.

⁶The infimum and supremum over the empty set are defined as plus and minus infinity, respectively.

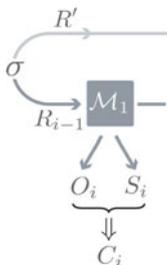


Fig. 9.2 A single step in the sequential process. The initial state is $\sigma_{R_{i-1}R'}$; $\sigma_{R_{i-1}}$ is the input of the map \mathcal{M}_i while R' acts as the environment register and is not affected by the map (similarly to E in Fig. 9.1). The map produces the registers O_i and S_i , from which C_i can be inferred

To comprehend these so called tradeoff functions, let us first discuss the set over which we perform the optimisations in Eqs. (9.5) and (9.6):

$$\Sigma(p) = \{ \sigma_{R_{i-1}R'} : \mathcal{M}_i(\sigma)_{C_i} = p \} . \tag{9.7}$$

$\mathcal{M}_i(\sigma_{R_{i-1}})$ is the output state of the map and $\mathcal{M}_i(\sigma)_{C_i}$ is its marginal over C_i . Recall that the classical registers C_i are used to collect statistics during the run of the considered protocol. Hence, $\mathcal{M}_i(\sigma)_{C_i}$ can be seen as a probability distribution over \mathcal{C} . The condition $\mathcal{M}_i(\sigma)_{C_i} = p$ therefore restricts the set of considered states— $\Sigma(p)$ only includes states σ that exhibit the statistics defined by the probability distribution p . If there are no such states then $\Sigma(p)$ is empty.

As an example, consider a protocol in which the CHSH game is being played in each round and C_i records whether the game was won ($C_i = 1$) or lost ($C_i = 0$). Denoting by ω the probability that σ wins the game, we can write

$$\mathcal{M}_i(\sigma)_{C_i} = \begin{pmatrix} \omega & 0 \\ 0 & 1 - \omega \end{pmatrix} . \tag{9.8}$$

$\Sigma(p)$ then includes all of the states for which $\omega = p(1)$. For a probability distribution with $p(1) = 1$, for example, the set $\Sigma(p)$ is empty, since there are no quantum states which can be used to play the CHSH game with probability $\omega = 1$.

Given the above, the infimum/supremum of $H(O_i|S_iR')$ over the set $\Sigma(p)$ describes the worst-case⁷ conditional von Neumann entropy in a *single round, restricted to states with the correct marginal over C_i* .

To get some intuition as to why the tradeoff functions in Definition 9.2 give an adequate way of quantifying the amount of entropy accumulated in a single step of the process, let us present two “alternative” definitions that one could try to use and refute them with the help of simple classical examples.

⁷By “worst-case” we mean lowest or largest, depending on whether we are working with min- or max- tradeoff functions. This will become clearer when discussing the EAT itself.

In both examples we consider classical processes in which each channel \mathcal{M}_i outputs a single bit O_i without any side information S_i about it; the system E is empty as well. Every bit O_i may depend on the ones produced previously. We would like to extract randomness out of \mathbf{O} and thus aim to calculate $H_{\min}^\varepsilon(\mathbf{O})$, which tightly describes the amount of extractable randomness [5, 6]. How much randomness does a single round contribute to the extractable randomness *given* that we already accounted for the randomness of the previous rounds?

One possible guess is the conditional von-Neumann entropy:

$$H(O_i | O_1, \dots, O_{i-1}) = -\mathbb{E}_{o_1, \dots, o_i} \log \Pr(o_i | o_1, \dots, o_{i-1}). \quad (9.9)$$

The von Neumann entropy fulfils the chain rule and so we have $\sum_i H(O_i | O_1, \dots, O_{i-1}) = H(\mathbf{O})$. Unfortunately, the smooth min-entropy $H_{\min}^\varepsilon(\mathbf{O})$ can be arbitrarily lower than $H(\mathbf{O})$. An example for a sequential process in which this is the case is as follows: O_1 is uniform while, for all $i \in [n] \setminus \{1\}$,

$$O_i = \begin{cases} 0 & O_1 = 0 \\ \text{uniform} & \text{otherwise.} \end{cases}$$

Direct calculation of $H(\mathbf{O})$ gives $H(\mathbf{O}) = 1 + (n-1)/2$. The min-entropy, however, depends on the most probable value of \mathbf{O} rather than its expected value. One can easily check that $H_{\min}(\mathbf{O}) = 1$, which implies that the extractable randomness is independent of n . Thus, $H(\mathbf{O})$ is too optimistic—it suggests that one can get arbitrarily more randomness than we can possibly extract from this process.

Let us try a worst-case version of the min-entropy instead:

$$H_{\min}^{w.c.} = -\log \max_{o_1, \dots, o_i} \Pr(o_i | o_1, \dots, o_{i-1}). \quad (9.10)$$

While this option at least does not result in a contradiction (in contrast to the one above), it is too pessimistic. To see this, consider IID RVs \mathbf{O} , where each O_i is a Bernoulli random variable with expectation $p < 1/2$. Then, Eq. (9.10) tells us that we can extract $-\log(1-p)$ randomness per round. However, it follows from the EAP that $h(p) > -\log(1-p)$ randomness can be extracted per round, for sufficiently large n .

The following quantity lies between those given in Eqs. (9.9) and (9.10):

$$\min_{o_1, \dots, o_{i-1}} H(O_i | O_1 = o_1, \dots, O_{i-1} = o_{i-1}).$$

This quantity describes the von Neumann entropy of O_i , evaluated for the worst case values of O_1, \dots, O_{i-1} . Going back to the two processes considered above, one can easily verify that this choice gives the “correct” amount of extractable randomness in both cases. The min-tradeoff function defined above is the quantum analogue of this.

The tradeoff functions are not uniquely defined by Eqs. (9.5) and (9.6). The equations merely pose a constraint on the functions. That is, a min-tradeoff function can be chosen to be any differentiable convex function satisfying the condition⁸ given in Eq. (9.5), i.e., it is upper-bounded by $\inf_{\Sigma(p)} H(O_i | S_i R')$. Similarly, a max-tradeoff function is any differentiable concave function lower-bounded by $\sup_{\Sigma(p)} H(O_i | S_i R')$. To get the tightest bounds on the smooth entropies using the EAT one should construct tradeoff functions in the tightest way possible, ideally matching the exact value of the worst-case von Neumann entropy given in Eqs. (9.5) and (9.6). In device-independent cryptographic protocols based on the CHSH game, for example, Lemma 5.3 can be used to construct a tight min-tradeoff function; this will be done in Chap. 11.

9.2.3 Statement of the Theorem

After presenting the prerequisites of the EAT, we are now ready to discuss the statement of the theorem.

9.2.3.1 Conditioning on Not Aborting

Consider a sequential protocol, i.e., one which proceeds in rounds; the development of the quantum state throughout the protocol can be described by a sequential process. In the end of the protocol, the honest parties can choose whether to abort the protocol or not. For example, if Alice and Bob run a device-independent cryptographic protocols and observe that the device does not win the game in sufficiently many games, they conclude that the device might be malicious and abort the protocol. Our goal is to bound the smooth entropies of the outputs when the protocol *does not abort*.

Whether the protocol aborts or not depends on the observed data produced during the execution of the protocol and, specifically, on the value assigned to \mathcal{C} . Thus, the event of not aborting the protocol, denoted by Ω , is defined to be a subset of \mathcal{C}^n . The most common way of choosing the set Ω is such that whether $\mathbf{c} = c_1, \dots, c_n \in \mathcal{C}^n$ belongs to Ω or not depends on its “frequencies”. Formally, for any $\mathbf{c} \in \mathcal{C}^n$, denote by $\text{freq}_{\mathbf{c}}$ the probability distribution over \mathcal{C} defined by

$$\text{freq}_{\mathbf{c}}(\tilde{c}) = \frac{|\{i | c_i = \tilde{c}\}|}{n} \quad (9.11)$$

for $\tilde{c} \in \mathcal{C}$. We define a set $\hat{\Omega}$ that includes all the probability distributions approved by the protocol, i.e., the desired frequencies $\text{freq}_{\mathbf{c}}$ for which the protocol does not abort. Then, we can write the event of not aborting in terms of the desired frequencies:

⁸The value of the functions at points p for which $\Sigma(p)$ is the empty set is unconstrained and can be chosen freely (while keeping the function differentiable and convex).

$$\Omega = \left\{ \mathbf{c} : \text{freq}_{\mathbf{c}} \in \hat{\Omega} \right\} \subseteq \mathcal{C}^n .$$

Note that one can also start by choosing the set Ω describing the event of not aborting the protocol. Then, $\hat{\Omega}$ can be chosen to be any set fulfilling⁹

$$\left\{ \text{freq}_{\mathbf{c}} : \mathbf{c} \in \Omega \right\} \subseteq \hat{\Omega} .$$

Focusing on permutation invariant sets Ω , in the sense that $\mathbf{c} \in \Omega$ if and only if $\pi(\mathbf{c}) \in \Omega$ for all permutations π of the n indices, defining $\hat{\Omega}$ via Ω is practically the same as defining Ω via $\hat{\Omega}$.

Let us present a simple example of the above definitions and sets. Let $\mathcal{C} = \{0, 1\}$ and consider, e.g.,

$$\mathbf{c} = 01101000110100111011 . \quad (9.12)$$

To write $\text{freq}_{\mathbf{c}}$ we count the number of zeros and ones in the above string and get, according to Eq. (9.11), the probability distribution over $\{0, 1\}$ defined by

$$\text{freq}_{\mathbf{c}}(0) = \frac{9}{20} \quad ; \quad \text{freq}_{\mathbf{c}}(1) = \frac{11}{20} .$$

We can now consider a protocol which *does not abort* whenever the observed statistics are such that the fraction of ones is greater than half. This leads to

$$\begin{aligned} \hat{\Omega} &= \left\{ p : p(1) > \frac{1}{2} \right\} \\ \Omega &= \left\{ \mathbf{c} : \text{freq}_{\mathbf{c}} \in \hat{\Omega} \right\} = \left\{ \mathbf{c} : \text{freq}_{\mathbf{c}}(1) > \frac{1}{2} \right\} \end{aligned}$$

and, in particular, for \mathbf{c} appearing in Eq. (9.12), $\mathbf{c} \in \Omega$.

9.2.3.2 The Theorem

We first give the formal statement of the EAT and then explain.

Theorem 9.3 (EAT) *Let $\mathcal{M}_i : R_{i-1} \rightarrow R_i O_i S_i C_i$ for $i \in [n]$ be EAT channels, ρ be the final state, Ω an event defined over \mathcal{C}^n , p_{Ω} the probability of Ω in ρ , and $\rho_{|\Omega}$ the final state conditioned on Ω . Let $\varepsilon \in (0, 1)$.*

⁹It will become clear from the statement of the EAT that one should choose a minimal convex set $\hat{\Omega}$ that includes the frequencies considered in Ω . It is perhaps instructive to observe that, for a finite n , Ω is a finite set; $\hat{\Omega}$, on the other hand, includes infinitely many probability distributions.

For $\hat{\Omega} = \{\text{freq}_c : c \in \Omega\}$ convex,¹⁰ f_{\min} a min-tradeoff function for $\{\mathcal{M}_i\}_{i \in [n]}$, and any $t \in \mathbb{R}$ such that $f_{\min}(\text{freq}_c) \geq t$ for any $\text{freq}_c \in \hat{\Omega}$,

$$H_{\min}^{\varepsilon}(\mathbf{O}|\mathbf{SE})_{\rho_{|\Omega}} > nt - \mu\sqrt{n}, \quad (9.13)$$

where

$$\mu = 2(\log(1 + 2d_O) + \lceil \|\nabla f_{\min}\|_{\infty} \rceil) \sqrt{1 - 2\log(\varepsilon \cdot p_{\Omega})}, \quad (9.14)$$

d_O the dimension of the systems O_i , and $\|\nabla f_{\min}\|_{\infty}$ is the infinity norm of the gradient of f_{\min} .

Similarly, for $\hat{\Omega} = \{\text{freq}_c : c \in \Omega\}$ convex, f_{\max} a max-tradeoff function and any $t \in \mathbb{R}$ such that $f_{\max}(\text{freq}_c) \leq t$ for any $\text{freq}_c \in \hat{\Omega}$,

$$H_{\max}^{\varepsilon}(\mathbf{O}|\mathbf{SE})_{\rho_{|\Omega}} < nt + \mu\sqrt{n}, \quad (9.15)$$

with

$$\mu = 2(\log(1 + 2d_O) + \lceil \|\nabla f_{\max}\|_{\infty} \rceil) \sqrt{1 - 2\log(\varepsilon \cdot p_{\Omega})}. \quad (9.16)$$

Let us parse the statement of the theorem while focusing on the smooth min-entropy for the moment. Equation (9.13) has exactly the form that we were aiming for: it gives a lower-bound on the conditional smooth min-entropy, evaluated on the state $\rho_{|\Omega}$ in the end of the protocol and conditioned on not aborting, where the first order term is linear in n and the second, describing finite statistic effects, scales like \sqrt{n} .

The constant t , governing the entropy rate $H_{\min}^{\varepsilon}(\mathbf{O}|\mathbf{SE})_{\rho_{|\Omega}}/n$ when $n \rightarrow \infty$, is defined via the min-tradeoff function in the following way. The min-tradeoff function f_{\min} assigns to each probability distribution p , or frequency, a number describing the minimal amount of conditional von Neumann entropy which is compatible with the probability distribution p (recall Definition 9.5). We now consider all frequencies freq_c (i.e., probability distributions) which are accepted by the protocol. The theorem asserts that t should be chosen as the minimal value of f_{\min} over this set of accepted frequencies. That is,¹¹

¹⁰We consider only convex sets $\hat{\Omega}$ (as was done in [3]). One can convince oneself that choosing a convex $\hat{\Omega}$ is the sensible thing to do. For example, a set $\hat{\Omega}$ including all frequencies, or probability distributions, for which $p(1) \in [a, b]$ for some constants $0 \leq a, b \leq 1$ is convex. If, nevertheless, one wishes to consider arbitrary sets Ω , which are not defined via a convex $\hat{\Omega}$, then this comes at the cost of considering only affine tradeoff functions (instead of convex/concave functions as in Definition 9.2); see [1] for the original claim. It is not clear that there are scenarios in which Ω cannot be defined with an underlying convex set $\hat{\Omega}$ and, at the same time, applying the EAT with adequate affine tradeoff functions does not result in a trivial statement. Hence, the convexity of $\hat{\Omega}$ should not be seen as a restriction.

¹¹The reader may be concerned that for finite n all frequencies belonging to $\text{freq}_c \in \hat{\Omega}$ actually lead to empty sets $\Sigma(\text{freq}_c)$, defined in Eq. (9.7) and hence t can be arbitrary. Note however that the tradeoff functions are defined over the set of *all* probability distributions, not only over the possible frequencies. Since the tradeoff functions must be differential convex/concave functions with a finite

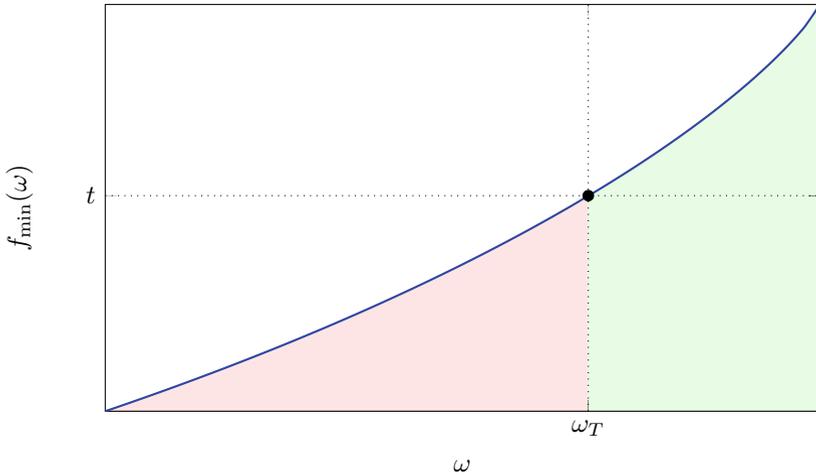


Fig. 9.3 First order term from the min-tradeoff function. We consider a protocol which does not abort if the fraction of games won is above ω_T . The value of t , appearing in Eq.(9.13), should be chosen to be the lowest value that the min-tradeoff function f_{\min} assigns to accepted winning probabilities, i.e., the black point in the plot

$$t = \inf \left\{ f_{\min}(\text{freq}_c) : \text{freq}_c \in \hat{\Omega} \right\} .$$

Since the min-tradeoff function is practically the worst-case conditional von Neumann entropy, we get that, asymptotically, $H_{\min}^{\varepsilon}(\mathbf{O}|\mathbf{SE})_{\rho_{i\Omega}}/n$ is equal to the lowest von Neumann entropy of a single-round (in the sense defined by the min-tradeoff function) that is compatible with the statistics observed in the protocol.

As an example, consider a device-independent protocol that assigns $C_i = 1$ when the i 'th game is won and $C_i = 0$ otherwise and which does not abort as long as the fraction of games won ($\sum_i C_i/n$) is above some threshold ω_T . The min-tradeoff function is defined over probability distributions p over $\mathcal{C} = \{0, 1\}$. Thus, we can also think of it as a function over winning probabilities ω via the relation $p(1) = \omega$ and $p(0) = 1 - \omega$ (see Eq.(9.8)). Assuming that the min-tradeoff function is increasing with ω (as expected to be; see Lemma 5.3), the lowest value that f_{\min} assigns to accepted frequencies is $f_{\min}(\omega_T)$ and hence this should be the value of t ; this is illustrated in Fig.9.3.

The above discussion refers to the first order term in Eq.(9.13). Let us now briefly discuss the second order term. Similarly to the AEP, the second order term scales like \sqrt{n} , which is the optimal scaling. The constant μ , defined in Eq.(9.14), depends on the different parameters and constants. In particular, it depends on the dimension of the output systems O_i and the gradient of the tradeoff functions. To get a good second order term one should therefore choose the possible values that can be assigned to

gradient, the points in which the functions are constrained by Eqs.(9.5) and (9.6) also constrain the values of the functions at the points $\text{freq}_c \in \hat{\Omega}$.

the registers O_i and the tradeoff functions such that the quantities of interest can be bounded in a good manner. In particular, to control the gradient of the tradeoff function one can “cut” when the gradient becomes too large and linearise the function at that point. An example is given in Sect. 11.2.2. Improved second order terms for the EAT were derived in [2].

A last remark is with regards to the statement of the EAT for the smooth max-entropy. When interested in the smooth max-entropy, the registers describing the environment, i.e., E in Eqs. (9.4) and (9.15) as well as R' in Eq. (9.6), can be dropped. The fact that R' can be dropped from Eq. (9.6) was already noted in [1, Remark 4.2]; the reason is that for the calculation of the supremum one can always assume that the system on R' is in product with the rest of the systems. To see that E can be dropped from Eqs. (2.21) and (9.15) note that the EAT must hold for any initial state $\rho_{R_0 E}^{\text{in}}$ and, hence, in particular to a tensor product state $\rho_{R_0 E}^{\text{in}} = \rho_{R_0} \otimes \rho_E$, for which the conditional smooth max-entropy is maximal.

To conclude this chapter, we summarise the reasons for why the EAT can be seen as an extension of the AEP for non-IID processes:

1. Similarly to the AEP, the EAT tells us that for large enough n the smooth entropies are equal to the *von Neumann entropy of a single-round times the number of rounds*.
2. The observed frequencies from the entire process are used when calculating the entropy accumulated in a single step of the process, *as if* all steps contribute equally and independently of each other. This is analogous to the analysis done in the IID case, as we saw in Sect. 7.3.
3. The asymptotic bounds on the smooth entropies derived using the EAT are equal to those derived using the AEP under the IID assumption. As a result, the bounds resulting from the application of the EAT are tight (assuming that the constructed tradeoff functions are tight). The second order terms of both theorems are also similar, as the scaling of both is \sqrt{n} .

All the above justifies the use of the EAT as a “reduction to IID” technique.

References

1. Dupuis F, Fawzi O, Renner R (2016) Entropy accumulation. [arXiv:1607.01796](https://arxiv.org/abs/1607.01796)
2. Dupuis F, Fawzi O (2019) Entropy accumulation with improved second-order term. *IEEE Trans Inf Theory* 65(11):7596–7612
3. Arnon-Friedman R, Dupuis F, Fawzi O, Renner R, Vidick T (2018) Practical device-independent quantum cryptography via entropy accumulation. *Nat Commun* 9(1):459
4. Arnon-Friedman R, Bancal J-D (2019) Device-independent certification of one-shot distillable entanglement. *New J Phys* 21(3):033010
5. Renner R (2008) Security of quantum key distribution. *Int J Quantum Inf* 6(01):1–127
6. Tomamichel M, Hayashi M (2013) A hierarchy of information quantities for finite block length analysis of quantum tasks. *IEEE Trans Inf Theory* 59(11):7693–7710

Chapter 10

Showcase: Non-signalling Parallel Repetition



In this chapter we consider the showcase of non-signalling parallel repetition, introduced in Sect. 4.1, and show how threshold theorems derived under the IID assumption can be extended to threshold theorems for general strategies, using a *reduction to IID*.

We focus on the case of non-signalling players.¹ That is, the *only* restriction on the players is that they are not allowed to communicate. Considering the non-signalling case is interesting for several reasons. A first reason is to minimise the set of assumptions to the mere necessary. Minimising the set of assumptions can be useful in cryptography when one wishes to get the strongest result possible, i.e., one where the attack strategies of malicious parties are only restricted minimally (as in [1–3] for example). In theoretical physics, non-signalling correlations enable the study of generalised theories possibly beyond quantum theory. It is also important to mention that, due to their linearity, the non-signalling constraints are often easier to analyse than the quantum or the classical constraints. Therefore, even if additional constraints hold, focusing on the non-signalling ones serves as a way to get first insights into a given problem.

Our theorem deals with complete-support games; these are game in which the distribution Q_{XY} over the questions has complete support, i.e., for all x, y , $Q_{XY}(x, y) > 0$.² The main result presented in this chapter can be informally stated as follows.³

¹Most steps of our proof can be used as is when considering classical and quantum players as well. There is one lemma, however, which we do not know how to modify so it can capture the classical and quantum case. We explain the difficulty later on.

²When considering games with only two players, the requirement for complete support can be dropped but, even though we focus on two-player games, we do not discuss this here; see [4] for the details.

³See Theorem 10.11 for the formal statement.

Theorem 10.1 (Informal) *For any complete-support game, a threshold theorem for general non-signalling strategies follows from a threshold theorem for non-signalling IID strategies. Furthermore, given a game with optimal non-signalling winning probability $1 - \alpha$, the resulting threshold theorem states that, for any $\beta > 0$, the probability to win more than a fraction $1 - \alpha + \beta$ of n games is exponentially small in $n\beta^2$, as in the IID case.*

The result previously appeared in [4]. We remark that while the non-signalling threshold theorem of [5] was known prior to [4], the dependence on β did not match that of the IID case, which is optimal (as follows from the optimal formulation of the Chernoff bound). Following [4], the work of [6] showed that parallel repetition does *not* hold for general games without complete-support in the non-signalling case. Thus, Theorem 10.1 is as general as it gets.

Proving parallel repetition via a reduction to IID has the advantage that, as in the IID case, the proof is oblivious to the number of players and structure of the considered game. This leads to a general theorem applicable to all games (with complete support) that is, arguably, simpler than other proof techniques.

When considering a strategy for the repeated game there is one type of symmetry which one can take advantage of—since the repeated game is permutation invariant the same symmetry can be assumed to hold for the optimal strategies, without loss of generality. Permutation-invariant strategies are strategies which are indifferent to the ordering of the questions given by the referee. That is, the probability of answering a specific set of questions correctly does not depend on the ordering of the questions (see Sect. 10.3 below for the formal definitions). Once we restrict our attention to permutation-invariant strategies, de Finetti theorems presents themselves as a natural tool to leverage for the analysis. Indeed, our proof builds on the de Finetti reduction discusses in Chap. 8, which acts as a reduction to IID in our analysis of parallel repetition.

The chapter is arranged as follows. In Sect. 10.1 we explain the main challenge when proving parallel repetition and threshold theorems using techniques employed by works predating [4] and why de Finetti theorems were not used in the context of parallel repetition before [4]. In Sect. 10.2 we give the main technical statements needed to prove our non-signalling threshold theorem. The different observations and lemmas of Sect. 10.2 may be of independent interest when analysing non-signalling parallel boxes and therefore we perform the analysis without referring to multi-player games. The explicit threshold theorem and its proof are given in Sect. 10.3. As in the rest of the thesis, we focus on the case of two parties for simplicity; we refer to [4] for the proofs in the case of more than two players as well as a couple of extensions of our theorem to games without complete support.

10.1 Main Challenge and Goal

The main difficulty in proving a parallel repetition result comes from the, almost arbitrary, correlations between the different questions-answers pairs in the players' strategy for the repeated $G_{1-\alpha+\beta}^n$: as the players get all the n questions together they can answer them in a correlated way. In most of the known parallel repetition results (e.g., [5, 7–9]) the main idea of the proof is to bound the winning probability for some of the questions, *conditioned* on winning the game in several other coordinates. However, as the strategy itself introduces correlations between the different tuples of questions, a large amount of technical work is devoted to dealing with the effect of conditioning on the event of winning the previous questions.

As mentioned above (and formally stated in Sect. 10.3), due to the permutation invariance of $G_{1-\alpha+\beta}^n$ one can study only permutation invariant strategies without loss of generality. Once we restrict our attention to permutation-invariant strategies, de Finetti theorems seem like a natural tool to leverage for the analysis. In the context of games and strategies, de Finetti theorems suggest one may be able to reduce the analysis of general permutation-invariant strategies to the analysis of a de Finetti strategy, i.e., a convex combination of IID strategies; recall Chap. 8. As presented in Sect. 7.3.1, the behaviour of IID strategies is trivial under parallel repetition. Hence, a reduction to IID using a de Finetti-type theorem could significantly simplify the analysis of parallel repetition theorems and threshold theorems.

Yet, de Finetti theorems were not used in this context prior to [4], and for a good reason. The many versions of quantum de Finetti theorems (e.g., [10, 11]) could not have been used as they depend on the dimension of the underlying quantum strategies, while in the quantum multi-player game setting one does not wish to restrict the dimension. Non-signalling de Finetti theorems, as in [12, 13], were also not applicable for non-signalling parallel repetition theorems, as they restrict almost completely the type of allowed correlations in the strategies for the repeated game by assuming very strict non-signalling constraints between the different repetitions, i.e., between the different questions-answers pairs.

In the proof presented in the next sections, we use the de Finetti reduction presented in Chap. 8, which imposes no assumptions at all regarding the structure of the strategies (apart from permutation invariance), and is therefore applicable in the context of parallel repetition. This allows us to devise a proof technique which is completely different from the proofs of parallel repetition results predating [4].⁴ In particular, at least in the non-signalling case presented here, the conditioning problem described above disappears completely and the number of players does not play a role in the proof structure.

As explained in Sect. 8.2, the de Finetti strategy that one ought to consider when using our de Finetti reduction assigns some weight to *signalling* IID strategies.⁵ Most of the effort is therefore directed to, informally, showing that when starting with a

⁴Following [4, 14] presented another, conceptually similar but technically different, proof of non-signalling parallel repetition based on de Finetti reductions.

⁵As discussed in Sect. 8.4, this is inevitable.

permutation-invariant non-signalling strategy the de Finetti strategy must assign only a small weight to signalling IID strategies. Formally, a similar in spirit but somewhat different statement is proven; see Theorem 10.2 below. Without further ado, let us get into the proof of the exact statements in the following section.

10.2 Approximately Non-signalling Marginals

Consider a parallel box $P_{AB|XY}$ (as defined in Sect. 6.1) and some complete-support distribution Q_{XY} , i.e., for all $\tilde{x} \in \mathcal{X}$ and $\tilde{y} \in \mathcal{Y}$, $Q(\tilde{x}, \tilde{y}) \neq 0$. Sample \mathbf{x} , \mathbf{y} , \mathbf{a} , and \mathbf{b} according to $Q_{XY}^{\otimes n} P_{AB|XY}$. We assume in this section that all possible inputs \tilde{x} , \tilde{y} appear in the observed data \mathbf{x} and \mathbf{y} (that is, there exists $i \in [n]$ for which $(x_i, y_i) = (\tilde{x}, \tilde{y})$). For a complete-support distribution Q_{XY} , the probability that this is *not* the case is exponentially small in n and we will account for it later.

Next, let $O_{ABXY}^{\text{freq}(a,b,x,y)}$ be the distribution derived from the frequencies in the observed data via

$$O_{ABXY}^{\text{freq}(a,b,x,y)}(\tilde{a}\tilde{b}\tilde{x}\tilde{y}) = \frac{\left| \left\{ i : (a_i, b_i, x_i, y_i) = (\tilde{a}, \tilde{b}, \tilde{x}, \tilde{y}) \right\} \right|}{n}$$

and define

$$O_{AB|XY}^{\text{freq}(a,b,x,y)} = \frac{O_{ABXY}^{\text{freq}(a,b,x,y)}}{Q_{XY}}. \quad (10.1)$$

Without the complete-support requirement on Q_{XY} it does not even make sense to talk about a fully defined $O_{AB|XY}^{\text{freq}(a,b,x,y)}$, i.e., a conditional probability distribution which is defined for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Indeed, $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ can only be defined for $x \in \text{Supp}(\mathcal{X})$ and $y \in \text{Supp}(\mathcal{Y})$ due to the estimation process (at least when assuming that all inputs appear in the observed data, which happens with high probability for large enough n). If one is willing to consider conditional probability distributions which are allowed to not assign values to certain inputs then $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ regains its meaning. In the context of non-signalling boxes, these conditional probability distributions were termed “sub-non-signalling” boxes in [14]; sub-non-signalling boxes fulfil the subset of the non-signalling conditions which apply for the defined inputs. In the case of two parties, it is known that there is always a way to “complete” a sub-non-signalling box to a non-signalling box, defined over all inputs [14, 15].

The current section deals with the following question: given that we start with a non-signalling box $P_{AB|XY}$, what is the probability that the *single-round box* $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ is signalling? Under the IID assumption, i.e., when $P_{AB|XY} = O_{AB|XY}^{\otimes n}$, this question is natural and can be easily answered. In that case, $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ can simply be seen as an estimation of the “real box”, or marginal, $O_{AB|XY}$. In particular, according to Sanov’s theorem (Lemma 2.2), as $n \rightarrow \infty$, we have $O_{AB|XY}^{\text{freq}(a,b,x,y)} =$

$O_{AB|XY}$ almost surely. Thus, for a non-signalling IID box $P_{AB|XY}$, $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ must be non-signalling as $n \rightarrow \infty$.

Our goal is to show that roughly the same is true for permutation invariant non-signalling parallel boxes $P_{AB|XY}$ when boxes such as $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ take the role of the marginals, which are not properly defined for parallel boxes. The theorem can be stated informally as follows⁶:

Theorem 10.2 (Informal) *Let $P_{AB|XY}$ be a permutation invariant non-signalling parallel box and $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ be the single-round box defined via the observed data sampled using $Q_{XY}^{\otimes n} P_{AB|XY}$, as in Eq. (10.1). Then, for sufficiently large n , $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ is close to a non-signalling single-round box with high probability. This also implies that the observed data can be seen as if, with high probability, it was sampled using an IID box $O_{AB|XY}^{\otimes n}$ with $O_{AB|XY}$ close to a non-signalling single-round box.*

To prove the theorem we utilise the concept of a test, discussed in Sect. 8.3.2. Roughly speaking, we define a *signalling test* \mathcal{T} , interacting with a parallel box, which *accepts* whenever the box $O_{AB|XY}^{\text{freq}(a,b,x,y)}$ is highly signalling and *rejects* whenever the box is close to a non-signalling box. With the aid of the variant of the de Finetti reduction phrased as Theorem 8.11 and a rather simple signalling game (defined in Sect. 10.2.3) we prove that the probability that the test accepts when interacting with a permutation invariant non-signalling parallel box is small. We follow this proof idea in the succeeding sections.

10.2.1 Single-Round Boxes from Frequencies

To ease notation we denote $\text{data} = a, b, x, y$ when it is clear from the context which a, b, x, y are considered. Every observed data is split into two non-overlapping parts, data_1 and data_2 . Specifically, let⁷

$$\begin{aligned} \text{data}_1 &= a_1, \dots, a_{\frac{n}{2}}, b_1, \dots, b_{\frac{n}{2}}, x_1, \dots, x_{\frac{n}{2}}, y_1, \dots, y_{\frac{n}{2}}, \\ \text{data}_2 &= a_{\frac{n}{2}+1}, \dots, a_n, b_{\frac{n}{2}+1}, \dots, b_n, x_{\frac{n}{2}+1}, \dots, x_n, y_{\frac{n}{2}+1}, \dots, y_n. \end{aligned} \tag{10.2}$$

data (and hence also data_1 and data_2) is sampled according to $Q_{XY}^{\otimes n} P_{AB|XY}$, where $P_{AB|XY}$ is a non-signalling permutation invariant parallel box. Note the following:

1. $P_{AB|XY}$ may be signalling between the different rounds $i \in [n]$ (i.e., for a given party, the output of one round may depend on the input of other rounds). Therefore, even though data_1 and data_2 are each defined only by part of the observed data, they may depend on the entire data.

⁶For the formal statement see Theorem 10.10.

⁷For simplicity we assume n is even; otherwise replace $n/2$ by $\lceil n/2 \rceil$ and modify everything else accordingly.

2. Due to permutation invariance, it does not matter which indices $i \in [n]$ belong to each part of the data. We could as well define data_1 and data_2 by splitting the data according to whether the index i is even or odd (for example). For any partition of the data, data_1 and data_2 are distributed in the same way. Hence, the choice of partition made in Eq. (10.2) is arbitrary and all other choices give rise to the same results.

We define two single-round boxes from the observed data, similarly to what was done in Eq. (10.1):

$$\begin{aligned} O_{ABXY}^{\text{freq}(\text{data}_1)}(\tilde{a}\tilde{b}\tilde{x}\tilde{y}) &= \frac{\left| \left\{ i \in [n/2] : (a_i, b_i, x_i, y_i) = (\tilde{a}, \tilde{b}, \tilde{x}, \tilde{y}) \right\} \right|}{n/2}, \\ O_{ABXY}^{\text{freq}(\text{data}_2)}(\tilde{a}\tilde{b}\tilde{x}\tilde{y}) &= \frac{\left| \left\{ i \in \{n/2 + 1, \dots, n\} : (a_i, b_i, x_i, y_i) = (\tilde{a}, \tilde{b}, \tilde{x}, \tilde{y}) \right\} \right|}{n/2} \end{aligned} \quad (10.3)$$

and, for $t \in \{0, 1\}$,

$$O_{AB|XY}^{\text{freq}(\text{data}_t)} = \frac{O_{ABXY}^{\text{freq}(\text{data}_t)}}{Q_{XY}}. \quad (10.4)$$

As mentioned above, for $O_{AB|XY}^{\text{freq}(\text{data}_1)}$ and $O_{AB|XY}^{\text{freq}(\text{data}_2)}$ to be defined for all inputs we assume that all inputs (x, y) appear in both data_1 and data_2 .

10.2.2 Signalling Test

Below we consider distributions $O_{ABXY} = Q_{XY}O_{AB|XY}$. One can then consider different marginals of O_{ABXY} . For example, O_{BY} is simply defined by $O_{BY}(b, y) = \sum_{a,x} O_{ABXY}(a, b, x, y)$. Note that the marginals of O_{ABXY} are all well-defined even if $O_{AB|XY}$ itself is signalling.

We wish to define a signalling test. To this end, let us first define a signalling *measure* over single-round boxes:

Definition 10.3 Let $O_{AB|XY}$ be a single-round box defined over $\mathcal{A}, \mathcal{B}, \mathcal{X}$, and \mathcal{Y} , Q_{XY} a distribution over the inputs of the single-round box, and $O_{ABXY} = Q_{XY}O_{AB|XY}$. The amount of signalling from Alice to Bob using the inputs (x, y) and Bob's output b is given by

$$\text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}) = O_{BY}(b, y) [O_{X|BY}(x|b, y) - Q_{X|Y}(x|y)].$$

Similarly, the amount of signalling from Bob to Alice using the inputs (x, y) , distributed according to Q_{XY} , and Alice's output a , distributed according to $O_{A|X=x, Y=y}$, is given by

$$\text{Sig}^{(B \rightarrow A, x, y, a)}(\mathcal{O}_{AB|XY}) = \mathcal{O}_{AX}(a, x) [\mathcal{O}_{Y|AX}(y|a, x) - \mathcal{Q}_{Y|X}(y|x)] .$$

The box $\mathcal{O}_{AB|XY}$ is non-signalling if and only if

$$\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}) = \text{Sig}^{(B \rightarrow A, x, y, a)}(\mathcal{O}_{AB|XY}) = 0 . \quad (10.5)$$

To see that the above definition makes sense as a signalling measure first notice that, when positive, $\mathcal{O}_{X|BY}(x|b, y) - \mathcal{Q}_{X|Y}(x|y)$ can be understood as quantifying Bob's advantage in guessing Alice's input x when observing b , compared to his prior information $\mathcal{Q}_{X|Y}(x|y)$ about her input. For a uniform distribution over $\mathcal{X} \times \mathcal{Y}$, $\mathcal{Q}_{X|Y}(x|y) = 0$ for all x and y . Then, the non-signalling requirement means that Bob cannot infer Alice's input from his output (as otherwise Alice could signal Bob), that is, $\mathcal{O}_{X|BY}(x|b, y) = 0$ as well. The above is a generalisation of this requirement to non-uniform distributions \mathcal{Q}_{XY} . On the more technical level—the non-signalling conditions (Definition 3.1) can be equivalently written as

$$\begin{aligned} \forall b, x, y \quad \sum_{\tilde{a}} \mathcal{O}_{AB|XY}(\tilde{a}, b|x, y) &= \sum_{\tilde{x}} \mathcal{Q}_{X|Y}(\tilde{x}|y) \sum_{\tilde{a}} \mathcal{O}_{AB|XY}(\tilde{a}, b|\tilde{x}, y) ; \\ \forall a, x, y \quad \sum_{\tilde{b}} \mathcal{O}_{AB|XY}(a, \tilde{b}|x, y) &= \sum_{\tilde{y}} \mathcal{Q}_{Y|X}(\tilde{y}|x) \sum_{\tilde{b}} \mathcal{O}_{AB|XY}(a, \tilde{b}|a, \tilde{y}) . \end{aligned} \quad (10.6)$$

One can verify that, for complete support \mathcal{Q}_{XY} , these conditions are equivalent to Eq. (10.5).

All statements proven below regarding our signalling measure hold for signalling in both directions, i.e., from Alice to Bob and from Bob to Alice. We present all the statements and proofs in terms of signalling from Alice to Bob; to derive the same statements for signalling from Bob to Alice one can simply replace the parties (their inputs and outputs) with one another.

We use the following definition to measure the distance between two single-round boxes.⁸

Definition 10.4 The distance between $\mathcal{K}_{AB|XY}$ and $\mathcal{R}_{AB|XY}$ is defined as

$$|\mathcal{K}_{AB|XY} - \mathcal{R}_{AB|XY}|_1 = \mathbb{E}_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} |\mathcal{K}_{AB|XY}(a, b|x, y) - \mathcal{R}_{AB|XY}(a, b|x, y)| .$$

The following lemma shows that our measure of signalling is continuous. That is, if two strategies are close to one another according to Definition 10.4 then their signalling values are also close. The proof is given in Appendix B.1.

Lemma 10.5 Let $\mathcal{O}_{AB|XY}^1$ and $\mathcal{O}_{AB|XY}^2$ be two single-round boxes such that

$$|\mathcal{O}_{AB|XY}^1 - \mathcal{O}_{AB|XY}^2|_1 \leq \epsilon .$$

⁸ More commonly in the literature, one considers a definition in which $\mathbb{E}_{(x,y)}$ is replaced by $\max_{x,y}$. We use Definition 10.4 since it allows us to apply Sanov's theorem later on.

Then, for all a, b, x , and y ,

$$\left| \text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}^1) - \text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}^2) \right| \leq 2\epsilon$$

We can now define our *signalling test*. The test interacts with the parallel box $\mathcal{P}_{AB|XY}$ by sampling **data** according to $\mathcal{Q}_{XY}^{\otimes n} \mathcal{P}_{AB|XY}$ and then checking whether $\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}$ is sufficiently signalling. Formally:

Definition 10.6 Let $\zeta, \epsilon > 0$ be parameters satisfying $\zeta \geq 7\epsilon$. For any x, y , and b , a signalling test is defined by⁹

$$\mathcal{T}^{(A \rightarrow B, x, y, b)}(\mathcal{P}_{AB|XY}) = \begin{cases} 1 & \text{if } \text{Sig}^{(A \rightarrow B, x, y, b)}\left(\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}\right) \geq \zeta - 2\epsilon \\ 0 & \text{otherwise,} \end{cases} \quad (10.7)$$

where $\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}$ is defined as in Eq. (10.4).

Let T denote the event that the signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ passes. The probability of the test passing when interacting with $\mathcal{P}_{AB|XY}$ is given by

$$\Pr_{\text{data} \sim \mathcal{P}_{ABXY}} [T] = \sum_{x, y} \mathcal{Q}_{XY}^{\otimes n}(x, y) \sum_{\substack{a, b: \\ \text{Sig}^{(A \rightarrow B, x, y, b)}\left(\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}\right) \geq \zeta - 2\epsilon}} \mathcal{P}_{AB|XY}(a, b|x, y).$$

The signalling test above is defined with $\text{Sig}^{(A \rightarrow B, x, y, b)}\left(\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}\right)$, rather than its absolute value, since this will be the only case relevant for our analysis; see Appendix B.2.

When considering IID boxes $\mathcal{O}_{AB|XY}^{\otimes n}$, the signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ is reliable— if $\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}) \geq \zeta$ the test will detect it with high probability, i.e. the test will accept with high probability, and if $\mathcal{O}_{AB|XY}$ is non-signalling then the test will reject with high probability. It follows, in particular, that if signalling is detected by the test in $\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)}$, $\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_1)}$ is also signalling with high probability. This holds also when considering the de Finetti box as in Definition 8.2.

To make the statement precise, let us define two sets of single-round boxes for every signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$. The first set is given by

$$\begin{aligned} \sigma^{(A \rightarrow B, x, y, b)} &= \left\{ \mathcal{O}_{AB|XY} : \forall \bar{\mathcal{O}}_{AB|XY} \text{ s.t. } |\mathcal{O}_{AB|XY} - \bar{\mathcal{O}}_{AB|XY}|_1 \leq \epsilon \right. \\ &\quad \left. \Rightarrow \text{Sig}^{(A \rightarrow B, x, y, b)}(\bar{\mathcal{O}}_{AB|XY}) \geq \zeta \right\}. \end{aligned}$$

Using the continuity of the signalling measure, Lemma 10.5, we observe that

$$\mathcal{O}_{AB|XY} \notin \sigma^{(A \rightarrow B, x, y, b)} \Rightarrow \text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}) \leq \zeta + 2\epsilon. \quad (10.8)$$

⁹If **data**₁ does not include an index in which the inputs are (x, y) then the test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ rejects by definition (recall Definition 10.3).

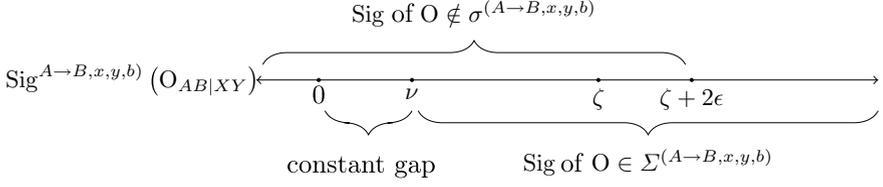


Fig. 10.1 Visualisation of the sets $\sigma^{(A \rightarrow B, x, y, b)}$ and $\Sigma^{(A \rightarrow B, x, y, b)}$

The second set is defined to be

$$\Sigma^{(A \rightarrow B, x, y, b)} = \left\{ O_{AB|XY} : \exists \bar{O}_{AB|XY} \text{ s.t. } |O_{AB|XY} - \bar{O}_{AB|XY}|_1 \leq \epsilon \right. \\ \left. \wedge \Pr_{\text{data} \sim \bar{O}_{AB|XY}^{\otimes n}} [T] > \delta \right\},$$

where $\delta = \delta\left(\frac{n}{2}, \epsilon\right) = \left(\frac{n}{2} + 1\right)^{|\mathcal{A}||\mathcal{B}||\mathcal{X}||\mathcal{Y}|-1} e^{-n\epsilon^2/4}$. Since the signalling test is reliable when acting on IID boxes, one can easily show that

$$O_{AB|XY} \in \Sigma^{(A \rightarrow B, x, y, b)} \Rightarrow \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}) > \nu \quad (10.9)$$

for any $0 < \nu < \zeta - 6\epsilon$. This is stated and proven as Lemma B.1 in Appendix B.1. The sets and the relevant constants are illustrated in Fig. 10.1.

We use below the following notation:

- in^σ denotes the event that $O_{AB|XY}^{\text{freq}(\text{data}_1)} \in \sigma^{(A \rightarrow B, x, y, b)}$.
- in^Σ denotes the event that $O_{AB|XY}^{\text{freq}(\text{data}_1)} \in \Sigma^{(A \rightarrow B, x, y, b)}$.
- “For all signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)} \dots$ ” should be understood as “for all x, y , and b , defining a signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}, \dots$ ” and similarly for other quantifiers.

Furthermore, to avoid confusion, we explicitly denote the probability distributions on which we evaluate the probability of the above events.

As shown in Appendix B.1, the following lemma holds for a de Finetti box:

Lemma 10.7 *Let $\tau_{ABXY} = Q_{XY}^{\otimes n} \tau_{AB|XY}$, where $\tau_{AB|XY}$ is a de Finetti box. For every signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$,*

1. $\Pr_{\text{data} \sim \tau_{ABXY}} [-in^\Sigma \wedge T] \leq \delta$
2. $\Pr_{\text{data} \sim \tau_{ABXY}} [in^\sigma \wedge \neg T] \leq \delta$,

where $\delta = \delta\left(\frac{n}{2}, \epsilon\right) = \left(\frac{n}{2} + 1\right)^{|\mathcal{A}||\mathcal{B}||\mathcal{X}||\mathcal{Y}|-1} e^{-n\epsilon^2/4}$.

A similar lemma can be proven for *permutation invariant parallel boxes*, using the de Finetti reduction of Theorem 8.3:

Lemma 10.8 *Given a parallel box $P_{AB|XY}$ let $P_{ABXY} = Q_{XY}^{\otimes n} P_{AB|XY}$. For every permutation-invariant box $P_{AB|XY}$ and every $\mathcal{T}^{(A \rightarrow B, x, y, b)}$:*

1. $\Pr_{\text{data} \sim P_{ABXY}} [-in^\Sigma \wedge T] \leq c\delta$

$$2. \Pr_{\text{data} \sim P_{ABXY}} [in^\sigma \wedge \neg T] \leq c\delta,$$

where $c = (n+1)^{|\mathcal{X}||\mathcal{Y}|(|\mathcal{A}||\mathcal{B}|-1)}$ and δ is as in Lemma 10.7.

Proof We prove both of the claims together. Denote the relevant event by $E(\text{data})$ and note that for both events we can write

$$\Pr_{\text{data} \sim P_{ABXY}} [E(\text{data}) = 1] = \sum_{\substack{\text{data} \\ E(\text{data})=1}} P_{ABXY}(\text{data}).$$

From Theorem 8.3 we get $P_{ABXY}(\text{data}) \leq c \cdot \tau_{ABXY}(\text{data})$ and therefore

$$\begin{aligned} \Pr_{\text{data} \sim P_{ABXY}} [E(\text{data}) = 1] &= \sum_{\substack{\text{data} \\ E(\text{data})=1}} P_{ABXY}(\text{data}) \\ &\leq c \cdot \sum_{\substack{\text{data} \\ E(\text{data})=1}} \tau_{ABXY}(\text{data}) \\ &= c \cdot \Pr_{\text{data} \sim \tau_{ABXY}} [E(\text{data}) = 1]. \end{aligned}$$

Combining this with Lemma 10.7 proves the lemma. \square

10.2.3 Guessing Game

The previous section discussed the relations between $O_{AB|XY}^{\text{freq}(\text{data}_1)}$ and $O_{AB|XY}^{\text{freq}(\text{data}_2)}$ in terms of the probabilities of certain events which depend on these averaged single-round boxes. All statements made so far were general, in the sense that they hold for any permutation-invariant parallel box $P_{AB|XY}$. In the current section we focus on permutation-invariant *non-signalling* parallel boxes $P_{AB|XY}$. Our goal is to show that for non-signalling boxes $P_{AB|XY}$, the averaged boxes $O_{AB|XY}^{\text{freq}(\text{data}_1)}$ and $O_{AB|XY}^{\text{freq}(\text{data}_2)}$ cannot be too signalling as we were set to prove.

To this end, we construct a guessing game for every signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$. In the game, a referee gives Alice and Bob $n/2$ questions, distributed according to $Q_{XY}^{\otimes n/2}$. Bob's goal is to output an index $j \in [n/2]$ for which $(x_j, y_j) = (x, y)$ (while Alice does not need to output anything). Alice and Bob are allowed to use any non-signalling box to win the game. Clearly, a non-signalling box should not allow Bob to learn anything about Alice's input from his outputs. Bob's best strategy is thus to guess an index j for which $y_j = y$. The probability that his guess is correct is $Q_{X|Y}(x|y)$. If Alice and Bob are able to win the game with higher probability than the used box must be signalling.¹⁰

¹⁰This motivates our signalling measure given in Definition 10.3.

The following lemma asserts that, for non-signalling $P_{AB|XY}$, conditioned on the signalling test detecting a lot of signalling in $O_{AB|XY}^{\text{freq}(\text{data}_2)}$, the probability that $O_{AB|XY}^{\text{freq}(\text{data}_1)}$ is highly signalling is bounded away from 1. Intuitively, we would have expected that if signalling is detected in data_2 then data_1 should exhibit signalling practically with certainty. The lemma (roughly) shows that, when starting with non-signalling boxes, this is not the case.

Before starting, we remind the reader that we assume in this section that all pairs of questions appear in data_1 and data_2 . For the lemmas and proofs below it is important to remember that all the probabilities are conditioned on data_1 . To ease notation we do not explicitly write it.

Lemma 10.9 *Let $\epsilon \in [0, 1]$ and n be such that*

$$\frac{n}{\ln(n)} > 20|\mathcal{X}||\mathcal{Y}||\mathcal{A}||\mathcal{B}|\frac{\ln(2/\epsilon)}{\epsilon^2}, \quad (10.10)$$

and $P_{AB|XY}$ a non-signalling parallel box. For any signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ denote by $P_{ABXY|\mathcal{T}=1}$ the probability distribution P_{ABXY} conditioned on the event $\mathcal{T}^{(A \rightarrow B, x, y, b)}(P_{AB|XY}) = 1$, whenever such a conditional probability distribution is defined. Then,

$$\Pr_{\text{data} \sim P_{ABXY|\mathcal{T}=1}}[in^\Sigma] < 1 - \sqrt{c\delta}, \quad (10.11)$$

for c and δ as in Lemma 10.8.

Proof We denote $\mathbf{x}_{\text{data}_1} = x_1, \dots, x_{n/2}$ and $\mathbf{y}_{\text{data}_1} = y_1, \dots, y_{n/2}$.

For every signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ and inputs for Bob $\mathbf{y}_{\text{data}_1}$ such that $\Pr_{\text{data} \sim P_{ABXY}}[T|\mathbf{y}_{\text{data}_1}] \neq 0$ we construct a guessing game. Our goal is to derive a contradiction by showing that if Eq. (10.11) is not true, then the guessing game can be won with probability higher than the optimal non-signalling winning probability.

The guessing game is as explained above. A referee gives Bob the inputs $\mathbf{y}_{\text{data}_1}$ and Alice gets $\mathbf{x}_{\text{data}_1}$ distributed according to $Q_{XY}(x|y)$. Bob's goal is to guess an index $j \in [n/2]$ such that $(x_j, y_j) = (x, y)$ (we assume that such exists).

If the parties share a non-signalling box $P_{AB|XY}$ then Bob's marginals are the same for all $\mathbf{x}_{\text{data}_1}$. Therefore, his outputs do not give him any information about the inputs that Alice got from the referee. The best non-signalling strategy for the guessing game is therefore to choose, uniformly at random, an index j for which $y_j = y$. The winning probability is then given by $W_{\text{ns}} = Q_{X|Y}(x|y) < 1$.¹¹

We now show that if the parties share $P_{AB|XY}$ for which

$$\Pr_{\text{data} \sim P_{ABXY|\mathcal{T}=1}}[in^\Sigma | \mathbf{y}_{\text{data}_1}] \geq 1 - \sqrt{c\delta} \quad (10.12)$$

then they can win the above guessing game with probability higher than the optimal non-signalling winning probability W_{ns} .

¹¹Note that while Bob's inputs, $\mathbf{y}_{\text{data}_1} = y_1, \dots, y_{n/2}$, are fixed in a specific instance of the guessing game, Alice's inputs are still distributed according to the prior $Q_{XY}(x|y)$.

The idea is as follows. The parties share many identical copies of $P_{AB|XY}$. They use the inputs given by the referee as x_{data_1} and y_{data_1} in all of the copies and choose, using shared randomness, the rest of the inputs, associated with data_2 , in each copy (i.e., there are different inputs for data_2 for each copy). They use the copies of $P_{AB|XY}$ with the described inputs. Bob then looks for the first copy of $P_{AB|XY}$ in which the event T holds—such a copy exists as long as¹² $\Pr_{\text{data} \sim P_{ABXY}} [T | y_{\text{data}_1}] \neq 0$; he can find it since he knows all the inputs in data_2 (as they were chosen using shared randomness).¹³ Alice does not need to know in which copy the test holds. Using the chosen copy, Bob chooses a random index $j \in [n/2]$ such that $y_j = y$ and $b_j = b$.

Let us show that, as long as $\Pr_{\text{data} \sim P_{ABXY}} [T | y_{\text{data}_1}] \neq 0$, this box achieves a winning probability which is higher than W_{ns} . For the chosen copy, the event T holds and hence data_1 can be seen as data distributed according to $n/2$ identical copies of $O_{AB|XY}^{\text{freq}(\text{data}_1)}$, which is with high probability in $\Sigma_{(i,b,x,y)}$ according to Eq. (10.12). Using Eq. (10.9) this implies

$$\Pr_{\text{data} \sim P_{ABXY} | T=1} \left[\text{Sig}^{(A \rightarrow B, b, x, y)} (O_{AB|XY}^{\text{freq}(\text{data}_1)}) > \nu | y_{\text{data}_1} \right] \geq 1 - \sqrt{c\delta}, \quad (10.13)$$

where $\nu > 0$ is any parameter satisfying $\nu < \zeta - 6\epsilon$ (recall Eq. (10.9)).

Using Definition 10.3 we know that if indeed

$$\text{Sig}^{(A \rightarrow B, b, x, y)} (O_{AB|XY}^{\text{freq}(\text{data}_1)}) > \nu$$

then $O_{BY}^{\text{freq}(\text{data}_1)}(b, y) > 0$ and

$$\begin{aligned} O_{X|BY}^{\text{freq}(\text{data}_1)}(x|b, y) &> \frac{\nu}{O_{BY}^{\text{freq}(\text{data}_1)}(b, y)} + Q_{X|Y}(x|y) \\ &= \frac{\nu}{O_{BY}^{\text{freq}(\text{data}_1)}(b, y)} + W_{\text{ns}}. \end{aligned} \quad (10.14)$$

That is, by choosing an index for which $b_j = b$ Bob increase the winning probability.

On the other hand, if $\text{Sig}^{(A \rightarrow B, b, x, y)} (O_{AB|XY}^{\text{freq}(\text{data}_1)}) \leq \nu$, which can happen with probability $\sqrt{c\delta}$, then Bob may decrease his winning probability. In the worst case the winning probability is 0. Therefore, for the chosen copy (for which the test passed) we get the following winning probability

$$W \geq (1 - \sqrt{c\delta}) \left(\frac{\nu}{O_{BY}^{\text{freq}(\text{data}_1)}(b, y)} + W_{\text{ns}} \right) + \sqrt{c\delta} \cdot 0. \quad (10.15)$$

¹²To see this note that since the box is non-signalling between Alice and Bob, Bob can check in which copy the test passes even before Alice uses her input. Therefore, the probability to pass the test is independent of Alice's inputs and hence must be non-zero for any of them.

¹³Recalling Definitions 10.3 and 10.6, we see that only $O_{BXY}^{\text{freq}(\text{data}_2)}$ is needed in order to check whether the signalling test passes or not. Thus, Bob indeed has all the relevant information and he can locally check whether the test passes or not.

Thus, $W > W_{\text{ns}}$ for

$$\nu > \frac{\sqrt{c\delta}}{1 - \sqrt{c\delta}} W_{\text{ns}} \geq \frac{\sqrt{c\delta}}{1 - \sqrt{c\delta}} W_{\text{ns}} \cdot \mathbf{O}_{BY}^{\text{freq}(\text{data}_1)}(b, y). \quad (10.16)$$

Using $W_{\text{ns}} \cdot \mathbf{O}_{BY}^{\text{freq}(\text{data}_1)}(b, y) \leq 1$ and $\sqrt{c\delta} \leq (n+1)^{|\mathcal{X}||\mathcal{Y}||\mathcal{A}||\mathcal{B}|} e^{-n\epsilon^2/8}$, we see that as long as $n/\ln(n) > 20|\mathcal{X}||\mathcal{Y}||\mathcal{A}||\mathcal{B}|\epsilon^{-2} \ln(2/\epsilon)$ we have

$$\frac{\sqrt{c\delta} W_{\text{ns}} \mathbf{O}_{BY}^{\text{freq}(\text{data}_1)}(b, y)}{1 - \sqrt{c\delta}} < \epsilon.$$

Assuming $\zeta \geq 7\epsilon$, there is a choice of ν that satisfies both Eq. (10.16) and the earlier condition that $\nu < \zeta - 6\epsilon$.

We get that Eq. (10.12) must not hold for any $\mathbf{y}_{\text{data}_1}$ and hence cannot hold also when we omit the conditioning on $\mathbf{y}_{\text{data}_1}$. The lemma therefore follows. \square

The bound given in Eq. (10.11) is weak for two reasons. First, data is distributed according to the conditional distribution $\mathbb{P}_{ABXY|\mathcal{T}=1}$ and not according to \mathbb{P}_{ABXY} itself. Second, it only tells us that $\Pr_{\text{data} \sim \mathbb{P}_{ABXY|\mathcal{T}=1}} [\mathbf{O}_{AB|XY}^{\text{freq}(\text{data}_1)} \notin \Sigma^{(A \rightarrow B, x, y, b)}] \geq \sqrt{c\delta}$, i.e., the probability that $\mathbf{O}_{AB|XY}^{\text{freq}(\text{data}_1)}$ has a small value of signalling is higher than $\sqrt{c\delta}$. We show how the statement can be amplified using Lemma 10.8, which utilised our de Finetti reduction.

Theorem 10.10 *Let $\mathbb{P}_{AB|XY}$ be a permutation-invariant non-signalling parallel box and n such that Eq. (10.16) is satisfied. Then for any signalling test $\mathcal{T}^{(A \rightarrow B, x, y, b)}$ such that $\mathbf{Q}_{XY}(x, y) \neq 0$ and $\mathbf{Q}_{X|Y}(x|y) \neq 1$ and conditioned on the event of all questions (x, y) appearing in data_1 and data_2 ,*

$$\Pr_{\text{data} \sim \mathbb{P}_{ABXY}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)} \left(\mathbf{O}_{AB|XY}^{\text{freq}(\text{data})} \right) > \zeta + 2\epsilon \right] \leq 4\sqrt{c\delta}. \quad (10.17)$$

Similarly, for any signalling test $\mathcal{T}^{(B \rightarrow A, x, y, a)}$,

$$\Pr_{\text{data} \sim \mathbb{P}_{ABXY}} \left[\text{Sig}^{(B \rightarrow A, x, y, b)} \left(\mathbf{O}_{AB|XY}^{\text{freq}(\text{data})} \right) > \zeta + 2\epsilon \right] \leq 4\sqrt{c\delta}. \quad (10.18)$$

Proof From Lemma 10.8 part 1 we get

$$\Pr_{\text{data} \sim \mathbb{P}_{ABXY}} [T] > \sqrt{c\delta} \Rightarrow \Pr_{\text{data} \sim \mathbb{P}_{ABXY|\mathcal{T}=1}} [\neg in^\Sigma] \leq \sqrt{c\delta}$$

and this can be rewritten as

$$\Pr_{\text{data} \sim \mathbb{P}_{ABXY}} [T] > \sqrt{c\delta} \Rightarrow \Pr_{\text{data} \sim \mathbb{P}_{ABXY|\mathcal{T}=1}} [in^\Sigma] \geq 1 - \sqrt{c\delta}.$$

According to Lemma 10.9, this implies

$$\Pr_{\text{data} \sim P_{ABXY}} [T] > \sqrt{c\delta} \Rightarrow P_{AB|XY} \text{ is signalling .}$$

Therefore it must be that

$$\Pr_{\text{data} \sim P_{ABXY}} [T] \leq \sqrt{c\delta} \quad (10.19)$$

or alternatively,

$$\Pr_{\text{data} \sim P_{ABXY}} [\neg T] \geq 1 - \sqrt{c\delta} \quad (10.20)$$

Next, combining Lemma 10.8 part 2 with Eq. (10.20) we get

$$\Pr_{\text{data} \sim P_{ABXY|T=0}} [in^\sigma] \leq \sqrt{c\delta} .$$

Using Eq. (10.19) we get

$$\Pr_{\text{data} \sim P_{ABXY}} [in^\sigma] \leq 2\sqrt{c\delta} .$$

Using the definition of the set $\sigma^{(A \rightarrow B, x, y, b)}$ and Eq. (10.8) we get that

$$\Pr_{\text{data} \sim P_{ABXY}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data}_1)} \right) > \zeta + 2\epsilon \right] \leq 2\sqrt{c\delta} .$$

Permutation invariance implies that data_1 and data_2 are distributed in the same way. Therefore, we also have

$$\Pr_{\text{data} \sim P_{ABXY}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data}_2)} \right) > \zeta + 2\epsilon \right] \leq 2\sqrt{c\delta} .$$

By definition,

$$O_{AB|XY}^{\text{freq}(\text{data})} = \frac{1}{2} O_{AB|XY}^{\text{freq}(\text{data}_1)} + \frac{1}{2} O_{AB|XY}^{\text{freq}(\text{data}_2)}$$

and, thus, using the linearity of the signalling measure, for any fixed observed data

$$\begin{aligned} \text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) &= \frac{1}{2} \text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data}_1)} \right) \\ &\quad + \frac{1}{2} \text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data}_2)} \right) . \end{aligned}$$

Equation (10.17) follows by combining the above equations. Switching Alice and Bob in all lemmas above, Eq. (10.18) follows in the exact same way. \square

Theorem 10.10 tells us that if $P_{AB|XY}$ is a permutation-invariant non-signalling parallel box then the probability that $O_{AB|XY}^{\text{freq}(\text{data})}$ is highly signalling (in any direction and using any inputs and outputs) is exponentially small in the number of games. De facto, this means that we can think of the observed data sampled from a non-signalling parallel box *as if* it came from an IID box defined via a single-round box that is approximately non-signalling, with high probability. This can be used to infer

properties of non-signalling parallel boxes. In Sect. 10.3 we show that Theorem 10.10 implies a threshold theorem almost directly.

Apart from the applications of Theorem 10.10, we see it as an abstract mathematical statement about the observed data produced by non-signalling parallel boxes. Deriving a similar statement for quantum boxes is also of interest (and, in particular, will imply a threshold theorem for all quantum games). The main difficulty in deriving a quantum analogue of Theorem 10.10 lies in finding a “non-quantumness” measure which, ideally, can be performed locally by one of the parties (as in our guessing game in the proof of Lemma 10.9).

10.3 Threshold Theorem

This section is devoted to deriving the following threshold theorem:

Theorem 10.11 *For any complete-support two-player game G whose optimal non-signalling winning probability is $w_{\text{ns}} = 1 - \alpha$, there exist $\mathcal{C}(G)$ such that for every $0 < \beta \leq \alpha$ and large enough n , the probability that non-signalling players win more than a fraction $1 - \alpha + \beta$ of the n questions in the threshold game $G_{1-\alpha+\beta}^n$ is at most $\exp[-\mathcal{C}(G)n\beta^2]$.*

That is, for sufficiently many repetitions the probability to win more than a fraction $1 - \alpha + \beta$ of the n games is exponentially small. A sufficient condition on the number of repetitions for the bound in the theorem to hold is stated in Eq. (10.10), and a choice of constants made around Eq. (10.28), for a more precise bound.

The proof builds on Theorem 10.10 and is rather simple. If $O_{AB|XY}^{\text{freq}(\text{data})}$ is not too signalling (for any signalling test), then its winning probability in a single game cannot be much higher than the winning probability of the optimal *non-signalling* strategy for G . Furthermore, the number of games won in any given observed data can be read directly from the winning probability of $O_{AB|XY}^{\text{freq}(\text{data})}$. Thus, by analysing $O_{AB|XY}^{\text{freq}(\text{data})}$ we are actually analysing the number of games won. The combination of these two observations gives the final theorem. We follow these steps below.

10.3.1 Winning Probability of Approximately Non-signalling Strategies

The linearity of the non-signalling conditions (Definition 3.1) and the winning probability in a game (Eq. (4.1)) allow us to phrase the optimisation problem of finding the optimal non-signalling winning probability in any game G as a linear program [16]. For complete-support game we can use the following linear program over the variables $O_{AB|XY}(a, b|x, y)$:

$$\max \sum_{a,b,x,y} Q_{XY}(xy)R(a,b,x,y)O_{AB|XY}(a,b|x,y) \quad (10.21a)$$

$$\text{s.t. } \text{Sig}^{(A \rightarrow B,x,y,b)}(O_{AB|XY}) = 0 \quad \forall x,y,b \quad (10.21b)$$

$$\text{Sig}^{(B \rightarrow A,x,y,b)}(O_{AB|XY}) = 0 \quad \forall x,y,a \quad (10.21c)$$

$$\sum_{a,b} O_{AB|XY}(a,b|x,y) = 1 \quad \forall x,y \quad (10.21d)$$

$$O_{AB|XY}(a,b|x,y) \geq 0 \quad \forall a,b,x,y \quad (10.21e)$$

The objective function, Eq. (10.21a), is exactly the winning probability in the game when using strategy $O_{AB|XY}$. Equations (10.21d) and (10.21e) are the normalisation and positivity constraints on the strategy $O_{AB|XY}$. In Eqs. (10.21b) and (10.21c) all the non-signalling constraints are listed, as follows for complete-support games from Eq. (10.5).¹⁴

Our goal is to upper-bound the winning probability of $O_{AB|XY}^{\text{freq}(\text{data})}$, which may be slightly signalling. The optimal winning probability of strategies which are slightly signalling can be written as a linear program similar to the one above, by relaxing the constraint in Eqs. (10.21b) and (10.21c) so that it allows for some signalling. Specifically, keeping in mind Eq. (10.17), we are interested in the following program:

$$\begin{aligned} \max \quad & \sum_{a,b,x,y} Q_{XY}(xy)R(a,b,x,y)O_{AB|XY}(a,b|x,y) \\ \text{s.t.} \quad & \text{Sig}^{(A \rightarrow B,x,y,b)}(O_{AB|XY}) \leq \zeta + 2\epsilon \quad \forall x,y,b \\ & \text{Sig}^{(B \rightarrow A,x,y,b)}(O_{AB|XY}) \leq \zeta + 2\epsilon \quad \forall x,y,a \quad (10.22) \\ & \sum_{a,b} O_{AB|XY}(a,b|x,y) = 1 \quad \forall x,y \\ & O_{AB|XY}(a,b|x,y) \geq 0 \quad \forall a,b,x,y \end{aligned}$$

Program (10.22) can be seen as a perturbation of Program (10.21). Their optimal values are therefore related to one another; the exact relation can be derived by studying how sensitive the objective function is to small modifications of the constraints. This process is called “sensitivity analysis of linear programs” and we follow it in Appendix B.2 for the programs of interest. As a result, we get that strategies $O_{AB|XY}$ with

$$\begin{aligned} \text{Sig}^{(A \rightarrow B,x,y,b)}(O_{AB|XY}) &\leq \zeta + 2\epsilon \\ \text{Sig}^{(B \rightarrow A,x,y,a)}(O_{AB|XY}) &\leq \zeta + 2\epsilon, \end{aligned}$$

for all x,y,a , and b achieve winning probability $w(O_{AB|XY})$ such that

$$w(O_{AB|XY}) \leq 1 - \alpha + (\zeta + 2\epsilon)d, \quad (10.23)$$

¹⁴Reference [4] includes an explanation of the implications of the linear program (10.21) to games with incomplete support.

where $1 - \alpha$ is the optimal winning probability of a *non-signalling* strategy, i.e., it is the solution of Program (10.21), and $d = |\mathcal{X}||\mathcal{Y}| (|\mathcal{A}| + |\mathcal{B}|)$ is the number of the non-signalling constraints in the linear programs above.

10.3.2 Final Result

The results of Sect. 10.2 are applicable when considering permutation invariant strategies $P_{AB|XY}$ (recall Definition 8.1). As the repeated game $G_{1-\alpha+\beta}^n$ is by itself permutation invariant we can restrict the strategies of the players to be permutation invariant without loss of generality.¹⁵ This is shown in the next lemma.

Lemma 10.12 *For every strategy $P_{AB|XY}$ for the repeated game $G_{1-\alpha+\beta}^n$ there exists a permutation-invariant strategy $\tilde{P}_{AB|XY}$ such that $w(P_{AB|XY}) = w(\tilde{P}_{AB|XY})$.*

Proof Given $P_{AB|XY}$ define its permutation-invariant version to be

$$\tilde{P}_{AB|XY} = \frac{1}{n!} \sum_{\pi} P_{AB|XY} \circ \pi .$$

The winning probability of the game is linear in the strategy, therefore we have

$$w(\tilde{P}_{AB|XY}) = w\left(\frac{1}{n!} \sum_{\pi} P_{AB|XY} \circ \pi\right) = \frac{1}{n!} \sum_{\pi} w(P_{AB|XY} \circ \pi) . \quad (10.24)$$

Since the questions in the repeated game are chosen in an IID manner and the winning condition is checked for each game separately, the winning probability is indifferent to the ordering of the questions-answers pairs. As π permutes the questions and answers together we have $w(P_{AB|XY} \circ \pi) = w(P_{AB|XY})$. Thus, we get $w(\tilde{P}_{AB|XY}) = w(P_{AB|XY})$. \square

We can now combine everything we have learned in the previous sections in order to derive the final results. As before, we denote by d the number of non-signalling conditions appearing in the linear programs above, i.e., $d = |\mathcal{X}||\mathcal{Y}| (|\mathcal{A}| + |\mathcal{B}|)$.

Lemma 10.13 *Let $w(G) = 1 - \alpha$ be the optimal winning probability of a non-signalling strategy in G . Let $0 < \beta \leq \alpha$ be some constant and n a sufficiently large integer such that Eq. (10.16) is satisfied. Then for any non-signalling strategy $P_{AB|XY}$ of the threshold game $G_{1-\alpha-\beta}^n$*

¹⁵This is not to say that all strategies are permutation invariant but only that the optimal strategy can be assumed to be permutation invariant. It is perhaps interesting to note that, more commonly, the optimal strategies are taken to be, without loss of generality, deterministic in proofs of classical parallel repetition and pure in proofs of quantum parallel repetition. Here we are choosing to focus on permutation invariant strategies instead.

$$\Pr_{\text{data} \sim P_{ABXY}} \left[w \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) > 1 - \alpha + \beta \right] \leq 6d\sqrt{c\delta}.$$

Proof We denote the event of all inputs appearing in data_1 and data_2 by aid . Furthermore, let $\zeta, \epsilon > 0$ be such that $d(\zeta + 2\epsilon) \leq \beta$, $\epsilon \leq \min_{x,y} Q_{XY}(xy)$ and $7\epsilon \leq \zeta \leq 1$.

If all questions (x, y) appear at least once in data_1 and data_2 , i.e., the event aid holds, then we can use Eq. (10.23) in combination with Theorem 10.10 and get

$$\begin{aligned} \Pr_{\text{data} \sim P_{ABXY}} \left[w \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) > 1 - \alpha + \beta \mid \text{aid} \right] \\ \leq \Pr_{\text{data} \sim P_{ABXY}} \left[\exists a, b, x, y \text{ s.t. } \text{Sig}^{(A \rightarrow B, x, y, b)} \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) \leq \zeta + 2\epsilon \right. \\ \left. \text{or } \text{Sig}^{(B \rightarrow A, x, y, a)} \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) \leq \zeta + 2\epsilon \mid \text{aid} \right] \\ \leq d \cdot 4\sqrt{c\delta}. \end{aligned}$$

The probability that the event aid does not hold is upper bounded by

$$\begin{aligned} 2|\mathcal{X}||\mathcal{Y}| \left(1 - \min_{x,y} Q_{XY}(x, y) \right)^{n/2} &\leq 2|\mathcal{X}||\mathcal{Y}| e^{-\min_{x,y} Q_{XY}(x,y)n/2} \\ &\leq 2|\mathcal{X}||\mathcal{Y}| e^{-\epsilon n/2} \\ &\leq 2d\delta \end{aligned}$$

and therefore all together we have

$$\Pr_{\text{data} \sim P_{ABXY}} \left[w \left(O_{AB|XY}^{\text{freq}(\text{data})} \right) > 1 - \alpha + \beta \right] \leq 6d\sqrt{c\delta}.$$

□

Our threshold theorem, Theorem 10.11, follows from Lemma 10.13:

Proof (Proof of Theorem 10.11) Let f denote the fraction of coordinates in which the players win the game in the observed data. Note that f is equals exactly $w \left(O_{AB|XY}^{\text{freq}(\text{data})} \right)$ by the definition of $O_{AB|XY}^{\text{freq}(\text{data})}$. Lemma 10.13 therefore implies

$$\Pr_{\text{data} \sim P_{ABXY}} [f > 1 - \alpha + \beta] \leq 6d\sqrt{c\delta}. \quad (10.25)$$

Plugging the values of the parameters d , c , and δ , we see that Eq. (10.25) can be written, for an appropriately defined $\hat{C}(G, n)$, as

$$\begin{aligned} \Pr_{\text{data} \sim P_{ABXY}} [f > 1 - \alpha + \beta] &\leq \hat{C}(G, n) \exp[-n\epsilon^2/8] \\ &= \text{poly}(n) \exp[-n\epsilon^2/8]. \end{aligned} \quad (10.26)$$

Our goal now is to show that, actually, it must be possible to replace $\hat{C}(G, n)$ with a constant smaller than 1 and by this drop the polynomial pre-factor. We do this using a step that appeared in [14, Proof of Theorem 3.1].¹⁶

To this end, denote by $\omega_{\text{opt}}(G_{1-\alpha+\beta}^n)$ the optimal winning probability in the threshold game $G_{1-\alpha+\beta}^n$. For any n , let \tilde{C}_n be the constant for which the tight bound

$$\omega_{\text{opt}}(G_{1-\alpha+\beta}^n) = \tilde{C}_n \exp[-n\epsilon^2/8] \quad (10.27)$$

holds. In particular, this means that there exists a strategy P_n that achieves the above winning probability.

Assume by contradiction that there exists N_0 such that $\tilde{C}_{N_0} > 1$. Thus, there exists a strategy P_{N_0} achieving $\tilde{C}_{N_0} \exp[-N_0\epsilon^2/8]$, with $\tilde{C}_{N_0} > 1$, in the game $G_{1-\alpha+\beta}^{N_0}$.

Let N_1 be sufficiently large, so that Eq. (10.10) holds for $n = N_0N_1$ and consider the threshold game $G_{1-\alpha+\beta}^{N_0N_1}$. On the one hand, using N_1 independent copies of P_{N_0} achieves winning probability of $(\tilde{C}_{N_0})^{N_1} \exp[-N_0N_1\epsilon^2/8]$ and thus

$$\omega_{\text{opt}}(G_{1-\alpha+\beta}^{N_0N_1}) \geq (\tilde{C}_{N_0})^{N_1} \exp[-N_0N_1\epsilon^2/8].$$

On the other hand, Eq. (10.26) must hold for $n = N_0N_1$:

$$\omega_{\text{opt}}(G_{1-\alpha+\beta}^{N_0N_1}) \leq \text{poly}(N_0N_1) \exp[-N_0N_1\epsilon^2/8].$$

To reconcile both bounds, we must have $(\tilde{C}_{N_0})^{N_1} \leq \text{poly}(N_0N_1)$ for all sufficiently large N_1 . Thus, $\tilde{C}_{N_0} \leq 1$, which leads to a contradiction.

We get that for *all* sufficiently large n , $\tilde{C}_n \leq 1$. In combination with Eq. (10.27) we therefore have

$$\omega_{\text{opt}}(G_{1-\alpha+\beta}^n) \leq \exp[-n\epsilon^2/8]. \quad \square$$

To get a better feeling of the result, without trying to optimise it, one can make the following choices. Let $\epsilon = \frac{\beta}{10d}$, $\zeta = 8\epsilon$ and $\nu = \epsilon$ (assuming $\min_{x,y} Q_{XY}(x, y) > \frac{\beta}{10d}$). Using these choices, our proof holds for n and β such that

$$\frac{n}{\ln(n)} > 20|\mathcal{A}||\mathcal{B}||\mathcal{X}||\mathcal{Y}| \frac{\ln(20d/\beta)}{(\beta/10d)^2}$$

¹⁶The part of the proof starting at this point onward did not appear in the proof of the threshold theorem of [4]. We follow here the last part of the proof of the threshold theorem presented in [14], which appeared after [4], and can be used to improve the result of [4].

with the following constants in Theorem 10.11:

$$\mathcal{C}(G) = (30d)^{-2} = (30|\mathcal{X}||\mathcal{Y}|(|\mathcal{A}| + |\mathcal{B}|))^{-2} . \quad (10.28)$$

The theorem then reads

$$\Pr_{\text{data} \sim P_{ABXY}} [f > 1 - \alpha + \beta] \leq \exp \left[-n\beta^2 (30|\mathcal{X}||\mathcal{Y}|(|\mathcal{A}| + |\mathcal{B}|))^{-2} \right] .$$

A different choice of parameters can improve the dependence of the constants on the game G .

10.4 Open Questions

In this chapter we considered the question of parallel repetition of games when the players are allowed to use any non-signalling strategy. The most interesting direction for future work is the development of a similar proof technique, based on de Finetti reductions or other forms of reductions to IID, for classical and quantum parallel repetition. In the case of classical games, parallel repetition results for general games with more than two parties are unknown. For quantum games, even the case of two-player games is not completely solved. (Recall Sect. 4.1.2 for further information). Since our proof captures all types of games and any number of players (see [4]), a similar proof technique for classical and quantum games will solve some open questions.

To understand what is the main challenge when trying to extend the proof to classical and quantum case, note the following. In the standard proofs of parallel repetition theorems, i.e., proofs following the approach of [7] most of the difficulties arise due to the effect of conditioning on the event of winning some of the game repetitions. As this event is one that depends on the structure of the game and strategy and we have no control over them, conditioning can introduce arbitrary correlations between the questions used in different repetitions of the game, a major source of difficulty for the remainder of the argument. In our proof we also need to analyse the effect of conditioning on a certain event, the event of the non-signalling test accepting, and this is done in Lemma 10.9. However, the key advantage of our approach is that the test has a very specific structure, and in particular conditioning on the test passing can be done locally by the players in a way that respects the non-signalling constraints. As a result it is almost trivial to deal with the conditioning in the remainder of the proof. This shift from conditioning on an uncontrolled event, success in the game, to a highly controlled one, a non-signalling test that we design ourselves, is a key simplification that we expect to play an important role in any extension of our method to classical or quantum strategies.

By finding appropriate “non-classicality” and “non-quantumness” measures which can replace our signalling measure in Definition 10.3 one may be able to adapt the proof to the multi-player classical and quantum cases as well. Unfortu-

nately, it is not clear which measure can be used by the players, preferably locally, to determine if their systems are classical or quantum. In other words, the main difficulty is finding a measure for which Lemma 10.9 can be proven. The rest of the proofs should follow easily for most “non-classicality” and “non-quantumness” measures of one-game strategies.

References

1. Hänggi E, Renner R, Wolf S (2010) Efficient device-independent quantum key distribution. *Advances in cryptology-EUROCRYPT 2010*. Springer, Berlin, pp 216–234
2. Masanes L (2009) Universally composable privacy amplification from causality constraints. *Phys Rev Lett* 102(14):140501
3. Masanes L, Renner R, Christandl M, Winter A, Barrett J (2014) Full security of quantum key distribution from no-signaling constraints. *IEEE Trans Inf Theory* 60(8):4973–4986
4. Arnon-Friedman R, Renner R, Vidick T (2016) Non-signaling parallel repetition using de finetti reductions. *IEEE Trans Inf Theory* 62(3):1440–1457
5. Buhrman H, Fehr S, Schaffner C (2013) On the parallel repetition of multi-player games: the no-signaling case. [arXiv:1312.7455](https://arxiv.org/abs/1312.7455)
6. Holmgren J, Yang L (2017) (a counterexample to) parallel repetition for non-signaling multi-player games. In: *Electronic colloquium on computational complexity (ECCC)*, vol 24, p 178
7. Raz R (1998) A parallel repetition theorem. *SIAM J Comput* 27(3):763–803
8. Holenstein T (2007) Parallel repetition: simplifications and the no-signaling case. In: *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pp 411–419. ACM
9. Rao A (2011) Parallel repetition in projection games and a concentration bound. *SIAM J Comput* 40(6):1871–1891
10. Renner R (2007) Symmetry of large physical systems implies independence of subsystems. *Nat Phys* 3(9):645–649
11. Christandl M, König R, Renner R (2009) Postselection technique for quantum channels with applications to quantum cryptography. *Phys Rev Lett* 102(2):020504
12. Barrett J, Leifer M (2009) The de Finetti theorem for test spaces. *New J Phys* 11(3):033024
13. Christandl M, Toner B (2009) Finite de Finetti theorem for conditional probability distributions describing physical theories. *J Math Phys* 50:042104
14. Lancien C, Winter A (2016) Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de finetti reduction. *Chic J Theor Comput Sci* (11)
15. Ito T (2010) Polynomial-space approximation of no-signaling provers. *Automata, Languages and Programming*. Springer, Berlin, pp 140–151
16. Schrijver A (1998) *Theory of linear and integer programming*. Wiley, New York

Chapter 11

Showcase: Device-Independent Quantum Cryptography



In this chapter we consider the showcase of device-independent quantum cryptography and show how the security proof of device-independent cryptographic protocols can be performed via a *reduction to IID*. We introduce a general framework for obtaining proofs of device-independent security for a broad range of cryptographic tasks. For the sake of explicitness, we focus in this chapter on the task of device-independent quantum key distribution (DIQKD).¹

The main result that we present can be phrased in the following informal way (the formal theorem is stated as Theorem 11.6):

Theorem 11.1 (Informal) *Security of DIQKD in the most general case follows from security under the IID assumption. Moreover, the dependence of the key rate on the number of rounds, n , is the same as the one in the IID case, up to terms that scale like $1/\sqrt{n}$.*

The theorem establishes the a priori surprising fact that general quantum adversaries are no stronger than an adversary restricted to IID attacks, even in the device-independent setting. This allows us to give simple and modular security proofs of DIQKD and to extend tight results known for DIQKD under the IID assumption to the most general setting thus deriving essentially optimal key rates and noise tolerance.²

Our technique takes advantage of the sequential nature of the protocol, as well as the specific way in which classical statistics are collected by users of the protocol,

¹Since the initial announcement of our work [1], our framework has already been applied to a variety of additional tasks, including conference key agreement [2], randomness expansion [1] and privatization [3], as well as randomness generation with sub-linear quantum resources [4].

²This is crucial for experimental implementations of device-independent protocols. Our quantitative results have been applied to the analysis of the first experimental implementation of a protocol for randomness generation in the fully device-independent framework [5].

and makes use of the entropy accumulation theorem (EAT), discussed as part of Chap. 9. The analysis and results of this chapter previously appeared in [1, 6].

The chapter is arranged as follows. We first explain in Sect. 11.1 what is the main challenge when proving the security of device-independent quantum cryptographic protocols, such as DIQKD. Section 11.2 deals with the analysis of the main subroutine of most device-independent protocols. Then, the security proof of DIQKD is given in Sect. 11.3. As we are about to encounter many parameters and variables throughout the proofs, we list them in Appendix C.3 for convenience.

11.1 Main Challenge and Goal

The central task when proving security of cryptographic protocols consists in bounding the information that an adversary, called Eve, may obtain about certain values generated by the protocol, which are supposed to be secret. In the case of QKD, for example, the relevant output of the protocol is the raw data K , and proving security is essentially equivalent³ to establishing a lower bound on the smooth conditional min-entropy $H_{\min}^{\varepsilon}(K|E)$, where E is Eve's quantum system, which can be initially correlated to the device producing K . The quantity $H_{\min}^{\varepsilon}(K|E)$ determines the maximal length of the secret key that can be created by the protocol. Hence, proving security amounts to establishing a lower bound on $H_{\min}^{\varepsilon}(K|E)$. Evaluating the smooth min-entropy $H_{\min}^{\varepsilon}(K|E)$ of a large system is often difficult, especially in the device-independent setting where not much is known about the way K is produced and the system E is out of our control.

The IID assumption, discussed in Chap. 7, is commonly used to simplify the task of calculating $H_{\min}^{\varepsilon}(K|E)$. The analysis of the smooth min-entropy under the IID assumption was sketched in Sect. 7.3.2; in that case the total smooth min-entropy can be easily related to the sum of the von Neumann entropies in each round separately, using the quantum asymptotic equipartition property (Sect. 7.2.2). A bound on the entropy accumulated in one round can usually be derived using the expected winning probability in the game played in that round (as appeared in Sect. 5.2), which in turn can be easily estimated during the protocol in the IID case using standard Chernoff-type bounds. A long line of works [7–16] considered the security of device-independent quantum and non-signalling cryptography under the IID assumption. Most relevant for our work are the results of [12], where security of a DIQKD protocol was proven in the asymptotic limit, i.e., when the device is used $n \rightarrow \infty$ times, and under the IID assumption. Their protocol is based on the CHSH inequality [17], and their analysis shows that it achieves the best possible rates under these assumptions.

Unfortunately, even though quite convenient for the analysis, the IID assumption is a very strong one in the DI scenario. In particular, under such an assumption the

³From that point onward standard classical post-processing steps, e.g., error correction and privacy amplification, suffice to prove the security of the protocol; recall Sect. 4.2.3.

device cannot use any internal memory (i.e., its actions in one round cannot depend on the previous rounds) or even display time-dependent behaviour (due to inevitable imperfections for example). Without this assumption, however, very little is known about the structure of the untrusted device and hence also about its output (as the device might correlate the different rounds in an almost arbitrary way). As a consequence, DIQKD security proofs [18–20] that estimated $H_{\min}^{\epsilon}(K|E)$ directly for the most general case had to use complicated techniques and statistical analysis compared to the IID case. This led to security statements which are of limited relevance for practical experimental implementations; they are applicable only in an unrealistic regime of parameters, e.g., small amount of tolerable noise and large number of signals.

To overcome the above difficulty we take the approach of *reductions to IID* in the analysis presented in the following sections. In particular, we leverage the sequential nature of our DIQKD protocol to prove its security by reducing the analysis of multi-round sequential boxes to that of IID boxes as discussed in Chap. 9. Specifically, we use the EAT presented in Sect. 9.2.3 to establish that entropy accumulates additively throughout the multiple rounds of the protocol and use it to bound the total amount of smooth min-entropy $H_{\min}^{\epsilon}(K|E)$.⁴

This results in a proof technique with several benefits. Firstly, since the analysis of the IID case is rather simple and modular (as it builds mainly on the analysis of a single-round box) a security proof via a reduction to IID ends up being simple and modular by itself. For example, if one wishes to consider a DIQKD protocol based on a game other than the CHSH game, the sole significant modification of the security proof is the analysis of a single-round box (see Sects. 5.2 and 11.2.2). Secondly, due to the optimality of the EAT (at least to first order in n), we are able to extend tight results known for, e.g., DIQKD, under the IID assumption, to the most general setting. This yields the best rates known for any protocol for a device-independent cryptographic task. Thirdly, performing a finite-size analysis is no harder than performing the asymptotic one as all dependency on n is either trivial or already incorporated in the EAT.

We are now ready to embark on the mission of proving the security of our DIQKD protocol, described in Sect. 4.2.2 (see also Protocol 11.2 below).

11.2 Device-Independent Entropy Accumulation

The current section is devoted to the analysis of the entropy accumulation protocol presented as Protocol 11.1. The entropy accumulation protocol acts as the main building block of many device-independent cryptographic protocols. It is used to

⁴The security proof presented in [20] is similar in spirit (but technically very different) to the one presented here. It bounds the total amount of smooth min-entropy generated in the protocol in a round-by-round fashion but the entropy accumulated in a single round is not the von Neumann entropy.

Protocol 11.1 CHSH-based entropy accumulation protocol

Arguments:

- D – untrusted device of two components that can play CHSH repeatedly
- $n \in \mathbb{N}_+$ – number of rounds
- $\gamma \in (0, 1]$ – expected fraction of test rounds
- ω_{exp} – expected winning probability in an honest implementation
- $\delta_{\text{est}} \in (0, 1)$ – width of the confidence interval for parameter estimation

- 1: For every round $i \in [n]$ do Steps 2-5:
 - 2: Alice and Bob choose a random $T_i \in \{0, 1\}$ such that $\Pr(T_i = 1) = \gamma$.
 - 3: If $T_i = 0$, Alice and Bob choose $(X_i, Y_i) = (0, 2)$ and otherwise $X_i, Y_i \in \{0, 1\}$ uniformly at random.
 - 4: Alice and Bob use D with X_i, Y_i and record their outputs as A_i and B_i respectively.
 - 5: If $T_i = 0$ then Bob updates B_i to $B_i = \perp$, and they set $W_i = \perp$. If $T_i = 1$ they set $W_i = w(A_i, B_i, X_i, Y_i)$.
 - 6: Alice and Bob abort if $\sum_{j: T_j=1} W_j < (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n$.
-

generate the raw data for Alice and Bob by playing a non-local game n times in sequence using an untrusted device D . We remark that even though we call the entropy accumulation protocol a “protocol”, one should see it more as a mathematical tool which allows us to use the machinery of the EAT rather than an actual protocol to be implemented.⁵ The relevance of the protocol stems from the fact that the final state at the end of the protocol, on which a smooth min-entropy is evaluated, is closely related to the final state in the actual protocol to be executed (e.g., our DIQKD protocol).

Our primary task is to lower-bound the amount of smooth min-entropy generated by playing the n games. This lower-bound can then be used as the starting point of security proofs of device-independent cryptographic protocols, such as DIQKD. The informal statement is given below (for the explicit formulation see Theorem 11.5):

Theorem 11.2 (Informal) *Fix a choice of parameters for Protocol 11.1. Then there exist constants $c_1, c_2 > 0$ such that the following holds. Let D be any device and $\rho_{|\Omega}$ the state generated using Protocol 11.1, conditioned on the protocol not aborting. Then for any $\varepsilon_1, \varepsilon_2 \in (0, 1)$, either the protocol aborts with probability greater than $1 - \varepsilon_1$ or*

$$H_{\min}^{\varepsilon_2}(\mathbf{AB}|\mathbf{XYTE})_{\rho_{|\Omega}} > c_1 n - c_2 \sqrt{n \log(1/\varepsilon_1 \varepsilon_2)}. \quad (11.1)$$

The registers \mathbf{AB} in Eq. (11.1) contain the classical outputs generated by the device during the protocol. The registers \mathbf{XYT} hold the classical information exchanged during the protocol, that may be leaked to the adversary. E is a quantum register that describes the adversary’s quantum system. Thus, Eq. (11.1) gives a precise bound on the amount of the smooth min-entropy present in the users’ outputs at the end of the protocol, conditioned on all information available to the adversary.

We give below explicit formulas for computing the constants c_1 and c_2 that appear in Eq. (11.1) as a function of the parameters of the protocol. Importantly, the con-

⁵In particular, in a setting with two distinct parties, Alice and Bob, communication is required to actually implement Protocol 11.1. We ignore this as it is not relevant for the analysis.

stant c_1 that governs the leading-order term equals the optimal constant, i.e., the same leading constant that would be obtained under the IID assumption, which by the asymptotic equipartition property (Theorem 7.3) is the von Neumann entropy accumulated in one round of the protocol. Furthermore, our analysis provides control over the constant c_2 in front of the second-order term. Such control is necessary for any application where finite values of n need to be considered, such as in quantum cryptography, where the values of n achieved in practice remain relatively small.⁶

As we show below, Theorem 11.2 can be proven by reducing the general sequential scenario to the IID one using the EAT. To use the EAT, we first need to construct the relevant objects, i.e., the EAT-channels and the min-tradeoff functions defined in Sect. 9.2.2. This is done in the following two sections. A lower-bound on the smooth min-entropy is then proven in Sect. 11.2.3.

11.2.1 EAT Channels

Protocol 11.1 proceeds in rounds and can therefore be presented by an application of a sequence of quantum channels (recall Sect. 9.1). In this section we define the considered channels and prove that they are EAT-channels, according to Definition 9.1. Note that one has some freedom in *choosing* the channels to work with (i.e., the channels are not completely defined by the protocol itself). We choose our particular channels so that all the prerequisites of the EAT are fulfilled and, at the same time, the final bound on the smooth min-entropy can be converted to a bound on the smooth min-entropy in our DIQKD protocol (see Sect. 11.3.2, and Lemma 11.8 in particular, for details).

Every EAT channel \mathcal{M}_i describes one round of the protocol, where one round includes Steps 2–5 of Protocol 11.1. For every $i \in \{0\} \cup [n]$, the (unknown) quantum state of the device D shared by Alice and Bob after round i of the protocol is denoted by $\rho_{Q_A Q_B}^i$. We denote the register holding this state by R_i . In particular, $R_0 = Q_A Q_B$ at the start of the protocol. At Step 4 in Protocol 11.1, the quantum state of the devices is changed from $\rho_{Q_A Q_B}^{i-1}$ in R_{i-1} to $\rho_{Q_A Q_B}^i$ in R_i by the use of the device. To be a bit more precise, the quantum state is changed in two stages. First, the relevant measurement of Step 4 is done (where it is assumed that the measurements of the different components are in tensor product). Then, after A_i and B_i are recorded, the different components of the device are allowed to communicate. Thus, some further changes can be made to the post-measurement state even based on the memory of all components together (recall Sect. 6.2.2).

In the notation of Chap. 9, we make the following choices:

⁶See e.g. Fig. 11.4, where one can see that finite-size effects can play an important role up to even moderately large values of $n \approx 10^{10}$.

$$\begin{aligned}
O_i &= A_i B_i \\
S_i &= X_i Y_i T_i \\
C_i &= W_i \\
R_i &= R_i \\
E &= E .
\end{aligned} \tag{11.2}$$

Our EAT channels are then

$$\mathcal{M}_i : R_{i-1} \rightarrow R_i A_i B_i X_i Y_i T_i W_i$$

defined by the CPTP map describing the i -th round of Protocol 11.1, as implemented by the untrusted device D . That is, the channel describes the random choices of T_i , X_i , and Y_i , the quantum operations made by the device, and the production of A_i , B_i , and W_i .

Since the operations of D are unknown, our EAT channels are not completely explicit. The important thing is merely that we know that some quantum channels describing the operation of the device *exist*. The lack of knowledge regarding the channels does not raise any problems when applying the EAT but it does make the task of deriving good min-tradeoff functions more challenging (compared to the scenario of a characterised device). This difficulty, however, is inherent to device-independent information processing tasks and has nothing to do with the proof technique; see Sect. 11.2.2 below for further details.

We prove that the described channels can act as our EAT channels.

Lemma 11.3 *The channels $\{\mathcal{M}_i : R_{i-1} \rightarrow R_i A_i B_i X_i Y_i T_i W_i\}_{i \in [n]}$ defined by the CPTP map describing the i -th round of Protocol 11.1, as implemented by the untrusted device D are EAT channels according to Definition 9.1 and the identification made in Eq. (11.2).*

Proof To prove that the constructed channels $\{\mathcal{M}_i\}_{i \in [n]}$ are EAT channels we need to show that the three conditions stated in Definition 9.1 are fulfilled.

1. $\{O_i\}_{i \in [n]} = \{A_i B_i\}_{i \in [n]}$, $\{S_i\}_{i \in [n]} = \{X_i Y_i T_i\}_{i \in [n]}$, and $\{C_i\}_{i \in [n]} = \{W_i\}_{i \in [n]}$ are all finite-dimensional classical systems. $\{R_i\}_{i \in [n]}$ are arbitrary quantum systems. Finally, we have $d_O = d_{A_i} \cdot d_{B_i} = 2 \cdot 3 = 6$.
2. For any $i \in [n]$ and any input state $\sigma_{R_{i-1}}$, W_i is a function of the classical values A_i , B_i , X_i , and Y_i . Hence, the marginal $\sigma_{O_i S_i} = \sigma_{A_i B_i X_i Y_i T_i}$ of the output state is unchanged when deriving W_i from it. (In other words, we can “measure” $\sigma_{A_i B_i X_i Y_i T_i}$ to get the value of W_i repeatedly without disturbing $\sigma_{A_i B_i X_i Y_i T_i}$).
3. For any initial state $\rho_{R_0 E}^{\text{in}}$ and the resulting final state $\rho_{OSCE} = \rho_{ABXYTWE}$, the Markov-chain conditions

$$(AB)_1, \dots, (AB)_{i-1} \leftrightarrow (XYT)_1, \dots, (XYT)_{i-1}, E \leftrightarrow (XYT)_i$$

trivially hold for all $i \in [n]$ since, according to Protocol 11.1, X_i , Y_i , and T_i are chosen independently from everything else. \square

When defining the EAT channels we identified O_i with $A_i B_i$. Looking ahead, this means that we are about to derive a bound on $H_{\min}^\varepsilon(\mathbf{AB}|\mathbf{XYTE})$. For the analysis of DIQKD, however, a bound on $H_{\min}^\varepsilon(A|\mathbf{XYTE})$ is needed. Why not set $O_i = A_i$ instead of $O_i = A_i B_i$? The reason is that the definition of an EAT channel requires that C_i can be derived from $O_i S_i$ alone. Thus, choosing $O_i = A_i$, $S_i = X_i Y_i T_i$, and $C_i = W_i$ would render the above lemma wrong. The reader may then ask—why not choose $O_i = A_i$ and $S_i = B_i X_i Y_i T_i$? With this choice, however, the required Markov-chain conditions read

$$A_1, \dots, A_{i-1} \leftrightarrow (BXYT)_1, \dots, (BXYT)_{i-1}, E \leftrightarrow (BXYT)_i$$

and these do not hold for an arbitrary initial states since nothing restricts B_i from being correlated to, e.g., A_{i-1} , when considering untrusted devices.⁷ Hence, the lemma would not hold for this choice as well. We therefore stick with the choices made in Eq. (11.2) and relate $H_{\min}^\varepsilon(\mathbf{AB}|\mathbf{XYTE})$ to $H_{\min}^\varepsilon(A|\mathbf{XYTE})$ in Sect. 11.3.2.

Now that our EAT channels are defined, the next step is to construct a min-tradeoff function for them. This is done in the next section.

11.2.2 Min-Tradeoff Function

When working with the EAT the most important task is to devise a good tradeoff function, as defined in Definition 9.2. As mentioned in Sect. 9.2.3, this is where the “physics kicks in”. We are aiming for a lower-bound on the smooth min-entropy and hence in need of a min-tradeoff function f_{\min} . That is, we need to construct a convex differential function for which, for all $i \in [n]$,

$$f_{\min}(p) \leq \inf_{\sigma_{R_{i-1}R'} : \mathcal{M}_i(\sigma)_{W_i} = p} H(A_i B_i | X_i Y_i T_i R')_{\mathcal{M}_i(\sigma)}, \quad (11.3)$$

where p is a probability distribution over $\mathcal{W} = \{\perp, 0, 1\}$ and $\{\mathcal{M}_i\}_{i \in [n]}$ are the EAT channels defined in the previous section.

To understand the task at hand, let us first focus on the set of states

$$\Sigma(p) = \{\sigma_{R_{i-1}R'} : \mathcal{M}_i(\sigma)_{W_i} = p\} .$$

on which the infimum is evaluated. First observe that, due to the structure of our channels, the distributions over X_i , Y_i , and T_i are *fixed* for any $\sigma \in \Sigma(p)$. That is, even though we take an infimum over many possible input states for the channels, and even though the actions of the untrusted device are not characterised, the values

⁷Consider for example a device in which the initial state $\rho_{Q_{A_1} Q_{B_1} Q_{A_2} Q_{B_2}} = |\Phi\rangle\langle\Phi|_{Q_{A_1} Q_{B_2}} \otimes |\Phi\rangle\langle\Phi|_{Q_{A_2} Q_{B_1}}$, i.e., the systems over Q_{A_1} and Q_{B_2} are entangled. Thus, A_1 and B_2 may be correlated even given $B_1 X_1 Y_1 T_1$. In this case the Markov-chain conditions do not hold since the side-information B_2 reveals information regarding the past output A_1 .

of X_i , Y_i , and T_i are always chosen according to Protocol 11.1.⁸ This implies, in particular, that for probability distributions p with $p(\perp) \neq 1 - \gamma$ the set $\Sigma(p)$ is empty.

Given the above, for any p over $\{\perp, 0, 1\}$, the set $\Sigma(p)$ includes only states σ for which

$$\mathcal{M}_i(\sigma)_{w_i} = \begin{pmatrix} p(\perp) & 0 & 0 \\ 0 & p(0) & 0 \\ 0 & 0 & p(1) \end{pmatrix} = \begin{pmatrix} 1 - \gamma & 0 & 0 \\ 0 & \gamma(1 - \omega) & 0 \\ 0 & 0 & \gamma\omega \end{pmatrix}, \quad (11.4)$$

where we identify ω with the winning probability of the state σ in the CHSH game (when using the measurements defining the channel \mathcal{M}_i). For this reason, we can, slightly informally,⁹ see the function f_{\min} as defined over a single variable $\omega \in [0, 1]$.

In total, we can understand the set $\Sigma(p) = \Sigma(p(\omega))$ as the set including all states σ that can be used to win the CHSH game with probability ω . It is this information about the relevant input states σ that allows us to construct a min-tradeoff function fulfilling Eq. (11.3). In fact, given the above observation, the main ingredient needed to construct a valid (and tight) min-tradeoff function is Lemma 5.3, which was discussed in the context of *single-round boxes* (Chap. 5). This clarifies why the presented proof technique can be seen as a *reduction to IID*.

We are ready to embark on the construction of the min-tradeoff function.

Lemma 11.4 *Let*¹⁰

$$g(p) = \begin{cases} 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\frac{p(1)}{\gamma}\left(\frac{p(1)}{\gamma} - 1\right) + 3}\right) & \frac{p(1)}{\gamma} \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right] \\ 1 & \frac{p(1)}{\gamma} \in \left[\frac{2+\sqrt{2}}{4}, 1\right] \end{cases},$$

and

⁸A different model for the sequential process could have been one in which the initial quantum state *itself* includes the registers X and Y and the channel is defined such that a measurement is performed on those registers to get the inputs (and then use the device in the protocol). When starting with maximally mixed states over XY the entire sequential process is exactly the same as the one described by our EAT channels. However, when coming to construct a min-tradeoff function with this (somewhat strange) alternative choice of channels, we see that the set $\Sigma(p)$ can include states in which, e.g., $X_i = 0$ with probability 1 (since we need to consider *all possible* input states). In the context of Bell inequalities, this is similar to dropping the “free choice assumption”. Clearly, if this had been the case, the only min-tradeoff function one could construct is the constant function $f_{\min}(p) = 0$ for all p , which is trivial and useless.

⁹Formally, we will need to extend the function to all probability distributions p (even those with $p(\perp) \neq 1 - \gamma$). We can extend the function in any way we wish, while keeping it convex and differentiable.

¹⁰We define the functions g and f_{\min} only in the regime in which the protocol does not abort, i.e., $p(1)/\gamma \geq 3/4$.

$$f_{\min}(p, p_{\text{cut}}) = \begin{cases} g(p) & p(1) \leq p_{\text{cut}}(1) \\ \frac{d}{dp(1)}g(p)\big|_{p_{\text{cut}}} \cdot p(1) & p(1) > p_{\text{cut}}(1) \\ + g(p_{\text{cut}}) - \frac{d}{dp(1)}g(p)\big|_{p_{\text{cut}}} \cdot p_{\text{cut}}(1) & \end{cases} . \quad (11.5)$$

Then, for any probability distribution p_{cut} over $\{\pm, 0, 1\}$, $f_{\min}(p, p_{\text{cut}})$ is a min-tradeoff functions for the EAT channels from Lemma 11.3.

Before proving the lemma, let us parse the above lengthy equations. The function g is basically the single-round bound presented in Lemma 5.3, where we replace ω with $\frac{p(1)}{\gamma}$ and trivially extend the function to the regime of winning probabilities above the optimal quantum winning probability $\frac{2+\sqrt{2}}{4}$. Notice that for an arbitrary p , the correct relation between p and ω is given by $\omega = \frac{p(1)}{p(0)+p(1)}$. However, as explained above, due to the definition of the channels \mathcal{M}_i , the set $\Sigma(p)$ is empty for p with $p(0) + p(1) \neq \gamma$. This implies that there are no constraints on the value of the min-tradeoff function for such p 's and we are free to define it as we wish in this regime. Thus, we are not going to run into problems even though the value of the function g does not seem to have any ‘‘physical meaning’’ for p with $p(0) + p(1) \neq \gamma$.¹¹

The function f_{\min} in Eq. (11.10) is governed by g and can be understood as follows. Fix a probability distribution $p_{\text{cut}} \in [0, 1]$. For p with $p(1) \leq p_{\text{cut}}(1)$, f_{\min} is identical to g . Otherwise, f_{\min} is a linear function (when restricting ourselves to a slice $p(0) + p(1) = \text{constant}$) defined via the value and the tangent of g at the point $p_{\text{cut}}(1)$. That is, we ‘‘cut and glue’’ the function at point p_{cut} . By doing so, we make sure that f_{\min} is a convex and differentiable function, as required by Definition 9.2 while restraining its gradient, which will later affect the bound on the smooth min-entropy (via Eq. (9.14)). This construction of f_{\min} is illustrated in Fig. 11.1.

Proof (Proof of Lemma 11.4) We start by using the chain rule of the von Neumann entropy,

$$H(A_i B_i | X_i Y_i T_i R')_{\mathcal{M}_i(\sigma)} \geq H(A_i | X_i Y_i T_i R')_{\mathcal{M}_i(\sigma)} .$$

Due to the bipartite requirement on the untrusted device D used to implement the protocol, the actions of Alice’s device are independent of Bob’s choice of Y_i as well as of and T_i .¹² We thus have

$$H(A_i | X_i Y_i T_i R')_{\mathcal{M}_i(\sigma)} = H(A_i | X_i R')_{\mathcal{M}_i(\sigma)} .$$

¹¹Alternatively, one could replace $p(1)/\gamma$ with $p(1)/(p(0) + p(1))$, which is more meaningful, in the definition of the function g . However, since the $df_{\min}/dp(1)$ will affect the final smooth min-entropy bound, using $p(1)/\gamma$ leads to better quantitative results.

¹²We assume that the value of T_i is exchanged over a classical authenticated channel to which the device D does not have access. In particular, Alice’s part of the device is independent from the value of T_i given X_i .

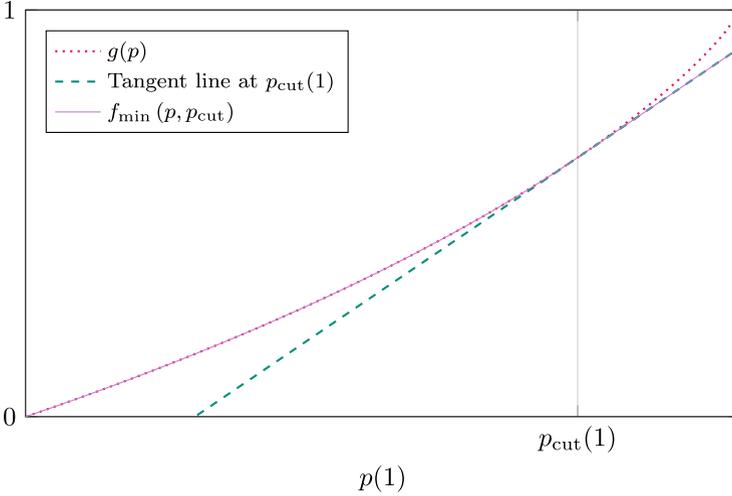


Fig. 11.1 The construction of the min-tradeoff function f_{\min} appearing in Eq. (11.10). The plot shows the values of the min-tradeoff function restricted to a slice $p(0) + p(1) = \text{constant}$. After [1, 6]

Combined with Lemma 5.3 we get that for any state σ with winning probability ω in the CHSH game,

$$H(A_i B_i | X_i Y_i T_i R') \geq 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\omega(\omega - 1) + 3}\right). \quad (11.6)$$

For probability distributions p with $p(0) + p(1) \neq \gamma$, the set of states fulfilling $\mathcal{M}_i(\sigma)_{W_i} = p$ is empty and the condition on the min-tradeoff function given in Eq. 11.3 becomes trivial. Hence, for the construction of the min-tradeoff function we can restrict our attention to p with $p(0) + p(1) = \gamma$. For such p 's one can write $\omega = \frac{p(1)}{p(0)+p(1)} = \frac{p(1)}{\gamma}$. All together we learn that for all p with $\frac{p(1)}{\gamma} \geq \frac{3}{4}$,

$$\inf_{\sigma_{R_{i-1}R'}} H(A_i B_i | X_i Y_i T_i R')_{\mathcal{M}_i(\sigma)} \geq 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16 \frac{p(1)}{\gamma} \left(\frac{p(1)}{\gamma} - 1\right) + 3}\right). \quad (11.7)$$

Define a function g by

$$g(p) = \begin{cases} 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16 \frac{p(1)}{\gamma} \left(\frac{p(1)}{\gamma} - 1\right) + 3}\right) & \frac{p(1)}{\gamma} \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right] \\ 1 & \frac{p(1)}{\gamma} \in \left[\frac{2+\sqrt{2}}{4}, 1\right]. \end{cases} \quad (11.8)$$

From Eq. (11.7) it follows that any choice of f_{\min} that is differentiable and satisfies $f_{\min}(p) \leq g(p)$ for all p will satisfy Eq. (11.3).

For $\frac{p(1)}{\gamma} = \frac{2+\sqrt{2}}{4}$ the derivative of g is infinite. Looking ahead, for the final bound on the smooth min-entropy derived using the EAT to be meaningful, f_{\min} should be chosen such that $\|\nabla f_{\min}\|_{\infty}$ is finite. To assure that this is the case we choose f_{\min} by “cutting” the function g and “gluing” it to a linear function at some point p_{cut} , while keeping the function differentiable. By doing this we ensure that the gradient of f_{\min} is bounded, at the cost of losing a bit of entropy for p with $p(1) > p_{\text{cut}}(1)$.¹³ Towards this, denote

$$a(p_{\text{cut}}) = \frac{d}{dp(1)}g(p)\Big|_{p_{\text{cut}}} \quad \text{and} \quad b(p_{\text{cut}}) = g(p_{\text{cut}}) - a(p_{\text{cut}}) \cdot p_{\text{cut}}(1). \quad (11.9)$$

We then make the following choice for the min-tradeoff function f_{\min} (see Fig. 11.1):

$$f_{\min}(p, p_{\text{cut}}) = \begin{cases} g(p) & p(1) \leq p_{\text{cut}}(1) \\ a(p_{\text{cut}}) \cdot p(1) + b(p_{\text{cut}}) & p(1) > p_{\text{cut}}(1) \end{cases} \quad (11.10)$$

From the definition of a and b in Eq. (11.9), this function is convex, differentiable, and fulfils the condition given in Eq. (11.3). f_{\min} can therefore be rightfully called a min-tradeoff function. Furthermore, by definition, for any choice of p_{cut} it holds that $\|\nabla f_{\min}(\cdot, p_{\text{cut}})\|_{\infty} \leq a(p_{\text{cut}})$. \square

11.2.3 Smooth Min-Entropy Rate

After constructing the EAT channels and min-tradeoff function in the previous sections, we are ready to apply Theorem 9.3 to derive our lower-bound on the conditional smooth min-entropy generated by the entropy accumulation protocol, Protocol 11.1.

We use the following notation. The event of *not aborting* the protocol is given by

$$\Omega = \left\{ \mathbf{w} : \sum_{j:T_j=1} w_j \geq (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n \right\}. \quad (11.11)$$

For any initial state $\rho_{Q_A Q_B E}^{\text{in}}$, the final state in the end of the protocol is denoted by $\rho = \rho_{\text{ABXYTWE}}$ and the final state conditioned on not aborting the entropy accumulation protocol is $\rho|_{\Omega}$.

As shown in Theorem 11.5 below, The smooth min-entropy rate is governed by the following functions, where h is the binary entropy and $\gamma, p(1) \in (0, 1]$:

¹³The point p_{cut} can later be chosen such that the derived smooth entropy bounds are optimised.

$$\begin{aligned}
g(p) &= \begin{cases} 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\frac{p(1)}{\gamma}\left(\frac{p(1)}{\gamma} - 1\right) + 3}\right) & \frac{p(1)}{\gamma} \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right] \\ 1 & \frac{p(1)}{\gamma} \in \left[\frac{2+\sqrt{2}}{4}, 1\right], \end{cases} \\
f_{\min}(p, p_{\text{cut}}) &= \begin{cases} g(p) & p(1) \leq p_{\text{cut}}(1) \\ \frac{d}{dp(1)}g(p)\Big|_{p_{\text{cut}}} \cdot p(1) + g(p_{\text{cut}}) - \frac{d}{dp(1)}g(p)\Big|_{p_{\text{cut}}} \cdot p_{\text{cut}}(1) & p(1) > p_{\text{cut}}(1), \end{cases} \\
\mu(p, p_{\text{cut}}, \varepsilon_s, \varepsilon_e) &= f_{\min}(p, p_{\text{cut}}) \\
&\quad - \frac{1}{\sqrt{n}}2\left(\log 13 + \frac{d}{dp(1)}g(p)\Big|_{p_{\text{cut}}}\right)\sqrt{1 - 2\log(\varepsilon_s \cdot \varepsilon_e)}, \\
\mu_{\text{opt}}(\varepsilon_s, \varepsilon_e) &= \max_{\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}} \mu(\omega_{\text{exp}}\gamma - \delta_{\text{est}}, p_{\text{cut}}, \varepsilon_s, \varepsilon_e). \tag{11.12}
\end{aligned}$$

Theorem 11.5 *Let D be any device, ρ the state generated by running Protocol 11.1, Ω the event that the protocol does not abort (as defined in Eq. (11.11)), and $\rho_{|\Omega}$ the state conditioned on Ω . Then, for any $\varepsilon_{\text{EA}}, \varepsilon_s \in (0, 1)$, either the protocol aborts with probability greater than $1 - \varepsilon_{\text{EA}}$ or*

$$H_{\min}^{\varepsilon_s}(\mathbf{AB}|\mathbf{XYTE})_{\rho_{|\Omega}} > n \cdot \mu_{\text{opt}}(\varepsilon_s, \varepsilon_{\text{EA}}), \tag{11.13}$$

where μ_{opt} is defined in Eq. (11.12).

Proof We wish to apply the EAT, stated as Theorem 9.3. To this end, denote by $\text{freq}_{\mathbf{w}}(\tilde{\mathbf{w}}) = \frac{|\{i|w_i=\tilde{w}_i\}|}{n}$ the frequency defined by the raw data \mathbf{w} (recall Eq. (9.11)) and observe the following:

1. The EAT channels $\{\mathcal{M}_i\}_{i \in [n]}$ constructed in Sect. 11.2.1 faithfully describe the protocol and the device D , in the sense that the final state of the protocol, ρ , can be written as

$$\rho_{\mathbf{ABXYTWE}} = (\text{Tr}_{R_n} \circ \mathcal{M}_n \circ \cdots \circ \mathcal{M}_1) \otimes \mathbb{I}_E \rho_{Q_A Q_B E}^{\text{in}}. \tag{11.14}$$

2. The set $\hat{\Omega} = \{p : p(1) \geq \omega_{\text{exp}}\gamma - \delta_{\text{est}}\}$ is convex and $\{\text{freq}_{\mathbf{w}} : \mathbf{w} \in \Omega\} \subseteq \hat{\Omega}$ (recall Sect. 9.2.3.1).
3. According to Lemma 11.4, $f_{\min}(p, p_{\text{cut}})$ is a min-tradeoff function for the considered EAT channels, for any p_{cut} with $\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}$.
4. For any p_{cut} with $\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}$, the value $t = f_{\min}(\omega_{\text{exp}}\gamma - \delta_{\text{est}}, p_{\text{cut}})$ satisfies $f_{\min}(\text{freq}_{\mathbf{w}}, p_{\text{cut}}) \geq t$ for any $\text{freq}_{\mathbf{w}} \in \hat{\Omega}$.
5. $d_O = d_{A_i B_i} = 6$ and $\|\nabla f_{\min}(\cdot, p_{\text{cut}})\|_{\infty} = a(p_{\text{cut}})$ for any p_{cut} with $\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}$.

Using the EAT (Theorem 9.3) in combination with the above observations we conclude that for any p_{cut} with $\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}$, either the protocol aborts with probability greater than $1 - \varepsilon_{\text{EA}}$, or

$$H_{\min}^{\varepsilon_s}(\mathbf{AB}|\mathbf{XYTE})_{\rho_{\Omega}} > n f_{\min}(\omega_{\text{exp}}\gamma - \delta_{\text{est}}, p_{\text{cut}}) - \sqrt{n}\zeta(p_{\text{cut}}), \quad (11.15)$$

for $\zeta(p_{\text{cut}}, \varepsilon_s, \varepsilon_{\text{EA}}) = 2(\log 13 + a(p_{\text{cut}}))\sqrt{1 - 2\log(\varepsilon_s \cdot \varepsilon_{\text{EA}})}$. To obtain the optimal rate we maximise $H_{\min}^{\varepsilon_s}(\mathbf{AB}|\mathbf{XYTE})_{\rho_{\Omega}}$ over p_{cut} . Denote $\mu(p, p_{\text{cut}}, \varepsilon_s, \varepsilon_{\text{EA}}) = f_{\min}(p, p_{\text{cut}}) - \frac{1}{\sqrt{n}}\zeta(p_{\text{cut}}, \varepsilon_s, \varepsilon_{\text{EA}})$ and let

$$\mu_{\text{opt}}(\varepsilon_s, \varepsilon_{\text{EA}}) = \max_{\frac{3}{4} < \frac{p_{\text{cut}}(1)}{\gamma} < \frac{2+\sqrt{2}}{4}} \mu(\omega_{\text{exp}}\gamma - \delta_{\text{est}}, p_{\text{cut}}, \varepsilon_s, \varepsilon_{\text{EA}}).$$

Plugging this into Eq. (11.15) the theorem follows. \square

Importantly, the theorem tells us that the first order term of the smooth min-entropy is linear in n . Moreover, asymptotically, the entropy rate is simply given by the min-tradeoff function $f_{\min}(p, p_{\text{cut}})$. This is why it was crucial to construct an optimal min-tradeoff function in Sect. 11.2.2.

The rate μ_{opt} is plotted in Fig. 11.2 as a function of the expected winning probability ω_{exp} in the CHSH game for $\gamma = 1$ and several choices of values for the parameters ε_{EA} , δ_{est} , and n (while optimising over all other parameters). For comparison, we also plot in Fig. 11.2 the asymptotic rate ($n \rightarrow \infty$) under the IID assumption. In this case, the quantum asymptotic equipartition property implies that the optimal rate is the Shannon entropy accumulated in one round of the protocol (recall Sect. 7.2.2). This

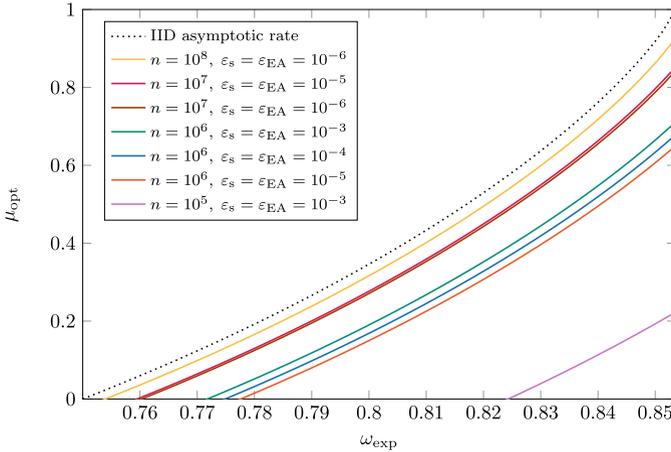


Fig. 11.2 $\mu_{\text{opt}}(\omega_{\text{exp}})$ for $\gamma = 1$ and several choices of n , ε_{EA} , and the smoothing parameter ε_s . $\delta_{\text{est}} = 10^{-2}$ in the curve with $n = 10^5$ and $\delta_{\text{est}} = 10^{-3}$ in all other curves. Note that for the errors of the protocols to be meaningful the number of rounds n should be at least of order $\delta_{\text{est}}^{-2} \cdot \varepsilon_{\text{EA}}$ and ε_s affect the soundness error in the DIQKD protocol considered in Sect. 11.3 and therefore should be chosen to be relatively small. The dashed line shows the optimal asymptotic ($n \rightarrow \infty$) rate under the IID assumption. After [1, 6]

rate, appearing as the dashed line in Fig. 11.2, is an upper bound on the smooth min-entropy that can be accumulated. One can see that as the number of rounds in the protocol increases the rate μ_{opt} approaches this optimal rate.

11.3 Device-Independent Quantum Key Distribution

Our DIQKD protocol is stated as Protocol 11.2. In the first part of the protocol Alice and Bob use their devices to produce the raw data, similarly to what is done in the entropy accumulation protocol, Protocol 11.1, analysed in the previous section. In the second part of the protocol Alice and Bob apply classical post-processing steps to produce their final keys from the raw data. The classical post-processing consists of error correction, parameter estimation, and privacy amplification; all discussed in detail in Sect. 4.2.2.

Apart from the classical post-processing, the main difference between the entropy accumulation protocol and the DIQKD protocol is the way we set Bob's outputs. In Protocol 11.1, Bob's outputs are being set to \perp in all rounds for which $T_i = 0$, i.e., in the generation rounds. In contrast, when dealing with QKD Bob needs to keep the outputs produced in the generation rounds so that he could create a key identical to Alice's key. To make the distinction explicit we denote Bob's outputs in Protocol 11.2 with a tilde, $\tilde{\mathbf{B}}$. We will get back to this point later and explain why the distinction is relevant for our analysis.

Our main goal in the following sections is to prove the security (according to Definition 4.5) of Protocol 11.2:

Theorem 11.6 *For any choice of parameters, the DIQKD protocol given in Protocol 11.2 is $(\varepsilon_{\text{QKD}}^s, \varepsilon_{\text{QKD}}^c)$ -secure according to Definition 4.5, with $\varepsilon_{\text{QKD}}^s \leq 2\varepsilon_{\text{EC}} + \varepsilon_{\text{PA}} + \varepsilon_s + \varepsilon_{\text{EA}}$, $\varepsilon_{\text{QKD}}^c \leq \varepsilon_{\text{EC}}^c + \varepsilon_{\text{EA}}^c + \varepsilon_{\text{EC}}$, and for key length*

$$\begin{aligned} \ell = & n \cdot \mu_{\text{opt}}(\varepsilon_s/4, \varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}) - \text{leak}_{\text{EC}} \\ & - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) - \gamma n \\ & - \sqrt{n} 2 \log(7) \sqrt{1 - 2 \log(\varepsilon_s/4 \cdot (\varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}))} - 2 \log(\varepsilon_{\text{PA}}^{-1}), \end{aligned} \quad (11.16)$$

where μ_{opt} is specified in Eq. (11.12).

The resulting key rates, ℓ/n , are discussed and plotted in Sect. 11.3.3 for different choices of parameters.

In the following sections we are set to prove Theorem 11.6. The theorem follows from the completeness of the protocol, stated as Lemmas 11.7, and its soundness, stated as Lemma 11.9.

Protocol 11.2 CHSH-based DIQKD protocol

Arguments:

- D – untrusted device of two components that can play CHSH repeatedly
- $n \in \mathbb{N}_+$ – number of rounds
- $\gamma \in (0, 1]$ – expected fraction of test rounds
- ω_{exp} – expected winning probability in an honest implementation
- $\delta_{\text{est}} \in (0, 1)$ – width of the confidence interval for parameter estimation
- EC – error correction protocol that leaks leak_{EC} bits and has completeness and soundness error probabilities $\varepsilon_{\text{EC}}^c$ and ε_{EC} respectively
- PA – privacy amplification protocol with error probability ε_{PA}

- 1: For every round $i \in [n]$ do Steps 2-4:
 - 2: Alice and Bob choose a random $T_i \in \{0, 1\}$ such that $\Pr(T_i = 1) = \gamma$.
 - 3: If $T_i = 0$, Alice and Bob choose $(X_i, Y_i) = (0, 2)$ and otherwise $X_i, Y_i \in \{0, 1\}$ uniformly at random.
 - 4: Alice and Bob use D with X_i, Y_i and record their outputs as A_i and \tilde{B}_i respectively.
 - 5: **Error correction:** Alice and Bob apply the error correction protocol EC. If EC aborts they abort the protocol. Otherwise, they obtain raw keys denoted by K_A and K_B .
 - 6: **Parameter estimation:** Using \tilde{B} and K_B , Bob sets $W_i = w_{\text{CHSH}}(K_{B_i}, \tilde{B}_i, X_i, Y_i)$ for the test rounds and $W_i = \perp$ otherwise. He aborts if $\sum_{j: T_j=1} W_j < (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n$.
 - 7: **Privacy amplification:** Alice and Bob apply the privacy amplification protocol PA on K_A and K_B to create their final keys \tilde{K}_A and \tilde{K}_B of length ℓ .
-

11.3.1 Completeness

We seek to prove that Protocol 11.2 is complete, i.e., that there exists an honest implementation of the device D that leads to a negligible probability of the protocol aborting. We remark that in order for the protocol to be relevant in practice, completeness has to be proven with respect to a *realistic* honest implementation that can be realised in experiments (or, at the least, believed to be feasible in the future). The honest implementation that we consider is the standard one and is described in Sect. 4.2.4. In short, the honest device makes IID measurements on an IID quantum state $\rho = \sigma^{\otimes n}$. The state and measurements are such that the winning probability achieved in the CHSH game in a single round is ω_{exp} that can be chosen freely.¹⁴

The following lemma gives the relation between the probability $\varepsilon_{\text{QKD}}^c$ that the protocol aborts for an honest implementation of the device D and the other parameters of the protocol.

Lemma 11.7 *Protocol 11.2 is complete with completeness error*

$$\varepsilon_{\text{QKD}}^c \leq \varepsilon_{\text{EC}}^c + \varepsilon_{\text{EC}} + \varepsilon_{\text{EA}}^c,$$

¹⁴For any ω_{exp} there are many devices that fit this description; an explicit example can be found in Sect. 4.2.4.

where $\varepsilon_{EA}^c \leq \exp(-2n\delta_{\text{est}}^2)$ and ε_{EC}^c and ε_{EC} are two independent parameters of the error correction protocol.

Proof We wish to upper-bound the probability that Protocol 11.2 aborts when running using the honest implementation. There are two steps in which Alice and Bob can abort Protocol 11.2:

1. The protocol may abort after the error correction step (Step 5). This happens with probability ε_{EC}^c .
2. Assuming the protocol did not abort in Step 5, it may abort after the parameter estimation step (Step 6). Recall that Bob performs parameter estimation using K_B and $\tilde{\mathbf{B}}$, i.e., he checks whether sufficiently many games were won when looking at his data K_B and $\tilde{\mathbf{B}}$. There are two scenarios which lead to the protocol aborting after parameter estimation:
 - (a) Error correction was successful, i.e., $K_B = K_A$, but not sufficiently many games were won when comparing K_A and $\tilde{\mathbf{B}}$. When utilising the honest implementation, W_i are IID RVs with $\mathbb{E}[W_i] = \omega_{\text{exp}}\gamma$. Therefore, we can use Hoeffding's inequality to bound the probability of such an event:

$$\varepsilon_{EA}^c = \Pr \left[\sum_{j: T_j \neq \perp} W_j \leq (\omega_{\text{exp}}\gamma - \delta_{\text{est}}) \cdot n \right] \leq \exp(-2n\delta_{\text{est}}^2). \quad (11.17)$$

- (b) Error correction was not successful, i.e., $K_B \neq K_A$ (but EC did not abort) and not sufficiently many games were won when comparing at K_B and $\tilde{\mathbf{B}}$. This happens with probability at most ε_{EC} .

The lemma follows by using the above in combination with the union bound. \square

11.3.2 Soundness

To establish soundness first note that, by definition, as long as Protocol 11.2 does not abort it produces a key of length ℓ . Therefore it remains to verify correctness (Definition 4.3), which depends on the error correction step, and security (Definition 4.4), which is based on the privacy amplification step.

To prove security we start by assuming that the error correction step is successful and lower-bound the smooth min-entropy of the quantum state shared between Alice and Bob right before the privacy amplification step. The main ingredient in the proof is the lower-bound on the smooth min-entropy established in Theorem 11.5. Most effort in proving security is devoted to relating the state considered in the entropy accumulation protocol (to which Theorem 11.5 refers) and the state in the end of the DIQKD protocol.

To be more precise, let $\tilde{\Omega}$ denote the event of Protocol 11.2 not aborting *and* the EC protocol being successful, and let $\tilde{\rho}_{A\tilde{B}XYTOE|\tilde{\Omega}}$ be the state at the end of the protocol,¹⁵ conditioned on this event. Success of the privacy amplification step relies on the smooth min-entropy $H_{\min}^{\varepsilon_s}(A|XYTOE)_{\tilde{\rho}_{|\tilde{\Omega}}}$ being sufficiently large. Lemma 11.8 connects this quantity to $H_{\min}^{\varepsilon_s/4}(AB|XYTE)_{\rho_{|\Omega}}$, on which a lower bound is provided by Theorem 11.5.

Lemma 11.8 *For any device D , let $\tilde{\rho}$ be the state generated in Protocol 11.2 right before the privacy amplification step, Step 7. Let $\tilde{\rho}_{|\tilde{\Omega}}$ be the state conditioned on not aborting the protocol and success of the EC protocol. Then, for any $\varepsilon_{EA}, \varepsilon_{EC}, \varepsilon_s \in (0, 1)$, either the protocol aborts with probability greater than $1 - \varepsilon_{EA} - \varepsilon_{EC}$ or*

$$\begin{aligned} H_{\min}^{\varepsilon_s}(A|XYTOE)_{\tilde{\rho}_{|\tilde{\Omega}}} &\geq n \cdot \mu_{\text{opt}}(\varepsilon_s/4, \varepsilon_{EA} + \varepsilon_{EC}) - \text{leak}_{\text{EC}} \\ &\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) - \gamma n \\ &\quad - \sqrt{n} 2 \log 7 \sqrt{1 - 2 \log(\varepsilon_s/4 \cdot (\varepsilon_{EA} + \varepsilon_{EC}))}. \end{aligned} \quad (11.18)$$

Proof Consider the following events:

1. Ω : the event of not aborting in the entropy accumulation protocol, Protocol 11.1. This happens when the Bell violation, calculated using Alice and Bob's outputs and inputs, is sufficiently high.
2. $\tilde{\Omega}$: Suppose Alice and Bob run Protocol 11.1, and then execute the EC protocol. The event $\tilde{\Omega}$ is defined by Ω and $K_B = \mathbf{A}$.
3. $\tilde{\tilde{\Omega}}$: the event of not aborting the DIQKD protocol, Protocol 11.2, and $K_B = \mathbf{A}$.

The state $\rho_{|\tilde{\tilde{\Omega}}}$ then denotes the state at the end of Protocol 11.1 conditioned on $\tilde{\tilde{\Omega}}$.

As we are only interested in the case where the EC protocol outputs the correct guess of Alice's bits, that is $K_B = \mathbf{A}$ (which happens with probability $1 - \varepsilon_{EC}$), we have $\tilde{\rho}_{AXYTE|\tilde{\tilde{\Omega}}} = \rho_{AXYTE|\tilde{\tilde{\Omega}}}$ (note that \tilde{B} and B were traced out from $\tilde{\rho}$ and ρ respectively). Hence,

$$H_{\min}^{\varepsilon_s}(A|XYTE)_{\tilde{\rho}_{|\tilde{\tilde{\Omega}}}} = H_{\min}^{\varepsilon_s}(A|XYTE)_{\rho_{|\tilde{\tilde{\Omega}}}}. \quad (11.19)$$

Using the chain rule given in [21, Lemma 6.8] together with Eq. (11.19) we get that

$$\begin{aligned} H_{\min}^{\varepsilon_s}(A|XYTOE)_{\tilde{\rho}_{|\tilde{\tilde{\Omega}}}} &\geq H_{\min}^{\varepsilon_s}(A|XYTE)_{\tilde{\rho}_{|\tilde{\tilde{\Omega}}}} - \text{leak}_{\text{EC}} \\ &= H_{\min}^{\varepsilon_s}(A|XYTE)_{\rho_{|\tilde{\tilde{\Omega}}}} - \text{leak}_{\text{EC}}, \end{aligned} \quad (11.20)$$

where leak_{EC} denotes the amount of information leaked during error correction.

¹⁵ O denotes the classical information sent from Alice to Bob during error correction; see Sect. 4.2.2.

To apply Theorem 11.5 it remains to relate $H_{\min}^{\varepsilon_s}(A|XYTE)_{\rho_{i\tilde{\Omega}}}$ to $H_{\min}^{\varepsilon'_s}(AB|XYTE)_{\rho_{i\tilde{\Omega}}}$ for some ε'_s . For this we first write

$$\begin{aligned} H_{\min}^{\varepsilon_s}(A|XYTE)_{\rho_{i\tilde{\Omega}}} &\geq H_{\min}^{\frac{\varepsilon_s}{4}}(AB|XYTE)_{\rho_{i\tilde{\Omega}}} - H_{\max}^{\frac{\varepsilon_s}{4}}(B|AXYTE)_{\rho_{i\tilde{\Omega}}} \\ &\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) \\ &\geq H_{\min}^{\frac{\varepsilon_s}{4}}(AB|XYTE)_{\rho_{i\tilde{\Omega}}} - H_{\max}^{\frac{\varepsilon_s}{4}}(B|TE)_{\rho_{i\tilde{\Omega}}} \\ &\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right), \end{aligned}$$

where the first inequality is due to the chain rule [21, Eq. (6.57)] and the second is due to strong sub-additivity of the smooth max-entropy.

One can now apply the EAT to upper bound $H_{\max}^{\frac{\varepsilon_s}{4}}(B|TE)_{\rho_{i\tilde{\Omega}}}$ in the following way. We use Theorem 9.3 with the replacements $\mathbf{O} \rightarrow \mathbf{B}$, $\mathbf{S} \rightarrow \mathbf{T}$, $\mathbf{E} \rightarrow \mathbf{E}$. The Markov conditions $B_{1,\dots,i-1} \leftrightarrow T_{1,\dots,i-1}E \leftrightarrow T_i$ then trivially hold and the condition on the max-tradeoff function reads

$$f_{\max}(p) \geq \sup_{\sigma_{R_{i-1}R'}: \mathcal{M}_i(\sigma)_{w_i}=p} H(B_i|T_i R')_{\mathcal{M}_i(\sigma)}.$$

By the definition of the EAT channels $\{\mathcal{M}_i\}_{i \in [n]}$, $B_i \neq \perp$ only for $T_i = 1$, which happens with probability γ .¹⁶ Hence, for any state $\sigma_{R_{i-1}R'}$ we have,

$$H(B_i|T_i R')_{\mathcal{M}_i(\sigma)} \leq H(B_i|T_i)_{\mathcal{M}_i(\sigma)} \leq \gamma$$

and the max-tradeoff function is simply $f_{\max}(p) = \gamma$ for any p (and thus $\|\nabla f_{\max}\|_{\infty} = 0$). Applying¹⁷ Theorem 9.3 with this choice of f_{\max} we get

$$H_{\max}^{\frac{\varepsilon_s}{4}}(B|TE)_{\rho_{i\tilde{\Omega}}} < \gamma n + \sqrt{n} 2 \log 7 \sqrt{1 - 2 \log(\varepsilon_s/4 \cdot (\varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}))}. \quad (11.21)$$

Combing the equations above we get that

$$\begin{aligned} H_{\min}^{\varepsilon_s}(A|XYTOE)_{\tilde{\rho}_{i\tilde{\Omega}}} &\geq H_{\min}^{\frac{\varepsilon_s}{4}}(AB|XYTE)_{\rho_{i\tilde{\Omega}}} - \text{leak}_{\text{EC}} \\ &\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) - \gamma n \\ &\quad - \sqrt{n} 2 \log 7 \sqrt{1 - 2 \log(\varepsilon_s/4 \cdot (\varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}))}. \end{aligned}$$

¹⁶This is why we made the distinction between B_i in the entropy accumulation protocol and \tilde{B}_i in the DIQKD protocol.

¹⁷Here a slightly more general version of the EAT than the one given in Sect. 9.2.3 is needed, in which the event Ω can be defined via A, B, X, Y and not only C ; see [22] for the details.

Finally, note that by applying the EAT on $\rho_{|\tilde{\Omega}}$, as in Theorem 11.5, we have that either $1 - \Pr(\tilde{\Omega}) \geq 1 - \varepsilon_{EA} - \varepsilon_{EC}$, or

$$H_{\min}^{\frac{\varepsilon_s}{4}}(\mathbf{AB}|\mathbf{XYTE})_{\rho_{|\tilde{\Omega}}} > n \cdot \mu_{\text{opt}}(\varepsilon_s/4, \varepsilon_{EA} + \varepsilon_{EC}) .$$

The last two equations together give us the desired bound on $H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYTOE})_{\tilde{\rho}_{|\tilde{\Omega}}}$: either the protocol aborts with probability greater than $1 - \varepsilon_{EA} - \varepsilon_{EC}$ or

$$\begin{aligned} H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYTOE})_{\tilde{\rho}_{|\tilde{\Omega}}} &\geq n \cdot \mu_{\text{opt}}(\varepsilon_s/4, \varepsilon_{EA} + \varepsilon_{EC}) - \text{leak}_{EC} \\ &\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) - \gamma n \\ &\quad - \sqrt{n} 2 \log 7 \sqrt{1 - 2 \log(\varepsilon_s/4 \cdot (\varepsilon_{EA} + \varepsilon_{EC}))} . \end{aligned}$$

□

Using Lemma 11.8, we prove that Protocol 11.2 is sound.

Lemma 11.9 *For any device D let $\tilde{\rho}$ be the state generated using Protocol 11.2. Then either the protocol aborts with probability greater than $1 - \varepsilon_{EA} - \varepsilon_{EC}$ or it is $(\varepsilon_{EC} + \varepsilon_{PA} + \varepsilon_s)$ -correct-and-secret while producing keys of length ℓ , as defined in Eq. (11.16).*

Proof Denote all the classical public communication during the protocol by $J = \mathbf{XYTOS}$ where S is the seed used in the privacy amplification protocol PA. Denote the final state of Alice, Bob, and Eve at the end of Protocol 11.2, conditioned on not aborting, by $\tilde{\rho}_{\tilde{K}_A \tilde{K}_B J E | \tilde{\Omega}}$.

We consider two cases. First assume that the EC protocol was not successful (but did not abort). Then Alice and Bob's final keys might not be identical. This happens with probability at most ε_{EC} .

Otherwise, assume the EC protocol was successful, i.e., $K_B = A$. In that case, Alice and Bob's keys must be identical also after the final privacy amplification step. That is, conditioned on $K_B = A$, $\tilde{K}_A = \tilde{K}_B$.

We continue to show that in this case the key is also secret. The secrecy depends only on the privacy amplification step, and for universal hashing a secure key is produced as long as

$$\ell = H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYTOE}) - 2 \log \frac{1}{\varepsilon_{PA}}$$

holds (recall Sect. 4.2.2). Hence, a uniform and independent key of length ℓ as in Eq. (11.16) is produced by the privacy amplification step unless the smooth min-entropy is not high enough (i.e., the bound in Eq. (11.18) does not hold) or the privacy amplification protocol was not successful, which happens with probability at most $\varepsilon_{PA} + \varepsilon_s$.

According to Lemma 11.8, either the protocol aborts with probability greater than $1 - \varepsilon_{\text{EA}} - \varepsilon_{\text{EC}}$, or the entropy is sufficiently high for us to have (recall Definition 4.6)

$$\|\tilde{\rho}_{\tilde{\kappa}_A J E|\hat{\Omega}} - \rho_{U_i} \otimes \tilde{\rho}_{J E}\|_1 \leq \varepsilon_{\text{PA}} + \varepsilon_s .$$

Combining both cases above the lemma follows. \square

11.3.3 Key Rate Analysis

Theorem 11.6 establishes a relation between the length ℓ of the secure key produced by our protocol and the different error terms. As this relation, given in Eq. (11.16), is somewhat hard to visualise, we analyse the key rate $r = \ell/n$ for some specific choices of parameters and compare it to the key rates achieved in device-dependent QKD with finite resources [23, 24] and DIQKD with infinite resources and a restricted set of attacks [12].

The key rate depends on the amount of leakage of information due to the error correction step, which in turn depends on the honest implementation of the protocol (recall Sect. 4.2.2). We use the honest IID implementation described in Sect. 4.2.4 and choose the honest state of each round to be the two-qubit Werner state $\rho_{Q_A Q_B} = (1 - \nu)|\phi^+\rangle\langle\phi^+| + \nu\mathbb{I}/4$ (and the measurements are as described in Sect. 4.2.4). The quantum bit error rate is then $Q = \frac{\nu}{2}$ and the expected winning probability is $\omega_{\text{exp}} = \frac{2 + \sqrt{2(1-2Q)}}{4}$.

We emphasise that this is only a choice of the *honest* implementation and it does not in any way restrict the actions of the adversary (and, in particular, the types of imperfections in the device). Furthermore, the analysis done below can be adapted to any other honest implementation of interest.

11.3.3.1 Leakage Due to Error Correction

To calculate the rates we first need to explicitly upper bound the leakage of information due to the error correction protocol, leak_{EC} . As shown in Eq. (4.2), this can be done by evaluating $H_0^{\varepsilon'_{\text{EC}}}(A|\tilde{\mathbf{B}}XYT)$ on Alice and Bob's state in an honest IID implementation of the protocol, described in Sect. 4.2.4.

For this we first use the following relation between H_0^{ε} and $H_{\text{max}}^{\varepsilon}$ [25, Lemma 18]:

$$H_0^{\varepsilon'_{\text{EC}}}(A|\tilde{\mathbf{B}}XYT) \leq H_{\text{max}}^{\frac{\varepsilon'_{\text{EC}}}{2}}(A|\tilde{\mathbf{B}}XYT) + \log(8/\varepsilon_{\text{EC}}^2 + 2/(2 - \varepsilon'_{\text{EC}})) .$$

The quantum asymptotic equipartition property, given as Theorem 7.3, tells us that

$$H_{\text{max}}^{\frac{\varepsilon'_{\text{EC}}}{2}}(A|\tilde{\mathbf{B}}XYT) \leq nH(A_i|\tilde{B}_i X_i Y_i T_i) + \sqrt{n}\delta(\varepsilon'_{\text{EC}}, \tau) ,$$

for $\tau = 2\sqrt{2^{H_{\max}(A_i|\tilde{B}_i X_i Y_i T_i)} + 1}$ and $\delta(\varepsilon'_{\text{EC}}, \tau) = 4 \log \tau \sqrt{2 \log(8/\varepsilon'^2_{\text{EC}})}$.

For the honest implementation of Protocol 11.2 $H_{\max}(A_i|\tilde{B}_i X_i Y_i T_i) = 1$ and

$$\begin{aligned} H(A_i|\tilde{B}_i X_i Y_i T_i) &= \Pr(T_i = 0) \cdot H(A_i|\tilde{B}_i X_i Y_i, T_i = 0) + \\ &\quad \Pr(T_i = 1) \cdot H(A_i|\tilde{B}_i X_i Y_i, T_i = 1) \\ &= (1 - \gamma) \cdot H(A_i|\tilde{B}_i X_i Y_i, T_i = 0) + \\ &\quad \gamma \cdot H(A_i|\tilde{B}_i X_i Y_i, T_i = 1) \\ &= (1 - \gamma) h(Q) + \gamma h(\omega_{\text{exp}}), \end{aligned}$$

where the first equality follows from the definition of conditional entropy and the second from the way T_i is chosen in Protocol 11.2. The last equality holds since for the generation rounds the error rate (i.e., the probability that A_i and \tilde{B}_i differ) in the honest case is Q and for the test rounds Bob can guess A_i with probability ω_{exp} given \tilde{B}_i , X_i , and Y_i .

We thus have

$$\begin{aligned} H_0^{\varepsilon'_{\text{EC}}}(A|\tilde{B}XYT) &\leq n [(1 - \gamma) h(Q) + \gamma h(\omega_{\text{exp}})] \\ &\quad + \sqrt{n} 4 \log(2\sqrt{2} + 1) \sqrt{2 \log(8/\varepsilon'^2_{\text{EC}})} \\ &\quad + \log(8/\varepsilon'^2_{\text{EC}} + 2/(2 - \varepsilon'_{\text{EC}})) . \end{aligned}$$

Plugging this into Eq. (4.2) we get

$$\begin{aligned} \text{leak}_{\text{EC}} &\leq n [(1 - \gamma) h(Q) + \gamma h(\omega_{\text{exp}})] \\ &\quad + \sqrt{n} 4 \log(2\sqrt{2} + 1) \sqrt{2 \log(8/\varepsilon'^2_{\text{EC}})} \\ &\quad + \log(8/\varepsilon'^2_{\text{EC}} + 2/(2 - \varepsilon'_{\text{EC}})) + \log\left(\frac{1}{\varepsilon_{\text{EC}}}\right) . \end{aligned} \tag{11.22}$$

11.3.3.2 Key Rate Curves

In Appendix C.2 a slightly modified protocol is considered in which, instead of fixing the number of rounds in the protocol, only the expected number of rounds is fixed. The completeness and soundness proofs follow the same lines as the proofs above, as detailed in Appendix C.2 and do not include additional crucial insights. The modification of the protocol improves the dependency of the key rate on the probability of a test round γ ¹⁸ The analysis presented in the appendix leads to the

¹⁸The second order term of the smooth min-entropy rate given in Lemma 11.8 scales with γ , roughly, as $1/\gamma$, while in Appendix C.2 the dependency is roughly $1/\sqrt{\gamma}$. The modified analysis can be seen as a ‘‘patch’’ used to overcome the non-optimal dependency of the EAT given in Theorem 9.3 on the

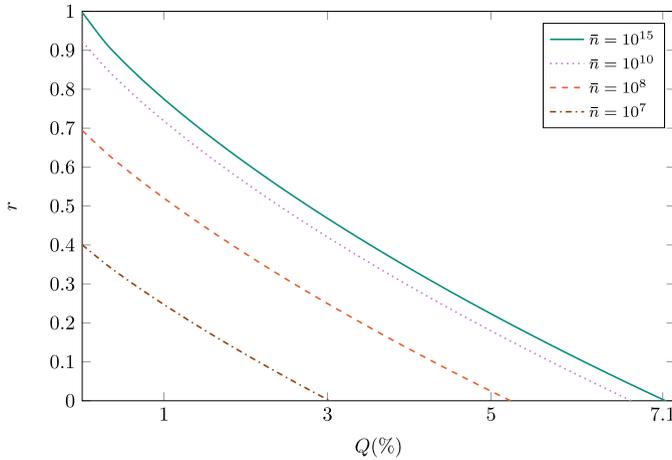


Fig. 11.3 The expected key rate $r = \ell/\bar{n}$ as a function of the quantum bit error rate Q for several values of the expected number of rounds \bar{n} (see the main text and Appendix C.2). For $\bar{n} = 10^{15}$ the curve essentially coincides with the curve for the IID asymptotic case [12, Eq. (12)]. The following values for the error terms were chosen: $\varepsilon_{\text{EC}} = 10^{-10}$, $\varepsilon_{\text{QKD}}^s = 10^{-5}$ and $\varepsilon_{\text{QKD}}^c = 10^{-2}$. After [1, 6]

key rates for the modified protocol and these are the rates presented here. Putting the technical details aside, the reader may simply think of \bar{n} below as taking the place of the number of rounds n used so far.

In an asymptotic analysis ($\bar{n} \rightarrow \infty$) it is well understood that the soundness and completeness errors $\varepsilon_{\text{QKD}}^s$, $\varepsilon_{\text{QKD}}^c$ should tend to zero as \bar{n} increases. However, in the non-asymptotic scenario considered here these errors are always finite. We therefore fix some values for them which are considered to be realistic and relevant for actual applications. We choose the parameters such that the security parameters are at least as good (and in general even better) as in [23], such that a fair comparison can be made. All other parameters are chosen in a consistent way while (roughly) optimising the key rate.

In Fig. 11.3 the expected key rate $r = \ell/\bar{n}$ is plotted as a function of the quantum bit error rate Q for several values of the expected number of rounds \bar{n} . For $\bar{n} = 10^{15}$ the curve essentially coincides with the rate achieved in the asymptotic IID case [12]. Since the latter was shown to be optimal [12] it provides an upper bound on the key rate and the amount of tolerable noise. Hence, for large enough \bar{n} our rates become optimal and the protocol can tolerate up to the maximal error rate $Q = 7.1\%$. For comparison, the previously established explicit rates [19] are well below the lowest curve presented in Fig. 11.3, even when the number of signals goes to infinity, with a maximal noise tolerance of 1.6%.

testing probability in the considered protocols. This issue was overcome in a more recent version of the EAT [26].

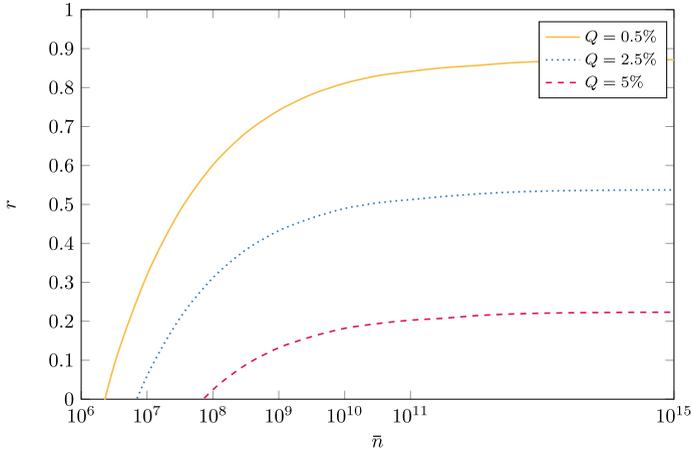


Fig. 11.4 The expected key rate $r = \ell/\bar{n}$ as a function of the expected number of rounds \bar{n} (see the main text and Appendix C.2) for several values of the quantum bit error rate Q . For $Q = 0.5\%$, 2.5% , and 5% the achieved key rates are approximately $r = 87\%$, 53% , and 22% respectively. The following values for the error terms were chosen: $\varepsilon_{\text{EC}} = 10^{-10}$, $\varepsilon_{\text{QKD}}^s = 10^{-5}$ and $\varepsilon_{\text{QKD}}^c = 10^{-2}$. After [1, 6]

In Fig. 11.4, r is plotted as a function of \bar{n} for several values of Q . As can be seen from the figure, the achieved rates are significantly higher than those achieved in previous works. Moreover, they are practically comparable to the key rates achieved in device-dependent QKD (see Fig. 1 in [23]). The main difference between the curves for the device-dependent case and the independent one is the minimal value of \bar{n} which is required for a positive key rate. (That is, for the protocols considered in [23] one can get a positive key rate with less rounds.)

11.4 Open Questions

To end the chapter, we list some future work directions and open questions specific for the showcase of quantum cryptography.

11.4.1 Experimental Realisations

The results presented in this chapter provide the theoretical groundwork for experimental implementations of device-independent cryptographic protocols. The quantitative results imply that the first proof of principle experiments, with small distances and small rates, are within reach with today's state-of-the-art technology, which recently enabled the violation of Bell inequalities in a loophole-free way [27–29] (a

necessity for device-independent cryptography). Indeed, Theorem 11.5 has already been applied to the analysis of the first experimental implementation of a protocol for randomness generation in the fully device-independent framework [5]. The next major challenge in experimental implementations is a field demonstration of a DIQKD protocol. This would provide the strongest cryptographic experiment ever realised.

As can be seen from Figs. 11.2, 11.3, and 11.4, implementing a DIQKD protocol is more challenging than implementing a randomness generation protocol—positive key rates require higher number of signals and lower noise levels. It therefore becomes increasingly relevant to achieve the best possible dependence of the rate curves on the number of rounds \bar{n} , even for very small values of \bar{n} . As can be seen from the figures our rate curves approach (and essentially coincide) with the optimal curves as the number of rounds increases. This is the case since our first-order term of the key rate is tight.

However, one thing that can perhaps still be further optimised is the dependency on the number of rounds, or in other words, how fast the curves approach the asymptotic curve. Although this seems like a minor issue, it can make actual implementations more feasible. The explicit dependency on \bar{n} given in Eq. (11.16) is already close to optimal and we did not try to optimise it. It can still be improved and several efforts were already made in the direction of achieving a better second order term.

1. A refined entropy accumulation theorem with a tighter second order term was developed in [26] following our work. Building on the improved theorem will lead directly to stronger quantitative results for finite and relatively small values of n .
2. In [30] a randomness expansion protocol was proven to be secure using our technique, while performing a more detailed analysis of the second order term. This, in particular, allowed for a more efficient experimental implementation of their protocol.
3. A different proof technique, termed “quantum probability estimators”, for bounding the smooth min-entropy was recently developed in [31, 32]. The entropy accumulation and the probability estimators approaches are closely related (see [31] for a discussion) but the exact relation between them is not fully understood to date. The technique of probability estimators has the potential of leading to better second order terms, possibly due to a more fine-tuned use of concentration bounds.

11.4.2 Possible Extensions

The optimality of our key rates is only with respect to the structure of the considered protocol, which is the standard (and only, as far as we are aware) DIQKD protocol studied in the literature. It is interesting to come up with new DIQKD protocols and see if they lead to key rates with higher first-order terms. Apart from the theoretical

curiosity, protocols with better asymptotic key rates can, of course, help us reach an experimental implementation.

Due to the modularity of our analysis, it can at large be directly applied to the analysis of other protocols. The main challenge is to come up with interesting protocols. We discuss several possible directions to consider.

On the “quantum side” of the protocol, one may modify the protocol by considering different Bell inequalities. Even more, one can construct protocols in which more information than the violation of a single inequality is used: Alice and Bob may use the collected statistics to evaluate several quantities and decide accordingly whether to abort or not; see for example the related work [33]. To apply our proof to other Bell inequalities and additional statistical information one should find a good bound on the min-tradeoff function, as done in Eq. (11.7) for the CHSH inequality. For many Bell inequalities such bounds are known, but for the min-entropy instead of the von Neumann entropy. In most cases using a bound on the min-entropy will result in far from optimal rate curves. Therefore, to adapt our protocol in this direction one should probably first bound the min-tradeoff function using the von Neumann entropy directly.¹⁹

On the “classical side” of the protocol, different classical post-processing steps can be considered. It is known that, asymptotically, considering protocols with other one-way classical post-processing cannot lead to an improvement over our protocol [34]. Hence, the interesting thing to check is whether there are protocols with either classical *pre*-processing of the data (i.e., prior to applying error correction and privacy amplification) and/or *two-way* classical post-processing protocols that lead to an improvement of the first-order term of the key rates. Two such protocols were suggested recently: [35, 36].

Of course, it is also interesting different protocols in which the entropy of a single round is certified by considering a more general setting than the one of the protocol analysed here. An example is the novel protocol of [37], where the randomness is produced from two different measurements settings rather than one. Once the single-round of the protocol is done [37], our framework is applied directly to get the full security proof.

An interesting related question is that of finding *upper*-bounds on the possible key rates that can be extracted from a given correlation. Such upper-bounds show us the limitations of a whole class of protocols all together, without needing to analyse each protocol separately and in a sense tell us what is the best that we can expect. This question was recently studied in [38–41].

¹⁹This should not be dismissed as can be seen from the following state of affairs. [33] reports an advantage in terms of the *min-entropy* when considering the full statistics instead of merely the violation of the CHSH inequality. Comparing the bound on the min-entropy from the full statistics to the bound on the von Neumann entropy from the violation alone, both evaluated on the quantum states produced by the honest implementations, we find that it is still better to use the bound on the von Neumann entropy as we do here. Thus, to truly see if an advantage can be gained by considering the full statistics, one should aim to a direct bound on the von Neumann entropy.

11.4.3 Bounding the von Neumann Entropy

As mentioned above, the main task to perform when modifying the protocol and the analysis is to lower-bound the min-tradeoff function. Getting a tight bound on the von Neumann entropy, and hence on the min-tradeoff function, does not seem to be an easy task. Is there a numerical technique that allows one to get good bounds on the von Neumann entropy for general Bell inequalities?

Interestingly, as it turns out, good numerical tools are known for a couple of similar quantities:

1. When considering the min-entropy instead of the von Neumann entropy one can use SDP hierarchies to get (not necessarily tight) lower-bounds [42].
2. In the device-*dependent* case, a recent development [43] presents a numerical technique to lower-bound the von Neumann entropy and by this derive better key rates for QKD.

We hope that it is possible to devise a general technique (numerical or analytical) to calculate good lower bounds on the von Neumann entropy relevant for our case. Two works that go in this direction are [44, 45]. Such a technique, in combination with our work, will allow us to “enumerate” over all possible protocols and calculate their key rates when looking for better protocols. If it is not possible to devise a general tool, it will at least be interesting to understand why this is the case. Is there a fundamental mathematical reason behind the complexity of the problem?

References

1. Arnon-Friedman R, Renner R, Vidick T (2019) Simple and tight device-independent security proofs. *SIAM J Comput* 48(1):181–225
2. Ribeiro J, Murta G, Wehner S (2017) Fully device independent conference key agreement. [arXiv:1708.00798](https://arxiv.org/abs/1708.00798)
3. Kessler M, Arnon-Friedman R (2017) Device-independent randomness amplification and privatization. [arXiv:1705.04148](https://arxiv.org/abs/1705.04148)
4. Bamps C, Massar S, Pironio S (2017) Device-independent randomness generation with sub-linear shared quantum resources. [arXiv:1704.02130](https://arxiv.org/abs/1704.02130)
5. Liu Y, Yuan X, Li M-H, Zhang W, Zhao Q, Zhong J, Cao Y, Li Y-H, Chen L-K, Li H et al. (2017) High speed self-testing quantum random number generation without detection loophole. In: *Frontiers in optics*, p FTh2E–1. Optical Society of America
6. Arnon-Friedman R, Dupuis F, Fawzi O, Renner R, Vidick T (2018) Practical device-independent quantum cryptography via entropy accumulation. *Nat Commun* 9(1):459
7. Acín A, Gisin N, Masanes L (2006) From Bell’s theorem to secure quantum key distribution. *Phys Rev Lett* 97(12):120405
8. Acín A, Massar S, Pironio S (2006) Efficient quantum key distribution secure against no-signalling eavesdroppers. *New J Phys* 8(8):126
9. Scarani V, Gisin N, Brunner N, Masanes L, Pino S, Acín A (2006) Secrecy extraction from no-signaling correlations. *Phys Rev A* 74(4):042339
10. Acín A, Brunner N, Gisin N, Massar S, Pironio S, Scarani V (2007) Device-independent security of quantum cryptography against collective attacks. *Phys Rev Lett* 98(23):230501

11. Masanes L (2009) Universally composable privacy amplification from causality constraints. *Phys Rev Lett* 102(14):140501
12. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021
13. Hänggi E, Renner R, Wolf S (2010) Efficient device-independent quantum key distribution. *Advances in cryptology-EUROCRYPT 2010*. Springer, Berlin, pp 216–234
14. Hänggi E, Renner R (2010) Device-independent quantum key distribution with commuting measurements. [arXiv:1009.1833](https://arxiv.org/abs/1009.1833)
15. Masanes L, Pironio S, Acín A (2011) Secure device-independent quantum key distribution with causally independent measurement devices. *Nat Commun* 2:238
16. Masanes L, Renner R, Christandl M, Winter A, Barrett J (2014) Full security of quantum key distribution from no-signaling constraints. *IEEE Trans Inf Theory* 60(8):4973–4986
17. Clauser JF, Horne MA, Shimony A, Holt RA (1969) Proposed experiment to test local hidden-variable theories. *Phys Rev Lett* 23(15):880
18. Reichardt BW, Unger F, Vazirani U (2013) Classical command of quantum systems. *Nature* 496(7446):456–460
19. Vazirani U, Vidick T (2014) Fully device-independent quantum key distribution. *Phys Rev Lett* 113(14):140501
20. Miller CA, Shi Y (2014) Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. In: *Proceedings of the 46th annual ACM symposium on theory of computing*, pp 417–426. ACM
21. Tomamichel M (2015) *Quantum information processing with finite resources: mathematical foundations*, vol 5. Springer, Berlin
22. Dupuis F, Fawzi O, Renner R (2016) Entropy accumulation. [arXiv:1607.01796](https://arxiv.org/abs/1607.01796)
23. Scarani V, Renner R (2008) Quantum cryptography with finite resources: unconditional security bound for discrete-variable protocols with one-way postprocessing. *Phys Rev Lett* 100(20):200501
24. Scarani V, Renner R (2008) Security bounds for quantum cryptography with finite resources. *Theory of quantum computation, communication, and cryptography*. Springer, Berlin, pp 83–95
25. Tomamichel M, Schaffner C, Smith A, Renner R (2011) Leftover hashing against quantum side information. *IEEE Trans Inf Theory* 57(8):5524–5535
26. Dupuis F, Fawzi O (2019) Entropy accumulation with improved second-order term. *IEEE Trans Inf Theory* 65(11):7596–7612
27. Hensen B, Bernien H, Dréau A, Reiserer A, Kalb N, Blok M, Ruitenber J, Vermeulen R, Schouten R, Abellán C et al (2015) Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* 526(7575):682–686
28. Shalm LK, Meyer-Scott E, Christensen BG, Bierhorst P, Wayne MA, Stevens MJ, Gerrits T, Glancy S, Hamel DR, Allman MS et al (2015) Strong loophole-free test of local realism. *Phys Rev Lett* 115(25):250402
29. Giustina M, Versteegh MA, Wengerowsky S, Handsteiner J, Hochrainer A, Phelan K, Steinlechner F, Kofler J, Larsson J-Å, Abellán C et al (2015) Significant-loophole-free test of Bell’s theorem with entangled photons. *Phys Rev Lett* 115(25):250401
30. Liu W-Z, Li M-H, Ragy S, Zhao S-R, Bai B, Liu Y, Brown PJ, Zhang J, Colbeck R, Fan J et al (2019) Device-independent randomness expansion against quantum side information. [arXiv:1912.11159](https://arxiv.org/abs/1912.11159)
31. Zhang Y, Fu H, Knill E (2020) Efficient randomness certification by quantum probability estimation. *Phys Rev Res* 2(1):013016
32. Zhang Y, Shalm LK, Bienfang JC, Stevens MJ, Mazurek MD, Nam SW, Abellán C, Amaya W, Mitchell MW, Fu H et al (2020) Experimental low-latency device-independent quantum randomness. *Phys Rev Lett* 124(1):010505
33. Nieto-Silleras O, Bamps C, Silman J, Pironio S (2018) Device-independent randomness generation from several bell estimators. *New J Phys* 20(2):023049
34. Devetak I, Winter A (2005) Distillation of secret key and entanglement from quantum states. In: *Proceedings of the Royal Society of London A: mathematical, physical and engineering sciences*, vol 461, pp 207–235. The Royal Society

35. Ho M, Sekatski P, Tan E-Z, Renner R, Bancal J-D, Sangouard N (2020) Noisy preprocessing facilitates a photonic realization of device-independent quantum key distribution. *Phys Rev Lett* 124(23):230502
36. Tan EY-Z, Lim CC-W, Renner R (2020) Advantage distillation for device-independent quantum key distribution. *Phys Rev Lett* 124(2):020502
37. Schwonnek R, Goh KT, Primateamaja IW, Tan EY-Z, Wolf R, Scarani V, Lim CC-W (2020) Robust device-independent quantum key distribution. [arXiv:2005.02691](https://arxiv.org/abs/2005.02691)
38. Kaur E, Wilde MM, Winter A (2020) Fundamental limits on key rates in device-independent quantum key distribution. *New J Phys* 22(2):023039
39. Arnon-Friedman R, Leditzky F (2020) Upper bounds on device-independent quantum key distribution rates and a revised peres conjecture. [arXiv:2005.12325](https://arxiv.org/abs/2005.12325)
40. Christandl M, Ferrara R, Horodecki K (2020) Upper bounds on the rate in device-independent quantum key distribution. [arXiv:2005.13511](https://arxiv.org/abs/2005.13511)
41. Winczewski M, Das T, Horodecki K (2019) Upper bounds on secure key against non-signaling adversary via non-signaling squashed secrecy monotones. [arXiv:1903.12154](https://arxiv.org/abs/1903.12154)
42. Navascués M, Pironio S, Acín A (2008) A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J Phys* 10(7):073013
43. Winick A, Lütkenhaus N, Coles PJ (2017) Reliable numerical key rates for quantum key distribution. [arXiv:1710.05511](https://arxiv.org/abs/1710.05511)
44. Wang Y, Primateamaja IW, Lavie E, Varvitsiotis A, Lim CCW (2019) Characterising the correlations of prepare-and-measure quantum networks. *npj Quantum Inf* 5(1):1–6
45. Tan EY-Z, Schwonnek R, Goh KT, Primateamaja IW, Lim CC-W (2019) Computing secure key rates for quantum key distribution with untrusted devices. [arXiv:1908.11372](https://arxiv.org/abs/1908.11372)

Chapter 12

Outlook



The development and application of the concept of reductions to IID, taking the form of de Finetti theorems, flourished in “standard” quantum information processing in the last decade and more. The tools used, unfortunately, were not applicable when considering device-independent information processing tasks, where the devices being analysed are uncharacterised. The reductions presented in the thesis, namely the de Finetti reduction (Chap. 8) and the entropy accumulation theorem (Chap. 9), are the first to be applicable in the device-independent setting. As such, they have opened the possibility of a significantly simpler analysis of device-independent information processing tasks.

Among the advantages of applying the approach of reductions to IID in the device-independent setting, compared to directly analysing the most general case, are tighter quantitative results and modular proofs. The thesis’ showcases, used to exemplify the usage of the reductions, indeed report such benefits. Our proof of non-signalling parallel repetition (Chap. 10) is automatically valid for any complete-support game with any number of players and achieves an exponential decrease that matches that of IID strategies. Our security proof for device-independent quantum key distribution (Chap. 11) achieves tight key rates, as under the IID assumption, that are significantly better than all prior results and can be easily adapted to other related protocols.

With this in mind, it is interesting to investigate how the presented reductions or variants thereof can be used in the analysis of other tasks. Let us discuss a partial list of questions and possible future work that we find intriguing.¹

¹We list here questions that are not directly related to the showcases considered in the thesis. For concrete open questions regarding parallel repetition (e.g., extensions of the results) and device-independent quantum key distribution (such as possible improvements and experimental implementations) see Sects. 10.4 and 11.4, respectively.

12.1 Two-Party Device-Independent Quantum Cryptography

In the cryptographic protocols discussed in the thesis we considered two honest and cooperating parties, Alice and Bob. Two-party cryptography, on the other hand, refers to cryptographic protocols in which Alice and Bob do not trust each other. When considering device-independent two-party cryptography the dishonest party (which can be either Alice or Bob) takes the role of the adversary and hence is allowed to prepare the device used to implement the protocol. References [1, 2] present examples for such protocols.

The above mentioned works study the security of the protocols under the IID assumption (or a closely related assumption). Clearly, it is interesting to see if the analysis can be extended to capture the most general adversarial scenario, which, in the case of these protocols, includes the use of sequential boxes. Applying a reduction to IID can be beneficial here. Unfortunately, it is not clear whether the entropy accumulation theorem, in its current form, can be of use in such protocols. The reason is that the Markov-chain conditions stated in Eq. (9.4) do not hold, at least when considering the most obvious choices of random variables.²

In some cases, one can overcome the problem by considering “imaginary” protocols, closely related to the “real” protocol, in which the Markov-chain conditions do hold. The idea is then to reduce the problem of proving the security of the real protocol to that of the imaginary one and perform the analysis of the imaginary protocol using the entropy accumulation theorem.

Such a proof technique is used in [3]. There, the protocol of interest is a device-independent entanglement certification protocol and its analysis requires an upper bound on the smooth max-entropy, rather than a lower bound on the smooth min-entropy as in cryptographic scenarios.³ Thus, the steps used in [3] are not directly applicable to two-party device-independent cryptography. It is interesting to see if similar ideas can be useful in cryptographic scenarios as well.

Alternatively, one could also try to prove a different variant of the entropy accumulation theorem in which the Markov-chain conditions are replaced by some other restrictions on the sequential process, which are fulfilled by two-party cryptographic protocols. (Finding such conditions is interesting by itself). As discussed in Sect. 9.2.1, some conditions on the process must appear in the theorem, since entropy does not accumulate in any sequential process. While the Markov-chain conditions are sufficient, we currently have no reason to believe that they are necessary; it might as well be that some weaker or incomparable conditions also suffice.

²In the case of two-party cryptography, the natural choice to make when trying to use the entropy accumulation theorem is one in which the \mathcal{O} systems belong to the honest party and the \mathcal{S} systems to the dishonest party. One can easily come up with boxes that do not fulfil Eq. (9.4) with these choices.

³In the considered scenarios the two quantities are not dual to one another; see [3] for the details.

12.2 Parallel Device-Independent Quantum Cryptography

Another type of cryptographic protocols to which the presented reductions to IID are not applicable in a trivial manner are ones in which the most general analysis should be done with quantum parallel boxes. An example is the parallel device-independent quantum key distribution protocol of [4], in which all the non-local games are played in parallel with the device (as in the parallel repetition question). While [4] includes a security proof that goes beyond the IID scenario, it achieves quantitatively weak results. This raises the fundamental question of whether parallel adversaries, i.e., adversaries that can create parallel boxes, are stronger than sequential and IID adversaries (which are proven to have the same strength by our work). To learn the answer to this question there is a need to supply tight key rates for parallel device-independent quantum key distribution protocols.

Utilising a reduction to IID instead of analysing the general case directly, as in [4], will almost surely lead to stronger, perhaps even tight, results. Alas, the known reductions are not directly applicable here. The entropy accumulation theorem is not useful in this case since it is restricted to sequential boxes and here one ought to analyse parallel boxes. The de Finetti reduction, while suitable for parallel boxes, is a priori not applicable here since the de Finetti box does not include the adversary and is not a quantum box; see Sect. 8.4. It is therefore interesting to investigate whether the analysis can somehow be manipulated so that the known techniques can be utilised to prove security of parallel device-independent quantum cryptography or, otherwise, whether other types of reductions, more adequate for such scenarios, can be developed.

12.3 Device-Independent Tomography

One of the applications of the “original” quantum de Finetti reduction (also called the post-selection technique) [5] is a technique for a reliable quantum state tomography [6]. The technique is said to be reliable since it reports not just an estimation of the quantum state but also a confidence region around the estimated state, which acts as a meaningful “error bar”. This is of crucial importance as the other more standard approaches, such as the maximum-likelihood optimisation and least-square-error estimation, suffer from systematic errors [7].

Recently, the device-independent equivalents of the maximum-likelihood optimisation and least-square-error estimation were considered in [8]. The goal of such device-independent tomographic techniques is to report an estimated quantum box from the observed finite statistics. Apart from systematic errors, device-independent tomographic procedures as above are also at risk of providing a non-quantum box, since up to date it is unknown how to perform optimisation problems over the set of quantum boxes. In analogy to [6], applying our de Finetti reductions to achieve reliable device-independent tomography can therefore be of interest.

References

1. Fu H, Miller CA (2018) Local randomness: examples and application. *Phys Rev A* 97(3):032324
2. Ribeiro J, Kaniewski J, Helsen J, Wehner S et al (2018) Device independence for two-party cryptography and position verification with memoryless devices. *Phys Rev A* 97(6):062307
3. Arnon-Friedman R, Bancal J-D (2019) Device-independent certification of one-shot distillable entanglement. *New J Phys* 21(3):033010
4. Jain R, Miller CA, Shi Y (2017) Parallel device-independent quantum key distribution. [arXiv:1703.05426](https://arxiv.org/abs/1703.05426)
5. Christandl M, König R, Renner R (2009) Postselection technique for quantum channels with applications to quantum cryptography. *Phys Rev Lett* 102(2):020504
6. Christandl M, Renner R (2012) Reliable quantum state tomography. *Phys Rev Lett* 109(12):120403
7. Schwemmer C, Knips L, Richart D, Weinfurter H, Moroder T, Kleinmann M, Gühne O (2015) Systematic errors in current quantum state tomography tools. *Phys Rev Lett* 114(8):080403
8. Lin P-S, Rosset D, Zhang Y, Bancal J-D, Liang Y-C (2018) Device-independent point estimation from finite data and its application to device-independent property estimation. *Phys Rev A* 97(3):032309

Appendix A

Additional Proofs: de Finetti Reductions

A.1 Bounding the de Finetti Box

We use the notation used in Sect. 8.2:

1. $|\mathcal{X}||\mathcal{Y}| = l$ and we identify each pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with a label $j \in [l]$ by writing $(x, y) = j$.
2. $|\mathcal{A}||\mathcal{B}| = m$ and we identify each pair $(a, b) \in \mathcal{A} \times \mathcal{B}$ with a label $k \in [m]$ by writing $(a, b) = k$.
3. For all $j \in [l]$ and $k \in [m]$, $p_k^j \in [0, 1]$ such that $\sum_k p_k^j = 1$ for all j .
4. For all $j \in [l]$ and $k \in [m]$, $c_k^j = 1 - \sum_{t < k} p_t^j$.
5. For all \mathbf{x}, \mathbf{y} , and $j \in [l]$, $n^j = |\{i : (x_i, y_i) = j\}|$, i.e., n^j denotes the number of indices of (\mathbf{x}, \mathbf{y}) in which the type of inputs is $(x, y) = j$.
6. For all $\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}$, $j \in [l]$, and $k \in [m]$, $n_k^j = |\{i : (x_i, y_i) = j \wedge (a_i, b_i) = k\}|$, i.e., n_k^j denotes the number of indices of $(\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b})$ in which the type of inputs is $(x, y) = j$ and the type of outputs is $(a, b) = k$.

and notice that:

1. For all $j \in [l]$ and $k \in [m - 1]$, $p_k^j \in [0, c_k^j]$ and $p_m^j = c_m^j$.
2. For all $j \in [l]$ and $k \in [m]$, $c_k^j = c_{k-1}^j - p_{k-1}^j$.
3. For all $j \in [l]$, $n_m^j = n^j - \sum_{k=1}^{m-1} n_k^j$.

Lemma 8.4 For all $\mathbf{a}, \mathbf{b}, \mathbf{x}$, and \mathbf{y} ,

$$\tau_{AB|XY}(\mathbf{ab}|\mathbf{xy}) \geq \prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1} \frac{1}{(n^j + 1)^{m-1}},$$

where $\tau_{AB|XY}$ is the de Finetti box defined by

$$\begin{aligned} \tau_{AB|XY}(\mathbf{a}, \mathbf{b} | \mathbf{x}, \mathbf{y}) &= \int \mathbf{O}_{AB|XY}^{\otimes n} d\mathbf{O}_{AB|XY} \\ &= \prod_{j=1}^l \left[\int_0^{c_1^j} \frac{dp_1^j}{c_1^j} (p_1^j)^{n_1^j} \right] \cdots \left[\int_0^{c_{m-1}^j} \frac{dp_{m-1}^j}{c_{m-1}^j} (p_{m-1}^j)^{n_{m-1}^j} \right] \cdot (p_m^j)^{n^j - \sum_{k=1}^{m-1} n_k^j}. \end{aligned}$$

In the proof of the lemma we use the following formula:

$$\begin{aligned} \forall c > 0 \forall n, n' \in \mathbb{N}, n' \leq n \\ \int_0^c \frac{dp}{c} p^{n'} (c-p)^{n-n'} &= c^n \int_0^1 q^{n'} (1-q)^{(n-n')} dq \\ &= c^n \mathbf{B}(n-n'+1, n'+1) \\ &= c^n \binom{n}{n'}^{-1} \frac{1}{n+1} \end{aligned} \tag{A.1}$$

where \mathbf{B} is the Beta function. We also need the following identity:

$$\binom{n - \sum_{t=1}^s n_t}{n_{s+1}} \cdot \binom{n}{n_1, \dots, n_s, n - \sum_{t=1}^s n_t} = \binom{n}{n_1, \dots, n_{s+1}, n - \sum_{t=1}^{s+1} n_t} \tag{A.2}$$

Proof Abusing notation we denote below, for $t \in [2, m-1]$,¹

$$\left[\int_0^{c_1^j} \frac{dp_1^j}{c_1^j} (p_1^j)^{n_1^j} \right] \cdots \left[\int_0^{c_t^j} \frac{dp_{m-1}^j}{c_t^j} (p_t^j)^{n_t^j} \right] = \prod_{k=1}^t \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right].$$

We start by proving the following, for all $j \in [l]$, by induction:

$$\begin{aligned} \prod_{k=1}^{m-1} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_{m-1}^j - p_{m-1}^j)^{n^j - \sum_{k=1}^{m-1} n_k^j} &\geq \\ \binom{n^j}{n_1^j, \dots, n_{m-1}^j, n^j - \sum_{k=1}^{m-1} n_k^j}^{-1} \frac{1}{(n^j + 1)^{m-1}} \end{aligned} \tag{A.3}$$

Base case, $m = 2$:

$$\int_0^{c_1^j} \frac{dp_1^j}{c_1^j} (p_1^j)^{n_1^j} [c_1^j - p_1^j]^{n^j - n_1^j} = \binom{n^j}{n_1^j}^{-1} \frac{1}{n^j + 1}$$

This follows from Eq. (A.1) while noting that for the first index we have $c_1^j = 1$ by definition.

¹This is just a notation and $\prod_{k=1}^l$ should not be understood as the product operation. In particular, the order of terms is relevant since the different parameters are not independent of one another.

Induction hypothesis for $m - 2$:

$$\prod_{k=1}^{m-2} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_{m-2}^j - p_{m-2}^j)^{n^j - \sum_{k=1}^{m-2} n_k^j} \geq \left(\binom{n^j}{n_1^j, \dots, n_{m-2}^j, n^j - \sum_{k=1}^{m-2} n_k^j} \right)^{-1} \frac{1}{(n^j + 1)^{m-2}} \quad (\text{A.4})$$

Inductive step:

$$\begin{aligned} & \prod_{k=1}^{m-1} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_{m-1}^j - p_{m-1}^j)^{n^j - \sum_{k=1}^{m-1} n_k^j} = \\ & \prod_{k=1}^{m-2} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] \int_0^{c_{m-1}^j} \frac{dp_{m-1}^j}{c_{m-1}^j} (p_{m-1}^j)^{n_{m-1}^j} (c_{m-1}^j - p_{m-1}^j)^{n^j - \sum_{k=1}^{m-2} n_k^j - n_{m-1}^j} = \end{aligned} \quad (\text{A.5})$$

$$\begin{aligned} & \prod_{k=1}^{m-2} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] \times \\ & \times (c_{m-1}^j)^{n^j - \sum_{k=1}^{m-2} n_k^j} \left(\binom{n^j - \sum_{k=1}^{m-2} n_k^j}{n_{m-1}^j} \right)^{-1} \frac{1}{n^j - \sum_{k=1}^{m-2} n_k^j + 1} = \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} & \left(\binom{n^j - \sum_{k=1}^{m-2} n_k^j}{n_{m-1}^j} \right)^{-1} \frac{1}{n^j - \sum_{k=1}^{m-2} n_k^j + 1} \times \\ & \times \prod_{k=1}^{m-2} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_{m-2}^j - p_{m-2}^j)^{n^j - \sum_{k=1}^{m-2} n_k^j} \geq \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} & \left(\binom{n^j - \sum_{k=1}^{m-2} n_k^j}{n_{m-1}^j} \right)^{-1} \frac{1}{n^j - \sum_{k=1}^{m-2} n_k^j + 1} \times \\ & \times \left(\binom{n^j}{n_1^j, \dots, n_{m-2}^j, n^j - \sum_{k=1}^{m-2} n_k^j} \right)^{-1} \frac{1}{(n^j + 1)^{m-2}} \geq \end{aligned} \quad (\text{A.8})$$

$$\left(\binom{n^j}{n_1^j, \dots, n_m^j} \right)^{-1} \frac{1}{(n^j + 1)^{m-1}}.$$

where we used Eq. (A.1) to get from (A.5) to (A.6), $c_{m-1}^j = c_{m-2}^j - p_{m-2}^j$ to get from (A.6) to (A.7), the induction hypothesis (A.4) to get from (A.7) to (A.8) and Eq. (A.2) as well as $n^j - \sum_k = 1^{m-2} + 1 \geq n^j + 1$ in the last line.

Finally, for any \mathbf{a} , \mathbf{b} , \mathbf{x} , and \mathbf{y} ,

$$\begin{aligned}
\tau_{AB|XY}(\mathbf{ab}|\mathbf{xy}) &= \\
\prod_{j=1}^l \prod_{k=1}^{m-1} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (p_m^j)^{n^j - \sum_{k=1}^{m-1} n_k^j} &= \\
\prod_{j=1}^l \prod_{k=1}^{m-1} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_m^j)^{n^j - \sum_{k=1}^{m-1} n_k^j} &= \\
\prod_{j=1}^l \prod_{k=1}^{m-1} \left[\int_0^{c_k^j} \frac{dp_k^j}{c_k^j} (p_k^j)^{n_k^j} \right] (c_{m-1}^j - p_{m-1}^j)^{n^j - \sum_{k=1}^{m-1} n_k^j} &\geq \\
\prod_{j=1}^l \binom{n^j}{n_1^j, \dots, n_m^j}^{-1} \frac{1}{(n^j + 1)^{m-1}} &
\end{aligned}$$

where we used Eq. (A.3) in the last step. \square

A.2 Diamond Norm Reduction

We prove Lemma 8.16

Lemma 8.16 *For every two permutation invariant channels $\mathcal{E}, \mathcal{F} : \mathcal{P} \rightarrow \mathcal{K}$ where $\mathbb{P}_{\mathcal{K}}$ is a probability distribution over $k \in \{0, 1\}^t$ for some $t > 0$, and all $\mathbb{P}_{ABC|XYZ}$,*

$$\|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1 \leq (n+1)^{l(m-1)} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC|XYZ}^{\mathbb{P}_{ABC|XYZ}})\|_1$$

where $\tau_{ABC|XYZ}^{\mathbb{P}_{ABC|XYZ}}$ is a non-signalling extension of $\tau_{AB|XY}$ which depends on the specific box $\mathbb{P}_{ABC|XYZ}$.

Proof First, as in the proof of Theorem 8.11, since the channels are permutation invariant it is sufficient to consider boxes $\mathbb{P}_{AB|XY}$ which are permutation invariant.

Given a specific box $\mathbb{P}_{ABC|XYZ}$ we can see this extension as a set of convex decompositions of $\mathbb{P}_{AB|XY}$, according to Lemma 8.9. That is, every possible input z induces a specific decomposition $\{(p_{c_z}, \mathbb{P}_{AB|XY}^{c_z})\}_{c_z}$ such that $p_{c_z} = \mathbb{P}_{C|Z}(c_z|z)$ and $\mathbb{P}_{AB|XY}^{c_z}(\mathbf{a}, \mathbf{b}|\mathbf{x}, \mathbf{y}) = \mathbb{P}_{ABC|XYZ}(\mathbf{a}, \mathbf{b}, c_z|\mathbf{x}, \mathbf{y}, z)$. Since this is a convex decomposition of $\mathbb{P}_{AB|XY}$ we also have

$$\forall z \quad \sum_c p_c \cdot \mathbb{P}_{AB|XY}^c = \mathbb{P}_{AB|XY} . \quad (\text{A.9})$$

We now use the set of decompositions of $\mathbb{P}_{AB|XY}$ to construct a set of decompositions of the de Finetti box $\tau_{AB|XY}$. Combining Lemmas 8.6, 8.8 and 8.9 together, we know that there exists a non-signalling box $\mathbb{R}_{AB|XY}$ such that

$$\begin{aligned}\tau_{AB|XY} &= \frac{1}{(n+1)^{l(m-1)}} \mathbb{P}_{AB|XY} + \left(1 - \frac{1}{(n+1)^{l(m-1)}}\right) \mathbb{R}_{AB|XY} \\ &= \frac{1}{(n+1)^{l(m-1)}} \sum_c p_c \cdot \mathbb{P}_{AB|XY}^c + \left(1 - \frac{1}{(n+1)^{l(m-1)}}\right) \mathbb{R}_{AB|XY},\end{aligned}$$

where the second equality is due to Eq. (A.9). For every z this defines a decomposition $\{(\frac{1}{(n+1)^{l(m-1)}} \cdot p_{c_z}, \mathbb{P}_{AB|XY}^{c_z})\}_{c_z} \cup \{(1 - \frac{1}{(n+1)^{l(m-1)}}, \mathbb{R}_{AB|XY})\}$ of $\tau_{AB|XY}$. That is, this defines an extension $\tau_{ABC'|XYZ}^{\mathbb{P}_{AB|XY}}$ of $\tau_{AB|XY}$ where $C' = C \cup \{c'\}$.

This connection between the extensions $\mathbb{P}_{ABC|XYZ}$ and $\tau_{ABC'|XYZ}^{\mathbb{P}_{AB|XY}}$ allow us to get the following bound on the trace distance, from which the lemma follows:

$$\|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC'|XYZ}^{\mathbb{P}_{AB|XY}})\|_1 \geq \frac{1}{(n+1)^{l(m-1)}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1. \quad (\text{A.10})$$

Equation (A.10) can be proven using the following sequence of steps. First, the diamond norm can be written in the following way.

$$\begin{aligned}\|\mathcal{E} - \mathcal{F}\|_\diamond &= \max_{\mathbb{P}_{ABC|XYZ}} \|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ})\|_1 \\ &= \max_{\mathbb{P}_{ABC|XYZ}} \|\mathbb{E}_{K|C} \cdot \mathbb{P}_{C|Z} - \mathbb{F}_{K|C} \cdot \mathbb{P}_{C|Z}\|_1 \\ &= \max_{\mathbb{P}_{ABC|XYZ}} \frac{1}{2} \sum_k \max_z \sum_c \mathbb{P}_{C|Z}(c|z) \left| \mathbb{E}_{K|C}(k|c) - \mathbb{F}_{K|C}(k|c) \right| \\ &= \max_{\mathbb{P}_{ABC|XYZ}} \frac{1}{2} \sum_k \max_z \sum_c \mathbb{P}_{C|Z}(c|z) \times \\ &\quad \times \left| \sum_{x,y} \Pr_{\mathcal{E}}(x, y) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \mathbb{P}_{AB|XYC}(a, b|x, y, c) - \right. \\ &\quad \left. \sum_{x,y} \Pr_{\mathcal{F}}(x, y) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \mathbb{P}_{AB|XYC}(a, b|x, y, c) \right|\end{aligned}$$

where the third equality is due to the explicit form of the trace distance previously given in [1, 2].

This can then be used to write

$$\begin{aligned}\|(\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\tau_{ABC'|XYZ}^{\mathbb{P}_{AB|XY}})\|_1 &= \frac{1}{2} \sum_k \max_z \sum_{c \in C'} \tau_{C'|Z}^{\mathbb{P}_{AB|XY}}(c|z) \left| \sum_{x,y} \Pr_{\mathcal{E}}(x, y) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\mathbb{P}_{AB|XY}}(ab|xy, c) \right. \\ &\quad \left. - \sum_{x,y} \Pr_{\mathcal{F}}(x, y) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\mathbb{P}_{AB|XY}}(ab|xy, c) \right|\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2} \sum_k \max_z \left[\sum_{c \in C} \tau_{C'|Z}^{\text{P}_{ABC|XYZ}}(c|z) \left| \sum_{x,y} \Pr_{\mathcal{E}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right. \right. \\
&\quad \left. \left. - \sum_{x,y} \Pr_{\mathcal{F}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right| \right. \\
&\quad \left. + \left(1 - \frac{1}{(n+1)^{l(m-1)}} \right) \left| \sum_{x,y} \Pr_{\mathcal{E}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}') \right. \right. \\
&\quad \left. \left. - \sum_{x,y} \Pr_{\mathcal{F}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}') \right| \right] \\
&\geq \frac{1}{2} \sum_k \max_z \sum_{c \in C} \tau_{C'|Z}^{\text{P}_{ABC|XYZ}}(c|z) \left| \sum_{x,y} \Pr_{\mathcal{E}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right. \\
&\quad \left. - \sum_{x,y} \Pr_{\mathcal{F}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \tau_{AB|XYC'}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right| \\
&= \frac{1}{2} \sum_k \max_z \sum_{c \in C} \frac{1}{(n+1)^{l(m-1)}} \cdot \mathbb{P}_{C|Z}(c|z) \\
&\quad \cdot \left| \sum_{x,y} \Pr_{\mathcal{E}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \mathbb{P}_{AB|XYC}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right. \\
&\quad \left. - \sum_{x,y} \Pr_{\mathcal{F}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \mathbb{P}_{AB|XYC}(\mathbf{ab}|\mathbf{x}y\mathbf{c}) \right| \\
&= \frac{1}{(n+1)^{l(m-1)}} \| (\mathcal{E} - \mathcal{F}) \otimes \mathbb{I}(\mathbb{P}_{ABC|XYZ}) \|_1 .
\end{aligned}$$

where in order to get the second equality we divide the sum over $C' = C \cup \{c'\}$ to the sum over C and then additional part of the partition c' . The next inequality is then correct since

$$\begin{aligned}
&\left(1 - \frac{1}{(n+1)^{l(m-1)}} \right) \left| \sum_{x,y} \Pr_{\mathcal{E}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{E}(a,b,x,y)=k}} \tau_{A|XC}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}') - \right. \\
&\quad \left. \sum_{x,y} \Pr_{\mathcal{F}}(\mathbf{x}, \mathbf{y}) \sum_{\substack{a,b: \\ \mathcal{F}(a,b,x,y)=k}} \tau_{A|XC}^{\text{P}_{ABC|XYZ}}(\mathbf{ab}|\mathbf{x}y\mathbf{c}') \right| \geq 0 ,
\end{aligned}$$

and the two last equalities are due to the specific decomposition of $\tau_{AB|XY}$ that we defined and the definition of the trace distance. This proves the lemma. \square

References

1. Masanes L (2009) Universally composable privacy amplification from causality constraints. *Phys Rev Lett* 102(14):140501
2. Hänggi E, Renner R, Wolf S (2010) Efficient device-independent quantum key distribution. *Advances in cryptology-EUROCRYPT 2010*. Springer, Berlin, pp 216–234

Appendix B

Additional Proofs: Non-signalling Parallel Repetition

B.1 Signalling Measure and Test

We present here proofs of lemmas relevant to Sect. 10.2. The proofs previously appeared in [1].

Lemma 10.5 *Let $O_{AB|XY}^1$ and $O_{AB|XY}^2$ be two single-round boxes such that*

$$|O_{AB|XY}^1 - O_{AB|XY}^2|_1 \leq \epsilon .$$

Then, for all $a, b, x,$ and $y,$

$$|\text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}^1) - \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}^2)| \leq 2\epsilon$$

Proof We prove a stronger result from which the lemma follows. We prove

$$\sum_{b, x, y} |\text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}^1) - \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}^2)| \leq 2\epsilon .$$

To do so first note the following,

$$\begin{aligned} |O_{AB|XY}^1 - O_{AB|XY}^2|_1 &= \mathbb{E}_{x, y} \sum_{a, b} |O_{AB|XY}^1(a, b|x, y) - O_{AB|XY}^2(a, b|x, y)| \\ &\geq \mathbb{E}_{x, y} \sum_b \left| \sum_a (O_{AB|XY}^1(a, b|x, y) - O_{AB|XY}^2(a, b|x, y)) \right| \\ &= \mathbb{E}_{x, y} \sum_b |O_{B|XY}^1(b|x, y) - O_{B|XY}^2(b|x, y)| \\ &= \sum_{b, x, y} Q_{XY}(x, y) |O_{B|XY}^1(b|x, y) - O_{B|XY}^2(b|x, y)| , \end{aligned}$$

therefore if $|\mathcal{O}_{AB|XY}^1 - \mathcal{O}_{AB|XY}^2|_1 \leq \epsilon$ then

$$\sum_{b,x,y} Q_{XY}(x,y) |\mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y)| \leq \epsilon. \quad (\text{B.1})$$

Next, using Definition 10.3 and the discussion following it,

$$\begin{aligned} & \sum_{b,x,y} |\text{Sig}^{(A \rightarrow B,x,y,b)}(\mathcal{O}_{AB|XY}^1) - \text{Sig}^{(A \rightarrow B,x,y,b)}(\mathcal{O}_{AB|XY}^2)| \\ &= \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) \mathcal{O}_{B|XY}^1(b|\tilde{x},y) \right. \\ & \quad \left. - \mathcal{O}_{B|XY}^2(b|x,y) + \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) \mathcal{O}_{B|XY}^2(b|\tilde{x},y) \right| \\ &= \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y) \right. \\ & \quad \left. + \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) (\mathcal{O}_{B|XY}^2(b|\tilde{x},y) - \mathcal{O}_{B|XY}^1(b|\tilde{x},y)) \right| \\ &\leq \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y) \right| \\ & \quad + \sum_{b,x,y} Q_{XY}(x,y) \left| \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) (\mathcal{O}_{B|XY}^2(b|\tilde{x},y) - \mathcal{O}_{B|XY}^1(b|\tilde{x},y)) \right| \\ &\leq \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y) \right| \\ & \quad + \sum_{b,x,y} \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^2(b|\tilde{x},y) - \mathcal{O}_{B|XY}^1(b|\tilde{x},y) \right| \\ &= \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y) \right| \\ & \quad + \sum_{b,y} \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) Q_Y(y) \left| \mathcal{O}_{B|XY}^2(b|\tilde{x},y) - \mathcal{O}_{B|XY}^1(b|\tilde{x},y) \right| \\ &= \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^1(b|x,y) - \mathcal{O}_{B|XY}^2(b|x,y) \right| \\ & \quad + \sum_{b,x,y} Q_{XY}(x,y) \left| \mathcal{O}_{B|XY}^2(b|x,y) - \mathcal{O}_{B|XY}^1(b|x,y) \right| \\ &\leq 2\epsilon \end{aligned}$$

where the last inequality follows from Eq. (B.1). \square

Lemma B.1 *Let $\nu > 0$ be any parameter such that $\nu < \zeta - 6\epsilon$. Then for every x , y , and b ,*

$$\forall \mathbf{O}_{AB|XY} \in \Sigma^{(A \rightarrow B, x, y, b)}, \quad \text{Sig}^{(A \rightarrow B, x, y, b)}(\mathbf{O}_{AB|XY}) > \nu.$$

Proof Assume by contradiction that there exists $\mathbf{O}_{AB|XY} \in \Sigma^{(A \rightarrow B, x, y, b)}$ such that $\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathbf{O}_{AB|XY}) \leq \nu$. Since $\mathbf{O}_{AB|XY} \in \Sigma^{(A \rightarrow B, x, y, b)}$ there exists $\bar{\mathbf{O}}_{AB|XY}$ such that $|\mathbf{O}_{AB|XY} - \bar{\mathbf{O}}_{AB|XY}|_1 \leq \epsilon$ and

$$\Pr_{\text{data} \sim \bar{\mathbf{O}}_{AB|XY}^{\otimes n}} [T] > \delta. \quad (\text{B.2})$$

Using Lemma 10.5 we get $\text{Sig}^{(A \rightarrow B, x, y, b)}(\bar{\mathbf{O}}_{AB|XY}) \leq \nu + 2\epsilon$. From Lemma 2.2 we know that $\Pr_{\text{data} \sim \bar{\mathbf{O}}_{AB|XY}^{\otimes n}} [|\bar{\mathbf{O}}_{AB|XY}^{\text{freq}(\text{data}_2)} - \bar{\mathbf{O}}_{AB|XY}|_1 > \epsilon] \leq \delta$ and therefore, using Lemma 10.5 again,

$$\Pr_{\text{data} \sim \bar{\mathbf{O}}_{AB|XY}^{\otimes n}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)}(\bar{\mathbf{O}}_{AB|XY}^{\text{freq}(\text{data}_2)}) > \nu + 4\epsilon \right] \leq \delta.$$

Since $\nu < \zeta - 6\epsilon$ this implies

$$\Pr_{\text{data} \sim \bar{\mathbf{O}}_{AB|XY}^{\otimes n}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)}(\bar{\mathbf{O}}_{AB|XY}^{\text{freq}(\text{data}_2)}) > \zeta - 2\epsilon \right] \leq \delta$$

and therefore, according to the definition of the test,

$$\Pr_{\text{data} \sim \bar{\mathbf{O}}_{AB|XY}^{\otimes n}} [T] \leq \delta,$$

which contradicts Eq. (B.2). \square

Next we would like to prove Lemma 10.7. To do so, we first prove the same statement but for IID boxes:

Lemma B.2 *Assume the players share an IID box $\mathbf{O}_{AB|XY}^{\otimes n}$ and let $\zeta, \epsilon > 0$ be the parameters defined as in Eq. (10.7). For every $\mathcal{T}^{(A \rightarrow B, x, y, b)}$,*

1. *If $\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathbf{O}_{AB|XY}) \geq \zeta$ then*

$$\Pr_{\text{data} \sim \mathbf{O}_{AB|XY}^{\otimes n}} [T] > 1 - \delta \quad (\text{B.3})$$

2. *If $\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathbf{O}_{AB|XY}) = 0$ then*

$$\Pr_{\text{data} \sim \mathbf{O}_{AB|XY}^{\otimes n}} [\neg T] > 1 - \delta \quad (\text{B.4})$$

where $\delta = \delta\left(\frac{n}{2}, \epsilon\right) = \left(\frac{n}{2} + 1\right)^{|\mathcal{A}| + |\mathcal{Q}| - 1} e^{-n\epsilon^2/4}$.

Proof For the first part of the lemma assume that $\text{Sig}^{(A \rightarrow B, x, y, b)}(\mathcal{O}_{AB|XY}) \geq \zeta$. Then

$$\begin{aligned} \Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[-T] &= \Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}} \left[\text{Sig}^{(A \rightarrow B, x, y, b)} \left(\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)} \right) < \zeta - 2\epsilon \right] \\ &\leq \Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}} \left[|\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_2)} - \mathcal{O}_{AB|XY}|_1 > \epsilon \right] \\ &\leq \delta \end{aligned}$$

where the first inequality is due to Lemma 10.5 and the second due to Lemma 2.2. This implies Eq. (B.3). Equation (B.4) can be proven in an analogous way. \square

Lemma 10.7 Let $\tau_{ABXY} = \mathcal{Q}_{XY}^{\otimes n} \tau_{AB|XY}$, where $\tau_{AB|XY}$ is a de Finetti box. For every $\mathcal{T}^{(A \rightarrow B, x, y, b)}$

1. $\Pr_{\text{data} \sim \tau_{ABXY}}[-in^\Sigma \wedge T] \leq \delta$
2. $\Pr_{\text{data} \sim \tau_{ABXY}}[in^\sigma \wedge \neg T] \leq \delta$,

where $\delta = \delta\left(\frac{n}{2}, \epsilon\right) = \left(\frac{n}{2} + 1\right)^{|\mathcal{A}||\mathcal{B}||\mathcal{X}||\mathcal{Y}|-1} e^{-n\epsilon^2/4}$.

Proof Since a de Finetti box is a convex combination of IID boxes, it is sufficient to prove this for IID boxes $\mathcal{O}_{AB|XY}^{\otimes n}$ and the lemma will follow. We start by proving the first part of the lemma.

If $\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[T] \leq \delta$ then we are done. Consider therefore single-round boxes $\mathcal{O}_{AB|XY}$ such that

$$\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[T] > \delta.$$

For such boxes

$$\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[-in^\Sigma] \leq \Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}} \left[|\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_1)} - \mathcal{O}_{AB|XY}|_1 > \epsilon \right] \leq \delta$$

where the first inequality follows from the definition of $\Sigma^{(A \rightarrow B, x, y, b)}$ and the second from Lemma 2.2. All together we get $\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[-in^\Sigma \wedge T] \leq \delta$ as required for the first part of the lemma.

We now proceed to the second part of the lemma. If $\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[in^\sigma] \leq \delta$ then we are done. Consider therefore boxes $\mathcal{O}_{AB|XY}$ such that

$$\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[in^\sigma] > \delta.$$

Using Lemma 2.2 we know that there exists a state $\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_1)} \in \Sigma^{(A \rightarrow B, x, y, b)}$ such that $|\mathcal{O}_{AB|XY}^{\text{freq}(\text{data}_1)} - \mathcal{O}_{AB|XY}|_1 \leq \epsilon$ and according to the definition of $\Sigma^{(A \rightarrow B, x, y, b)}$ this implies that $\mathcal{O}_{AB|XY}$ is ζ signalling or more. Therefore, according to Lemma B.2, $\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[-T] \leq \delta$. All together we get

$$\Pr_{\text{data} \sim \mathcal{O}_{AB|XY}^{\otimes n}}[in^\sigma \wedge \neg T] \leq \delta. \quad \square$$

B.2 Sensitivity Analysis

Linear programs (see, e.g., [2]) are a useful tool when considering non-signalling games, as the non-signalling constraints are linear. The following general results regarding the sensitivity of linear programs will be of use for us.

Lemma B.3 (Sensitivity analysis of linear programs, [2] Sect.10.4). *Let $\max\{c^T x | Ax \leq b\}$ be a primal linear program and $\min\{b^T y | A^T y = c, y \geq 0\}$ its dual. Denote the optimal value of the programs by w and the optimal dual solution by y^* . Then the optimal value of the perturbed program $w_e = \max\{c^T x | Ax \leq b + e\}$ for some perturbation e is bounded by $w_e \leq w + e^T y^*$.*

Lemma B.4 (Dual optimal solution bound, [2] Sect.10.4). *Let A be an $r_1 \times r_2$ -matrix and let Δ be such that for each non-singular sub-matrix B of A all entries of B^{-1} are at most Δ in absolute value. Let c be a row vector of dimension r_2 and let y^* be the optimal dual solution of $\min\{b^T y | A^T y = c, y \geq 0\}$. Then*

$$\kappa = \sum_{j=1}^{r_1} |y_j^*| \leq r_2 \Delta \sum_{j=1}^{r_2} |c_j|.$$

We start with the following program from Sect. 10.3.1:

$$\max \sum_{a,b,x,y} Q_{XY}(xy) R(a, b, x, y) O_{AB|XY}(ab|xy)$$

$$\text{s.t. } \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}(ab|xy)) = 0 \quad \forall x, y, b \quad (\text{B.5a})$$

$$\text{Sig}^{(B \rightarrow A, x, y, b)}(O_{AB|XY}(ab|xy)) = 0 \quad \forall x, y, a \quad (\text{B.5b})$$

$$\sum_{a,b} O_{AB|XY}(ab|xy) = 1 \quad \forall x, y$$

$$O_{AB|XY}(ab|xy) \geq 0 \quad \forall a, b, x, y$$

To apply Lemma B.3 we first need to write the program in the form $\max\{c^T x | Ax \leq b\}$. For this purpose, one can relax the linear program (B.5) to the following:

$$\max \sum_{a,b,x,y} Q_{XY}(xy) R(a, b, x, y) O_{AB|XY}(ab|xy)$$

$$\text{s.t. } \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}(ab|xy)) \leq 0 \quad \forall x, y, b \quad (\text{B.6a})$$

$$\text{Sig}^{(B \rightarrow A, x, y, b)}(O_{AB|XY}(ab|xy)) \leq 0 \quad \forall x, y, a \quad (\text{B.6b})$$

$$\sum_{a,b} O_{AB|XY}(ab|xy) = 1 \quad \forall x, y$$

$$O_{AB|XY}(ab|xy) \geq 0 \quad \forall a, b, x, y$$

To see that the relaxation of the non-signalling constraints (B.5a) and (10.21c) to the constraints (B.6a) and (B.6b) does not change the program, i.e., does not change the value of the optimal solution, recall Eq. (10.6) and assume there exists x, y, b for which

$$Q_{XY}(x, y) \left[O_{B|XY}(b|x, y) - \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) O_{B|XY}(b|\tilde{x}, y) \right] < 0 .$$

That is, $O_{B|XY}(b|x, y)$ is smaller than the average $\sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) O_{B|XY}(b|\tilde{x}, y)$, and therefore there must be some x' for which $O_{B|XY}(b|x', y)$ is larger than the average, meaning,

$$Q_{XY}(x', y) \left[O_{B|XY}(b|x', y) - \sum_{\tilde{x}} Q_{X|Y}(\tilde{x}|y) O_{B|XY}(b|\tilde{x}, y) \right] > 0 ,$$

but this contradicts the constraints (B.6a) and (B.6b).

The dual program of the primal (B.6) is given below.

$$\begin{aligned} \min \quad & \sum_{x,y} z(x, y) \\ \text{s.t.} \quad & z(x, y) + y_A(x, y, b) Q_{XY}(x, y) + y_B(x, y, a) Q_{XY}(x, y) \\ & - \sum_{\tilde{x}} y_A(\tilde{x}, y, b) Q_{XY}(\tilde{x}, y) Q_{X|Y}(x|y) \\ & - \sum_{\tilde{y}} y_B(x, \tilde{y}, a) Q_{XY}(x, \tilde{y}) Q_{Y|X}(y|x) \\ & \geq Q_{XY}(x, y) R(a, b, x, y) \quad \forall a, b, x, y \\ & y_A(x, y, b) \geq 0 \quad \forall x, y, b \\ & y_B(x, y, a) \geq 0 \quad \forall x, y, a \end{aligned} \tag{B.7a}$$

Lemma B.5 Let $\kappa = \sum_{j=1}^d |y_j^*|$ where d is the number of signalling tests and y^* is an optimal solution of the dual program (B.7). Let $O_{AB|XY}$ be a strategy such that the following holds for all a, b, x, y

$$\begin{aligned} \text{Sig}^{(A \rightarrow B, x, y, b)}(O_{AB|XY}) &\leq \zeta + 2\epsilon \\ \text{Sig}^{(B \rightarrow A, x, y, a)}(O_{AB|XY}) &\leq \zeta + 2\epsilon . \end{aligned} \tag{B.8}$$

Then $w(O_{AB|XY}) \leq 1 - \alpha + (\zeta + 2\epsilon) d$.

Proof The fact that $O_{AB|XY}$ is not “too signalling” in any direction can be used to bound its winning probability in the game G.

The following linear program describes the optimal winning probability of a strategy $O_{AB|XY}$ which fulfils Eq. (B.8):

$$\begin{aligned}
\max \quad & \sum_{a,b,x,y} Q_{XY}(xy) R(a,b,x,y) O_{AB|XY}(ab|xy) \\
\text{s.t.} \quad & \text{Sig}^{(A \rightarrow B,x,y,b)}(O_{AB|XY}(ab|xy)) \leq \zeta + 2\epsilon \quad \forall x,y,b \\
& \text{Sig}^{(B \rightarrow A,x,y,b)}(O_{AB|XY}(ab|xy)) \leq \zeta + 2\epsilon \quad \forall x,y,a \\
& \sum_a O_{AB|XY}(ab|xy) = 1 \quad \forall x,y \\
& O_{AB|XY}(ab|xy) \geq 0 \quad \forall a,b,x,y
\end{aligned} \tag{B.9}$$

Program (B.9) can be seen as a perturbation of the linear program (B.6), we can therefore bound its optimal value by using known tools for sensitivity analysis of linear programs, stated in Lemmas B.3 and B.4.

Denote by y^* an optimal solution of the dual program² (B.7) and let $\kappa = \sum_{j=1}^d |y_j^*|$ where d is the number of signalling tests. That is, κ is the sum of all the dual variables which are associated to the non-signalling constraints.

According to Lemma B.3 the perturbed winning probability is then bounded by

$$w_e \leq 1 - \alpha + (\zeta + 2\epsilon) \kappa.$$

In the case of a game with two players, using [3, Sect. 4], one can show that $\kappa \leq d$ where d is the number of different signalling tests, i.e., $d = |\mathcal{X}||\mathcal{Y}|(|\mathcal{A}| + |\mathcal{B}|)$. \square

References

1. Arnon-Friedman R, Renner R, Vidick T (2016b) Non-signaling parallel repetition using de finetti reductions. IEEE Trans Inf Theory 62(3):1440–1457
2. Schrijver A (1998) Theory of linear and integer programming. Wiley, New York
3. Ito T (2010) Polynomial-space approximation of no-signaling provers. Automata, languages and programming. Springer, Berlin, pp 140–151

²We are only interested in the value of y^* as z^* will not affect the bound.

Appendix C

Additional Proofs: Device-Independent Quantum Cryptography

This appendix is devoted to presenting the technical proofs of the statements made in this thesis related to the device-independent cryptography. The proofs previously appeared in [1].

C.1 Single-Round Statement

As mentioned in Sect. 5.2, Lemma 5.3 (restated below) follows, almost directly, from [2, 3]. We show here how the bound derived in these works can be manipulated in a simple way to get the bound used in this thesis.

Lemma 5.3 *For any quantum single-round box $P_{AB|XY}$ with winning probability $\omega \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right]$ in the CHSH game,*

$$H(A|XYE) \geq 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\omega(\omega-1)+3}\right),$$

where E denotes the quantum side-information belonging to the adversary and $h(\cdot)$ is the binary entropy function.

Proof Our starting point is the result of [4]. There, a quantum single-round box $P_{\hat{A}\hat{B}|XY}$ with symmetric marginals on \hat{A} and \hat{B} was considered (i.e., \hat{A} and \hat{B} are uniformly distributed). To derive a bound which holds for any $P_{AB|XY}$ we do the following:

1. Symmetrisation of $P_{AB|XY}$ —Alice chooses a bit F uniformly at random and communicates it to Bob. They then symmetrise their marginals by setting $\hat{A} = A \oplus F$ and $\hat{B} = B \oplus F$.

2. Use [4] to lower-bound $H(\hat{A}|XYFE)$.
3. Derive a bound on $H(A|XYE)$ from $H(\hat{A}|XYFE)$.

Let us follow the above steps. After applying $\hat{A} = A \oplus F$ and $\hat{B} = B \oplus F$, with F uniformly distributed, A and B are unbiased. In our notation, [4] considered the following Holevo quantity

$$\chi(\hat{A} : FE|X = 0) = H(FE|X = 0) - H(FE|\hat{A}, X = 0) .$$

and showed that for states leading to a CHSH violation of $\beta \in [2, 2\sqrt{2}]$, related to the winning probability via $\omega = 1/2 + \beta/8$, the following tight bound holds [4, Sect. 2.3]:

$$\chi(\hat{A} : FE|X = 0) \leq h\left(\frac{1}{2} + \frac{1}{2}\sqrt{\frac{\beta^2}{4} - 1}\right) .$$

Rewriting the bound in terms of the winning probability ω one gets that for all $\omega \in \left[\frac{3}{4}, \frac{2+\sqrt{2}}{4}\right]$ (i.e. a winning probability in the quantum regime)

$$\chi(\hat{A} : FE|X = 0) \leq h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\omega(\omega - 1) + 3}\right) . \quad (\text{C.1})$$

We now wish to related the above Holevo quantities to our von Neumann entropy. Using the definition of the conditional von Neumann entropy one can rewrite $H(A|FE, X = 0)$ (σ) as follows:

$$\begin{aligned} H(\hat{A}|FE, X = 0) &= H(\hat{A}FE|X = 0) - H(FE|X = 0) \\ &= H(\hat{A}|X = 0) + H(FE|\hat{A}, X = 0) - H(FE|X = 0) \\ &= H(\hat{A}|X = 0) - \chi(\hat{A} : FE|X = 0) \\ &= 1 - \chi(\hat{A} : FE|X = 0) , \end{aligned} \quad (\text{C.2})$$

where the last equality holds since A is uniformly random due to the symmetrisation step (note that we do not condition on F in $H(A|X = 0)$). Furthermore,

$$\begin{aligned} H(\hat{A}|XYFE)_{\mathcal{M}_i(\sigma)} &= \Pr[X_i = 0] \cdot H(\hat{A}|YFE, X_i = 0)_{\mathcal{M}_i(\sigma)} \\ &\quad + \Pr[X_i = 1] \cdot H(\hat{A}|YFE, X_i = 1)_{\mathcal{M}_i(\sigma)} . \end{aligned} \quad (\text{C.3})$$

Combining Eqs. (C.1) and (C.2) while noting that, due to the symmetry between the cases $X = 0$ and $X = 1$, the same relations can be written for $X = 1$ and plugging

the bounds in Eq. (C.3) we get

$$H\left(\hat{A}|XYFE\right) \geq 1 - h\left(\frac{1}{2} + \frac{1}{2}\sqrt{16\omega(\omega-1)+3}\right). \quad (\text{C.4})$$

The only think left to do is to use the above to get our bound for the original box $P_{AB|XY}$. For this simply observe that

$$H\left(\hat{A}|XYFE\right) = H(A|XYFE) = H(A|XYE),$$

where the first equality holds since $\hat{A} = A \oplus F$ and the second follows since F is independent of $A, X, Y,$ and E . The lemma therefore follows. \square

C.2 An Improved Dependency on the Test Probability

In this section we show how the EAT can be used in a slightly different way than what was done in the main text. This results in an entropy rate which has a better dependency on the probability of a test round γ , compared to the entropy rate given in Eq. (11.12). The improved entropy rate derived here is the one used for calculating the key rates of the DIQKD protocol in Sect. 11.3.3.2.

C.2.1 Modified Entropy Accumulation Protocol

We use a different entropy accumulation protocol, given as Protocol C.1. In this modified protocol instead of considering each round separately we consider blocks of rounds. A block is defined by a sequence of rounds: in each round a test is carried out with probability γ (and otherwise the round is a generation round). The block ends when a test round is being performed and then the next block begins. If for s_{\max} rounds there was no test, the block ends without performing a test and the next begins. Thus, the blocks can be of different length, but they all consist at most s_{\max} rounds.

In this setting, instead of fixing the number of rounds n in the beginning of the protocol, we fix the number of blocks m . The expected length of block is

$$\begin{aligned} \bar{s} &= \sum_{s \in [s_{\max}]} [s(1-\gamma)^{(s-1)}\gamma] + s_{\max}(1-\gamma)^{s_{\max}} = \frac{1 - (1-\gamma)^{s_{\max}}}{\gamma} \\ &= \sum_{s \in [s_{\max}]} [(1-\gamma)^{(s-1)}]. \end{aligned} \quad (\text{C.5})$$

The expected number of rounds is denoted by $\bar{n} = m \cdot \bar{s}$.

Compared to the main text, we now have a RV $\tilde{W}_j \in \{0, 1, \perp\}$ for each block, instead of each round. Alice and Bob set \tilde{W}_j to be 0 or 1 depending on the result of the game in the block's test round (i.e., the last round of the block), or $\tilde{W}_j = \perp$ if a test round was not carried out in the block. By the definition of the blocks we have $\Pr[\tilde{W}_j = \perp] = (1 - \gamma)^{s_{\max}}$.

Protocol C.1 Modified entropy accumulation protocol

Arguments:

G —two-player non-local game

$\mathcal{X}_g, \mathcal{X}_t \subset \mathcal{X}$ —generation and test inputs for Alice

$\mathcal{Y}_g, \mathcal{Y}_t \subset \mathcal{Y}$ —generation and test inputs for Bob

D —untrusted device of (at least) two components that can play G repeatedly

$m \in \mathbb{N}_+$ —number of blocks

$s_{\max} \in \mathbb{N}_+$ —maximal length of a block

$\gamma \in (0, 1]$ —probability of a test round

ω_{exp} —expected winning probability in G for an honest (perhaps noisy) implementation

$\delta_{\text{est}} \in (0, 1)$ —width of the statistical confidence interval for the estimation test

- 1: For every block $j \in [m]$ do Steps 2-9:
 - 2: Set $i = 0$ and $W_j = \perp$.
 - 3: If $i \leq s_{\max}$:
 - 4: Set $i = i + 1$.
 - 5: Alice and Bob choose $T_i \in \{0, 1\}$ at random such that $\Pr(T_i = 1) = \gamma$.
 - 6: If $T_i = 0$ Alice and Bob choose inputs $X_i \in \mathcal{X}_g$ and $Y_i \in \mathcal{Y}_g$ respectively. If $T_i = 1$ they choose inputs $X_i \in \mathcal{X}_t$ and $Y_i \in \mathcal{Y}_t$.
 - 7: Alice and Bob use D with X_i, Y_i and record their outputs as A_i and B_i respectively.
 - 8: If $T_i = 0$ Bob updates B_i to $B_i = \perp$.
 - 9: If $T_i = 1$ they set $\tilde{W}_j = w(A_i, B_i, X_i, Y_i)$ and $i = s_{\max} + 1$.
 - 10: Alice and Bob abort if $\sum_{j \in [m]} \tilde{W}_j < [\omega_{\text{exp}}(1 - (1 - \gamma)^{s_{\max}}) - \delta_{\text{est}}] \cdot m$.
-

C.2.2 Modified Min-tradeoff Function

Below, we apply the EAT on blocks of outputs instead of single rounds directly. Let \mathcal{M}_j denote the EAT channels defined by the actions of Steps 2-9 in Protocol C.1 together with the behaviour of the device. It is easy to verify that \mathcal{M}_j fulfil the necessary conditions given in Definition 9.1.

We now construct a min-tradeoff function for \mathcal{M}_j . Let \tilde{p} be a probability distribution over $\{0, 1, \perp\}$. Our goal is to find F_{\min} such that

$$\forall j \in [m] \quad F_{\min}(\tilde{p}) \leq \inf_{\sigma_{R_{j-1}R'}} \inf_{\mathcal{M}_j(\sigma)_{\tilde{W}_j} = \tilde{p}} H\left(\vec{A}_j \vec{B}_j | \vec{X}_j \vec{Y}_j \vec{T}_j R'\right)_{\mathcal{M}_j(\sigma)}, \quad (\text{C.6})$$

where \vec{A}_j is a vector of varying length (but at most s_{\max}). We use $A_{j,i}$ to denote the i 'th entry of \vec{A}_j and $A_{j,1}^{j,i-1} = A_{j,1} \dots A_{j,i-1}$. Since we will only be interested in the entropy of \vec{A}_j we can also describe it as a vector of length s_{\max} which is initialised to be all \perp . For every actual round being performed in the block the value of $A_{j,i}$ is updated. Thus, the entries of \vec{A}_j which correspond to rounds which were not performed do not contribute to the entropy. We use similar notation for the other vectors of RVs.

To lower-bound the right-hand side of Eq. (C.6) we first use the chain rule

$$H\left(\vec{A}_j \vec{B}_j | \vec{X}_j \vec{Y}_j \vec{T}_j R'\right) = \sum_{i \in [s_{\max}]} H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j \vec{T}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1}). \quad (\text{C.7})$$

Next, for every $i \in [s_{\max}]$,

$$\begin{aligned} & H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j \vec{T}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1}) = \\ & \Pr[T_{j,1}^{j,i-1} = \vec{0}] H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1} T_{j,i}^{j,s_{\max}} T_{j,1}^{j,i-1} = \vec{0}) \\ & + \Pr[T_{j,1}^{j,i-1} \neq \vec{0}] H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1} T_{j,i}^{j,s_{\max}} T_{j,1}^{j,i-1} \neq \vec{0}) \\ = & (1 - \gamma)^{(i-1)} H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1} T_{j,i}^{j,s_{\max}} T_{j,1}^{j,i-1} = \vec{0}) \end{aligned}$$

since the entropy is not zero only if the i 'th round is being performed in the block, i.e., if a test was not performed before that round. Plugging this into Eq. (C.7) we get

$$\begin{aligned} & H\left(\vec{A}_j \vec{B}_j | \vec{X}_j \vec{Y}_j \vec{T}_j R'\right) \\ = & \sum_{i \in [s_{\max}]} (1 - \gamma)^{(i-1)} H(A_{j,i} B_{j,i} | \vec{X}_j \vec{Y}_j R' A_{j,1}^{j,i-1} B_{j,1}^{j,i-1} T_{j,i}^{j,s_{\max}} T_{j,1}^{j,i-1} = \vec{0}). \end{aligned}$$

Each term in the sum can now be identified as the entropy of a single round. We can therefore use the bound derived in the main text, as given in Eq. (11.6). For this we denote by ω_i the winning probability in the i 'th round (given that a test was not performed before). Then it holds that

$$H\left(\vec{A}_j \vec{B}_j | \vec{X}_j \vec{Y}_j \vec{T}_j R'\right) \geq \sum_{i \in [s_{\max}]} (1 - \gamma)^{(i-1)} \left[1 - h\left(\frac{1}{2} + \frac{1}{2} \sqrt{16\omega_i (\omega_i - 1) + 3}\right) \right], \quad (\text{C.8})$$

where, by the actions of the EAT channel \mathcal{M}_j , the ω_i 's must fulfil the constraint

$$\tilde{p}(1) = \sum_{i \in [s_{\max}]} \gamma (1 - \gamma)^{(i-1)} \omega_i. \quad (\text{C.9})$$

Note that, similarly to what was done in the main text, we only need to consider \tilde{p} for which $\tilde{p}(1) + \tilde{p}(0) = 1 - (1 - \gamma)^{s_{\max}}$ (otherwise the condition on the min-tradeoff function is trivial, as the infimum is over an empty set).

To find the min-tradeoff function $F_{\min}(\tilde{p})$ we therefore need to minimise Eq. (C.8) under the constraint of Eq. (C.9). The following lemma shows that the minimum is achieved when all ω_i are equal.

Lemma C.1 *The minimum of the function given on the righthand-side of Eq. (C.8) over ω_i constrained by Eq. (C.9) is achieved for $\omega_i^* = \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}}$ for all $i \in [s_{\max}]$.*

Proof Let $\vec{\omega} = \omega_1, \dots, \omega_{s_{\max}}$ and

$$f(\vec{\omega}) \equiv \sum_{i \in [s_{\max}]} (1-\gamma)^{(i-1)} \left[1 - h \left(\frac{1}{2} + \frac{1}{2} \sqrt{16\omega_i (\omega_i - 1) + 3} \right) \right];$$

$$g(\vec{\omega}) \equiv \sum_{i \in [s_{\max}]} \gamma(1-\gamma)^{(i-1)} \omega_i - \tilde{p}(1).$$

Using the method of Lagrange multipliers, we should look for $\vec{\omega}^*$ such that $g(\vec{\omega}^*) = 0$ and $\nabla f(\vec{\omega}^*) = -\lambda \nabla g(\vec{\omega}^*)$ for some constant λ . $\nabla f(\vec{\omega}^*) = -\lambda \nabla g(\vec{\omega}^*)$ implies that for any i ,

$$(1-\gamma)^{(i-1)} \frac{d}{d\omega_i} \left[1 - h \left(\frac{1}{2} + \frac{1}{2} \sqrt{16\omega_i (\omega_i - 1) + 3} \right) \right] \Big|_{\omega_i^*} = -\lambda \gamma (1-\gamma)^{(i-1)}$$

and therefore

$$\frac{d}{d\omega_i} \left[1 - h \left(\frac{1}{2} + \frac{1}{2} \sqrt{16\omega_i (\omega_i - 1) + 3} \right) \right] \Big|_{\omega_i^*} = -\lambda \gamma.$$

The function on the left-hand side of the above equation is strictly increasing. Hence, it must be that all ω_i^* are equal to some constant ω^* .

Lastly, we must have $g(\vec{\omega}^*) = 0$. Thus,

$$\sum_{i \in [s_{\max}]} \gamma(1-\gamma)^{(i-1)} \omega^* - \tilde{p}(1) = 0$$

which means

$$\omega^* = \frac{\tilde{p}(1)}{\sum_{i \in [s_{\max}]} \gamma(1-\gamma)^{(i-1)}} = \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}}.$$

□

Plugging the minimal values of ω_i into Eq. (C.8) we get that

$$\begin{aligned} & H\left(\vec{A}_j \vec{B}_j | \vec{X}_j \vec{Y}_j \vec{T}_j; R'\right) \\ & \geq \sum_{i \in [s_{\max}]} (1-\gamma)^{i-1} \left[1 - h\left(\frac{1}{2} + \frac{1}{2} \sqrt{16 \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} \left(\frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} - 1\right) + 3}\right)} \right] \\ & = \bar{s} \left[1 - h\left(\frac{1}{2} + \frac{1}{2} \sqrt{16 \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} \left(\frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} - 1\right) + 3}\right)} \right], \end{aligned}$$

where we used Eq. (C.5) to get the last equality.

From this point we can follow the same steps as in Sect. 11.2.2 (cutting and gluing the function etc.). The resulting min-tradeoff function is given by

$$\begin{aligned} g(\tilde{p}) &= \tag{C.10} \\ & \begin{cases} \bar{s} \left[1 - h\left(\frac{1}{2} + \frac{1}{2} \sqrt{16 \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} \left(\frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} - 1\right) + 3}\right)} \right] & \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} \in \left[0, \frac{2+\sqrt{2}}{4}\right] \\ \bar{s} & \frac{\tilde{p}(1)}{1-(1-\gamma)^{s_{\max}}} \in \left[\frac{2+\sqrt{2}}{4}, 1\right], \end{cases} \\ F_{\min}(\tilde{p}, \tilde{p}_t) &= \begin{cases} g(\tilde{p}) & \tilde{p}(1) \leq \tilde{p}_t(1) \\ \frac{d}{d\tilde{p}(1)} g(\tilde{p})|_{\tilde{p}_t} \cdot \tilde{p}(1) + \left(g(\tilde{p}_t) - \frac{d}{d\tilde{p}(1)} g(\tilde{p})|_{\tilde{p}_t} \cdot \tilde{p}_t(1)\right) & \tilde{p}(1) > \tilde{p}_t(1). \end{cases} \end{aligned}$$

The min-tradeoff function given above is effectively identical to the one derived in the main text; although it gives us a bound on the von Neumann entropy in a block, instead of a single round, this bound is exactly the expected length of a block, \bar{s} , times the entropy in one round. For $s_{\max} = 1$ the min-tradeoff function constructed in the main text is retrieved.

C.2.3 Modified Entropy Rate

Since we apply the EAT on the blocks, the entropy rate is now defined to be the entropy *per block*. We therefore get

$$\begin{aligned} \mu(\tilde{p}, \tilde{p}_t, \varepsilon_s, \varepsilon_c) &= F_{\min}(\tilde{p}, \tilde{p}_t) \\ & \quad - \frac{1}{\sqrt{m}} 2 \left(\log(1 + 2 \cdot 2^{s_{\max}} 3^{s_{\max}}) + \left\| \frac{d}{d\tilde{p}(1)} g(\tilde{p}) \right\|_{\infty} \right) \sqrt{1 - 2 \log(\varepsilon_s \cdot \varepsilon_c)}, \\ \mu_{\text{opt}}(\varepsilon_s, \varepsilon_c) &= \max_{\frac{3}{4} < \tilde{p}_t(1) < \frac{2+\sqrt{2}}{4}} \mu(\omega_{\text{exp}}[1 - (1-\gamma)^{s_{\max}}] - \delta_{\text{est}}, \tilde{p}_t, \varepsilon_s, \varepsilon_c), \end{aligned}$$

and the total amount of entropy is given by

$$H_{\min}^{\varepsilon_s}(\mathbf{AB} | \mathbf{XYTE})_{\rho_{\Omega}} > m \cdot \mu_{\text{opt}}(\varepsilon_s, \varepsilon_{\text{EA}}) = \frac{\bar{n}}{\bar{s}} \cdot \mu_{\text{opt}}(\varepsilon_s, \varepsilon_{\text{EA}}). \tag{C.11}$$

By choosing $s_{\max} = \lceil \frac{1}{\gamma} \rceil$ the scaling of the entropy rate with γ is better than the rate derived in the main text. In particular, a short calculation reveals that the second order term scales, roughly, as $\sqrt{\bar{n}}/\gamma$ instead of \sqrt{n}/γ .

C.2.4 Modified Key Rate

To get the final key rate we need to repeat the same steps from the main text, but this time applied to random variables of varying length.

For this we first observe that, with high probability, the actual number of rounds, n , cannot be much larger than the expected number of rounds \bar{n} . Let S_i be the RV describing the length of block i , for $i \in [m]$, and N the RV describing the total number of rounds. Then $N = S_1 + \dots + S_m$. Since all the S_i are independent, identical, and have values in $\left[1, \frac{1}{\gamma}\right]$ we have

$$\Pr[N \geq \bar{n} + t] \leq \exp\left[-\frac{2t^2\gamma^2}{m(1-\gamma)^2}\right].$$

Let $\varepsilon_t = \exp\left[-\frac{2t^2\gamma^2}{m(1-\gamma)^2}\right]$ then

$$t = \sqrt{-\frac{m(1-\gamma)^2 \log \varepsilon_t}{2\gamma^2}}.$$

The first step in the derivation of the key rate which needs to be changed is the one given in Eq.(11.21). The quantity that needs to be upper bounded is $H_{\max}^{\frac{\varepsilon_s}{4}}(\mathbf{B}|\mathbf{T}EN)_{\rho_{|\hat{\Omega}}}$; N can be included in the entropy since its value is fixed by \mathbf{T} . By the definition of the smooth max-entropy we have

$$H_{\max}^{\frac{\varepsilon_s}{4}}(\mathbf{B}|\mathbf{T}EN) \leq H_{\max}^{\frac{\varepsilon_s}{4} - \sqrt{\varepsilon_t}}(\mathbf{B}|\mathbf{T}EN, N \leq \bar{n} + t).$$

Following the same steps as in the proof of Lemma 11.8 we have

$$H_{\max}^{\frac{\varepsilon_s}{4} - \sqrt{\varepsilon_t}}(\mathbf{B}|\mathbf{T}EN, N \leq \bar{n} + t)_{\rho_{|\hat{\Omega}}} < \gamma(\bar{n} + t) + \sqrt{\bar{n} + t} 2 \log 7 \sqrt{1 - 2 \log\left((\varepsilon_s/4 - \sqrt{\varepsilon_t}) \cdot (\varepsilon_{\text{EA}} + \varepsilon_{\text{EC}})\right)}.$$

With this modification and the modified entropy rate given in Eq.(C.11) we get

$$\begin{aligned}
H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYT}OE)_{\tilde{\rho}_{|\Omega}} &\geq \frac{\bar{n}}{s} \cdot \mu_{\text{opt}}(\varepsilon_s/4, \varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}) - \text{leak}_{\text{EC}} \\
&\quad - 3 \log \left(1 - \sqrt{1 - (\varepsilon_s/4)^2} \right) - \gamma(\bar{n} + t) \\
&\quad - \sqrt{\bar{n} + t} 2 \log 7 \sqrt{1 - 2 \log \left((\varepsilon_s/4 - \sqrt{\varepsilon_t}) \cdot (\varepsilon_{\text{EA}} + \varepsilon_{\text{EC}}) \right)}.
\end{aligned}$$

Similarly, the amount of leakage due to the error correction step leak_{EC} should be modified as well. Following the steps in Sect. 11.3.3.1, the quantity to be upper bounded is $H_{\max}^{\varepsilon'_{\text{EC}}}(\mathbf{A}|\tilde{\mathbf{B}}\mathbf{XYTN})$. Here as well we have

$$H_{\max}^{\varepsilon'_{\text{EC}}}(\mathbf{A}|\tilde{\mathbf{B}}\mathbf{XYTN}) \leq H_{\max}^{\varepsilon'_{\text{EC}} - \sqrt{\varepsilon_t}}(\mathbf{A}|\tilde{\mathbf{B}}\mathbf{XYTN}, N \leq \bar{n} + t).$$

The asymptotic equipartition property can be used with the maximal length $\bar{n} + t$ to get

$$\begin{aligned}
&H_{\max}^{\varepsilon'_{\text{EC}} - \sqrt{\varepsilon_t}}(\mathbf{A}|\tilde{\mathbf{B}}\mathbf{XYTN}, N \leq \bar{n} + t) \\
&\leq (\bar{n} + t) \cdot H(\mathbf{A}_i|\tilde{\mathbf{B}}_i X_i Y_i T_i) + \sqrt{\bar{n} + t} \delta(\varepsilon'_{\text{EC}} - 2\sqrt{\varepsilon_t}, \tau),
\end{aligned}$$

for

$$\begin{aligned}
\tau &= 2\sqrt{2H_{\max}(\mathbf{A}_i|\tilde{\mathbf{B}}_i X_i Y_i T_i)} + 1 \\
\delta(\varepsilon'_{\text{EC}} - 2\sqrt{\varepsilon_t}, \tau) &= 4 \log \tau \sqrt{2 \log \left(8/(\varepsilon'_{\text{EC}} - 2\sqrt{\varepsilon_t})^2 \right)}.
\end{aligned}$$

Continuing exactly as in Sect. 11.3.3.1 we get

$$\begin{aligned}
\text{leak}_{\text{EC}} &\leq (\bar{n} + t) \cdot \left[(1 - \gamma) h(\mathcal{Q}) + \gamma h(\omega_{\text{exp}}) \right] \\
&\quad + \sqrt{\bar{n} + t} 4 \log \left(2\sqrt{2} + 1 \right) \sqrt{2 \log \left(8/(\varepsilon'_{\text{EC}} - 2\sqrt{\varepsilon_t})^2 \right)} \\
&\quad + \log \left(8/\varepsilon_{\text{EC}}^2 + 2/(2 - \varepsilon'_{\text{EC}}) \right) + \log \left(\frac{1}{\varepsilon_{\text{EC}}} \right).
\end{aligned}$$

The parameter ε_t should be chosen such that the key rate is optimised. The resulting key rates are shown in Figures 11.3 and 11.4 in the main text.

C.3 Summary of Parameters and Variables

For convenience, all the parameters and variables are listed in Tables C.1 and C.2.

Table C.1 Parameters used in Chap. 11

Symbol	Meaning
$n \in \mathbb{N}_+$	Number of rounds
$\gamma \in (0, 1]$	Expected fraction of Bell violation estimation rounds
$\omega_{\text{exp}} \in [0, 1]$	Expected winning probability in an honest (perhaps noisy) implementation
$\delta_{\text{est}} \in (0, 1)$	Width of the statistical confidence interval for the Bell violation estimation test
ε_s	Smoothing parameter
ε_{EA}^c	Completeness error of the entropy accumulation protocol
ε_{EA}	The error probability of the entropy accumulation protocol
leak_{EC}	The leakage of the error correction protocol
$\varepsilon_{EC}, \varepsilon'_{EC}$	Error probabilities of the error correction protocol
ε_{EC}^c	Completeness error of the error correction protocol
ε_{PE}^c	Completeness error of the parameter estimation step
ε_{PA}	Error probability of the privacy amplification protocol
ℓ	Final key length in the DIQKD protocol
ε_{QKD}^c	Completeness error of the DIQKD protocol
ε_{QKD}^s	Soundness error of the DIQKD protocol

Table C.2 Random variables and quantum systems used in Chap. 11

Random variables and systems	Meaning
$X_i \in \{0, 1\}$	Alice's input in round $i \in [n]$
$Y_i \in \{0, 1\}$	Bob's input in round $i \in [n]$
$A_i \in \{0, 1\}$	Alice's output in round $i \in [n]$
$B_i \in \{0, 1, \perp\}, \tilde{B}_i \in \{0, 1\}$	Bob's output in round $i \in [n]$
$T_i \in \{0, 1\}$	Indicator of the estimation test in round i : $T_i = \begin{cases} 0 & i\text{'th round is not a test round} \\ 1 & i\text{'th round is a test round} \end{cases}$
$W_i \in \{\perp, 0, 1\}$	Indicator of the correlation in the test rounds: $W_i = \begin{cases} \perp & T_i = 0 \\ 0 & T_i = 1 \text{ and the test fails} \\ 1 & T_i = 1 \text{ and the test succeeds} \end{cases}$
E	Register of Eve's quantum state
R_i	Register of the (unknown) quantum state $\rho_{Q_A Q_B}^i$ of Alice and Bob's devices after step i of the protocol, for $i \in \{0\} \cup [n]$

References

1. Arnon-Friedman R, Renner R, Vidick T (2019) Simple and tight device-independent security proofs. *SIAM J Comput* 48(1):181–225
2. Pironio S, Acín A, Massar S, de La Giroday AB, Matsukevich DN, Maunz P, Olmschenk S, Hayes D, Luo L, Manning TA et al (2010) Random numbers certified by Bell’s theorem. *Nature* 464(7291):1021–1024
3. Acín A, Massar S, Pironio S (2012) Randomness versus nonlocality and entanglement. *Phys Rev Lett* 108(10):100402
4. Pironio S, Acín A, Brunner N, Gisin N, Massar S, Scarani V (2009) Device-independent quantum key distribution secure against collective attacks. *New J Phys* 11(4):045021