# 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.

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

R. Arnon-Friedman, *Device-Independent Quantum Information Processing*, Springer Theses, https://doi.org/10.1007/978-3-030-60231-4_1

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

[1] Clearly, 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.

[2] Notice 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.

[3] Consider for example the case of device-independent QKD. Classical devices can always be preprogrammed 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

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

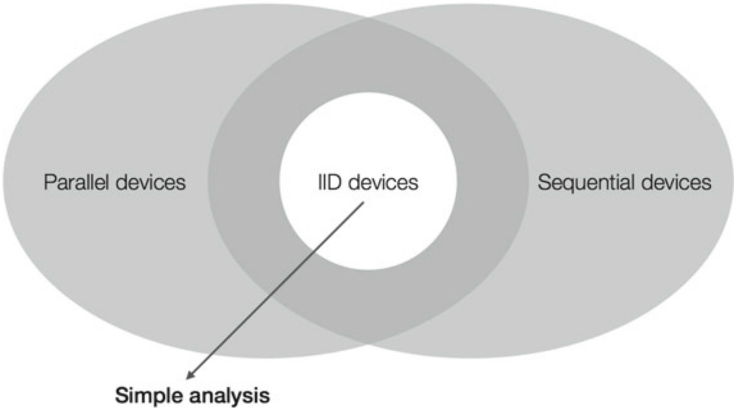

**Fig. 1.1** The relation between the different sets of devices. The intersection of the sets of sequential and parallel devices includes the set of IID devices. The analysis of IID devices, i.e., that done under the IID assumption, is rather simple

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 conceived[5] 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.

[5] Perhaps 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.

[6] In 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 quantitively strong results.[8]

The importance of quantitively 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 quantitive 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 quantitively 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

[7] The 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.

[8] This 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.

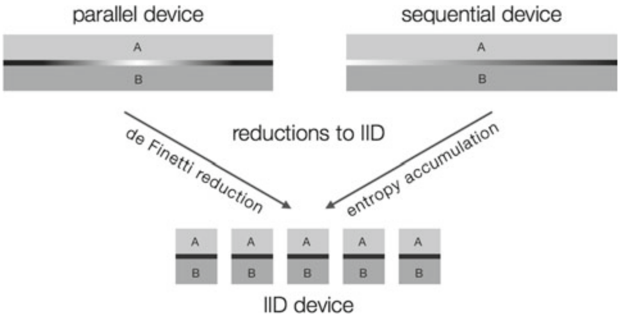

**Fig. 1.2** Reductions to IID in device-independent information processing. de Finetti reductions can be used to reduce the study of parallel devices to IID device (see Chaps. 8 and 10), while the entropy accumulation theorem can be used when dealing with sequential devices (Chaps. 9 and 11)

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 know 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]

[9] 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, 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 quantitive 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.

[10] To 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

[11] This 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.

[12] When 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 $\mathrm{P}_{\boldsymbol{AB}|\boldsymbol{XY}}$, where $\boldsymbol{A} = A_1, \ldots, A_n$ is the random variable describing Alice's answers in the threshold game ($A_i$ being her answer in the $i$'th game) and, similarly, $\boldsymbol{B} = B_1, \ldots, B_n$ describes Bob's answers, and $\boldsymbol{X} = X_1, \ldots, X_n$ and $\boldsymbol{Y} = Y_1, \ldots, Y_n$ 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

[13] We 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 $\mathrm{P}_{\boldsymbol{A}|\boldsymbol{X}}$ and $\mathrm{P}_{\boldsymbol{B}|\boldsymbol{Y}}$ are proper conditional probability distributions, objects such as $\mathrm{P}_{A_1|X_1}$ are not well-defined.

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

$$\mathrm{O}_{ABXY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}(\tilde{a}\tilde{b}\tilde{x}\tilde{y}) = \frac{\left|\left\{i : (a_i, b_i, x_i, y_i) = (\tilde{a}, \tilde{b}, \tilde{x}, \tilde{y})\right\}\right|}{n}$$

and define

$$\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})} = \frac{\mathrm{O}_{ABXY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}}{\mathrm{Q}_{XY}} . \tag{1.1}$$

$\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ 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 of[15] $\mathrm{P}_{\boldsymbol{AB}|\boldsymbol{XY}} = \mathrm{O}_{AB|XY}^{\otimes n}$, it holds that if the device $\mathrm{O}_{AB|XY}$ is non-signalling then $\mathrm{P}_{\boldsymbol{AB}|\boldsymbol{XY}}$ is non-signalling and vice versa. This also implies that, for sufficiently large $n$, $\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ is non-signalling with high probability.

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

**Theorem 1.2** (Informal) *Let $\mathrm{P}_{\boldsymbol{AB}|\boldsymbol{XY}}$ be a non-signalling permutation invariant parallel device and $\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ as in Eq. (1.1). Then, for sufficiently large n, $\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ 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 $\mathrm{O}_{AB|XY}^{\otimes n}$ in which every single device $\mathrm{O}_{AB|XY}$ 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 $\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ and that if $\mathrm{O}_{AB|XY}^{\mathrm{freq}(\boldsymbol{a},\boldsymbol{b},\boldsymbol{x},\boldsymbol{y})}$ is close to a non-signalling device then its winning probability cannot be too far from the optimal non-signalling winning probability $1 - \alpha$.

[14] In 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 $i \neq j$).

[15] An IID device is illustrated in the bottom of Fig. 1.2. We can then think of each copy $\mathrm{O}_{AB|XY}$ as describing a single copy of the smaller boxes in the figure, while $\mathrm{P}_{\boldsymbol{AB}|\boldsymbol{XY}} = \mathrm{O}_{AB|XY}^{\otimes n}$ 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, $\omega_{\text{exp}}$, 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 $\boldsymbol{A}$ and $\boldsymbol{B}$.
```

### 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 $\omega_{\text{exp}}$ 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 $\boldsymbol{A} = A_1, \ldots, A_n$ 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 $H^{\varepsilon}_{\text{min}}(\boldsymbol{A}|E)$, 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 $H^{\varepsilon}_{\text{min}}(\boldsymbol{A}|E)$ determines the maximal length of the secret key that can be created by the protocol. Hence, proving security amounts to lower-bounding $H^{\varepsilon}_{\text{min}}(\boldsymbol{A}|E)$. Evaluating the smooth min-entropy $H^{\varepsilon}_{\text{min}}(\boldsymbol{A}|E)$ 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 $H^{\varepsilon}_{\min}(\boldsymbol{A}|E)$. 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 $E = E_1, \ldots, E_n$, 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 $H^{\varepsilon}_{\min}(\boldsymbol{A}|E)$ 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 form preparing more complex devices.

On the quantitive 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 quantitive 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 Phy 56(5):052203
12. Dupuis F, Fawzi O, Renner R (2016) Entropy accumulation. arXiv: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