ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Challenges in information-theoretic secret-key agreement

Master Thesis

João Miguel Lourenço Ribeiro

November 2017

Advisors: Prof. Dr. Ueli Maurer, Daniel Jost
Department of Computer Science, ETH Zürich

Abstract

This thesis is concerned with *information-theoretic secret-key agreement*: Alice and Bob want to agree, through public communication, on a shared secret-key about which a computationally-unbounded adversary, Eve, learns almost no information. We work in the source model, where Alice, Bob, and Eve have access to sources of randomness. Each source emits independent and identically distributed realizations of a random variable to its corresponding party. The three sources are correlated, in the sense that their underlying random variables are jointly distributed according to some probability distribution. To such a distribution we can assign its *secret-key rate*, which is the best rate (per number of realizations) at which Alice and Bob can create secret-key bits, while ensuring that Eve learns almost nothing about them. Our main focus will be the study of this fundamental quantity for different types of distributions.

We first consider the *satellite setting*, where a random bit is sampled (e.g. by a satellite), and then sent to the three parties via different binary symmetric channels. We start by surveying the main known results about the secret-key rate in this setting. Then, we show that the main ideas behind the parity-check protocol for advantage distillation can be extended in a natural way to yield an improved protocol. Finally, we study what happens when the satellite is allowed to choose the quality of the channels to Alice, Bob, and Eve, as long as Eve's channel is always D times better than Alice's and Bob's. Of particular importance, we establish the exact asymptotic behavior of the secret-key rate *per time unit* in this setting as a function of the quality ratio D, which also settles a conjecture proposed by Gander and Maurer.

In the second part, we study the secret-key rate for general distributions. First, we survey the best known upper-bounds on the secret-key rate and *bound information*: Do there exist distributions which require secret common information between Alice and Bob to be created, but from which no secret-key can be extracted? Finding a distribution where one of these upper-bounds is zero while another one is positive implies that this conjecture is true. Our results show that such a separation is unlikely to exist between the most tractable upper bounds considered, and that, if it does exist, then proving its existence ought to be a very complex task. Next, we analyze a class of distributions which is believed to have bound information and, in particular, provide further evidence for this via connections to *non-interactive correlation distillation* and *entanglement distillation*. To conclude, we discuss relaxed notions of secret-key rate, in particular the *deterministic secret-key rate*, for which Alice and Bob can only use deterministic protocols. We study the limit of a technique of Ozols, Smith, and Smolin for finding distributions with positive secret-key rate but zero deterministic secret-key rate, and give evidence for the existence of a candidate distribution for the separation which is beyond this limit through a connection to the *communication for omniscience by a neutral observer* problem.

Contents

Acknowledgements

I would like to start by thanking my supervisors, Daniel and Ueli, for the many discussions where interesting ideas arose, and for giving me the freedom to pursue my own ideas as well. I also want to thank Ueli for guiding me throughout my time in Zurich, and for always keeping the door open for me. His guidance has turned me into a much more confident researcher.

I had my first taste of research during my undergraduate studies in Lisbon. I thank Paulo Mateus and André Souto for introducing me to theoretical computer science, and I thank André especially for his support throughout the past few years.

I thank my closest friends, both in Lisbon and in Zurich, for the countless great times we shared, and for enduring my constant hops between these two cities.

This thesis and all that came before it would not have been possible without the unconditional support of my family. They always encouraged me to focus on what I enjoy.

Finally, I thank Sara for all we have shared – I could not have asked for better. This thesis is dedicated to the connection drops during our video calls.

Chapter 1

Introduction

Alice is preparing a surprise birthday party for Eve. She wants to tell Eve's friend, Bob, about all the details. There is a problem, though: Eve is nearby, and so she might hear Alice speaking to Bob. The surprise would be ruined!

In order to avoid this, Alice wants to encrypt her message to Bob in a way that Eve learns only a negligible amount of information about its content. If Alice and Bob share a secret-key, with the same length as the message, about which Eve has almost no information, then Alice can use the one-time pad [43] to transmit her message in a secure manner to Bob over an authenticated noiseless channel (a channel Eve can listen in to, but cannot tamper). Therefore, Alice and Bob only need to worry about *secret-key agreement*: They want to agree, through an authenticated noiseless channel, on a shared secret-key about which Eve learns only negligible information.

Alice and Bob could use the Diffie-Hellman secret-key agreement protocol [8] to generate a secret-key, which is widely used in practice. It works roughly as follows: Alice and Bob select a finite cyclic group G (a finite set with an operation having some nice properties) and a generator g (every element of G can be written as g^k for some integer k). Then, Alice and Bob pick secret integers a and b, respectively. Alice computes g^a and sends it to Bob over the channel, and Bob computes g^b and sends it to Alice over the channel. The fulcral observation is that Alice and Bob can now compute g^{ab} – this is their shared secret-key. Eve, on the other hand, sees only g^a and g^b. In order for the Diffie-Hellman protocol to be secure, it must be the case that Eve learns very little about g^{ab} from g^a and g^b.

Eve really hates surprises, though, and she is a very influential individual. Therefore, she is more than willing and able to amass whatever computational power she needs to recover the secret-key, if possible. Indeed, Eve can just compute all powers of g (comparing them with g^a and g^b, which she observed), find a and b, and then finally compute g^{ab}. This means that

we can only hope for security against computationally-bounded adversaries who are only allowed to run efficient (polynomial-time) algorithms. This corresponds to the notion of *computational security*. Unfortunately, even the assumption that a computationally-bounded Eve learns almost nothing about g^{ab} in the Diffie-Hellman protocol has still not been proven. She could have access to an efficient algorithm for breaking Diffie-Hellman, for all we know. This is a general phenomenon: Cryptographic schemes for computational security all rely on some currently unproven hardness assumption.

Due to all of the above, Alice and Bob are, naturally, not very confortable with a secret-key agreement protocol for computational security. Even if the hardness assumption turns out to be correct, Eve can still obtain their key if she uses enough computational power. They would be much happier if they could agree on a shared secret-key through public communication while ensuring that a computationally-unbounded adversary learns only arbitrarily little amount of information about the secret-key, also without relying on unproven hardness assumptions. This is called *information-theoretic secret-key agreement*, and it is the main focus of this thesis. It is a subset of *information-theoretic security*, which concerns itself with the general study of cryptographic schemes secure against computationally-unbounded adversaries. No matter how much computational power Eve directs towards breaking such a scheme, she will not succeed.

The Diffie-Hellman protocol only requires that an authenticated noiseless channel exists between Alice and Bob (besides the unproven hardness assumption mentioned above). Naturally, one may wonder whether such a channel is also enough for information-theoretic secret-key agreement. Unfortunately, this is not the case. In his ground-breaking paper which sparked the study of information-theoretic security, Shannon [39] proved the following: If there are no a priori assumptions about the secret-key and the message (other than their length), then Alice and Bob, using an authenticated noiseless channel, can only ensure that Eve learns nothing about the message from the ciphertext if they already start with a secret-key at least as long as the message. This implies that information-theoretic secret-key agreement is impossible using only an authenticated noiseless channel if we require that Eve learns nothing about the generated key. Furthermore, a result of Maurer [22] implies that information-theoretic secret-key agreement is impossible even if we only require that Eve learns sufficiently little information about the key. Therefore, we need some additional assumption to make it possible.

To avoid the impossibility results, we assume that Alice, Bob, and Eve have access to correlated sources of randomness. The idea here is that each party receives probabilistic information from its corresponding source, in particular some information about the other two parties. Exactly how much and

what kind of information each party receives about the others is specified by the correlations between sources. This means that, depending on the sources under consideration, some parties may start with incomplete information about the others. Note that this is always the case in practice, due to the physical limits of information transmission [45], and so our assumption is reasonable. We then obtain the *source model* for information-theoretic secret-key agreement, introduced by Maurer [21][22], which we will be working on. In this model, there exists a two-way authenticated noiseless channel between Alice and Bob. Moreover, Alice, Bob, and Eve receive several independent and identically distributed realizations of random variables X, Y, and Z, respectively, jointly distributed according to a fixed probability distribution P_{XYZ}. Obviously, the addition of randomness sources is realistic if sources exhibiting the distribution under consideration can be found in practice. Chapter 3 is dedicated to the study of a realistic class of distributions. Practicality need not be the main motivation behind such a model, though. The study of general distributions, and pathological ones, has led to exciting problems, concepts, and techniques in (classical and quantum) information theory and cryptography, as we shall see in Chapter 4.

We will focus on a well-known fundamental quantity associated to each probability distribution P_{XYZ}, called the *secret-key rate*. It was originally defined in [22][23] (these two definitions were later shown to be equivalent [27]). Informally, it is described in [23] as the optimal number of secret-key bits Alice and Bob can generate per realization of X, Y, and Z in the source model while ensuring that Eve learns arbitrarily little information about them. This quantity is difficult to compute, and it is still unknown for many classes of distributions.

A more detailed account of the historical context, early developments, and the most significant models of information-theoretic secret-key agreement can be found in [45] and [24], on which the exposition above is based.

Fairly recently, results from information-theoretic secret-key agreement, in particular ones obtained in models where Alice, Bob, and Eve are connected through noisy channels in some way, such as the influential wire-tap channel model of Wyner [46], and in the source model described above led to the emergence of *physical layer security*. This field focuses on the design and implementation of practical methods for secure information transmission in wireless networks, which is of particular importance in an age where the number of sensitive devices connected to each other is rapidly increasing. A detailed, recent survey of this area can be found in [33].

The bulk of this thesis (after Chapter 2, which goes over the required technical background) is divided into two parts with clearly distinct motivations. In each part, we study an important challenge in information-theoretic secret-key agreement in the source model. The main difference is that the

first challenge is *concrete*, in the sense that we focus on very specific settings, while the second challenge is *general*, because we make almost no assumptions about the settings under consideration.

First, in Chapter 3, we study the *satellite setting*, introduced in [21][22]. Suppose there is a satellite which samples random bits and broadcasts them to Earth. On the ground, Alice, Bob, and Eve are listening to the satellite with their respective antennae, and each party receives the random bit with a certain error probability depending on the quality of its antenna. Then, the random variables X, Y, and Z correspond to the bits Alice, Bob, and Eve receive, respectively. This is a conceptually simple setting where computing the secret-key rate remains an open problem. Moreover, it exemplifies the power of interaction between Alice and Bob. If only messages from Alice to Bob are allowed, then secret-key agreement is possible only when Eve's error probability is larger than Alice's and Bob's. On the other hand, if interaction is allowed, then secret-key agreement is possible even in cases where Eve's error probability is smaller than Alice's and Bob's. The concrete challenge is to obtain better estimates of the secret-key rate as a function of the error probabilities, better explicit secret-key agreement protocols, and to attempt to translate the satellite setting into practice.

We begin by surveying fundamental protocols for advantage distillation, which can then be used to achieve secret-key agreement, namely the repeater-code protocol [21][22] and the parity-check protocol [21][11]. Then, we show that the ideas behind the parity-check protocol can be exploited further to obtain an improved secret-key agreement protocol. The key observation here is that Alice and Bob throw away many bits during the parity-check protocol which are still good, in the sense that we can still obtain some extra secret-key rate from them. Finally, we investigate how the satellite setting translates into practice. As part of this, we study the secret-key rate per time unit (instead of per realization) when the satellite is allowed to choose the error probabilities of all parties under the constraint that Eve's capacity is D times larger than Alice's and Bob's. We determine the asymptotic behavior of this notion of secret-key rate as a function of D (it behaves like $1/D$), and show that the parity-check protocol is already very good in practice (its rate also behaves like $1/D$). In the process, we also show that the secret-key rate per realization of the parity-check protocol in this alternative setting behaves like $1/D^2$, settling a conjecture of Gander and Maurer [11].

In Chapter 4, we study a conjectured criterion for the positivity of the secret-key rate for general distributions and the associated notion of *bound information*. Maurer and Wolf [26] conjectured that the intrinsic mutual information, an upper bound on the secret-key rate which is relatively tractable, is positive if and only if the secret-key rate is also positive. The general challenge is then to verify the validity of this conjecture. It is widely believed that

this conjecture is false, and so that there exist distributions with bound information [12], i.e. distributions with positive intrinsic mutual information but zero secret-key rate. It has proved very challenging to find such a distribution.

We start by introducing some of the best upper bounds on the secret-key rate, namely the intrinsic mutual information [26], the reduced intrinsic mutual information [35], and a bound due to Gohari and Anantharam [13] (in increasing order of tightness), and fundamental results about bound information. Then, we consider the existence of distributions with positive intrinsic mutual information but for which one of the other two bounds above is zero. This implies that these distributions have bound information. Known results imply that such distributions do not exist when X, Y, and Z are finite random variables. We extend this result with respect to the reduced intrinsic mutual information to the case where only one of X or Y must be finite. We also show that if there exists a distribution P_{XYZ} having finite X and Y but unbounded Z with positive intrinsic mutual information and for which the bound of Gohari and Anantharam is zero, then this should be very hard to prove, in the sense that the ranges of a sequence of random variables one must produce in the proof grow very fast, or are infinite. Next, we analyze a class of distributions, parameterized by a parameter a, which are believed to have bound information for large enough a. We show that the only clear possible strategy for secret-key agreement does not work when a is large enough. We also showcase a connection to non-interactive correlation distillation [28][47] which gives more evidence for bound information. Finally, we investigate relaxed notions of secret-key rate, namely the secret-key rate by public discussion [30] (all auxiliary random variables in a protocol must be public) and the deterministic secret-key rate (protocols are deterministic), where the latter is a weaker notion than the former. We show that the techniques from [30] for finding distributions with positive secret-key rate but zero secret-key rate by public discussion do not work when X and Y are binary random variables. With this in mind, we provide a candidate distribution where X and Y are binary with positive secret-key rate which we believe has zero deterministic secret-key rate. This belief comes from a connection between the deterministic secret-key rate and the problem of communication for omniscience by a neutral observer [13].

There are two goals in each part described above. The first goal is to provide a comprehensive survey of the topic in question, such that a dedicated reader, armed only with mathematical maturity and the concepts of Chapter 2, can learn the fundamental results of the topic and start his/her own research. The second goal is to present the results obtained independently while working on the relevant topic for the thesis. These can be found mainly in Sections 3.4, 3.5, 3.6, 4.2, 4.3, and 4.4.

Chapter 2

Notation and background

In this chapter, we introduce some notation and the basic concepts and results used throughout the thesis.

2.1 Notation

In this section, we fix some notation.

We denote random variables by uppercase letters like X, Y, and Z, and their ranges are denoted by calligraphic letters $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$. In fact, sets are denoted by calligraphic letters, except for a few special cases. We denote the set of natural numbers (starting at 0) by $\mathbb{N}$. A set $\mathcal{X}$ is countable if there is a surjection $f : \mathbb{N} \to \mathcal{X}$. The set of real numbers is denoted by $\mathbb{R}$. The real closed line $\mathbb{R} \cup \{-\infty, \infty\}$ is denoted by $\overline{\mathbb{R}}$. When enclosing an interval of real numbers, parentheses mean that the interval is open on that side, while square brackets mean that the interval is closed. For example, [0,1) is closed on the left and open on the right. We denote the set of finite binary strings by $\{0,1\}^*$. For the cardinality of a set $\mathcal{X}$, we use the symbol $|\mathcal{X}|$. For $x \in \mathbb{R}$, $|x| = \max(x, -x)$. The logarithm to the base 2 is denoted by log, and the natural logarithm is denoted by ln.

Given a string $x \in \{0,1\}^*$, its Hamming weight $w(x)$ is defined as

$$w(x) := |\{i : x_i \neq 0\}|.$$

An N-tuple $(x_1, \ldots, x_N)$ may be denoted by x^N. Collections of random variables $(X_1, \ldots, X_N)$ may also be denoted by X^N.

We denote sequences $a_1, a_2, \ldots$ as (a_i). We denote the fact that $\lim_{i\to\infty} a_i = c$ by $a_i \to c$. We say that $a_i \sim b_i$ if $\lim_{i\to\infty} \frac{a_i}{b_i} = 1$. We say a limit *exists* if it is in $\overline{\mathbb{R}}$, and we say it is *finite* if it is in $\mathbb{R}$, which we may also denote by $L < \infty$, where L stands for a limit. We use lim inf to denote the limit inferior of a sequence.

Given two functions $f, g : \mathbb{N} \to \mathbb{R}$, we say that $f = O(g)$ if there exists $c > 0$ such that $f(n) \leq cg(n)$ for all large enough n. We say that $f = \Theta(g)$ if there exist constants $c_1, c_2 > 0$ such that $c_1 g(n) \leq f(n) \leq c_2 g(n)$ for all large enough n, i.e. if $f = O(g)$ and $g = O(f)$. We say that $f = o(g)$ if for every $c > 0$ there exists $n_0 \in \mathbb{N}$ such that $f(n) \leq cg(n)$ for all $n \geq n_0$, i.e. if

$$\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0.$$

2.2 Probability theory

Information-theoretic security is heavily based on probability theory. In this section, we cover basic concepts and results from probability theory that will be useful in the following chapters. A detailed and rigorous introduction to probability theory can be found in [10].

We assume some very basic familiarity with the concept of a discrete random variable (a random variable taking on countably many values), and with events in discrete probability spaces (a set of outcomes in a countable sample space). A concise introduction to these concepts can be found in [44, Sections 2.2.1 and 2.2.2]. As a matter of uniformity, we use similar notation to [44, Section 2.2]. Our exposition here will also have some features in common with, and parts inspired by, that section, since the concepts and results we require for this thesis overlap those covered in [44, Section 2.2] to a certain degree.

Given an event A, we denote its probability by $\Pr[A]$, i.e. $\Pr[A]$ is the sum of the probabilities of the outcomes in A. The probability that events A and B both happen is denoted by $\Pr[A, B]$. The probability that at least one of A and B happens is denoted by $\Pr[A \vee B]$. Both notations are extended in the obvious way to a larger set of events.

One of the most useful and basic inequalities in probability theory is the union bound. It implies that if the probability of each event is small, and there are not too many events, then the probability that at least one of them happens is also small.

Lemma 2.1 (Union bound) *If $A_1, \ldots, A_N$ is a sequence of events, then*

$$\Pr[A_1 \vee \cdots \vee A_N] \leq \sum_{i=1}^{N} \Pr[A_i].$$

We move on to random variables. A *discrete random variable* X is induced by a function

$$P_X : \mathcal{X} \to [0, 1],$$

where $\mathcal{X}$ is a countable set, satisfying $\sum_{x\in\mathcal{X}} P_X(x) = 1$, which we call the *probability distribution of* X, where $P_X(x)$ is the probability that X is equal to x. The *range of* X is $\{x \in \mathcal{X} : P_X(x) > 0\}$, i.e. the set of values X can take. If the range of X is finite, we say that X is a *finite* random variable. Throughout this thesis, random variables will always be discrete, unless stated otherwise.

Indicator variables are simple, but very relevant, examples of discrete random variables that we will use later on. Given an event A, the *indicator variable of* A, denoted by 1_A, is the random variable with range $\{0,1\}$ and probability distribution P such that $P(1) = \Pr[A]$ and $P(0) = 1 - \Pr[A]$.

There are two fundamental quantities associated to a discrete random variable X defined over $\mathcal{X} \subseteq \mathbb{R}$. The *expected value of* X, denoted by $\mathrm{E}[X]$, is defined as

$$\mathrm{E}[X] := \sum_{x\in\mathcal{X}} P_X(x) \cdot x.$$

Note that the expected value need not be finite, or even exist. However, if $\sum_{x\in\mathcal{X}} P_X(x) \cdot |x| < \infty$, then the expected value is well-behaved, and X is said to be *integrable*. When it is not clear over which random variable we are taking the expected value, we may denote the expected value of Y with respect to X, where Y is a function $Y(X)$ of X, as $\mathrm{E}_X(Y) := \sum_{x\in\mathcal{X}} P_X(x) \cdot Y(x)$. Intuitively, the expected value of X is the value around which X is concentrated.

The *variance of* X, denoted by $\mathrm{Var}[X]$, is defined as

$$\mathrm{Var}[X] := \mathrm{E}[(X - \mathrm{E}[X])^2].$$

Intuitively, the variance of X tells us how tightly X is concentrated around its expected value. Indeed, small variance leads to tight concentration.

The relationship between the expected value and the variance mentioned above can be made rigorous by the following inequality, due to Chebyshev.

Lemma 2.2 (Chebyshev's inequality) *Let X be a random variable with range contained in $\mathbb{R}$. Then*

$$\Pr[|X - \mathrm{E}[X]| \geq a] \leq \frac{\mathrm{Var}[X]}{a^2}$$

for all $a > 0$.

We will be making heavy use of the important *Jensen's inequality*, which relates the image of the expected value of a random variable to the expected value of its image for a special class of functions. A function $f : \mathcal{X} \to \mathbb{R}$ with $\mathcal{X} \subseteq \mathbb{R}$ is said to be *convex* if, for all $\lambda \in [0,1]$ and $x, y \in \mathcal{X}$,

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda) f(y).$$

Intuitively, a function is convex if we can connect two points on its graph by a line segment that always stays above the graph. It is possible to check that a function f is convex by verifying that its second derivative is non-negative, if it exists. A function f is said to be *concave* if $-f$ is convex.

Lemma 2.3 (Jensen's inequality) *Let $f : \mathcal{X} \to \mathbb{R}$ be a convex function with $\mathcal{X} \subseteq \mathbb{R}$, and let X be an integrable discrete random variable defined over $\mathcal{X}$ such that $\mathrm{E}[X] \in \mathcal{X}$. Then*

$$f(\mathrm{E}[X]) \leq \mathrm{E}[f(X)],$$

provided that $f(X)$ is also integrable.

Jensen's inequality is very useful for obtaining lower and upper bounds, for example when f is a logarithm or a square root.

Two random variables X and Y defined over $\mathcal{X}$ and $\mathcal{Y}$, respectively, have a joint probability distribution P_{XY} defined over $\mathcal{X} \times \mathcal{Y}$. In particular, this distribution encodes how much information X gives about Y, and vice-versa. The distributions of X and Y are induced by P_{XY} by setting $P_X(x) = \sum_{y \in \mathcal{Y}} P_{XY}(x, y)$. The random variables X and Y are said to be *independent* if $P_{XY}(x, y) = P_X(x)P_Y(y)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Intuitively, X and Y are independent if they give no information about each other. Two important consequences, based on what we have covered, are that, whenever X and Y are independent, $\mathrm{E}[XY] = \mathrm{E}[X] \cdot \mathrm{E}[Y]$ and $\mathrm{Var}[X + Y] = \mathrm{Var}[X] + \mathrm{Var}[Y]$.

When we know an event A happened, our knowledge about a random variable X may change. This is modelled by updating the distribution of X. For an event A with $\Pr[A] > 0$, the *conditional distribution of X given A*, denoted by $P_{X|A}$, is given by

$$P_{X|A}(x) := \frac{\Pr[X = x, A]}{\Pr[A]}.$$

An equivalent characterization of independence between two random variables X and Y is that $P_{X|Y=y}(x) = P_X(x)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. We may also write $P_{X|Y}(x, y)$ for $P_{X|Y=y}(x)$. Two random variables are said to be *conditionally independent given A* if $P_{XY|A}(x, y) = P_{X|A}(x)P_{Y|A}(y)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Note that if X and Y are independent, it is not necessarily true that they are conditionally independent given any event A. We define the *expected value of X given A*, denoted by $\mathrm{E}[X|A]$, as $\mathrm{E}[X|A] := \sum_x P_{X|A}(x) \cdot x$.

Sometimes, we will be dealing with a sequence of random variables $X_1, \ldots, X_N$ such that X_i is obtained by applying some randomized function to X_{i-1}. In this case, X_i only depends on $X_1, \ldots, X_{i-2}$ through X_{i-1}, and so X_i is conditionally independent of $X_1, \ldots, X_{i-2}$ given X_{i-1}. In other words,

$$P_{X_i|X_1,\ldots,X_{i-1}} = P_{X_i|X_{i-1}}$$

holds for all i. Such a sequence is called a *Markov chain*, and is denoted by $X_1 \to X_2 \to \cdots \to X_N$. As a special case, a sequence of random variables

$X_1, \ldots, X_N$ is said to be *independent and identically distributed*, or i.i.d. for short, if the probability distributions of all the X_i are the same, and all the X_i are independent.

We can define a natural distance between discrete probability distributions. Given two discrete probability distributions P and Q over the same set $\mathcal{X}$, the *statistical distance between P and Q*, denoted by $\Delta(P, Q)$, is defined as

$$\Delta(P, Q) := \sum_{x \in \mathcal{X}} |P(x) - Q(x)|.$$

We remark that the above expression is usually multiplied by $1/2$ to normalize the statistical distance, but we opt to drop this factor as it is not relevant to us. The following properties hold for the statistical distance:

1. $\Delta(P, Q) = 0$ if and only if $P = Q$;
2. $\Delta(P, Q) = \Delta(Q, P)$;
3. $0 \leq \Delta(P, Q) \leq 2$;
4. $\Delta(P, Q) \leq \Delta(P, R) + \Delta(R, Q)$ for all discrete probability distributions R.

The statistical distance can also be interpreted as the L^1 distance in the space of probability distributions. Moreover, an important property of the statistical distance is that, if $\Delta(P, Q) \leq \varepsilon$, the probabilities P and Q assign to an event never differ by more than $\varepsilon/2$.

We conclude with a result that is not purely probability theoretic. When bounding some probabilities later on, it will be useful to know the asymptotic behavior of $n!$.

Lemma 2.4 (Stirling's approximation) *We have*

$$n! \sim \sqrt{2\pi n} \left(\frac{n}{e}\right)^n.$$

2.3 Information theory

In this section, we provide an overview of basic results and concepts from information theory which will be useful in this thesis. A very nice and complete overview of both introductory and advanced topics in information theory, on which we base the results we cover, can be found in [6]. A more general and advanced approach to information theory can be found in [14].

The study of information theory was initiated by Shannon [37][38]. Since then, it has become an extremely vibrant and influential field.

All the results stated in this section are well-known and hold for finite random variables, which will be the main focus of our thesis. However, we

will work with discrete random variables with infinite ranges in Sections 4.1 and 4.2. Therefore, we present the results in as general form as we will need them in these sections. We include a detailed proof whenever an extension is required but it is not straightforward from previous properties.

Fix a discrete random variable X with range $\mathcal{X}$. The *Shannon entropy of* X, or entropy for short, denoted by $H(X)$, is defined as

$$H(X) := -\sum_{x\in\mathcal{X}} P_X(x)\log P_X(x) = \mathrm{E}_X[-\log P_X(X)].$$

The entropy is a property of the distribution P_X, so we may also write $H(P_X)$ for the entropy of X. When considering a conditional distribution $P_{X|A}$ for some event A, we write $H(P_{X|A})$ as $H(X|A)$. Intuitively, the entropy of X measures the uncertainty one has about the value X takes. Since $H(X)$ is a sum of non-negative real numbers, it follows that $H(X) \in \overline{\mathbb{R}}$, and $H(X) \geq 0$. Note that there exist discrete random variables X with $H(X) = \infty$. For finite random variables, the entropy satisfies the following inequalities.

Lemma 2.5 *For any finite random variable X with range $\mathcal{X}$,*

$$0 \leq H(X) \leq \log|\mathcal{X}|.$$

The case where X has range $\{0,1\}$ gives rise to a very important function. Suppose that $P_X(1) = p$. Then $H(X) = h(p)$, where $h : [0,1] \to \mathbb{R}$ is the *binary entropy function*, defined as

$$h(p) := -p\log(p) - (1-p)\log(1-p).$$

The binary entropy function is concave, and so, by Jensen's inequality,

$$\mathrm{E}[h(X)] \leq h(\mathrm{E}[X])$$

for all random variables with range contained in $[0,1]$. Furthermore, h is symmetric around $1/2$, which means that $h(p) = h(1-p)$ for all $p \in [0,1]$. Another useful property is that, when $p \leq 1/2$, we have

$$p \leq h(p) \leq 2p\log(1/p).$$

All these properties will turn out to be extremely useful later in the thesis.

The *conditional entropy of X given Y*, denoted by $H(X|Y)$, is defined as

$$H(X|Y) := \sum_y P_Y(y)H(X|Y=y).$$

Intuitively, $H(X|Y)$ is the expected uncertainty we have about X when we know Y.

A useful concept in information theory and probability theory is the *Kullback-Leibler divergence* of two probability distributions P and Q defined over the same set $\mathcal{X}$, denoted by $D(P||Q)$, and defined as

$$D(P||Q) := \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)},$$

with the conventions that, if there is x such that $P(x) > 0$ and $Q(x) = 0$, then $D(P||Q) = \infty$, and that $0 \log(0/q) = 0$ for all $q \in [0,1]$. It is possible to prove, using Jensen's inequality, that $D(P||Q) \geq 0$ whenever $D(P||Q)$ exists, i.e. whenever $D(P||Q) \in \overline{\mathbb{R}}$. Note also that $D(P||Q) = 0$ if and only if $P = Q$.

There exists a well-known inequality in information theory that relates the Kullback-Leibler divergence and the statistical distance of two distributions, called Pinsker's inequality [32]. Informally, if two distributions have small Kullback-Leibler divergence, then they also have small statistical distance. We present a version of Pinsker's inequality with an improved constant.

Lemma 2.6 (Pinsker's inequality [6, Lemma 11.6.1]) *For any two finite probability distributions P and Q, we have*

$$\frac{\log(e)}{2} \Delta(P,Q)^2 \leq D(P||Q).$$

The *mutual information between X and Y*, denoted by $I(X;Y)$, is the amount of information we gain about X when we know Y, and is defined as

$$I(X;Y) := D(P_{XY}||P_X P_Y).$$

The mutual information has some useful basic properties, as stated in the following lemma.

Lemma 2.7 *We have*

1. $I(X;Y) \geq 0$ *if it exists;*
2. $I(X;Y) = 0$ *if and only if X and Y are independent;*
3. $I(X;Y) = I(Y;X)$*;*
4. *If $H(X)$ is finite, $I(X;Y)$ is also finite, and $I(X;Y) = H(X) - H(X|Y)$;*
5. *If $H(X)$ and $H(Y)$ are finite, $I(X;Y) = H(X) + H(Y) - H(XY)$.*

Proof We prove only property 4. Properties 1-3 are straightforward and property 5 follows in an analogous way to property 4. Suppose $H(X)$ is

finite. We have

$$I(X;Y) = \sum_{x,y} P_{XY}(x,y) \log \frac{P_{XY}(x,y)}{P_X(x)P_Y(y)} = \\ = \sum_{x,y} \left[P_Y(y) P_{X|Y}(x,y) \log P_{X|Y}(x,y) - P_{XY}(x,y) \log P_X(x) \right].$$

Note that

$$\sum_{x,y} -P_{XY}(x,y) \log P_X(x) = H(X) < \infty.$$

Thus,

$$I(X;Y) = H(X) + \sum_{x,y} P_Y(y) P_{X|Y}(x,y) \log P_{X|Y}(x,y).$$

The infinite series in the right hand side is a sum of non-positive terms only. Therefore, $I(X;Y) \in \overline{\mathbb{R}}$. This implies that $I(X;Y) \geq 0$ since the Kullback-Leibler divergence is always non-negative whenever it exists. Then,

$$-\sum_{x,y} P_Y(y) P_{X|Y}(x,y) \log P_{X|Y}(x,y) \leq H(X) < \infty,$$

and the desired property follows by noting that

$$-\sum_{x,y} P_Y(y) P_{X|Y}(x,y) \log P_{X|Y}(x,y) = H(X|Y).$$

□

The following lemma states that additional knowledge can only reduce our uncertainty about a given random variable.

Lemma 2.8 (Conditioning reduces entropy) *For any two discrete random variables X and Y,*

$$H(X) \geq H(X|Y).$$

Proof Follows directly from the properties 1 and 4 of Lemma 2.7. □

For an event A, we define $I(X;Y|A) := D(P_{XY|A} || P_{X|A} P_{Y|A})$. The *conditional mutual information of X and Y given Z*, denoted by $I(X;Y|Z)$, is defined as

$$I(X;Y|Z) := \sum_z P_Z(z) I(X;Y|Z=z).$$

We have the following lemma.

Lemma 2.9 *If $H(X)$ is finite, then $I(X;Y|Z)$ is also finite, and*

$$I(X;Y|Z) = H(X|Z) - H(X|YZ).$$

Moreover, if $H(Y)$ is also finite,

$$I(X;Y|Z) = H(X|Z) + H(Y|Z) - H(XY|Z).$$

Proof Note that $H(X|Z = z)$ is finite for all z, since $H(X|Z) \leq H(X)$. It follows that

$$I(X;Y|Z = z) = H(X|Z = z) - H(X|Y, Z = z),$$

for all z, by Lemma 2.7. Then,

$$I(X;Y|Z) = \sum_z P_Z(z)(H(X|Z = z) - H(X|Y, Z = z)) = H(X|Z) - H(X|YZ),$$

which is finite, as desired. The second equality follows in a similar manner.□

The following alternative expression for the conditional mutual information will be useful,

$$I(X;Y|Z) = H(XZ) + H(YZ) - H(XYZ) - H(Z),$$

which holds whenever $H(X)$, $H(Y)$, and $H(Z)$ are finite.

The entropy and the mutual information satisfy *chain rules*, as stated in the following lemmas.

Lemma 2.10 (Chain rule for entropy) *Given a sequence of discrete random variables X^N, we have*

$$H(X^N|Z) = \sum_{i=1}^{N} H(X_i|X^{i-1}Z),$$

whenever $H(X^N)$ is finite, for all discrete random variables Z.

Proof The equality $H(X_1X_2|Z) = H(X_1|Z) + H(X_2|X_1Z)$ is a direct consequence of Lemma 2.9 when $H(X_1X_2)$ is finite, and the general case follows easily by induction. □

Lemma 2.11 (Chain rule for mutual information) *Given a discrete random variable X with $H(X)$ finite, and a sequence of discrete random variables Y^N, we have*

$$I(X;Y^N|Z) = \sum_{i=1}^{N} I(X;Y_i|Y^{i-1}Z)$$

for all discrete random variables Z.

Proof Direct consequence of Lemma 2.9 when $H(X)$ is finite. □

The conditional mutual information has two more useful properties. Both are obtained as easy consequences of the chain rule for mutual information and the properties of Lemma 2.7. First, the conditional mutual information is additive in a sense.

Lemma 2.12 *Suppose X^N and Y^N are sequences of random variables such that the pairs (X_i, Y_i) are all conditionally independent given Z. Then,*

$$I(X^N; Y^N|Z) = \sum_{i=1}^{N} I(X_i; Y_i|Z)$$

whenever either $H(X^N)$ or $H(Y^N)$ is finite.

Proof It suffices to note that

$$\begin{aligned} I(X^N; Y^N|Z) = H(X^N|Z) - H(X^N|Y^N Z) = \\ = \sum_{i=1}^{N} H(X_i|X^{i-1}Z) - H(X_i|X^{i-1}Y^N Z) = \\ = \sum_{i=1}^{N} H(X_i|Z) - H(X_i|Y_i Z) = \sum_{i=1}^{N} I(X_i; Y_i|Z). \end{aligned}$$

where the first and fourth equalities follow from Lemma 2.9 (because $H(X^N)$ and $H(X_i)$ are finite), the second equality follows from the chain rule, and the third equality follows from the conditional independence hypothesis. □

Second, the conditional mutual information satisfies some continuity property with respect to side information.

Lemma 2.13 *For any discrete random variable U,*

$$|I(X; Y|ZU) - I(X; Y|Z)| \leq H(U)$$

whenever either $H(X)$ or $H(Y)$ is finite.

Proof We prove only that $I(X; Y|ZU) \leq I(X; Y|Z) + H(U)$. The lower bound is analogous. Since the case $H(U) = \infty$ is trivial, suppose $H(U) < \infty$. Then,

$$\begin{aligned} I(X; Y|ZU) = H(X|ZU) - H(X|YZU) \leq \\ \leq H(X|Z) - H(XU|YZ) + H(U|YZ) \leq I(X; Y|Z) + H(U), \end{aligned}$$

where the equality follows from Lemma 2.9, the first inequality holds because conditioning reduces entropy, and the second inequality follows by applying the chain rule to $H(XU|YZ)$. □

Channels are ubiquitous objects in information theory. A *discrete memoryless channel* receives a discrete random variable X over a set $\mathcal{X}$ as input, and outputs a discrete random variable W over a set $\mathcal{W}$ according to a distribution which may depend on the value of X. Therefore, a channel is characterized by a conditional probability distribution $P_{W|X}$. The channel does not

keep state between inputs, i.e. the output of the channel, W, is conditionally independent of everything else given its corresponding input, X.

We will be dealing with the binary symmetric channel with error probability ε (BSC_ε), which is one of the simplest channels. Both the input and output alphabets are $\{0,1\}$, and it is defined by the following conditional probability distribution,

$$P_{W|X}(b,b) = 1-\varepsilon, \quad P_{W|X}(1-b,b) = \varepsilon,$$

where $b \in \{0,1\}$. Intuitively, the BSC_ε receives a bit and flips it with probability ε.

A fundamental quantity associated to every channel is its capacity. The *capacity of* $P_{W|X}$, denoted by $C(P_{W|X})$, is the supremum of all rates at which one can communicate through the channel with error probability converging to 0. It was proved by Shannon [37][38] that

$$C(P_{W|X}) = \sup_{P_X} I(X;W).$$

In particular, the capacity of the BSC_ε equals $1-h(\varepsilon)$, where h is the binary entropy function.

There are other entropy measures beyond Shannon entropy. We will make use of the so-called *min-entropy of* X, denoted by $H_\infty(X)$, which is defined as

$$H_\infty(X) := -\log \max_x P_X(x).$$

Random variables X with $H_\infty(X) \leq k$ satisfy $P_X(x) \geq 2^{-k}$ for some x. Note that $H_\infty(X)$ is defined for all discrete random variables, even if their range is infinite. This is because, for a discrete random variable X with infinite range, we must have $P_X(i) \to 0$ as $i \to \infty$, and so there are only finitely many values i satisfying $P_X(i) \geq a$ for each $a \in (0,1]$. It holds that $H(X) \geq H_\infty(X)$ for every discrete random variable X.

2.4 Information-theoretic secret-key agreement

In this section, we introduce the basic concepts and results of information-theoretic secret-key agreement in Maurer's source model, which we will be focusing on in the next chapters, and which we already discussed briefly in the introduction. This topic has gained significant attention both from the cryptography and information theory communities. A detailed introduction to information-theoretic secrecy can be found in [9]. Surveys covering historical context, early results and influential models for information-theoretic secret-key agreement can be found in [45] and [24].

Maurer first introduced and studied the *source model* in 1990 [21][22]. In this model, Alice and Bob are connected by a two-way noiseless authenticated

channel which Eve has full access to. Furthermore, Alice, Bob, and Eve receive i.i.d. realizations of discrete random variables X, Y, and Z, respectively, with joint probability distribution P_{XYZ}. Ahlswede and Csiszár [2] studied the case where only messages from Alice to Bob are allowed in the source model.

Throughout this section, we will assume that $H(XYZ)$ is finite. This is weaker than assuming that X, Y, and Z have finite ranges, and it will prove to be useful when we have to deal with certain classes of distributions P_{XYZ} with infinite ranges in parts of Sections 4.1 and 4.2. When the original proof of a result does not extend directly to distributions such that $H(XYZ)$ is finite, we provide such an extension.

As previously discussed, a fundamental quantity associated with every discrete probability distribution P_{XYZ} is its secret-key rate, i.e. the optimal rate (with respect to the number of realizations of X and Y required) at which Alice and Bob can generate secret-key bits while keeping them secret from Eve. More precisely, we have the following definition.

Definition 2.14 ([23, Definition 2]) *Given a discrete probability distribution P_{XYZ}, the* secret-key rate of X and Y given Z, *denoted by $S(X;Y||Z)$, is the supremum of all real numbers $R \geq 0$ such that, for all $\varepsilon > 0$ and all N large enough there exists a communication protocol for Alice and Bob, who start with N i.i.d. realizations of X and Y, respectively, with communication $C = (C_1, C_2, \ldots, C_M)$, where the number of messages M is also a random variable, such that*

$$H(C_i|C^{i-1}, X^N, R_A) = 0,$$

for i odd, where R_A is the random variable corresponding to Alice's local randomness, and

$$H(C_i|C^{i-1}, Y^N, R_B) = 0,$$

for i even, where R_B is the random variable corresponding to Bob's local randomness. In other words, the protocol proceeds in an alternating style, with Alice sending the odd-numbered messages, and Bob the even-numbered ones for a finite number of steps. We further require that $H(R_A)$ and $H(R_B)$ be finite. Eve, on the other hand, has access to N i.i.d. realizations of Z, and observes the communication C. At the end of the protocol, Alice and Bob must produce finite random variables S and S', respectively, with S having range $\mathcal{S}$, such that

$$H(S|C, X^N, R_A) = 0,$$

and

$$H(S'|C, Y^N, R_B) = 0.$$

Furthermore, S and S' must satisfy the following properties:

1. $H(S) \geq N(R - \varepsilon)$;

2. $H(S) \geq \log |\mathcal{S}| - \varepsilon$;

3. $\Pr[S = S'] \geq 1 - \varepsilon$;

4. $I(S; Z^N C) \leq \varepsilon$.

Note that Definition 2.14 does not coincide with the original definition of the secret-key rate in [22, Definition 2], and in fact appears to be much stronger. Originally, only the rate at which Eve learns information about the key was bounded, i.e. property 4 in Definition 2.14 was instead $I(S; Z^N C) \leq \varepsilon N$. From a cryptographic point of view, bounding the rate is not enough, since Eve could still learn a significant amount of information about the key [22]. Furthermore, the condition for almost-uniformity (property 2 in Definition 2.14) was not enforced.

This gave rise to an initial distinction between the *weak* secret-key rate, corresponding to the definition in [22, Definition 2], and the *strong* secret-key rate, first introduced in [23, Definition 2] and corresponding to Definition 2.14. It turns out that such a distinction is not necessary, since Maurer and Wolf [27, Theorem 1] proved that the two rates are the same. In other words, the secret-key rate is not affected by requiring that the stronger properties (2 and 4) hold, and so we can always work with the stronger notion of secrecy.

The following lemma presents lower and upper bounds for the secret-key rate. The lower bound was proved by Maurer in [22] (originally for the weak secret-key rate) and in [23] (for the strong secret-key rate), with a simplified proof by Wolf [44], and also by Ahlswede and Csiszár [2] (initially only for the weak secret-key rate). The upper bound was proved by Maurer [22].

Lemma 2.15 ([22, Theorem 2] and [23, Theorem 4]) *We have*

$$S(X;Y||Z) \geq \max(I(X;Y) - I(X;Z), I(X;Y) - I(Y;Z))$$

and

$$S(X;Y||Z) \leq \min(I(X;Y), I(X;Y|Z)).$$

Proof In [36, Theorems 1 and 2], Renner and Wolf state that the bounds of this lemma hold for all discrete probability distributions. We prove this by examining and/or extending the original proofs.

For the upper bound, the proof by Maurer [22, Theorem 2] can be seen to hold whenever $H(X)$ and $H(Y)$ are finite. Alternatively, we can also take the bound when X, Y, and Z is finite and extend it to the case where only one of X and Y is necessarily finite (and the other random variables can have unbounded range) without making any assumptions about finite entropies. In fact, the results we obtain in Section 4.2 are directed at distributions P_{XYZ} where one of X or Y is finite, and so this alternative would also be enough.

For the lower bound, the proofs by Maurer [23, Theorem 4] and Wolf [44, Theorem 4.7] hold whenever X, Y, and Z are finite random variables. We first extend the bound to the case where Z has unbounded range and X, Y are finite. Then, we extend it to the case where X, Y, and Z all have unbounded ranges.

Suppose X and Y are finite random variables, and Z has unbounded range. Fix $\alpha > 0$ and k_α such that $\Pr[Z \leq k_\alpha] > 1 - \alpha$. Let Z_α be the random variable such that $Z_\alpha = Z$ if $Z \leq k_\alpha$ and $Z_\alpha = \perp$ otherwise. Furthermore, consider U_α such that $U_\alpha = \perp$ if $Z_\alpha \neq \perp$ and $U_\alpha = (X, Y)$ otherwise. Then,

$$S(X;Y||Z) \geq S(X;Y||Z_\alpha U_\alpha),$$

because if Eve knows $U_\alpha = (X, Y)$ and $Z_\alpha = \perp$ then she can simulate Z perfectly and discard U_α, which would leave her in the original situation. Since $Z_\alpha U_\alpha$ is a finite random variable, we have

$$S(X;Y||Z_\alpha U_\alpha) \geq I(X;Y) - I(X;Z_\alpha U_\alpha) = H(X|Z_\alpha U_\alpha) - H(X|Y).$$

Now it suffices to note that

$$H(X|Z_\alpha U_\alpha) = \sum_{z \leq k_\alpha} \Pr[Z = z] H(X|Z = z),$$

and so $H(X|Z_\alpha U_\alpha) \to H(X|Z)$ as $\alpha \to 0$. Therefore,

$$S(X;Y||Z) \geq I(X;Y) - I(X;Z),$$

and the symmetric bound follows in the same way.

Suppose now that X, Y, and Z have unbounded ranges. Fix $\alpha > 0$ and k_α such that

$$\Pr[X \leq k_\alpha] > 1 - \alpha \quad \text{and} \quad \Pr[Y \leq k_\alpha] > 1 - \alpha.$$

Let A_α be the indicator variable of the event that X and Y are in $\{0, \ldots, k_\alpha\}$. Furthermore, let X_α, Y_α, and Z_α denote X, Y, and Z conditioned on $A_\alpha = 1$. By the union bound we have $\Pr[A_\alpha = 1] > 1 - 2\alpha$. If we consider the case where Alice and Bob keep realizations of X and Y only if they lie in $\{0, \ldots, k_\alpha\}$, we obtain

$$S(X;Y||Z) > (1 - 2\alpha) S(X_\alpha; Y_\alpha||Z_\alpha) \geq (1 - 2\alpha)(I(X_\alpha; Y_\alpha) - I(X_\alpha; Z_\alpha)). \quad (2.1)$$

We can rewrite Inequality 2.1 as

$$S(X;Y||Z) > (1 - 2\alpha)(I(X;Y|A_\alpha = 1) - I(X;Z|A_\alpha = 1)).$$

Then it suffices to note that, when $\alpha \to 0$,

$$I(X;Y|A_\alpha = 1) \to I(X;Y) \quad (2.2)$$

and

$$I(X;Z|A_\alpha = 1) \to I(X;Z). \tag{2.3}$$

To see that 2.2 holds, write

$$I(X;Y|A_\alpha = 1) = H(X|A_\alpha = 1) + H(Y|A_\alpha = 1) - H(XY|A_\alpha = 1).$$

Then we can compute the limit for each quantity in the right-hand side. We show only that $H(X|A_\alpha = 1) \to H(X)$. The same reasoning can be applied in an analogous manner to the other quantities. We have

$$H(X|A_\alpha = 1) \leq \frac{H(X)}{1-2\alpha} \to H(X)$$

and

$$\begin{aligned} H(X|A_\alpha = 1) &= \sum_{x \leq k_\alpha} P_{X|A_\alpha=1}(x) \log \frac{1}{P_{X|A_\alpha=1}(x)} = \\ &= \log(\Pr[A_\alpha = 1]) + \sum_{x \leq k_\alpha} P_{X|A_\alpha=1}(x) \log \frac{1}{\Pr[X = x, A_\alpha = 1]} \geq \\ &\geq \log(1-2\alpha) + \sum_{x \leq k_\alpha} \Pr[X = x, A_\alpha = 1] \log \frac{1}{P_X(x)}. \end{aligned}$$

It holds that

$$\Pr[X = x, A_\alpha = 1] = \sum_{y \leq k_\alpha} P_{XY}(x,y) \to P_X(x)$$

as $\alpha \to 0$. Thus,

$$\sum_{x \leq k_\alpha} \Pr[X = x, A_\alpha = 1] \log \frac{1}{P_X(x)} \to H(X),$$

and so $H(X|A_\alpha = 1) \to H(X)$ too. The limit in 2.3 follows by the same reasoning.

Combining 2.1, 2.2, and 2.3 we conclude that

$$S(X;Y||Z) \geq I(X;Y) - I(X;Z).$$

The symmetric lower bound follows analogously. □

There exist some details which are ignored when proving the upper bound of Lemma 2.15, but which we feel deserve a careful discussion. The proof in question (see the proof of [22, Theorems 1 and 2]) considers Alice's and Bob's local randomness R_A and R_B as part of X^N and Y^N, respectively. In order for the proof to go through, it is necessary to enforce that $H(R_A)$ and $H(R_B)$ are finite. This is one of the reasons why such a condition is added

in the definition of the secret-key rate. Another reason is concerned with practicality, since in a sensible application a party would not extract infinite entropy from a randomness source.

In general, the bounds of Lemma 2.15 are not tight, and they do not necessarily tell us whether $S(X;Y||Z) > 0$. In fact, we can have $S(X;Y||Z) > 0$ while the left-hand side is negative, and $S(X;Y||Z) = 0$ while the right-hand side is positive.

As we have discussed before, Ahlswede and Csiszár [2] studied secret-key agreement in the source model when only communication from Alice to Bob is allowed.

Definition 2.16 ([2, Definition 2.1]) *Given a discrete probability distribution P_{XYZ}, the* one-way secret-key rate of X and Y given Z, *denoted by $S_{\mathrm{ow}}(X;Y||Z)$, is defined like $S(X;Y||Z)$, except that we only allow protocols where $C_i = \perp$ with probability 1 for all even i, i.e. we consider only communication protocols with one-way communication from Alice to Bob.*

It follows immediately from the definition that $S_{\mathrm{ow}}(X;Y||Z) \leq S(X;Y||Z)$. There exists a single-letter characterization of the one-way secret-key rate, shown by Ahlswede and Csiszár.

Lemma 2.17 ([2, Theorem 1]) *For every finite distribution P_{XYZ} we have*

$$S_{\mathrm{ow}}(X;Y||Z) = \sup_{U \to T \to X \to YZ} [I(T;Y|U) - I(T;Z|U)],$$

where the supremum is taken over all finite random variables U and T such that $U \to T \to X \to YZ$ holds. Moreover, the supremum is actually a maximum over finite random variables U and T with ranges $\mathcal{U}$ and $\mathcal{T}$ such that $|\mathcal{U}|, |\mathcal{T}| \leq |\mathcal{X}|$.

2.5 Entanglement distillation

In this section we introduce some results related to entanglement distillation, specialized for our needs in Section 4.4. We assume familiarity with some very basic concepts of quantum computation, which can be found in [29, Chapter 2]. The following exposition is partly based on [30], [12, Section 2.3], and [29, Section 12.5.3]

Suppose Alice, Bob, and Eve have access to their own subsystem of a tripartite pure quantum state $|\varphi\rangle$. The goal of entanglement distillation is for Alice and Bob to distill some maximally entangled bipartite quantum states, that are completely disentangled from Eve's environment, from several copies of $|\varphi\rangle$, solely through local (quantum) operations and classical communication between themselves. More precisely, Alice and Bob want to obtain states of the form

$$|\psi^-\rangle := \frac{|01\rangle - |10\rangle}{\sqrt{2}},$$

which are not entangled with Eve's environment.

A concept which is intimately connected with entanglement distillation is the *trace over Eve's environment* of $|\varphi\rangle$. For our purpose of analyzing the results of [30], $|\varphi\rangle$ will be of the form

$$|\varphi\rangle = |\varphi_{XYZ}\rangle := \sum_{x,y,z} \sqrt{P_{XYZ}(x,y,z)}|xyz\rangle,$$

for some finite distribution P_{XYZ}. Then, in our case, the trace over Eve's environment of $|\varphi\rangle$, which we denote by ρ_{XY}, is the density matrix

$$\rho_{XY} := \sum_{z} \sum_{x,x',y,y'} \sqrt{P_{XYZ}(x,y,z)P_{XYZ}(x',y',z)}|xy\rangle\langle x'y'|.$$

Note that ρ_{XY} can be represented by a $|\mathcal{X}||\mathcal{Y}| \times |\mathcal{X}||\mathcal{Y}|$ matrix, where $\mathcal{X}$ and $\mathcal{Y}$ are the ranges of X and Y, respectively. The partial transpose with respect to Bob of $|xy\rangle\langle x'y'|$, denoted by $(|xy\rangle\langle x'y'|)^{\top_B}$, is the linear operator satisfying

$$(|xy\rangle\langle x'y'|)^{\top_B} := |xy'\rangle\langle x'y|.$$

We then have the following definition.

Definition 2.18 ([29, Adaptation of definition in Section 12.5.3]) *The* entanglement distillation rate of a tripartite quantum state $|\varphi\rangle$, *denoted by* $D(|\varphi\rangle)$, *is the supremum of all rates at which Alice and Bob can create* $|\psi^-\rangle$ *states from several instances of* ρ_{XY} *through local operations and classical communication.*

The following theorem is a specialization of several results regarding entanglement distillation and separability criteria, most notably the so-called PPT-condition of Horodecki and Peres [17][31].

Theorem 2.19 (First statement follows from [31], equivalence proved in [18]) *Given a finite distribution* P_{XYZ}, *we have that* $D(|\varphi_{XYZ}\rangle) = 0$ *if* $\rho_{XY}^{\top_B}$ *has no negative eigenvalues. Moreover, if* $|\mathcal{X}| = |\mathcal{Y}| = 2$, *this is an equivalence.*

Chapter 3

A concrete challenge – the satellite setting

In this chapter we study one of the most basic settings for information-theoretic secret-key agreement in the source model – the satellite setting, which was introduced by Maurer [21][22]. While there exist nice results in this setting, there is still much progress to be done. For example, we know when the secret-key rate is positive in the satellite setting, but a formula for computing the exact secret-key rate (or even a decent approximation!) is still unknown. Studying and improving known protocols for secret-key agreement in this setting helps us understand the secret-key rate better and get closer to an exact characterization.

We introduce the satellite setting and advantage distillation in Section 3.1. In Sections 3.2 and 3.3 we provide a detailed analysis of two protocols for advantage distillation which will influence the rest of the chapter. In Section 3.4 we provide an example of a setting where the parity-check protocol with block-length 3 and 2 rounds is superior to the parity-check protocol with block-length 2 and *any* number of rounds. In Section 3.5, motivated by Section 3.4, we present a natural modification of the parity-check protocol with block-length $k > 2$ which strictly improves on the original parity-check protocol. Finally, in Section 3.6, we translate the satellite setting to practice. We determine the asymptotic behavior of the secret-key rate *per time unit* when the satellite is allowed to choose the error probabilities subject to a constraint on the ratio of the channel capacities.

3.1 The satellite setting and advantage distillation

Consider the following setting with three parties Alice, Bob, and Eve. Alice and Bob share an authenticated noiseless channel which Eve can listen

to. Furthermore, Alice, Bob, and Eve receive i.i.d. realizations of random variables X, Y, and Z, respectively, which are generated as follows:

1. A bit R is sampled uniformly at random from $\{0,1\}$;
2. The bit R is sent to Alice through a $\mathrm{BSC}_{\varepsilon_A}$, to Bob through a $\mathrm{BSC}_{\varepsilon_B}$, and to Eve through a $\mathrm{BSC}_{\varepsilon_E}$. The random variables X, Y, and Z correspond to the output of each of these channels, respectively.

We can assume without loss of generality that $\varepsilon_A, \varepsilon_B, \varepsilon_E \in [0, 1/2]$, since otherwise Alice, Bob, or Eve can just flip the bit they receive.

This is the so-called satellite setting, which was introduced and first studied by Maurer [21][22]. Its name derives from the fact that a possible physical realization of such a setting would be to have a satellite in space sample the bit R, and then broadcast it to Earth [22]. Then, the reliability with which Alice, Bob, and Eve receive R depends on the quality of their respective antennae. The better the antenna, the smaller the error probability when receiving R.

The significance of this setting comes from the fact that it is a simple example where one-way secret-key agreement is impossible under some conditions, while secret-key agreement with interaction is possible except in some trivial cases. More precisely, the following theorem, proved by Maurer [22], holds.

Theorem 3.1 ([22, Theorem 5], rewritten) *In the satellite setting, it holds that* $S_{\mathrm{ow}}(X;Y||Z) = 0$ *if* $\varepsilon_E \leq \varepsilon_A$ *and* $\varepsilon_E \leq \varepsilon_B$. *On the other hand, we have* $S(X;Y||Z) > 0$ *if and only if* $\varepsilon_E > 0$, $\varepsilon_A < 1/2$, *and* $\varepsilon_B < 1/2$.

Recall the following lower bound on the secret-key rate from Lemma 2.15,

$$S(X;Y||Z) \geq \max(I(X;Y) - I(X;Z), I(X;Y) - I(Y;Z)). \tag{3.1}$$

If $\varepsilon_E \leq \varepsilon_A$ and $\varepsilon_E \leq \varepsilon_B$, the lower bound in Inequality 3.1 vanishes. Therefore, in this situation it is not a priori obvious how to prove that the secret-key rate is positive.

Suppose $\varepsilon_A \leq \varepsilon_B$, $\varepsilon_E \leq \varepsilon_A$, and $\varepsilon_E \leq \varepsilon_B$. A possible way of achieving a positive secret-key rate in this situation is to have Alice and Bob run an interactive protocol such that, at the end, the three parties are in a situation where Bob has more knowledge about Alice's information than Eve does. This discrepancy between advantages can then be exploited to generate a secret-key. Such a protocol is called a *protocol for advantage distillation*. More precisely, starting with N realizations X^N and Y^N each, for fixed N large enough, Alice and Bob agree, through public discussion, on random variables $\hat{X}$ and $\hat{Y}$, respectively, such that, if $\hat{Z}$ denotes Eve's total information about $\hat{X}$ and $\hat{Y}$,

$$I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}) > 0,$$

i.e. such that Bob has more information about $\hat{X}$ (via $\hat{Y}$) than Eve does (via $\hat{Z}$). It then follows that

$$S(X;Y||Z) \geq \frac{S(\hat{X};\hat{Y}||\hat{Z})}{N} \geq \frac{I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z})}{N} > 0.$$

While a protocol for advantage distillation which yields a positive lower-bound for the secret-key rate is enough to show that secret-key agreement is possible in the given setting, it is also of interest to find increasingly more efficient protocols, since they yield improved lower bounds for the secret-key rate. Another goal is to find explicit protocols, and not merely prove that they exist.

3.2 The repeater-code protocol

In this section we present the repeater-code protocol for advantage distillation, which was used by Maurer [22] and Maurer and Wolf [25] to show that the secret-key rate in the satellite setting is positive even if Eve has a lower error probability than Alice and Bob.

Fix an integer N, and suppose Alice, Bob, and Eve have access to N i.i.d. realizations X^N, Y^N, and Z^N, respectively. The repeater-code protocol works as follows:

1. Alice samples a bit $C_A \in \{0,1\}$ uniformly at random, and sends $C' = C_A \oplus X^N$ to Bob over the authenticated channel.
2. Bob accepts publicly if the bits of $C' \oplus Y^N$ all coincide (let C_B be such that $C' \oplus Y^N = (C_B, \ldots, C_B)$). Otherwise, Bob rejects publicly.
3. If Bob accepts, Alice sets $\hat{X} = C_A$, and Bob sets $\hat{Y} = C_B$. If Bob rejects, $\hat{X}$ and $\hat{Y}$ are both set to $\perp$.

Another way of looking at this protocol, which gives it its name, is that Alice picks a codeword of a repeater-code uniformly at random, and sends each bit to Bob through a binary symmetric channel with error probability $\Pr[X \neq Y]$. Bob accepts if he sees a codeword of the repeater-code. Note that, given that Bob accepts, he decodes the codeword correctly if $X_i = Y_i$ for all i, and incorrectly if $X_i \neq Y_i$ for all i.

Intuitively, as discussed in [22], the repeater-code protocol should work for large enough N because Bob only accepts when he is very confident about C_A, as the probability that $X_i = Y_i$ for all i is much larger than the probability that $X_i \neq Y_i$ for all i. Moreover, the event of Bob accepting is independent of how well Z^N estimates X^N. Therefore, while Eve has more confidence about C_A when averaging over all possible outcomes, it makes sense that this does

not hold if we average only over outcomes where Bob accepts, for N large enough.

In the remainder of this section, we will take a closer look at the technical details of the repeater-code protocol. Consider $\hat{X}$ and $\hat{Y}$ as defined above, fix a number of realizations N, and let A be the indicator variable of the event that Bob accepts a run of the repeater-code protocol with N realizations, i.e. $A := 1_{\{\hat{X} \neq \perp\}}$. We define $\varepsilon := \Pr[X \neq Y] = \varepsilon_A(1-\varepsilon_B) + (1-\varepsilon_A)\varepsilon_B$. Furthermore, we denote the probability that Bob accepts a run of the repeater-code protocol with N realizations as $p_{a,N}$, which satisfies

$$p_{a,N} = \Pr[A = 1] = \varepsilon^N + (1-\varepsilon)^N.$$

Many quantities in this discussion will depend on N, ε_A, ε_B, and ε_E. For simplicity, we do not usually make this dependence explicit in their names, although they can be easily discerned via their definitions. Such parameters will remain fixed throughout the discussion.

We first carefully derive the formula for $I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z})$ determined in [22], where $\hat{Z} = [Z^N, C_A \oplus X^N, A]$ is Eve's total information, for various reasons: First, it helps us understand the inner workings of the protocol, and in particular it helps us determine the optimal strategy for Eve, which Maurer and Wolf [25] use to prove that the protocol indeed achieves advantage distillation. Second, this derivation (and the optimal strategy) will be a basis for the analysis of the protocols in Sections 3.3 and 3.5.

It came to our attention in the final stages of this thesis that a similar derivation can already be found in [3, Propositions 4.3 and 4.4]. We opt to still include our derivation here. This is mainly for completeness of our exposition, since it is a well-known result with a natural proof, and because, as already mentioned, it will be an important building block in the following sections.

First, note that

$$I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}) = I(\hat{X};\hat{Y}|A) - I(\hat{X};\hat{Z}|A), \tag{3.2}$$

since

$$I(\hat{X};\hat{Y}|A) = H(\hat{X}A) + H(\hat{Y}A) - H(\hat{X}\hat{Y}A) - H(A) = I(\hat{X};\hat{Y}) - H(A),$$

as A is completely determined from one of $\hat{X}$, $\hat{Y}$, and $\hat{Z}$. Analogously, $I(\hat{X};\hat{Z}|A) = I(\hat{X};\hat{Z}) - H(A)$. Furthermore,

$$I(\hat{X};\hat{Y}|A=0) = I(\hat{X};\hat{Z}|A=0) = 0,$$

since, conditioned on $A = 0$ (Bob rejects), $\hat{X} = \hat{Y} = \perp$. Therefore,

$$I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}) = p_{a,N}(I(\hat{X};\hat{Y}|A=1) - I(\hat{X};\hat{Z}|A=1)).$$

For ease of notation, we set $I_B(N) := I(\hat{X};\hat{Y}|A=1)$, for $\hat{X}$ and $\hat{Y}$ obtained from running the repeater-code protocol with N realizations, and analogously we set $I_E(N) := I(\hat{X};\hat{Z}|A=1)$.

Note that $I_B(N) = H(\hat{X}|A=1) - H(\hat{X}|\hat{Y},A=1)$. To compute $H(\hat{X}|A=1)$, note that $H(\hat{X}|A=1) = H(C_A|A=1) = 1$. Furthermore, note that

$$\beta_N := \Pr[\hat{X} \neq \hat{Y}|A=1] = \frac{\varepsilon^N}{p_{a,N}},$$

and so $H(\hat{X}|\hat{Y},A=1) = h(\beta_N)$. Therefore,

$$I_B(N) = 1 - h(\beta_N).$$

It remains to find an expression for $I_E(N)$. We have

$$I_E(N) = H(\hat{X}|A=1) - H(\hat{X}|\hat{Z},A=1) = 1 - H(\hat{X}|\hat{Z},A=1).$$

Therefore, we only need to compute $H(\hat{X}|\hat{Z},A=1)$. Given that Bob accepts, Eve's total information about $\hat{X}$ consists of $\hat{Z} = [Z^N, C_A \oplus X^N, A=1]$. The next lemma states that Eve can compute $C_A \oplus X^N \oplus Z^N$ and then discard Z^N and $C_A \oplus X^N$ without any loss of information about $\hat{X}$.

Lemma 3.2 *We have*

$$H(\hat{X}|Z^N, C_A \oplus X^N, A=1) = H(\hat{X}|C_A \oplus X^N \oplus Z^N, A=1).$$

Proof This proof is inspired by part of the proof of [22, Proposition 1]. Let $U := C_A \oplus X^N \oplus Z^N$. Note that

$$\begin{aligned} H(\hat{X}|C_A \oplus X^N, Z^N, A=1) &= \\ &= H(C_A, C_A \oplus X^N|U, A=1) - H(C_A \oplus X^N|U, A=1) = \\ &= H(C_A|U,A=1) + H(C_A \oplus X^N|C_A,U,A=1) - H(C_A \oplus X^N|U,A=1), \end{aligned}$$

where the first equality follows from the chain rule for entropy and from the fact that $[C_A \oplus X^N, U]$ completely determines $[C_A \oplus X^N, Z^N]$ and vice-versa. The second equality follows from the chain rule for entropy.

Therefore, it suffices to show that

$$H(C_A \oplus X^N|C_A,U,A=1) = H(C_A \oplus X^N|U,A=1).$$

Since $C_A \oplus X^N$ is uniformly distributed and independent of $[C_A, U]$ given $A=1$, it follows that

$$H(C_A \oplus X^N|C_A,U,A=1) = N.$$

Finally, note that $H(C_A \oplus X^N|U,A=1) \geq H(C_A \oplus X^N|C_A,U,A=1)$, and thus $H(C_A \oplus X^N|U,A=1) = N$ also. □

By Lemma 3.2, we have

$$H(\hat{X}|\hat{Z}, A=1) = H(\hat{X}|U, A=1) = \\ = \sum_{u\in\{0,1\}^N} \Pr[U=u|A=1] H(\hat{X}|U=u, A=1), \quad (3.3)$$

where $U := C_A \oplus X^N \oplus Z^N$. Thus, we must compute $\Pr[U=u|A=1]$ and $\Pr[\hat{X}=0|U=u, A=1]$ for every $u \in \{0,1\}^N$, since

$$H(\hat{X}|U=u, A=1) = h(\Pr[\hat{X}=0|U=u, A=1]). \quad (3.4)$$

Note that

$$\Pr[\hat{X}=0|U=u, A=1] = \frac{\Pr[\hat{X}=0, X^N \oplus Z^N = u, A=1]}{\Pr[U=u, A=1]}. \quad (3.5)$$

In view of this, let p_w be the probability of the event that $X^N \oplus Z^N$ is a particular word of weight w and Bob accepts. This probability is the same for all words of weight w because $A=1$ holds if and only if $X_i = Y_i$ for all i or $X_i \neq Y_i$ for all i, and in both cases the bits of $X^N \oplus Z^N$ are independent and identically distributed. In order to find a formula for p_w, let α_{rs} be the probability that $X \oplus Y = r$ and $X \oplus Z = s$, for $r, s \in \{0,1\}$. Then

$$\begin{aligned}
\alpha_{00} &= \varepsilon_A\varepsilon_B\varepsilon_E + (1-\varepsilon_A)(1-\varepsilon_B)(1-\varepsilon_E) \\
\alpha_{01} &= \varepsilon_A\varepsilon_B(1-\varepsilon_E) + (1-\varepsilon_A)(1-\varepsilon_B)\varepsilon_E \\
\alpha_{10} &= \varepsilon_A(1-\varepsilon_B)\varepsilon_E + (1-\varepsilon_A)\varepsilon_B(1-\varepsilon_E) \\
\alpha_{11} &= \varepsilon_A(1-\varepsilon_B)(1-\varepsilon_E) + (1-\varepsilon_A)\varepsilon_B\varepsilon_E.
\end{aligned}$$

As a result, we have

$$p_w = \alpha_{00}^{N-w}\alpha_{01}^{w} + \alpha_{10}^{N-w}\alpha_{11}^{w}.$$

Thus,

$$\Pr[\hat{X}=0|U=u, A=1] = \frac{p_w}{p_w + p_{N-w}} \quad (3.6)$$

if u has weight w. This follows from Equation 3.5, the fact that

$$\Pr[\hat{X}=0, X^N \oplus Z^N = u, A=1] = \frac{1}{2}\Pr[X^N \oplus Z^N = u, A=1] = \frac{p_w}{2},$$

and finally that

$$\Pr[U=u, A=1] = \\ = \Pr[\hat{X}=0, X^N \oplus Z^N = u, A=1] + \Pr[\hat{X}=1, X^N \oplus Z^N = 1 \oplus u, A=1] = \\ = \frac{1}{2}(p_w + p_{N-w}). \quad (3.7)$$

Combining 3.3, 3.4, 3.6, and 3.7 yields

$$I_E(N) = 1 - \sum_{u \in \{0,1\}^N} \frac{p_{w(u)} + p_{N-w(u)}}{2p_{a,N}} \cdot h\left(\frac{p_{w(u)}}{p_{w(u)} + p_{N-w(u)}}\right), \tag{3.8}$$

where $w(u)$ is the weight of u. Since there is a one-to-one correspondence between words of weight w and words of weight $N - w$, and since $h(p) = h(1-p)$, it follows that

$$I_E(N) = 1 - \sum_{u \in \{0,1\}^N} \frac{p_{w(u)}}{p_{a,N}} \cdot h\left(\frac{p_{w(u)}}{p_{w(u)} + p_{N-w(u)}}\right).$$

Clustering words with the same weight together finally leads to

$$I_E(N) = 1 - \sum_{w=0}^{N} \binom{N}{w} \frac{p_w}{p_{a,N}} \cdot h\left(\frac{p_w}{p_w + p_{N-w}}\right).$$

This train of thought, besides yielding a tractable formula for both $I_B(N)$ and $I_E(N)$, also gives some additional insight about the optimal guessing strategy for Eve. Lemma 3.2 and Equation 3.6 tell us that, conditioned on $A = 1$, the maximum likelihood estimate of $\hat{X}$ is 0 whenever the weight of $C_A \oplus X^L \oplus Z^L$ is smaller than $N/2$, and 1 when the weight is larger than $N/2$, since $p_w > p_{N-w}$ if $w < N/2$, as $\varepsilon_A < 1/2$ and $\varepsilon_B < 1/2$. Therefore, Eve's optimal guessing strategy for $\hat{X}$ given that $A = 1$ in terms of average error probability, which we denote by γ_N, is to take the majority of $C_A \oplus X^N \oplus Z^N$ (with arbitrary tie-breaking).

Maurer and Wolf [25] use these connections to the optimal guessing strategy to prove that, for large enough N (depending on ε_A, ε_B, and ε_E), we have $I_B(N) - I_E(N) > 0$.

Theorem 3.3 ([25, Consequence of Lemmas 3 and 4]) *For any $\varepsilon_A < 1/2$, $\varepsilon_B < 1/2$, and $\varepsilon_E > 0$, we have $I_B(N) - I_E(N) > 0$ for large enough even N.*

We do not present a full proof of Theorem 3.3 here, because the proof of the main result of Section 3.5 already uses similar techniques, and so we believe it is preferable to showcase only the general, intuitive idea of the proof in [25] first.

In order to prove Theorem 3.3, we first find constants $b < c$ such that Eve's optimal average error probability γ_N and Bob's error probability β_N satisfy

$$\gamma_N \geq c^N > b^N \geq \beta_N \tag{3.9}$$

for large enough even N. It turns out upper bounding β_N is straightforward, but the lower bound for γ_N requires a bit more care. Such a lower bound

can be obtained by focusing on the cases where U has weight $N/2$, in which case taking the majority does not help. Finally, we make use of the following lemma proved by Maurer and Wolf [25]. We reproduce its proof here for two reasons. First, we need it to show that its hypotheses are too strong, and in fact a weaker separation between the error probabilities is enough. Second, it is going to be used again in Section 3.5.

Lemma 3.4 ([25, Lemma 4]) *Suppose Alice and Bob run an advantage distillation protocol with N realizations and, with positive probability, end up with bits $\hat{X}$ and $\hat{Y}$ such that Bob's error probability between $\hat{X}$ and $\hat{Y}$, denoted β_N, and Eve's optimal average error probability between her total information $\hat{Z}$ and $\hat{X}$, denoted γ_N, satisfy the inequalities in 3.9 for large enough N. Let E be the event that Alice and Bob end up with such bits. Then*

$$I(\hat{X};\hat{Y}|E) - I(\hat{X};\hat{Z}|E) > 0$$

holds for large enough N.

Proof (Reproduction of proof in [25]) Note that

$$I(\hat{X};\hat{Y}|E) - I(\hat{X};\hat{Z}|E) = H(\hat{X}|\hat{Z},E) - H(\hat{X}|\hat{Y},E).$$

We can now use the inequalities in 3.9 to bound both terms. First, we have

$$H(\hat{X}|\hat{Y},E) = h(\beta_N) \leq h(b^N),$$

since h is increasing in $[0,1/2]$. Moreover,

$$h(b^N) = b^N \log(1/b^N) + (1-b^N)\log(1/(1-b^N)) \leq 2b^N \log(1/b^N)$$

for large enough N, since $p\log(1/p) \geq (1-p)\log(1/(1-p))$ for $p \leq 1/2$. This is because

$$f(p) := p\log(1/p) - (1-p)\log(1/(1-p))$$

is a concave function in p (it suffices to compute its second derivative), and $f(0) = f(1/2) = 0$. Now it suffices to note that

$$2b^N \log(1/b^N) = b^N(2N\log(1/b)) < c^N$$

for large enough N.

It remains to bound $H(\hat{X}|\hat{Z},E)$. We have

$$H(\hat{X}|\hat{Z},E) = \mathrm{E}_{\hat{Z}}[h(p_{e,N}(\hat{Z}))|E],$$

where $p_{e,N}(\hat{z})$ is Eve's error probability given that $\hat{Z} = \hat{z}$ and E both hold. Then, since Eve's guessing strategy is optimal, $p_{e,N}(\hat{z}) \leq 1/2$ always holds, and so

$$\mathrm{E}_{\hat{Z}}[h(p_{e,N}(\hat{Z}))|E] \geq \mathrm{E}_{\hat{Z}}[p_{e,N}(\hat{Z})|E] = \gamma_N \geq c^N,$$

as $h(p) \geq p$ for $p \leq 1/2$, since $\log(1/p) \geq 1$ for $p \leq 1/2$. Therefore,

$$I(\hat{X};\hat{Y}|E) - I(\hat{X};\hat{Z}|E) > c^N - c^N = 0$$

for large enough N, which concludes the proof. □

Note that Lemma 3.4 can be applied in a more general setting than just the repeater-code protocol, as long as there is an event with positive probability (which the parties can check that it occurred) under which we can find a large enough asymptotic separation between Bob's and Eve's error probabilities.

If we analyze the previous proof, it is also clear that we do not require an exponential separation between the error probabilities. In fact, a superlinear separation suffices! Suppose we find constants $c > 0$ and $\delta > 0$ such that

$$\gamma_N \geq c^N > \frac{c^N}{N^{1+\delta}} \geq \beta_N$$

for large enough N. Then we can proceed exactly like in the proof to obtain

$$h(\beta_N) \leq h\left(\frac{c^N}{N^{1+\delta}}\right) \leq 2\frac{c^N}{N^{1+\delta}} \log\left(\frac{N^{1+\delta}}{c^N}\right) \tag{3.10}$$

for large enough N. Expanding the right side of Inequality 3.10, we get

$$2\frac{c^N}{N^{1+\delta}} \log\left(\frac{N^{1+\delta}}{c^N}\right) = 2\frac{c^N}{N^{1+\delta}}((1+\delta)\log(N) + N\log(1/c)),$$

which we can bound further as

$$2\frac{c^N}{N^{1+\delta}}((1+\delta)\log(N) + N\log(1/c)) \leq c^N \frac{4(1+\delta)\log(1/c)}{N^\delta} < c^N$$

for large enough N, since

$$N\log(1/c) > (1+\delta)\log(N)$$

and

$$N^\delta > 4(1+\delta)\log(1/c)$$

for large enough N. The desired result then follows immediately.

We do not make use of this observation in this chapter, but we think it deserves mention, and might make the analysis of more complex protocols for advantage distillation much more approachable (for example, potential extensions of the protocol proposed in Section 3.5).

As a final observation, note that Theorem 3.3 tells us we can find a large enough *even* N such that running the repeater-code protocol with N realizations yields an advantage to Bob against Eve. While this is enough for the

repeater-code protocol and the goals of this section, it will be required that Theorem 3.3 actually holds for *all* large enough N in Section 3.3, since we will have cases where the proposed protocol "works like" the repeater-code protocol with an odd number of realizations. Such an extension is relatively straightforward. If one follows the proof of Theorem 3.3, it is sufficient to note that, for constants $b < c$ as determined in the proof,

$$\gamma_{N+1} \geq \binom{N+1}{N/2+1} \frac{\alpha_{00}^{N/2-1}\alpha_{01}^{N/2+1}}{\varepsilon^{N+1}+(1-\varepsilon)^{N+1}} \geq$$
$$\geq \frac{\alpha_{01}}{(1-\varepsilon)\alpha_{00}} \binom{N}{N/2} \frac{\alpha_{00}^{N/2}\alpha_{01}^{N/2}}{\varepsilon^{N}+(1-\varepsilon)^{N}} \geq d \cdot c^{N+1},$$

for d constant and large enough even N. Then,

$$d \cdot c^{N+1} \geq (c-\delta)^{N+1},$$

for some $\delta \to 0$ as N grows. Thus we can pick $c_2 := c - \delta > b$ for N large enough, since $c > b$ for N large enough, and thus $\gamma_N \geq c_2^N > b^N \geq \beta_N$ for *all* N large enough.

The repeater-code protocol we analyzed in this section meets its goal of showing that secret-key agreement is possible in the satellite setting under very general conditions, but only implies the lower bound

$$\frac{p_{a,N}}{N}(I_B(N) - I_E(N)),$$

where $p_{a,N}$ decreases roughly like $(1-\varepsilon)^N$. Furthermore, the protocol requires local randomness on Alice's part. A natural question arises: is there a more efficient protocol (in terms of the resulting secret-key rate) which preferably does not use local randomness? In Section 3.3 we present the parity-check protocol, which answers this question positively.

3.3 The parity-check protocol

In this section we present the parity-check protocol, first introduced by Maurer [21] and later studied by Gander and Maurer [11]. This protocol yields more rate than the repeater-code protocol presented in Section 3.2, and also has the advantage of not requiring any local randomness.

We showcase a detailed technical analysis of the protocol. Furthermore, we opt to analyze the general case of the protocol, where there may be k parity-checks, for some $k \geq 2$, instead of only two as in [11], since this general analysis will be useful in Sections 3.4 and 3.5.

Consider the usual satellite setting of Section 3.1, and suppose Alice, Bob, and Eve have error probabilities $\varepsilon_A < 1/2$, $\varepsilon_B < 1/2$, and $\varepsilon_E > 0$, and

furthermore have access to N i.i.d. realizations of X, Y, and Z, respectively. The parity-check protocol with block-length k, for $k \geq 2$, and ℓ rounds works as follows:

1. Alice and Bob have initially empty strings S and $\hat{S}$, respectively.
2. Alice and Bob divide X^N and Y^N into tuples $(X_{ki+1}, \ldots, X_{ki+k})$ and $(Y_{ki+1}, \ldots, Y_{ki+k})$, respectively, for $i = 0, \ldots, \lfloor N/k \rfloor$.
3. For each i, Alice sends $X_{ki+1} \oplus X_{ki+j}$, for $j = 2, \ldots, k$, to Bob via the noiseless authenticated channel.
4. Bob computes the tuple

$$(Y_{ki+1}, X_{ki+1} \oplus X_{ki+2} \oplus Y_{ki+2}, \ldots, X_{ki+1} \oplus X_{ki+k} \oplus Y_{ki+k}),$$

   and accepts publicly if and only if all entries have the same value. Otherwise, Bob rejects publicly.
5. If Bob accepts, Alice adds X_{ki+1} to her string S, and Bob adds Y_{ki+1} to his string $\hat{S}$, and they discard the remaining bits. If Bob rejects, Alice and Bob discard all the relevant bits.
6. If $\ell = 1$, then Alice and Bob stop the protocol. Alice sets $\hat{X} = S$ and Bob sets $\hat{Y} = \hat{S}$.
7. If $\ell > 1$ and $|S| \geq k^{\ell-1}$, Alice and Bob run the parity-check protocol with block-length k and $\ell - 1$ rounds on the strings S and $\hat{S}$. Otherwise, if $|S| < k^{\ell-1}$, then Alice and Bob abort the protocol, and set $\hat{X} = \perp$ and $\hat{Y} = \perp$, respectively.

Intuitively, the parity-check protocol works by applying the repeater-code protocol to small substrings over many rounds, and discarding the unsuccessful runs. In a sense, this yields an "adaptive" version of the repeater-code protocol in the following sense: in the original repeater-code protocol, a bad bit (i.e. $X_i \neq Y_i$) in the middle of many good bits (i.e. $X_i = Y_i$) makes the protocol fail, while in the parity-check protocol these bad bits can be successfully removed (if there are not too many), and the protocol is able to carry on with a smaller string.

We now proceed to the analysis of the protocol. Suppose that Alice and Bob run the parity-check protocol with block-length k and ℓ rounds, and end up with $\hat{X}$ and $\hat{Y}$. Note that $\hat{X}$ and $\hat{Y}$ may consist of several bits. An important observation is that the bits of $\hat{X}$ are independent and identically distributed. This is because every bit of $\hat{X}$ depends on $L := k^\ell$ bits of X^N, which are all independent and identically distributed, and because no two bits of $\hat{X}$ depend on the same bit of X^N. Furthermore, the same holds for the bits of $\hat{Y}$. Finally, Eve's total information $\hat{Z}_i$ about a bit $\hat{X}_i$ of $\hat{X}$ which depends on bits $X_{i_1}, \ldots, X_{i_L}$ of X^N, for some sequence (i_j), consists of $Z_{i_1}, \ldots, Z_{i_L}$, of

the parities $\hat{X}_i \oplus X_{i_j}$ for $j = 1, \ldots, L$ (since the parities Eve actually observes during a run of the protocol can be reconstructed from these parities, and vice-versa), and the fact that $\hat{X}_i$ is part of the final string $\hat{X}$. Therefore, it follows that all the $\hat{Z}_i$ are also independent and identically distributed.

Just like in Section 3.2, we want to study the behavior of $I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z})$ as a function of N. Let $T_{k,\ell}$ be the random variable corresponding to the length of $\hat{X}$ after Alice and Bob run the parity-check protocol with block-length k and ℓ rounds. Then, as in Equation 3.2 in Section 3.2, we have

$$I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}) = I(\hat{X};\hat{Y}|T_{k,\ell}) - I(\hat{X};\hat{Z}|T_{k,\ell}),$$

since $T_{k,\ell}$ is fully determined from one of $\hat{X}$, $\hat{Y}$, and $\hat{Z}$. Noting that $\hat{X}$ has length at most N/L, we have

$$I(\hat{X};\hat{Y}|T_{k,\ell}) = \sum_{t=0}^{N/L} \Pr[T_{k,\ell} = t] I(\hat{X};\hat{Y}|T_{k,\ell} = t)$$

and

$$I(\hat{X};\hat{Z}|T_{k,\ell}) = \sum_{t=0}^{N/L} \Pr[T_{k,\ell} = t] I(\hat{X};\hat{Z}|T_{k,\ell} = t).$$

Since we have seen above that both the pairs $(\hat{X}_i, \hat{Y}_i)$ and $(\hat{X}_i, \hat{Z}_i)$ are i.i.d., it follows that

$$I(\hat{X};\hat{Y}|T_{k,\ell} = t) = t \cdot I(\hat{X}_1;\hat{Y}_1),$$

and

$$I(\hat{X};\hat{Z}|T_{k,\ell} = t) = t \cdot I(\hat{X}_1;\hat{Z}_1).$$

Therefore,

$$I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}) = \mathrm{E}[T_{k,\ell}](I(\hat{X}_1;\hat{Y}_1) - I(\hat{X}_1;\hat{Z}_1)).$$

We can then turn our attention to a single bit $\hat{X}_1$, the corresponding prediction $\hat{Y}_1$ by Bob, and Eve's total information about $\hat{X}_1$, $\hat{Z}_1$. We can assume without loss of generality that $\hat{X}_1$ depends on the L bits $X^L := (X_1, \ldots, X_L)$, i.e. that $\hat{X}_1$ is the result of running the parity-check protocol on the first L bits of X^N. In this case, we have $\hat{X}_1 = X_1$, and X_i is still present in round $2 \leq r \leq \ell$ of the parity-check protocol if and only if

$$i - 1 = 0 \mod k^{r-1}.$$

Therefore, in the first round we have $X_1, \ldots, X_L$, in the second round we have $X_1, X_{k+1}, X_{2k+1}, \ldots$, and so on. The following lemma provides evidence that the parity-check protocol behaves like the repeater-code protocol in this case.

Lemma 3.5 *Suppose Alice and Bob run the parity-check protocol with block-length k and ℓ rounds on X^L and obtain $\hat{X} = X_1$. Then either $X_i = Y_i$ for all $i = 1, \ldots, L$ if $\hat{Y} = \hat{X}$, or $X_i \neq Y_i$ for all $i = 1, \ldots, L$ if $\hat{Y} \neq \hat{X}$.*

Proof In the case where $\ell = 1$, the parity-check protocol coincides with the repeater-code protocol of length k, and so the desired result follows. If $\ell > 1$, we can apply the reasoning for $\ell = 1$ iteratively, starting with the last round. Then, we have that the result holds for the bits present in round $\ell - 1$, and applying the result again to every bit implies that the result holds for the bits of round $\ell - 2$ (this is because either $X_i = Y_i$ for all i in round $\ell - 1$, or $X_i \neq Y_i$ for all such i), and so on. □

Lemma 3.5 implies that Bob's error probability about X_1 after the parity-check protocol coincides with the one obtained from running the repeater-code protocol on X^L. In fact, let A be the indicator variable of the event that X_1 is part of $\hat{X}$ after running the parity-check protocol on X^L for ℓ rounds. Then

$$\beta_\ell := \Pr[X_1 \neq \hat{Y}_1 | A = 1] = \frac{\varepsilon^L}{\varepsilon^L + (1-\varepsilon)^L},$$

where $\varepsilon := \Pr[X \neq Y] = \varepsilon_A(1-\varepsilon_B) + (1-\varepsilon_A)\varepsilon_B$, as in Section 3.2. Therefore,

$$I(\hat{X}_1; \hat{Y}_1) = 1 - h(\beta_\ell) = I_B(L), \tag{3.11}$$

with $I_B(\cdot)$ defined as in Section 3.2.

It remains to compute $I(\hat{X}_1; \hat{Z}_1)$. The following lemma is very similar to Lemma 3.2 in Section 3.2, and the proof goes through in the same way.

Lemma 3.6 *We have*

$$H(X_1 | Z^L, X_1 \oplus X^L, A = 1) = H(X_1 | X_1 \oplus X^L \oplus Z^L, A = 1).$$

Therefore, Eve can simply compute $X_1 \oplus X^L \oplus Z^L$, and then discard $X_1 \oplus X^L$ and Z^L without losing any information about X_1. Notice that, due to Lemma 3.5 and Lemma 3.6, we are in the exact same setting as after a successful run of the repeater-code protocol. Therefore, Eve's optimal strategy is to use the majority of $X_1 \oplus X^L \oplus Z^L$ as a guess for X_1 (again with arbitrary tie-breaking), and we have

$$I(\hat{X}_1; \hat{Z}_1) = I_E(L). \tag{3.12}$$

Combining Equations 3.11 and 3.12 with Theorem 3.3 (which, as we have seen, actually holds for all large enough N), it follows that, for any suitable choice of the error probabilities, there is L large enough so that

$$I(\hat{X}_1; \hat{Y}_1) - I(\hat{X}_1; \hat{Z}_1) = I_B(L) - I_E(L) > 0.$$

Finally, we need to compute $\mathrm{E}[T_{k,\ell}]$. It turns out that finding an exact expression for $\mathrm{E}[T_{k,\ell}]$ is complex, but in our case it suffices to provide suitable lower and upper bounds, which are given in the following lemma.

Lemma 3.7 *If Alice and Bob run the parity-check protocol with block-length k and ℓ rounds on N realizations, we have*

$$-\frac{c}{N}+\frac{1}{k^\ell}\prod_{i=0}^{\ell-1}\left(\beta_i^k+(1-\beta_i)^k\right)\leq\frac{\mathrm{E}[T_{k,\ell}]}{N}\leq\frac{1}{k^\ell}\prod_{i=0}^{\ell-1}\left(\beta_i^k+(1-\beta_i)^k\right), \quad (3.13)$$

for some $c>0$ which does not depend on N.

Proof We start by noting that, at the beginning of the protocol, Alice and Bob have N realizations each. Therefore we can set $T_{k,0}=N$ with probability 1 and so $\mathrm{E}[T_{k,0}]=N$.

In order to estimate $\mathrm{E}[T_{k,1}]$, note that if N is not divisible by k we must throw at most $k-1$ realizations away. The remaining realizations are split into k-tuples, from which a single bit is selected with probability $\varepsilon^k+(1-\varepsilon)^k=\beta_0^k+(1-\beta_0)^k$. Therefore,

$$\frac{\beta_0^k+(1-\beta_0)^k}{k}N-\frac{k-1}{k}\leq\frac{\beta_0^k+(1-\beta_0)^k}{k}(N-k+1)\leq$$
$$\leq\mathrm{E}[T_{k,1}]\leq\frac{\beta_0^k+(1-\beta_0)^k}{k}N.$$

We can follow a similar reasoning for $\mathrm{E}[T_{k,\ell}]$, with $\ell\geq 2$. In fact,

$$\mathrm{E}[T_{k,\ell}]=\sum_{t=0}^{N/k^{\ell-1}}\Pr[T_{k,\ell-1}=t]\cdot\mathrm{E}[T_{k,\ell}|T_{k,\ell-1}=t].$$

As before, $\mathrm{E}[T_{k,\ell}|T_{k,\ell-1}=ki+j]=\mathrm{E}[T_{k,\ell}|T_{k,\ell-1}=ki]$ for all i and $j=1,\ldots,k-1$, since we must throw away the leftover realizations. Therefore,

$$\mathrm{E}[T_{k,\ell}]=\sum_{i=0}^{N/k^\ell}\sum_{j=0}^{k-1}\Pr[T_{k,\ell-1}=ki+j]\cdot\mathrm{E}[T_{k,\ell}|T_{k,\ell-1}=ki],$$

and, since

$$\mathrm{E}[T_{k,\ell}|T_{k,\ell-1}=ki]=\frac{\beta_{\ell-1}^k+(1-\beta_{\ell-1})^k}{k}\cdot ki,$$

we have

$$\mathrm{E}[T_{k,\ell}]=\frac{\beta_{\ell-1}^k+(1-\beta_{\ell-1})^k}{k}\sum_{i=0}^{N/k^\ell}\sum_{j=0}^{k-1}\Pr[T_{k,\ell-1}=ki+j]\cdot ki. \quad (3.14)$$

It follows, by replacing ki by $ki + j - j$ in Equation 3.14, that

$$\begin{aligned}\mathrm{E}[T_{k,\ell}] = \frac{\beta_{\ell-1}^k + (1-\beta_{\ell-1})^k}{k}\mathrm{E}[T_{k,\ell-1}] - \\ - \frac{\beta_{\ell-1}^k + (1-\beta_{\ell-1})^k}{k}\sum_{i=0}^{N/k^\ell}\sum_{j=0}^{k-1}\Pr[T_{k,\ell-1} = ki+j]\cdot j \geq \\ \geq \frac{\beta_{\ell-1}^k + (1-\beta_{\ell-1})^k}{k}\mathrm{E}[T_{k,\ell-1}] - \frac{1}{k}\sum_{j=0}^{k-1} j. \quad (3.15)\end{aligned}$$

In view of this, set $c' := \frac{1}{k}\sum_{j=0}^{k-1} j$, which is of the order of k.

From 3.15, it follows that

$$\frac{\beta_{\ell-1}^k + (1-\beta_{\ell-1})^k}{k}\cdot \mathrm{E}[T_{k,\ell-1}] - c' \leq \mathrm{E}[T_{k,\ell}] \leq \frac{\beta_{\ell-1}^k + (1-\beta_{\ell-1})^k}{k}\cdot \mathrm{E}[T_{k,\ell-1}].$$

This recurrence yields

$$-c + \frac{N}{k^\ell}\prod_{i=0}^{\ell-1}\left(\beta_i^k + (1-\beta_i)^k\right) \leq \mathrm{E}[T_{k,\ell}] \leq \frac{N}{k^\ell}\prod_{i=0}^{\ell-1}\left(\beta_i^k + (1-\beta_i)^k\right),$$

for all N and ℓ with $k^\ell \leq N$, and for some $c \leq \sum_{i=0}^{\ell-1} c'/k^i \leq 2c'$, which concludes the proof. □

Let $R_{N,k,\ell}$ and $R_{k,\ell}$ be the left hand and right hand sides of 3.13, respectively. It immediately follows that

$$R_{N,k,\ell}(I_B(L) - I_E(L)) \leq \frac{\mathrm{E}[T_{k,\ell}](I(\hat{X};\hat{Y}) - I(\hat{X};\hat{Z}))}{N} \leq R_{k,\ell}(I_B(L) - I_E(L)),$$

and so, since $R_{N,k,\ell} \to R_{k,\ell}$ as $N \to \infty$, we conclude that the parity-check protocol with block-length k and ℓ rounds yields a lower bound on the secret-key rate of

$$R_{k,\ell}(I_B(k^\ell) - I_E(k^\ell)),$$

with

$$R_{k,\ell} = \frac{1}{k^\ell}\prod_{i=0}^{\ell-1}\left(\beta_i^k + (1-\beta_i)^k\right).$$

While the protocol presented in this section is a good improvement on the repeater-code protocol, questions arise: Do we need to consider block-lengths $k > 2$, or is the $k = 2$ version already optimal, at least among all parity-check protocols with different block-lengths? Furthermore, can one modify the parity-check protocol in order to improve its efficiency?

In Section 3.4, we show that the parity-check protocol with block-length 2 is not the optimal protocol for all choices of error probabilities. This will then motivate the analysis of a modified parity-check protocol in Section 3.5, which further exploits the ideas behind the parity-check protocol in a natural way.

3.4 Block-length 2 is not always optimal for the parity-check protocol

In [11], Gander and Maurer conjectured that the secret-key rate of the parity-check protocol with $k = 2$ is close to optimal for all choices of the error probabilities in the satellite setting. In this section, we show that such a protocol is, in fact, not optimal. We give (uncountably many) examples of choices of the error probabilities for which running the parity-check protocol with $k = 3$ and 1 round is strictly better than running the parity-check protocol with $k = 2$ and *any* number of rounds.

We make use of the formulas for $I_B(\cdot)$ and $I_E(\cdot)$ derived in Section 3.2, and also use the formula for $R_{k,\ell}$ derived in Section 3.3. The results and plots displayed in this section were obtained from Wolfram Mathematica 10. We analyzed the following example.

Example 3.8 *Consider the satellite setting with $\varepsilon_A = \varepsilon_B = 0.1$. Figure 3.1 showcases the secret-key rates of different versions of the parity-check protocol as functions of ε_E, which we vary between 0 and 0.1. Note that, for ε_E roughly between 0.6 and 0.7, the parity-check protocol with $k = 3$ and $\ell = 1$ is superior to the protocols with $k = 2$. Consider the setting with $\varepsilon_A = \varepsilon_B = 0.1$ and $\varepsilon_E = 0.065$. Table 3.1 confirms the parity-check protocol with $k = 3$ and $\ell = 1$ is superior to any protocol with $k = 2$.*

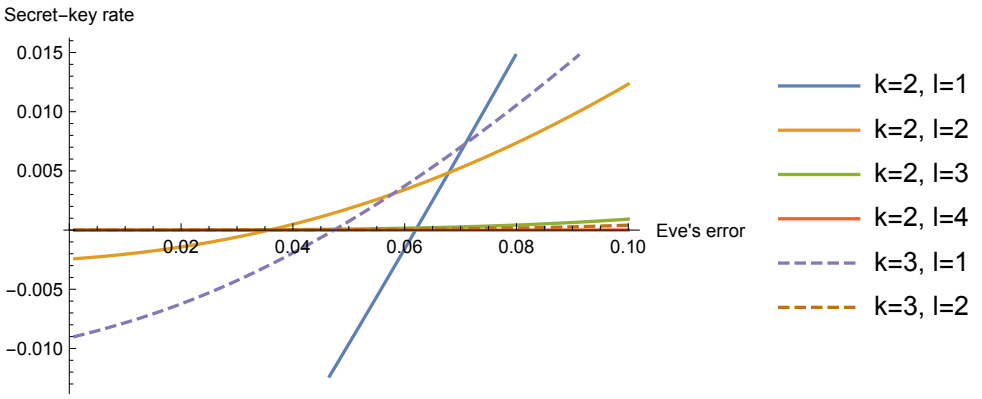

Figure 3.1: Plot of secret-key rates as a function of Eve's error probability ε_E with $\varepsilon_A = \varepsilon_B = 0.1$ for different block-lengths k and number of rounds ℓ of the parity-check protocol.

block-length k	number of rounds ℓ	secret-key rate
2	1	0.00249
2	2	0.00430
2	3	0.00021
3	1	0.00533
3	2	0.00008

Table 3.1: Comparison of secret-key rates of different parity-check protocols (parameterized by block-length and number of rounds) in the satellite setting with $\varepsilon_A = \varepsilon_B = 0.1$ and $\varepsilon_E = 0.065$.

Since we know that block-length $k = 2$ is not always optimal, we can turn our attention to potential improvements to the parity-check protocol when $k > 2$. The reason why finding such improvements for $k > 2$ is more plausible than for $k = 2$ is due to the additional degrees of freedom provided by a larger block-length: for $k > 2$, Bob can potentially relax the acceptance condition, and obtain some additional good bits (in the sense that they can be used to increase the secret-key rate) from bits whose parity-checks initially had some, but few, errors.

Note that relaxing the acceptance condition for the parity-check protocol with $k = 2$ is of no help. In fact, suppose Bob receives $X_1 \oplus X_2$, computes $(Y_1, X_1 \oplus X_2 \oplus Y_2)$, and notices that there is one error (so the string Bob computes is either $(1,0)$ or $(0,1)$). This happens if and only if $X_1 \oplus X_2 \neq Y_1 \oplus Y_2$. In this case, we have

$$\Pr[X_1 \neq Y_1 | X_1 \oplus X_2 \neq Y_1 \oplus Y_2] = \frac{\varepsilon(1-\varepsilon)}{2\varepsilon(1-\varepsilon)} = \frac{1}{2}.$$

Therefore, according to Bob's updated view about X^N, both X_1 and X_2 now look like uniformly random bits which are independent of Y_1 and Y_2, respectively. Furthermore, Eve now knows $X_1 \oplus X_2$ and that $X_1 \oplus X_2 \neq Y_1 \oplus Y_2$, which also severely limits what common information we can still extract from those bits.

Nevertheless, we can use this observation to obtain an improved advantage distillation protocol by considering $k > 2$, as we shall see in Section 3.5.

3.5 Modifying the parity-check protocol

In this section, we present a modification of the parity-check protocol for block-lengths $k > 2$, and then prove and analyze a condition (depending on the error probabilities and block-length) which implies that this modified protocol yields a strict improvement over the original one with corresponding block-length.

The intuition behind this modification was discussed at the end of Section 3.4. For $k > 2$, we can try to relax Bob's acceptance condition for each

k-tuple whose parities he receives. Consider the case where Bob receives some parities $X_1 \oplus X_j$, for $j = 2, \ldots, k$, computes

$$(Y_1, X_1 \oplus X_2 \oplus Y_2, \ldots, X_1 \oplus X_k \oplus Y_k),$$

and notices that $k-1$ entries coincide, and one differs. It seems a sub-optimal decision for Bob to reject this k-tuple. In fact, he has a higher confidence about the value of X_1 than before, albeit slightly less than in the case where all entries coincide. Therefore, it is plausible that Bob, instead of outright rejecting such k-tuples, can move the first bit into a separate "bucket", on which he can possibly run the original parity-check protocol for a large enough number of rounds to extract some extra secret-key rate, as long as the number of errors reported is not too large.

Consider the satellite setting where Alice, Bob, and Eve have error probabilities $\varepsilon_A < 1/2$, $\varepsilon_B < 1/2$, and $\varepsilon_E > 0$, respectively. Furthermore, suppose the three parties have access to N i.i.d. realizations of X, Y, and Z, respectively. Taking into account this intuition, the modified parity-check protocol with block-length k, error bound $D < k/2$, and $(\ell_0, \ell_1, \ldots, \ell_D)$ rounds works as follows:

1. Alice and Bob start with empty ordered sets (we shall call them buckets) B_j, for $j = 0, \ldots, D$.
2. Alice and Bob divide X^N and Y^N into tuples $(X_{ki+1}, \ldots, X_{ki+k})$ and $(Y_{ki+1}, \ldots, Y_{ki+k})$, respectively, for $i = 0, \ldots, \lfloor N/k \rfloor$.
3. For each i, Alice sends $X_{ki+1} \oplus X_{ki+j}$, for $j = 2, \ldots, k$, to Bob via the noiseless authenticated channel.
4. Bob computes the tuple

   $$W = (Y_{ki+1}, X_{ki+1} \oplus X_{ki+2} \oplus Y_{ki+2}, \ldots, X_{ki+1} \oplus X_{ki+k} \oplus Y_{ki+k}).$$

   Let m be the majority of W. Bob reports $d := |\{i : W_i \neq m\}|$ publicly.
5. If $d \leq D$, Alice puts X_{ki+1} at the end of her ordered set B_d, and Bob puts m at the end of his ordered set B_d. Otherwise, Alice and Bob discard the k-tuple.
6. For every bucket B_d, Alice and Bob run the original parity-check protocol with block-length k and $\ell_d - 1$ rounds on the bits of B_d.
7. At the end of the protocol, Alice and Bob hold strings resulting from running the parity-check protocol on each bucket, which they can concatenate and set to $\hat{X}$ and $\hat{Y}$, respectively. They can also compute, for each $\hat{X}_i$ and without interacting, a string $\mathrm{path}(\hat{X}_i)$ whose j-th entry is the index of the bucket where $\hat{X}_i$ was in round j. Let P be a tuple such that $P_i = \mathrm{path}(\hat{X}_i)$ for each i. The final outcome of the protocol for Alice and Bob is thus $(\hat{X}, P)$ and $(\hat{Y}, P)$, respectively.

For a bit $\hat{X}_i$ of the final string $\hat{X}$ which was put in B_d in the first round and then survived an application of the parity-check protocol with $\ell_d - 1$ rounds, we have $\text{path}(\hat{X}_i) = d0^{\ell_d - 1}$. Note that the definition of $\text{path}(\hat{X}_i)$ seems to be a bit redundant. The reason why we chose this particular definition is that it generalizes very easily to potential extensions of this protocol, where we can separate bits into different buckets beyond the first round.

Note that bits $\hat{X}_i$ such that $\text{path}(\hat{X}_i) = 0^{\ell_0}$ went through the original parity-check protocol. Therefore, the modified parity-check protocol with block-length k strictly improves on the original protocol with the same block-length, provided we can show that we can successfully apply the original protocol to the buckets B_d with $d > 0$ if we choose ℓ_d large enough.

Let $\hat{Z}$ denote Eve's total information about $\hat{X}$, and let $\hat{Z}_i$ denote Eve's total information about $\hat{X}_i$. Note that, similarly to the original parity-check protocol of Section 3.3, the triple $(\hat{X}_i, \hat{Y}_i, \hat{Z}_i)$ is conditionally independent of the tuples $(\hat{X}_j, \hat{Y}_j, \hat{Z}_j, \text{path}(\hat{X}_j))$ given $\text{path}(\hat{X}_i)$, for $j \neq i$, since the sets of bits of X^N, Y^N, and Z^N they depend on are disjoint. Moreover, for every i and j, triples $(\hat{X}_i, \hat{Y}_i, \hat{Z}_i)$ and $(\hat{X}_j, \hat{Y}_j, \hat{Z}_j)$ are identically distributed given that $\text{path}(\hat{X}_i) = \text{path}(\hat{X}_j)$.

The secret-key rate of the protocol is

$$\frac{1}{N}(I(\hat{X}P; \hat{Y}P) - I(\hat{X}P; \hat{Z})).$$

Since P is fully determined by one of $(\hat{X}, P)$, $(\hat{Y}, P)$, and $\hat{Z}$, it follows that

$$I(\hat{X}P; \hat{Y}P) - I(\hat{X}P; \hat{Z}) = I(\hat{X}; \hat{Y}|P) - I(\hat{X}; \hat{Z}|P).$$

Let T_{N,k,d,ℓ_d} be the random variable corresponding to the size of Alice's final string generated by running the original parity-check protocol with block-length k and $\ell_d - 1$ rounds on the set B_d, starting from N realizations, and let

$$R_{k,d,\ell_d} := \lim_{N \to \infty} \frac{\mathrm{E}[T_{N,k,d,\ell_d}]}{N}$$

be the corresponding compression ratio. Note that T_{N,k,d,ℓ_d} also corresponds to the number of entries i of P such that $P_i = d0^{\ell_d - 1}$, or, equivalently, the number of $\hat{X}_i$ such that $\text{path}(\hat{X}_i) = d0^{\ell_d - 1}$. By the previous observations, and similarly to the parity-check protocol of Section 3.3, we have, for a fixed N,

$$\begin{aligned} I(\hat{X}; \hat{Y}|P) - I(\hat{X}; \hat{Z}|P) = \\ = \sum_{d=0}^{D} \mathrm{E}[T_{N,k,d,\ell_d}](I(\hat{X}_1; \hat{Y}_1|\text{path}(\hat{X}_1) = d0^{\ell_d - 1}) \\ - I(\hat{X}_1; \hat{Z}_1|\text{path}(\hat{X}_1) = d0^{\ell_d - 1})), \end{aligned}$$

and therefore the modified parity-check protocol generates a secret-key rate of

$$\sum_{d=0}^{D} R_{k,d,\ell_d}(I(\hat{X}_1;\hat{Y}_1|\text{path}(\hat{X}_1) = d0^{\ell_d-1}) - I(\hat{X}_1;\hat{Z}_1|\text{path}(\hat{X}_1) = d0^{\ell_d-1})). \tag{3.16}$$

An advantage of the reasoning above is that it generalizes directly to the natural extensions of our protocol which use buckets in more (or all) of their rounds.

We wish to prove that 3.16 is strictly larger than the rate of the corresponding original parity-check protocol, under suitable conditions on the error probabilities and error bound. The following theorem provides an easily checkable condition which implies that Alice and Bob can extract good bits from B_d by choosing a large enough number of rounds ℓ_d for the modified parity-check protocol. Recall that $\varepsilon = \Pr[X \neq Y]$ and $\alpha_{rs} = \Pr[X \oplus Y = r, X \oplus Z = s]$.

Theorem 3.9 *Let $k > 2$, $d < k/2$, $\varepsilon_A = \varepsilon_B = \alpha$, and ε_E be such that*

$$(2\sqrt{\alpha_{00}\alpha_{01}})^{1-d/k}\,(1-\varepsilon)^{-d/k} > \varepsilon^{1-2d/k},$$

and suppose Alice holds $\hat{X}_1$ with $\text{path}(\hat{X}_1) = d0^{\ell_d-1}$. Furthermore, let $\hat{Y}_1$ be Bob's prediction of $\hat{X}_1$, and let $\hat{Z}_1$ be Eve's information about $\hat{X}_1$. Then, for ℓ_d large enough, Bob has more information about $\hat{X}_1$ than Eve does, i.e.

$$I(\hat{X}_1;\hat{Y}_1|\text{path}(\hat{X}_1) = d0^{\ell_d-1}) - I(\hat{X}_1;\hat{Z}_1|\text{path}(\hat{X}_1) = d0^{\ell_d-1}) > 0.$$

Before we proceed to the proof of Theorem 3.9, we make some comments, and then introduce some auxiliary lemmas and notation.

To begin, note that if the condition of Theorem 3.9 holds for some d, then it also holds for all $d' < d$ (assuming all other parameters are fixed). In fact, if

$$(2\sqrt{\alpha_{00}\alpha_{01}})^{1-d/k}\,(1-\varepsilon)^{-d/k} > \varepsilon^{1-2d/k}$$

holds for some d, then the condition for $d-1$,

$$(2\sqrt{\alpha_{00}\alpha_{01}})^{1-(d-1)/k}\,(1-\varepsilon)^{-(d-1)/k} > \varepsilon^{1-2(d-1)/k},$$

is satisfied whenever

$$2\sqrt{\alpha_{00}\alpha_{01}}(1-\varepsilon) \geq \varepsilon^2,$$

which always holds, as $2\sqrt{\alpha_{00}\alpha_{01}} \geq \varepsilon$ and $1-\varepsilon \geq \varepsilon$ are both true. Thus, it suffices to verify the condition for the largest number of errors allowed, i.e. for the error bound D. An interesting case arises when we set $D = k/3$, since, as we will see, this choice of D already yields protocols which perform better

than *any* original parity-check protocol in certain settings. The condition of Theorem 3.9 holds for $d = k/3$ if and only if

$$0 \leq \alpha < \frac{1}{2} \quad \text{and} \quad \frac{1}{2}\left(1 - \sqrt{1 - 2\alpha + 2\alpha^2}\right) < \varepsilon_E \leq \frac{1}{2}.$$

Therefore, if the error probabilities satisfy these two inequalities, then Alice and Bob can extract good bits from all buckets B_d with $d \leq D = k/3$, for all k. We thus have the following corollary.

Corollary 3.10 *Fix $k > 2$. Let $\varepsilon_A = \varepsilon_B = \alpha < 1/2$ and ε_E be such that*

$$\varepsilon_E > \frac{1}{2}\left(1 - \sqrt{1 - 2\alpha + 2\alpha^2}\right).$$

Then there exist large enough $(\ell_1, \ldots, \ell_{\lfloor k/3 \rfloor})$ such that the modified parity-check protocol with block-length k, error bound $D = k/3$, and $(\ell_0, \ell_1, \ldots, \ell_{\lfloor k/3 \rfloor})$ rounds is a strict improvement (in terms of secret-key rate) on the parity-check protocol with block-length k and ℓ_0 rounds.

Note that there exist pairs (α, ε_E) satisfying the conditions of Corollary 3.10 such that $\varepsilon_E < \alpha$. Figure 3.2 showcases the region of interest for Corollary 3.10, i.e. the set of pairs satisfying the conditions of the condition of the corollary and also $\varepsilon_E < \alpha$. Recall that Figure 3.1 of Section 3.4 show-

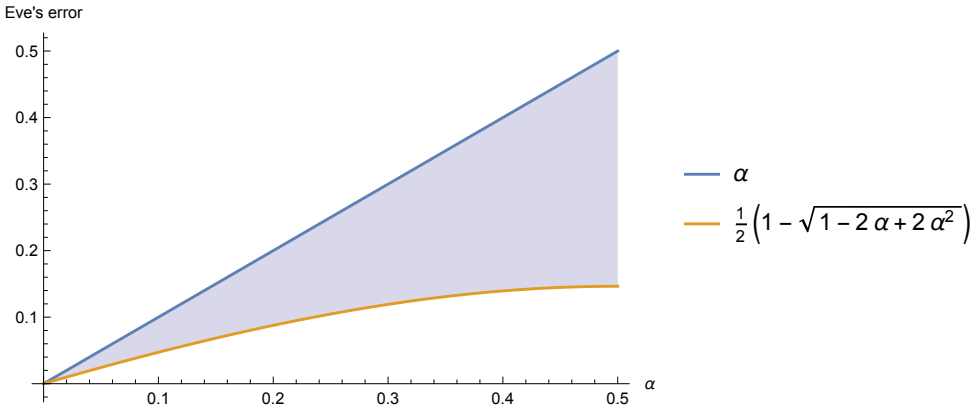

Figure 3.2: Region where the condition of Corollary 3.10 is satisfied by pairs (α, ε_E) with $\varepsilon_E < \alpha$. Ideally, one would like to cover the whole region below the blue line (and also cases where $\varepsilon_A \neq \varepsilon_B$).

cases a choice $(\alpha, \varepsilon_E) = (0.1, 0.065)$ where the parity-check protocol with block-length 2 is not optimal (since block-length 3 performs better). Setting $\alpha = 0.1$, it follows by Corollary 3.10 that the modified parity-check protocol with error bound $D = k/3$ works for $\varepsilon_E > 0.048$, whenever the number of rounds is large enough. Then, if we run the modified parity-check protocol with any block-length $k > 2$, we can extract good bits from the 1-error

bucket B_1, and so we obtain a larger secret-key rate than by simply running the original parity-check protocol with *any* block-length and number of rounds.

Corollary 3.11 *There exist choices of α and ε_E for which the original parity-check protocol does not yield the optimal secret-key rate for any block-length and number of rounds.*

The statement of Theorem 3.9 has deeper consequences than Corollaries 3.10 and 3.11. Given a block-length k and error probabilities α, ε_E, it provides an easy way of checking what is the largest error bound D we can select that ensures we can extract good bits from all buckets B_d for $d \leq D$, provided we choose a large enough number of rounds for each d. If we choose a pair outside the region displayed in Figure 3.2, then it no longer holds that we can provably use buckets B_d for $d \leq D = k/3$. Nevertheless, not all is lost: We can still successfully run the modified parity-check protocol if we decrease the error bound D. It is interesting to consider which pairs (α, ε_E) allow us to run the modified parity-check protocol with block-length k and buckets B_d for $d \leq D = c \cdot k$, for some $c < 1/2$. Obtaining a lower for ε_E as a function of α and D, as we did for $D = k/3$, seems to not be viable in general. One can instead opt to fix α and investigate how the lower bound on ε_E varies with D as we decrease c, with the help of a computer algebra system. Table 3.2 showcases the behavior of the lower bound on ε_E when $\alpha = 0.25$, as a function of $D = c \cdot k$. As expected, the lower bound on ε_E decreases as D becomes smaller in relation to k.

D	lower bound on ε_E
$9k/20$	0.2427
$2k/5$	0.1642
$k/3$	0.1047
$k/4$	0.0607
$k/5$	0.0427
$k/10$	0.0172
$k/15$	0.0108

Table 3.2: Lower bounds on ε_E as a function of the error bound D which guarantee that we can successfully run the modified parity-check protocol with error bound D and $\alpha = 0.25$, obtained through the condition of Theorem 3.9.

Note that we can fix $D > k/3$ if we accept that both error probabilities will need to be larger, and that the gap between thems will need to be smaller for the modified parity-check protocol to work for larger D. As an example, consider $D = 9k/20$. Figure 3.3 showcases the relationship required between α and ε_E in order to ensure the modified parity-check protocol is successful for this value of D. It is interesting to note that the current condition does

not guarantee success in this case when α is too small (smaller than 0.2, for example), if we require $\varepsilon_E < \alpha$.

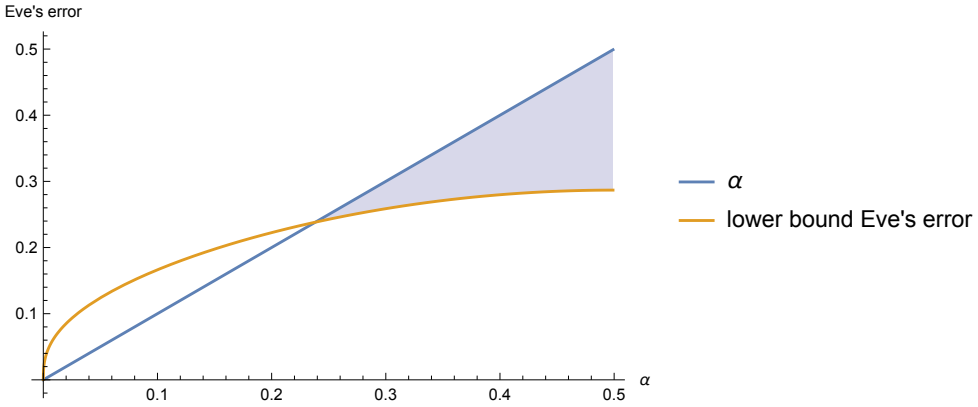

Figure 3.3: Region where the condition of Theorem 3.9 is satisfied by pairs (α, ε_E) with $\varepsilon_E < \alpha$ and $D = 9k/20$. Ideally, one would like to cover the whole region below the blue line (and also cases where $\varepsilon_A \neq \varepsilon_B$).

In order to obtain a strict improvement on the original parity-check protocol with block-length k, it is enough that the condition of Theorem 3.9 is satisfied for $d = 1$, i.e. it is enough that

$$(2\sqrt{\alpha_{00}\alpha_{01}})^{1-1/k}(1-\varepsilon)^{-1/k} > \varepsilon^{1-2/k}.$$

The following lemma, proved by Maurer and Wolf [25], will be helpful.

Lemma 3.12 ([25, part of Lemma 3]) *Let $\varepsilon_A = \varepsilon_B = \alpha < 1/2$ and $\varepsilon_E > 0$. Then*

$$2\sqrt{\alpha_{00}\alpha_{01}} > \varepsilon = 2\alpha(1-\alpha).$$

Proof It suffices to show that

$$\alpha_{00}\alpha_{01} > \alpha^2(1-\alpha)^2. \tag{3.17}$$

Note that equality holds in 3.17 when $\varepsilon_E = 0$, and that $\alpha_{00}\alpha_{01}$ is strictly increasing as a function of ε_E, for fixed α. □

Let $\delta := \frac{2\sqrt{\alpha_{00}\alpha_{01}}}{\varepsilon}$. We have $\delta > 1$ due to Lemma 3.12. Then, the condition of Theorem 3.9 is satisfied for $d = 1$ if and only if

$$\delta^{1-1/k} > \left(\frac{1-\varepsilon}{\varepsilon}\right)^{1/k}. \tag{3.18}$$

If $k \to \infty$, we have that the left-hand side of Inequality 3.18 converges to δ, while the right-hand side converges to 1. Therefore, we have the following corollary.

Corollary 3.13 *For any pair (α, ε_E) with $\alpha < 1/2$ and $\varepsilon_E > 0$, the modified parity-check protocol with block-length k and error bound $D = 1$ is a strict improvement on the original parity-check protocol with block-length k for large enough k.*

The proof of Theorem 3.9 is similar in spirit to the proof of Theorem 3.3. The idea will be to upper bound Bob's error probability and lower bound Eve's optimal error probability so that we get an exponential separation between the two terms. We can then apply Lemma 3.4 to obtain the desired result.

Fix some d satisfying $0 \leq d < k/2$. Suppose, without loss of generality, that, after running the modified parity-check protocol, Alice has a final bit $\hat{X}_1$ such that $\text{path}(\hat{X}_1) = d0^{\ell_d - 1}$, which depends on $X_1, \ldots, X_{L_d}$, for $L_d := k^{\ell_d}$. Thus, $\hat{X}_1 = X_1$. Let $\hat{Y}_1$ be Bob's prediction about $\hat{X}_1$. Eve's total information about $\hat{X}_1$, which we call $\hat{Z}_1$, consists of

$$\hat{Z}_1 = (Z^{L_d}, X_1 \oplus X^{L_d}, \text{path}(\hat{X}_1)).$$

The following lemma states, analogously to Lemma 3.2 and Lemma 3.6, that Eve can compute $X_1 \oplus X^{L_d} \oplus Z^{L_d}$, and then discard Z^{L_d} and $X_1 \oplus X^{L_d}$ without losing any information about $\hat{X}_1$. Its proof follows the same steps as the proof of Lemma 3.2.

Lemma 3.14 *We have*

$$H(X_1|Z^{L_d}, X_1 \oplus X^{L_d}, \text{path}(X_1)) = H(X_1|X_1 \oplus X^{L_d} \oplus Z^{L_d}, \text{path}(X_1)).$$

Suppose $\hat{X}_1 = \hat{Y}_1$. Note that $\hat{X}_1$ was obtained by running the parity-check protocol on bits X_{ki+1}, for $i = 0, \ldots, L_d/k - 1$, which were all put into B_d, since $\text{path}(\hat{X}_1) = d0^{\ell_d - 1}$. Let $\hat{Y}_{1,ki+1}$ be Bob's prediction of X_{ki+1} after Step 5 of the modified parity-check protocol. By Lemma 3.5, it follows that $\hat{Y}_{1,ki+1} = X_{ki+1}$ for all i, i.e. all of Bob's predictions are correct. This happens if and only if every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight d. An analogous reasoning when $\hat{X}_1 \neq \hat{Y}_1$ shows that this happens if and only if every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight $k - d$. Therefore, we have the following lemma.

Lemma 3.15 *Suppose Alice and Bob run the modified parity-check protocol with block-length k and ℓ_d rounds on X^{L_d}, and obtain $\hat{X}_1 = X_1$ with $\text{path}(\hat{X}_1) = d0^{\ell_d - 1}$. Then either every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight d if $\hat{X}_1 = \hat{Y}_1$, or every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight $k - d$ if $\hat{X}_1 \neq \hat{Y}_1$.*

Lemma 3.15 allows us to determine Bob's error probability on $\hat{X}_1$ easily, which we do in the next lemma.

Lemma 3.16 *Suppose $\hat{X}_1$ is such that $\text{path}(\hat{X}_1) = d0^{\ell_d - 1}$, and let $\varepsilon := \Pr[X \neq Y]$. Then*

$$\beta_d := \Pr[\hat{X}_1 \neq \hat{Y}_1 | \text{path}(\hat{X}_1) = d0^{\ell_d - 1}] = \frac{\varepsilon^{L_d(1-2d/k)}}{\varepsilon^{L_d(1-2d/k)} + (1-\varepsilon)^{L_d(1-2d/k)}}.$$

Proof Let E_d be the event that every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight d, and let E_{k-d} be the event that every consecutive block of k bits of $X^{L_d} \oplus Y^{L_d}$ has weight $k-d$. Then

$$\Pr[E_d] = \left(\binom{k}{d} \varepsilon^d (1-\varepsilon)^{k-d} \right)^{L_d/k}, \tag{3.19}$$

and

$$\Pr[E_{k-d}] = \left(\binom{k}{d} \varepsilon^{k-d} (1-\varepsilon)^d \right)^{L_d/k}. \tag{3.20}$$

Furthermore, we know that

$$\beta_d := \Pr[\hat{X}_1 \neq \hat{Y}_1 | \text{path}(\hat{X}_1) = d0^{\ell_d - 1}] = \frac{\Pr[E_{k-d}]}{\Pr[E_d] + \Pr[E_{k-d}]}. \tag{3.21}$$

Plugging in Equation 3.19 and Equation 3.20 into Equation 3.21 and then simplifying yields the desired result. □

We require some notation and a lemma, which will turn out to be useful when we try to lower bound Eve's optimal guessing probability. Recall, from Section 3.2, that α_{rs} denotes the probability that $X \oplus Y = r$ and $X \oplus Z = s$, for $r, s \in \{0, 1\}$. Suppose that $\varepsilon_A = \varepsilon_B = \alpha < 1/2$. Then

$$\begin{aligned} \alpha_{00} &= \alpha^2 \varepsilon_E + (1-\alpha)^2 (1-\varepsilon_E), \\ \alpha_{01} &= \alpha^2 (1-\varepsilon_E) + (1-\alpha)^2 \varepsilon_E. \end{aligned}$$

It will be of our interest to have a decent lower bound on $\binom{N}{k}$ when k is very close to $N/2$ and N is large enough. This can be accomplished via sharp asymptotic estimates on binomial coefficients, stemming from Sterling's approximation.

Lemma 3.17 (First inequality proved in [25, Lemma 3]) *For sufficiently large even N,*

$$\binom{N}{N/2} \geq \frac{2^N}{\sqrt{2\pi N}},$$

and, for sufficiently large odd N,

$$\binom{N}{\frac{N+1}{2}} \geq \frac{2^N}{\sqrt{2\pi N}}.$$

Proof Both bounds can be obtained as corollaries of the fact that, for k a constant independent of N,

$$\binom{N}{k} \sim \frac{2^N e^{-d^2/(2N)}}{\sqrt{\pi N/2}},$$

where $d := N - 2k$, which can be found in [40, Expression 5.41]. □

We can now proceed to the proof of Theorem 3.9.

Proof (Theorem 3.9) Fix $k > 2$, d such that $0 < d < k/2$, $\ell > 0$, $\varepsilon_A = \varepsilon_B = \alpha < 1/2$, and $\varepsilon_E > 0$. Furthermore, define $\varepsilon := \Pr[X \neq Y]$, and $L := k^{\ell_d}$. Suppose Alice and Bob run the modified parity-check protocol and at the end Alice holds $\hat{X}_1$ with $\text{path}(\hat{X}_1) = d0^{\ell-1}$. Without loss of generality, suppose $\hat{X}_1$ depends solely on $X_1 \ldots, X_{L_d}$, and so, in particular, $\hat{X}_1 = X_1$.

By Lemma 3.16, Bob's error probability about X_1 at the end of the protocol, β_L, satisfies

$$\beta_L = \frac{\varepsilon^{L(1-2d/k)}}{\varepsilon^{L(1-2d/k)} + (1-\varepsilon)^{L(1-2d/k)}} \leq \left(\frac{\varepsilon}{1-\varepsilon}\right)^{L(1-2d/k)}. \tag{3.22}$$

Let γ_L be Eve's optimal average error probability about X_1 at the end of the protocol, conditioned on $\text{path}(X_1) = d0^{\ell_d-1}$. By Lemma 3.14, we can assume Eve's optimal guess is computed from $X_1 \oplus X^L \oplus Z^L$ and the fact that $\text{path}(X_1) = d0^{\ell_d-1}$. Note that if $d = 0$, then Eve's optimal guessing strategy for X_1 is to take the majority of $X_1 \oplus X^L \oplus Z^L$. This is the main observation towards lower bounding Eve's average error probability in the proof of Theorem 3.3 in [25]. Unfortunately, the same does not hold for $d > 0$, since, conditioned on $\text{path}(X_1) = d0^{\ell_d-1}$, there are entries i and j such that $X_i = Y_i$ and $X_j \neq Y_j$ simultaneously in each consecutive block of k bits. Thus, Eve's optimal guess may also depend on how the weight of $X_1 \oplus X^L \oplus Z^L$ is spread over the L entries.

We will now find a lower bound for γ_L. Note that, since $\varepsilon_A = \varepsilon_B = \alpha$, it follows that

$$\Pr[X \neq Z | X \neq Y] = \frac{\alpha_{11}}{2\alpha(1-\alpha)} = \frac{1}{2}. \tag{3.23}$$

Let $\mathcal{I} := \{i : X_i = Y_i\}$. Consider the alternative experiment where Eve, besides having access to $X_1 \oplus X^L \oplus Z^L$ and $\text{path}(X_1) = d0^{\ell-1}$ to guess X_1, also has access to $\mathcal{I}$. Let γ'_L denote the optimal average error probability for Eve in this case. It follows immediately that $\gamma_L \geq \gamma'_L$, since Eve could just choose to ignore $\mathcal{I}$.

We analyze Eve's error probability in this alternative case. Note that now Eve knows which entries i of $X_1 \oplus X^L \oplus Z^L$ are such that $X_i \neq Y_i$. By Equation 3.23, it follows that Eve can simply discard such entries, and thus keep only entries $i \in \mathcal{I}$, since $X \oplus Z$ is uniform when conditioned on $X \neq Y$. Let $W := (X^L \oplus Z^L)_{\mathcal{I}}$, i.e. W is equal to $X^L \oplus Z^L$ restricted to entries in $\mathcal{I}$. Since all the bits of W are independent and identically distributed, we have that, similarly to the previous sections, the optimal strategy for Eve (in terms of average error probability) is to take the majority of $(X_1 \oplus X^L \oplus Z^L)_{\mathcal{I}} = X_1 \oplus W$ as the guess for X_1, because it coincides with the maximum likelihood estimate of X_1.

Recall that, by Lemma 3.15, $X_1 = \hat{Y}_1$ holds if and only if every consecutive block of k bits of $X^L \oplus Y^L$ has weight d. Therefore, $\mathcal{I}$ has cardinality $L(1-d/k)$ in this case. Fix d such that $L(1-d/k)$ is even and let A be the event that W has weight exactly $\frac{L(1-d/k)}{2}$. By the previous observations, if A happens, then Eve has no information about X_1, since any possible outcome of $X_1 \oplus X^L \oplus Z^L$, conditioned on A and $\text{path}(X_1) = d0^{\ell-1}$, is equally likely whether $X_1 = 0$ or $X_1 = 1$. Therefore, Eve's optimal average error probability in the original setting satisfies

$$\gamma_L \geq \gamma'_L \geq \frac{1}{2}\Pr[A|\text{path}(X_1) = d0^{\ell_d-1}].$$

It suffices now to find a suitable lower bound for $\Pr[A|\text{path}(X_1) = d0^{\ell_d-1}]$. Consider the following set $\mathcal{G}$, defined as

$$\mathcal{G} := \{u \in \{0,1\}^L : u \text{ has weight } d \text{ in every consecutive block of } k \text{ bits}\}.$$

By a simple counting argument, we have that

$$|\mathcal{G}| = \binom{k}{d}^{L/k}. \tag{3.24}$$

Furthermore, for every $u \in \mathcal{G}$ it holds that

$$\Pr[A, X^L \oplus Y^L = u] = \binom{L(1-d/k)}{L(1-d/k)/2} \sqrt{\alpha_{00}\alpha_{01}}^{L(1-d/k)} \varepsilon^{dL/k}.$$

Therefore,

$$\begin{aligned}\Pr[A|\text{path}(X_1) = d0^{\ell-1}] &= \sum_{u\in\mathcal{G}} \frac{\Pr[A, X^L \oplus Y^L = u]}{\Pr[\text{path}(X_1) = d0^{\ell-1}]} \\ &= \frac{|\mathcal{G}|\binom{L(1-d/k)}{L(1-d/k)/2}\sqrt{\alpha_{00}\alpha_{01}}^{L(1-d/k)}\varepsilon^{dL/k}}{\Pr[\text{path}(X_1) = d0^{\ell-1}]}. \end{aligned} \tag{3.25}$$

In order to simplify the expression in Equation 3.25, note that

$$\Pr[\text{path}(X_1) = d0^{\ell-1}] = \left(\binom{k}{d}\varepsilon^d(1-\varepsilon)^{k-d}\right)^{L/k} + \left(\binom{k}{d}\varepsilon^{k-d}(1-\varepsilon)^d\right)^{L/k}, \tag{3.26}$$

and that, by Lemma 3.17, for large enough L,

$$\binom{L(1-d/k)}{L(1-d/k)/2} \geq \frac{2^{L(1-d/k)}}{\sqrt{2\pi L(1-d/k)}}. \tag{3.27}$$

Combining 3.24, 3.26, and 3.27, we obtain

$$\Pr[A|\text{path}(X_1) = d0^{\ell-1}] \geq$$
$$\geq \frac{1}{\sqrt{2\pi L(1-d/k)}} \frac{(2\sqrt{\alpha_{00}\alpha_{01}})^{L(1-d/k)}\varepsilon^{dL/k}}{\varepsilon^{dL/k}(1-\varepsilon)^{L(1-d/k)} + \varepsilon^{L(1-d/k)}(1-\varepsilon)^{dL/k}}. \quad (3.28)$$

Our goal now is to obtain a denominator as in the right hand side of Inequality 3.22. Note that

$$\varepsilon^{dL/k}(1-\varepsilon)^{L(1-d/k)} + \varepsilon^{L(1-d/k)}(1-\varepsilon)^{dL/k} =$$
$$= (\varepsilon(1-\varepsilon))^{dL/k}((1-\varepsilon)^{L(1-2d/k)} + \varepsilon^{L(1-2d/k)}) \leq$$
$$\leq 2(\varepsilon(1-\varepsilon))^{dL/k}(1-\varepsilon)^{L(1-2d/k)},$$

since $\varepsilon \leq 1-\varepsilon$. Therefore, it follows that

$$\frac{(2\sqrt{\alpha_{00}\alpha_{01}})^{L(1-d/k)}\varepsilon^{dL/k}}{\varepsilon^{dL/k}(1-\varepsilon)^{L(1-d/k)} + \varepsilon^{L(1-d/k)}(1-\varepsilon)^{dL/k}} \geq \frac{(2\sqrt{\alpha_{00}\alpha_{01}})^{L(1-d/k)}(1-\varepsilon)^{-dL/k}}{2(1-\varepsilon)^{L(1-2d/k)}}.$$

We then have

$$\Pr[A|\text{path}(X_1) = d0^{\ell-1}] \geq C(L)\frac{(2\sqrt{\alpha_{00}\alpha_{01}})^{L(1-d/k)}\,(1-\varepsilon)^{-dL/k}}{(1-\varepsilon)^{L(1-2d/k)}}, \quad (3.29)$$

for $C(L) = O(1/\sqrt{L})$. By hypothesis, the pair (α, ε_E) satisfies the condition

$$(2\sqrt{\alpha_{00}\alpha_{01}})^{L(1-d/k)}\,(1-\varepsilon)^{-dL/k} > \varepsilon^{L(1-2d/k)}. \quad (3.30)$$

We can then set

$$b := \frac{\varepsilon}{1-\varepsilon},$$

and

$$\sigma := \frac{(2\sqrt{\alpha_{00}\alpha_{01}})^{(1-d/k)/(1-2d/k)}\,(1-\varepsilon)^{-d/(k-2d)}}{1-\varepsilon},$$

which is obtained by raising the left hand side of 3.30 to the power of $\frac{1}{L(1-2d/k)}$ and dividing it by $1-\varepsilon$. This immediately yields

$$\sigma > \frac{\varepsilon}{1-\varepsilon} = b,$$

as α and ε_E satisfy condition 3.30 (recall also that $d < k/2$). Then

$$\gamma_L \geq C(L)\cdot\sigma^{L(1-2d/k)} \geq (\sigma - \delta(L))^{L(1-2d/k)},$$

since $C(L) = O(1/\sqrt{L})$, for some function $\delta(\cdot)$ such that $\delta(L) \to 0$ as $L \to \infty$. Since $\sigma > b$, we can find L^* large enough such that $\sigma - \delta(L^*) > b$, and thus we can set

$$c := \sigma - \delta(L^*).$$

Therefore, for L large enough, we have

$$\gamma_L \geq c^{L(1-2d/k)} > b^{L(1-2d/k)} \geq \beta_L,$$

and so we can apply Lemma 3.4 to obtain the desired result.

Recall that the previous reasoning works only for $d > 0$ such that $L(1 - d/k)$ is even. For $d > 0$ such that $L(1 - d/k)$ is odd, the proof is in the same spirit of the extension of Theorem 3.3 to large odd N. It suffices to consider the event that W has weight exactly $\frac{L(1-d/k)+1}{2}$ and follow a reasoning analogous to the case where $L(1 - d/k)$ is even above, making use of the second inequality in Lemma 3.17. □

Note that one can recover the argument used by Maurer and Wolf [25] to prove Theorem 3.3 by setting $d = 0$ in the proof of Theorem 3.9 above and invoking Lemma 3.12. It is interesting that this argument extends to $d > 0$ in a relatively natural way.

It is plausible that the condition of Theorem 3.9 can be considerably relaxed in order to allow a much larger selection of error probabilities and error bounds. Furthermore, it should also be possible to keep dividing the selected bits into buckets depending on the number of errors reported by Bob in further rounds (possibly indefinitely), instead of restricting the buckets to the first round of the protocol and then applying the original parity-check protocol, as it stands. The main obstacle towards this extension lies in the fact that it is not clear how to find an exponential separation between Eve's and Bob's error probability in those cases. A possible way to overcome this barrier is to consider the fact that, as already remarked in Section 3.2, a suitable superlinear separation between Eve's and Bob's error probabilities is enough.

We remark that Liu et al. [20] presented an alternative protocol for advantage distillation, which might also improve on the parity-check protocol. Nevertheless, they do not prove that such an improvement holds, as bounding Eve's optimal error probability in their case seems hopeless. The claim stems from simulations where Eve uses a not necessarily optimal strategy for decoding. Furthermore, their protocol uses large quantities of local randomness. In contrast, the protocol we presented in this section uses no local randomness, is as conceptually simple as the parity-check protocol, yields a provable improvement, and is easily generalizable.

We believe the main message of our work in this section is that the ideas behind the parity-check protocol can still be exploited in a deeper way to yield a better understanding of the secret-key rate in the satellite setting.

In Section 3.6 we investigate a practical translation of the satellite setting. We study the secret-key rate when the satellite is allowed to choose the error probabilities according to some channel quality ratio constraint between

Alice and Eve. In particular, we will be interested in determining the asymptotic behavior of this alternative notion of secret-key rate.

3.6 The secret-key rate per time unit under a channel quality constraint

In the previous sections, we studied the quantity of secret-key bits Alice and Bob can obtain per number of realizations of X, Y, and Z in the satellite setting. While the results surveyed are fundamental to understanding the secret-key rate, it is also the case that the model considered does not perfectly capture a real-world satellite setting.

In fact, suppose Alice, Bob, and Eve each possess an antenna listening in to broadcasts from a satellite. Furthermore, suppose Alice and Bob have antennae of the same type. The satellite broadcasts a random bit over a time period of its choice: the larger the time period, the smaller the error probabilities at the receiving ends of the broadcast. Thus, we may think of the satellite as choosing Alice and Bob's error probability α, and Eve's error probability γ, or, equivalently, as choosing the capacities of the binary symmetric channels to Alice and Bob, and to Eve. There exists, though, a relationship between these capacities. One may assume that Eve's antenna is more powerful than Alice's and Bob's (e.g. it may have a larger surface area). This is modelled by requiring the existence of a fixed $D \geq 1$ such that Eve's capacity is always D times larger than Alice's and Bob's capacities, i.e. the satellite is only allowed to choose pairs $(\alpha, \gamma) \in [0, 1/2]^2$ such that

$$\frac{1 - h(\gamma)}{1 - h(\alpha)} = D.$$

Additionally, there is a tradeoff between the time period allocated to transmit a single bit and how many bits the satellite can transmit per some fixed time unit. In fact, if the time period for transmitting a bit shortens, then the satellite can transmit more bits in a given time period than originally possible, although the channels to Alice, Bob, and Eve now have smaller capacities. Therefore, it makes sense in practice to analyze the secret-key *per time unit*. If the satellite transmits two bits per time unit, then the number of secret-key bits extracted per time unit is double whatever number of secret-key bits we extract per number of realizations, but, as we mentioned before, the quality of each bit is worse than if we only transmit one bit per time unit. As a natural approximation, we assume there is an inverse linear relationship between the number of bits transmitted per time unit and Alice's capacity $1 - h(\alpha)$. We can now define our quantity of interest. We will interested in studying its asymptotic behavior as a function of the capacity ratio constraint D.

Definition 3.18 *For $D \geq 1$, the* secret-key rate per time-unit under a channel quality constraint *D, denoted by $S^*(D)$, is defined as*

$$S^*(D) := \sup_{\alpha,\gamma:\frac{1-h(\gamma)}{1-h(\alpha)}=D} \frac{S(\alpha;\gamma)}{1-h(\alpha)},$$

where $S(\alpha;\gamma)$ denotes the secret-key rate per number of realizations in the satellite setting when Alice and Bob have error probability α, and Eve has error probability γ.

Note that we do not know how to compute $S(\alpha;\gamma)$ in general. Therefore, in order to get worst-case guarantees for the behavior of $S^*(D)$, we replace $S(\alpha;\gamma)$ by the rate of some advantage distillation protocol. A first choice would be the repeater-code protocol of Section 3.2. The rate achieved by this protocol is exponentially small, and so we instead opt for the parity-check protocol of Section 3.3.

Recall that the parity-check protocol with block-length 2 achieves a rate (per number of realizations) of

$$R(\ell,\alpha,\gamma) := 2^{-\ell}\prod_{i=0}^{\ell-1}[\beta_i^2+(1-\beta_i)^2](I_B-I_E),$$

where

$$I_B = 1-h(\beta_\ell),$$

and

$$I_E = 1-\sum_{w=0}^{L}\binom{L}{w}\frac{p_w}{\varepsilon^L+(1-\varepsilon)^L}\cdot h\left(\frac{p_w}{p_w+p_{L-w}}\right).$$

In these expressions, α is Alice's and Bob's error probability, γ is Eve's error probability, ℓ is the number of rounds, $L=2^\ell$, $\varepsilon = 2\alpha(1-\alpha)$ is the initial error probability between Alice's and Bob's bits,

$$\beta_i = \frac{\varepsilon^{2^i}}{\varepsilon^{2^i}+(1-\varepsilon)^{2^i}}$$

is the error probability between Alice's and Bob's bits after i rounds, and

$$p_w = \alpha_{00}^{L-w}\alpha_{01}^{w}+\alpha_{10}^{L-w}\alpha_{11}^{w}$$

is the probability that $X^L\oplus Z^L$ is a particular codeword of weight w, and either $X^L\oplus Y^L = 0^L$ or $X^L\oplus Y^L = 1^L$, where $\alpha_{rs} = \Pr[X\oplus Y = r, X\oplus Z = s]$ satisfy

$$\begin{aligned}
\alpha_{00} &= \alpha^2\gamma+(1-\alpha)^2(1-\gamma)\\
\alpha_{01} &= \alpha^2(1-\gamma)+(1-\alpha)^2\gamma\\
\alpha_{10} &= \alpha(1-\alpha)\\
\alpha_{11} &= \alpha(1-\alpha).
\end{aligned}$$

We will focus on the worst-case behavior of the quantity

$$R^*(D) := \sup_{\ell,\alpha,\gamma:\frac{1-h(\gamma)}{1-h(\alpha)}=D} \frac{R(\ell,\alpha,\gamma)}{1-h(\alpha)}.$$

Clearly, $S^*(D) \geq R^*(D)$. It is easy to see that $R^*(D)$ is a decreasing function of D. Indeed, fix $D < D'$. For each choice (α', γ') for $R^*(D')$, we can obtain a choice (α, γ) for $R^*(D)$ by setting $\alpha = \alpha'$ and $\gamma > \gamma'$. It follows immediately that running the parity-check protocol with the same parameters for the new choices (α, γ) yields a larger secret-key rate, and so $R^*(D) \geq R^*(D')$.

We show that, in our practical setting, the parity-check protocol already yields a good secret-key rate per time unit, in the sense that this protocol guarantees that $R^*(D)$ (and thus $S^*(D)$) decreases very mildly with D. In fact, we have the following theorem.

Theorem 3.19 *There exists a constant $c > 0$ such that*

$$R^*(D) \geq \frac{c}{D}$$

holds for all $D \geq 1$.

Before we prove Theorem 3.19, we need some additional definitions and lemmas. Let p'_w be the probability that $X^L \oplus Z^L$ is a particular codeword of weight w and that $X^L = Y^L$. Then, in our case

$$p'_w = \alpha_{00}^{L-w}\alpha_{01}^{w},$$

and

$$p_w = (\alpha(1-\alpha))^L + p'_w \geq p'_w.$$

Lemma 3.20 *We have*

$$h\left(\frac{p_w}{p_w + p_{L-w}}\right) \geq h\left(\frac{p'_w}{p'_w + p'_{L-w}}\right)$$

for all L and w.

Proof This lemma is a consequence of the fact that, for $a, b, x > 0$,

$$\frac{a+x}{a+b+2x} \leq \frac{a}{a+b}$$

holds if and only if $a \geq b$.

Fix $w \leq L/2$. Then

$$\frac{p_w}{p_w + p_{L-w}} \geq \frac{1}{2}$$

since $p_w \geq p_{L-w}$. Furthermore, it holds that

$$\frac{p'_w}{p'_w + p'_{L-w}} \geq \frac{(\alpha(1-\alpha))^L + p'_w}{2(\alpha(1-\alpha))^L + p'_w + p'_{L-w}} = \frac{p_w}{p_w + p_{L-w}} \geq \frac{1}{2},$$

On the other hand, if $w > L/2$, then $p_w < p_{L-w}$ holds, and so

$$\frac{p'_w}{p'_w + p'_{L-w}} < \frac{p_w}{p_w + p_{L-w}} < 1/2.$$

This implies the desired result. □

Lemma 3.21 *Suppose $w = L(1/2 - \delta)$ for some $\delta > 0$. Then*

$$\frac{p'_{L-w}}{p'_w} = \left(\frac{\alpha_{01}}{\alpha_{00}}\right)^{2\delta L}.$$

Proof It suffices to note that

$$\frac{p'_{L-w}}{p'_w} = \frac{\alpha_{00}^w \alpha_{01}^{L-w}}{\alpha_{00}^{L-w} \alpha_{01}^w} = \left(\frac{\alpha_{01}}{\alpha_{00}}\right)^{L-2w},$$

and that $L - 2w = L - 2L(1/2 - \delta) = 2\delta L$. □

Lemma 3.22 ([4, Theorem 2.2]) *If $p = 1/2 - \delta$, we have*

$$\frac{2\delta^2}{\ln(2)} \leq 1 - h(p) \leq 4\delta^2.$$

We can now proceed to the proof of Theorem 3.19.

Proof (Theorem 3.19) The idea of the proof is as follows: We define a sequence of triples $(\ell_n, \alpha_n, \gamma_n)$ and lower-bound its asymptotic parity-check rate $R(\ell_n, \alpha_n, \gamma_n)$. We show that

$$R(\ell_n, \alpha_n, \gamma_n) \geq \frac{c_1}{n^4}$$

for all n, for some positive constant $c_1 > 0$. The sequence itself is also chosen such that both $(1 - h(\gamma_n))/(1 - h(\alpha_n))$ and $1 - h(\alpha_n)$ grow like n^2, which implies the desired result.

For each n we fix $\alpha_n = 1/2 - 1/n$, $\ell_n = 2\log(n)$, and $\gamma_n = 0.4$ for all n. We run the parity-check protocol with block-length 2 using these parameters.

From now on we drop all subscripts for simplicity. First, note that $2^{-\ell} = 1/n^2$. Second, we have

$$\prod_{i=0}^{\ell-1} [\beta_i^2 + (1-\beta_i)^2] \geq \prod_{i=0}^{\ell-1} \frac{1}{2} = \frac{1}{n^2},$$

since $p^2 + (1-p)^2 \geq 1/2$ holds for all $p \in [0,1]$. Therefore,

$$R^*\left(\frac{1-h(\gamma)}{1-h(\alpha)}\right) \geq \frac{1}{(1-h(\alpha))n^4}\left[\sum_{w=0}^{n^2}\binom{n^2}{w}\frac{p_w}{\varepsilon^{n^2}+(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p_w}{p_w+p_{n^2-w}}\right) - h\left(\beta_{n^2}\right)\right],$$

where

$$\beta_{n^2} = \frac{\varepsilon^{n^2}}{\varepsilon^{n^2}+(1-\varepsilon)^{n^2}}.$$

It also holds that

$$\lim_{n\to\infty} h\left(\beta_{n^2}\right) = h\left(\lim_{n\to\infty}\frac{1}{1+(1+8/n^2)^{n^2}}\right) = h\left(\frac{1}{1+e^8}\right) < 0.0044. \quad (3.31)$$

Our goal now is to obtain a lower bound for the weight-sum in the expression of $R^*(D)$ for large enough n that is strictly larger than the upper bound in Inequality 3.31. We have

$$\begin{aligned}
&\sum_{w=0}^{n^2}\binom{n^2}{w}\frac{p_w}{\varepsilon^{n^2}+(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p_w}{p_w+p_{n^2-w}}\right) \geq \\
&\quad\geq \sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)}\frac{p_w}{\varepsilon^{n^2}+(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p_w}{p_w+p_{n^2-w}}\right) \geq \\
&\quad\geq \frac{1}{2}\sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)}\binom{n^2}{w}\frac{p_w}{(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p_w}{p_w+p_{n^2-w}}\right) \geq \\
&\quad\geq \frac{1}{2}\sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)}\binom{n^2}{w}\frac{p'_w}{(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p_w}{p_w+p_{n^2-w}}\right) \geq \\
&\quad\geq \frac{1}{2}\sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)}\binom{n^2}{w}\frac{p'_w}{(1-\varepsilon)^{n^2}}\cdot h\left(\frac{p'_w}{p'_w+p'_{n^2-w}}\right) \geq \\
&\quad\geq \frac{1}{2}\sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)}\binom{n^2}{w}\frac{p'_w}{(1-\varepsilon)^{n^2}}\cdot h\left(\frac{1}{1+\left(\frac{\alpha_{01}}{\alpha_{00}}\right)^{4n}}\right) \quad (3.32)
\end{aligned}$$

for large enough n, where the first inequality holds for $n \geq 5$ since all terms in the sum are positive, the second inequality holds because

$$\varepsilon^{n^2} + (1-\varepsilon)^{n^2} \leq 2(1-\varepsilon)^{n^2}$$

is true for large enough n, the third inequality follows because $p'_w < p_w$, the fourth inequality follows from Lemma 3.20, and the fifth inequality follows from Lemma 3.21 with $L = n^2$ and $\delta = 2/n$, and from the fact that

$$h\left(\frac{p'_w}{p'_w+p'_{L-w}}\right) \leq h\left(\frac{p'_{w+1}}{p'_{w+1}+p'_{L-w-1}}\right)$$

for all $w < L/2$. Furthermore, it holds that

$$\lim_{n\to\infty} h\left(\frac{1}{1+\left(\frac{\alpha_{01}}{\alpha_{00}}\right)^{4n}}\right) = h\left(\frac{1}{1+e^{-32/5}}\right) > 0.0177, \tag{3.33}$$

since

$$\lim_{n\to\infty}\left(\frac{\alpha_{01}}{\alpha_{00}}\right)^{4n} = \lim_{n\to\infty}\left(1-\frac{8}{5n}\right)^{4n} = e^{-32/5}.$$

Let $W := (w(X^{n^2} \oplus Z^{n^2})|X^{n^2} = Y^{n^2})$. Then

$$\sum_{w=n^2(1/2-2/n)}^{n^2(1/2+2/n)} \binom{n^2}{w} \frac{p'_w}{(1-\varepsilon)^{n^2}} = \Pr[|W - n^2/2| \le 2n]. \tag{3.34}$$

It suffices now to find a suitable lower bound for $\Pr[|W - n^2/2| \le 2n]$. In order to do that, we will apply Chebyshev's inequality. First, note that

$$\mathrm{E}[W] = n^2 \cdot \frac{\alpha_{01}}{\alpha_{00}+\alpha_{01}} = n^2 \cdot \frac{\alpha^2(1-\gamma)+(1-\alpha)^2\gamma}{\alpha^2+(1-\alpha)^2} \le \frac{n^2}{2}.$$

Second, algebraic manipulation yields

$$\frac{n^2/2 - \mathrm{E}[W]}{n} = \frac{1}{2.5 + 10/n^2} \le \frac{2}{5},$$

which implies that

$$n^2/2 - 2n/5 \le \mathrm{E}[W] \le n^2/2$$

for all n. Thus, we have

$$n^2(1/2 - 2/n) = n^2/2 - 2n \le \mathrm{E}[W] - n,$$

and

$$n^2(1/2 + 2/n) = n^2/2 + 2n \ge \mathrm{E}[W] + n.$$

Therefore,

$$\Pr[|W - n^2/2| \le 2n] \ge \Pr[|W - \mathrm{E}[W]| \le n] \ge 1 - \frac{\mathrm{Var}[W]}{n^2} \ge \frac{3}{4}, \tag{3.35}$$

where the second inequality follows from Chebyshev's inequality, and the third inequality follows from the fact that

$$\mathrm{Var}[W] = n^2 \cdot \Pr[X \ne Z|X = Y](1 - \Pr[X \ne Z|X = Y]) \le \frac{n^2}{4}.$$

Combining 3.32, 3.33, 3.34, and 3.35 yields

$$\sum_{w=0}^{n^2} \binom{n^2}{w} \frac{p_w}{\varepsilon^{n^2} + (1-\varepsilon)^{n^2}} \cdot h\left(\frac{p_w}{p_w + p_{n^2-w}}\right) > \frac{1}{2} \cdot \frac{3}{4} \cdot 0.0177 > 0.0044$$

for large enough n.

Finally, Lemma 3.22 yields

$$\frac{1-h(\gamma)}{1-h(\alpha)} > \frac{n^2}{200},$$

and

$$1-h(\alpha) < 200(1-h(2/5))/n^2 < 6/n^2$$

for large enough n. Therefore, we have the inequality

$$R^*\left(\frac{n^2}{200}\right) \geq R^*\left(\frac{1-h(\gamma)}{1-h(\alpha)}\right) \geq \frac{c_1}{n^2} \tag{3.36}$$

for some constant $c_1 > 0$ and $n = 2^j$ for some integer j, since $R^*(\cdot)$ is a decreasing function. It follows that we can rewrite Inequality 3.36 as

$$R^*(D) \geq \frac{c_2}{D}$$

for $c_2 = c_1/200 > 0$ and $D = 4^j/200$ for some integer j. This inequality can be extended to all values of D by noting that, for each $D \in [1, \infty)$, there is an integer j such that

$$D \leq \frac{4^j}{200} \leq 6D.$$

In fact, if j is such that $D \leq 4^j \leq 4D$, which we know exists, then

$$D \leq \frac{4^{j+4}}{200} = \frac{256 \cdot 4^j}{200} \leq 6D.$$

Thus,

$$R^*(D) \geq R^*(4^{j+4}/200) \geq \frac{c_2}{6D} = \frac{c_3}{D}$$

for large enough D, where $c_3 = c_2/6 > 0$ is a positive constant independent of D, which yields the desired result. □

It is also possible to show that we cannot do better than $S^*(D) \geq c/D$ for some positive constant $c > 0$. In fact, the following lemma holds.

Theorem 3.23 *We have*

$$S^*(D) \leq \frac{4\ln(2)^2}{D}$$

for all D.

Proof Fix some D and $\alpha, \gamma \in [0, 1/2]$ satisfying $\frac{1-h(\gamma)}{1-h(\alpha)} = D$. Note that

$$\frac{S(\alpha;\gamma)}{1-h(\alpha)} \leq \frac{I(X;Y)}{1-h(\alpha)} = \frac{1-h(\varepsilon)}{1-h(\alpha)},$$

where the first inequality follows from Lemma 2.15, and $\varepsilon = 2\alpha(1-\alpha)$. We claim that

$$\frac{1-h(\varepsilon)}{1-h(\alpha)} \leq \frac{4\ln(2)^2}{D}.$$

Let $\delta := 1/2 - \alpha$. Then, by Lemma 3.22, we have $1 - h(\alpha) \geq 2\delta^2/\ln(2)$. Furthermore, since $\varepsilon = 2\alpha(1-\alpha) = 1/2 - 2\delta^2$, it follows that $1 - h(\varepsilon) \leq 16\delta^4$. Therefore,

$$\frac{1-h(\varepsilon)}{1-h(\alpha)} \leq 8\ln(2)\delta^2.$$

It remains to bound δ^2 by a function of D. Note that

$$\frac{2\delta^2}{\ln(2)} \leq 1 - h(\alpha) = \frac{1-h(\gamma)}{D} \leq \frac{1}{D}.$$

It follows that $\delta^2 \leq \ln(2)/(2D)$, and so

$$\frac{1-h(\varepsilon)}{1-h(\alpha)} \leq \frac{4\ln(2)^2}{D}$$

holds, as desired. Since the choice of α and γ was arbitrary, we have

$$S^*(D) \leq \frac{4\ln(2)^2}{D}$$

for all D. □

Theorems 3.19 and 3.23 establish the correct asymptotic behavior of $S^*(D)$, as showcased in the following corollary.

Corollary 3.24 *We have $S^*(D) = \Theta(1/D)$.*

As mentioned before, the proofs of Theorems 3.19 and 3.23 also yield a proof for a conjecture of Gander and Maurer [11], stated in the following theorem.

Theorem 3.25 *We have*

$$\sup_{\ell,\alpha,\gamma:\,\frac{1-h(\gamma)}{1-h(\alpha)}=D} R(\ell,\alpha,\gamma) = \Theta(1/D^2)$$

and

$$\sup_{\alpha,\gamma:\,\frac{1-h(\gamma)}{1-h(\alpha)}=D} S(\alpha;\gamma) = \Theta(1/D^2).$$

In other words, the (unnormalized) secret-key rate decreases like $1/D^2$ with the channel quality ratio D.

We conclude by proposing an open question. Theorem 3.19 and Lemma 3.23 show that there exist constants $c, c' > 0$ such that $c/D \leq R^*(D) \leq c'/D$ for all $D \geq 1$. Our current techniques allow us to obtain $c' < 2$ and $c > 10^{-5}$ for large enough D, and thus there exists a large gap between the constants in the lower and upper bounds. Furthermore, it is not clear how to significantly improve these bounds without using different techniques. A natural open question arises: Is it possible to bridge the gap between c and c' in the lower and upper bounds for $S^*(D)$?

Chapter 4

A general challenge – when is secret-key agreement possible?

In this chapter, we address one of the most significant open problems in classical information-theoretic secret-key agreement: Deciding when the secret-key rate is positive for general distributions. This leads to the fundamental question of whether distributions with bound information exist. These distributions require secret bits to be created, but have zero secret-key rate, and thus no secret bits can be extracted from them. Studying these questions has given rise to a large body of work, and has uncovered interesting connections with quantum information theory. It is widely believed that bound information exists.

In Section 4.1, we introduce improved upper bounds for the secret-key rate. In Section 4.2, we discuss the information of formation of a distribution, which allows us to grasp the significance of bound information, and study the existence of separations between the bounds of Section 4.1, which is a possible strategy for establishing the existence of bound information. In Section 4.3, we analyze a distribution which may have bound information, and we provide some evidence that this is indeed the case. Finally, in Section 4.4, we study notions of secret-key rate arising from restricting the types of protocols allowed, and study separations between them.

4.1 Better upper bounds for the secret-key rate

In this section, we discuss known improved upper bounds on the secret-key rate, along with their main properties, which will also be useful in the following sections. Throughout this section, we will assume that $H(XYZ)$ is finite.

We start by introducing the *intrinsic mutual information*, an information-theoretic quantity introduced by Maurer and Wolf [26] which, besides being a better

upper bound on the secret-key rate than the bound of Lemma 2.15, is also a lower bound on the number of secret-key bits shared by Alice and Bob they require to create a distribution like P_{XYZ} through an authenticated channel, which is called the information of formation [36], as we will see in Section 4.2. This property gives the intrinsic mutual information a fundamental character in information-theoretic secret-key agreement.

We present the idea Maurer and Wolf [26] used to derive the intrinsic mutual information. For a given distribution P_{XYZ}, Eve can use a strategy that minimizes the upper bound $I(X;Y|Z)$, instead of a strategy that attempts to minimize the secret-key rate of Alice and Bob. While this is a worse idea for Eve than minimizing the secret-key rate directly, it is a very good idea for us, because, unlike for the secret-key rate, it should be easy to compute the new conditional mutual information after Eve applies a simple strategy.

A possible simple strategy for Eve is to transform Z into a new discrete random variable $\overline{Z}$, and then discard Z. This can be thought of as Eve sending Z through a channel defined by $P_{\overline{Z}|Z}$. We thus have that $XY \to Z \to \overline{Z}$ holds. It follows that

$$S(X;Y||Z) \leq S(X;Y||\overline{Z}) \leq I(X;Y|\overline{Z}),$$

where the first inequality is true because we are restricting the set of possible strategies for Eve, and the second inequality follows from Lemma 2.15. Since $P_{\overline{Z}|Z}$ was an arbitrary channel, we can take the infimum over all such channels, and so we obtain

$$S(X;Y||Z) \leq \inf_{XY\to Z\to\overline{Z}} I(X;Y|\overline{Z}),$$

which gives rise to the following definition.

Definition 4.1 ([26, Definition 2]) *The* intrinsic mutual information of X and Y given Z, *denoted by* $I(X;Y\downarrow Z)$, *is defined as*

$$I(X;Y\downarrow Z) := \inf_{XY\to Z\to\overline{Z}} I(X;Y|\overline{Z}),$$

where the infimum is understood to be taken over all discrete random variables $\overline{Z}$ *such that* $XY \to Z \to \overline{Z}$ *holds.*

The following lemma states that the intrinsic mutual information is a strictly better upper bound on the secret-key rate than the conditional mutual information.

Lemma 4.2 ([26, Theorem 2]) *We have*

$$S(X;Y||Z) \leq I(X;Y\downarrow Z) \leq I(X;Y|Z),$$

and there exists a probability distribution for which

$$I(X;Y\downarrow Z) < I(X;Y|Z).$$

Computing the intrinsic mutual information seems hard at first sight, since it involves an infimum. There exists a sequence $(P_{\overline{Z}_i|Z})$ of channels such that

$$I(X;Y|\overline{Z}_i) \to I(X;Y \downarrow Z),$$

but it can happen that no channel $P_{\overline{Z}|Z}$ satisfies $I(X;Y|\overline{Z}) = I(X;Y \downarrow Z)$. Then, in order to compute $I(X;Y \downarrow Z)$, we would need to find and analyze the sequence $(P_{\overline{Z}_i|Z})$. This may be very complex, since it is an infinite sequence of random variables which can behave in a very arbitrary way.

The following lemma, which was proved by Christandl, Renner, and Wolf [5], significantly simplifies the computation of the intrinsic mutual information when Z is a finite random variable, by showing that there exists a channel achieving the infimum in this case.

Lemma 4.3 ([5, Theorem 1]) *Fix a probability distribution P_{XYZ} where Z has finite range $\mathcal{Z}$. Then, there exists a random variable $\overline{Z}$ such that $XY \to Z \to \overline{Z}$ and*

$$I(X;Y \downarrow Z) = I(X;Y|\overline{Z}).$$

Moreover, the range of $\overline{Z}$, denoted by $\overline{\mathcal{Z}}$, satisfies $|\overline{\mathcal{Z}}| \leq |\mathcal{Z}|$.

In particular, Lemma 4.3 implies that, if one wishes to prove that $I(X;Y \downarrow Z) = 0$, it suffices to look for a channel $P_{\overline{Z}|Z}$ such that $I(X;Y|\overline{Z}) = 0$, i.e. such that X and Y are conditionally independent given $\overline{Z}$. This observation will be useful in the following sections.

In order to improve on this upper bound, we look into another property of the secret-key rate. Suppose Alice, Bob, and Eve receive i.i.d. realizations of X, Y, and Z, respectively. The main question is now the following: How is the secret-key rate affected when Eve additionally receives some side information U such that $XYZ \to U$ holds? The following lemma, proved by Renner and Wolf [36], states that the secret-key rate does not decrease by more than the entropy of the side information.

Lemma 4.4 ([36, Theorem 3]) *For any distribution P_{XYZ} and discrete random variable U,*

$$S(X;Y||ZU) \geq S(X;Y||Z) - H(U).$$

Lemma 4.4 also implies that computing an estimate of the decrease of the secret-key rate after adding some side information U is easy, if U is a reasonable random variable. Recall that we could easily compute the conditional mutual information after Eve applied a simple strategy to Z, which was the main motivation behind Definition 4.1. We can now combine both steps, as done in [35, Section 1], by first giving Eve some side information U, and then letting her run a simple strategy on ZU. In order to obtain another upper

bound from these observations, we rearrange the inequality of Lemma 4.4 to get, for any discrete random variable U,

$$S(X;Y||Z) \leq S(X;Y||ZU) + H(U) \leq I(X;Y \downarrow ZU) + H(U).$$

We then have the following definition.

Definition 4.5 ([35, Section 1]) *The* reduced intrinsic mutual information of X and Y given Z, *denoted by* $I(X;Y \downarrow\downarrow Z)$, *is defined as*

$$I(X;Y \downarrow\downarrow Z) := \inf_{XYZ \to U} [I(X;Y \downarrow ZU) + H(U)],$$

where the infimum is taken over all discrete random variables U such that $XYZ \to U$ holds.

Note that we recover the intrinsic mutual information by setting $U = \perp$ with probability 1, as expected from our previous intuition. The following lemma implies that the reduced intrinsic mutual information is a strictly better upper bound on the secret-key rate than the intrinsic mutual information.

Lemma 4.6 ([35, Theorems 1 and 3]) *We have*

$$S(X;Y||Z) \leq I(X;Y \downarrow\downarrow Z) \leq I(X;Y \downarrow Z).$$

Moreover, there exists a probability distribution satisfying

$$I(X;Y \downarrow\downarrow Z) < I(X;Y \downarrow Z).$$

The reduced intrinsic mutual information, unlike the intrinsic mutual information, satisfies a property analogous to the one of Lemma 4.4 for the secret-key rate: For any random variable U,

$$I(X;Y \downarrow\downarrow ZU) \geq I(X;Y \downarrow\downarrow Z) - H(U).$$

It is thus reasonable to ask whether

$$S(X;Y||Z) \stackrel{?}{=} I(X;Y \downarrow\downarrow Z),$$

as was done in [36, Section 5]. More recently, Gohari and Anthanaram [13] proved that there exists a probability distribution for which

$$S(X;Y||Z) < I(X;Y \downarrow\downarrow Z).$$

This is done by obtaining the following improved upper bound on the secret-key rate for finite distributions,

$$GA(X;Y|Z) := \inf_{J} [I(X;Y|J) + I(XY;J|Z)],$$

where the infimum is taken over all discrete random variables J, and finding a distribution such that

$$GA(X;Y|Z) < I(X;Y \downarrow\downarrow Z).$$

This bound was obtained by trial and error, following a streamlined process for deriving upper bounds on the secret-key rate described in [13, Section IV A].

We will now see that $GA(X;Y|Z)$ also satisfies a property akin to the one of Lemma 4.3 for the intrinsic mutual information, and we use the results developed in [5] to prove Lemma 4.3.

Lemma 4.7 *Let X, Y, and Z be random variables with finite ranges $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$, respectively. Then, there exists a finite random variable J with range $\mathcal{J} \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ such that*

$$GA(X;Y|Z) = I(X;Y|J) + I(XY;J|Z).$$

Gohari and Anantharam [13] had already mentioned that such a result could be obtained through measure-theoretic methods (without a proof). Before we prove Lemma 4.7, we need the following lemma, which was proved by Christandl, Renner, and Wolf [5].

Lemma 4.8 ([5, Lemma 6 with $Z^i = W$ for all i]) *Let W be random variable with finite range $\mathcal{W}$, and let (J_i) be a sequence of discrete random variables. Furthermore, fix a continuous function $f : \mathcal{P}(\mathcal{W}) \to \mathbb{R}$, where $\mathcal{P}(\mathcal{W})$ denotes the space of all probability distributions over $\mathcal{W}$. Then, there exists a random variable J with range $\mathcal{J} \subseteq \mathcal{W}$ such that*

$$\mathrm{E}_J[f(P_{W|J}(\cdot, J))] \leq \lim_{i\to\infty} \mathrm{E}_{J_i}[f(P_{W|J_i}(\cdot, J_i))].$$

We can now proceed to the proof of Lemma 4.7.

Proof (Lemma 4.7) Fix random variables X, Y, and Z with finite ranges $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$, respectively. Then, there exists a sequence (J_i) of discrete random variables such that

$$I(X;Y|J_i) + I(XY;J_i|Z) \to GA(X;Y|Z).$$

Consider the following real-valued function, defined over the set of distributions on $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$,

$$f(P_{X'Y'Z'}) = H(X') + H(X'|Y') - H(X'Y'|Z') + c,$$

where $c = H(XY|Z)$ for X, Y, and Z as fixed above (c is a constant). Note that $P_{X'}$, $P_{Y'}$, $P_{Z'}$, $P_{X'|Y'}$, $P_{X'Y'|Z'}$ can all be computed from $P_{X'Y'Z'}$, and so all

the quantities in the expression for f can be computed from $P_{X'Y'Z'}$. Then, we can instantiate $P_{X'Y'Z'} = P_{XYZ|J}(\cdot,\cdot,\cdot,j)$ for some random variable J and obtain

$$f(P_{XYZ|J}(\cdot,\cdot,\cdot,j)) = H(X|J=j) + H(X|Y,J=j) - H(XY|Z,J=j) + H(XY|Z).$$

It follows that

$$\mathrm{E}_J[f(P_{XYZ|J}(\cdot,\cdot,\cdot,J))] = I(X;Y|J) + I(XY;J|Z).$$

Furthermore, f is continuous. This can be seen as a corollary of Lemma 4.21, which we will use in Section 4.2.

Applying Lemma 4.8 with $W = XYZ$ as fixed at the start of the proof, it follows that there is a random variable J with range $\mathcal{J} \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ such that

$$I(X;Y|J) + I(XY;J|Z) = \mathrm{E}_J[f(P_{XYZ|J})] \leq \lim_{i\to\infty} \mathrm{E}_{J_i}[f(P_{XYZ|J_i})] = GA(X;Y|Z),$$

which concludes the proof. □

Note that, unlike in the statement of Lemma 4.3, we require all of X, Y, and Z to have finite range, instead of only Z. This is because XY, Z, and J do not necessarily form a Markov chain. If $XY \to Z \to J$ holds, as is the case for the intrinsic mutual information, then we can represent $I(X;Y|J=j)$ as a continuous function of $P_{Z|J}(\cdot,j)$, since

$$P_{XY|J}(\cdot,\cdot,j) = \sum_z P_{Z|J}(z,j) P_{XY|Z}(\cdot,\cdot,z).$$

Therefore, in order to apply Lemma 4.8 to the intrinsic mutual information, we only require that the range of Z be finite.

The following lemma shows that $GA(X;Y|Z)$ is in general at least as tight as the reduced intrinsic mutual information.

Lemma 4.9 ([13, Corollary 2 and its proof]) *For all finite probability distributions P_{XYZ}, we have*

$$S(X;Y||Z) \leq GA(X;Y|Z) \leq I(X;Y \downarrow\downarrow Z).$$

The upper bound $GA(X;Y|Z)$ was obtained by Gohari and Anantharam as a corollary of the following lemma, also proved in [13].

Lemma 4.10 ([13, Corollary 2]) *Fix a finite probability distribution P_{XYZ}. Then*

$$S(X;Y||Z) \leq \inf_J [S(X;Y||J) + S_{\mathrm{ow}}(XY;J||Z)],$$

where the infimum is taken over all finite random variables J.

The inequality in Lemma 4.10 allows us to obtain more refined upper bounds. By replacing $S(X;Y||J)$ with the corresponding upper-bound $I(X;Y|J)$, and $S_{\text{ow}}(XY;J||Z)$ with the upper bound $I(XY;J|Z)$, we recover $GA(X;Y|Z)$ as an upper bound on the secret-key rate.

We can also obtain even better upper bounds than $GA(X;Y|Z)$ by using the intrinsic or reduced intrinsic mutual informations instead of the conditional mutual information in Lemma 4.10. In fact, we know a single-letter characterization for $S_{\text{ow}}(XY;J||Z)$ (Lemma 2.17), albeit it is not very tractable in general. We could even nest the bounds obtained, e.g. by replacing $S(X;Y||J)$ by $GA(X;Y||J)$ and so on, giving rise to an infinite sequence of potentially increasingly refined upper bounds on the secret-key rate. Unfortunately, bounds beyond $GA(X;Y|Z)$ are generally intractable with our current knowledge, and so we focus our discussion on $GA(X;Y|Z)$ only.

Finally, we remark that Tyagi and Watanabe [41][42] showcase interesting connections between the secret-key rate and binary hypothesis testing, which allows them to obtain upper-bounds for the single-shot secret-key rate (where the parties receive only one realization of their respective source), even when the error probability and secrecy constraint are not asymptotically zero. Applying their techniques to the i.i.d. model we focus on, they recover known upper bounds (namely $I(X;Y \downarrow Z)$ and $GA(X;Y||Z)$).

In the following section, we will connect these improved upper bounds to some fundamental conjectures and concepts in classical information-theoretic secret-key agreement.

4.2 The positivity conjecture and bound information

In this section, we first motivate and discuss two important conjectures in information-theoretic secret-key agreement. Afterwards, we show connections between the most tractable upper bounds discussed in Section 4.1, which imply that they are unlikely to help us settle these conjectures.

Throughout this section, we assume that $H(XYZ)$ is finite.

As we have seen, computing the secret-key rate of a distribution seems to be intractable in general. A more approachable, but still fundamental question is as follows [26]:

Is there a simple, general characterization of the probability distributions P_{XYZ} that satisfy $S(X;Y||Z) > 0$?

Maurer and Wolf studied several settings for information-theoretic secret-key agreement (see [25] [26] [44]), and, in most of those settings, were able to prove that distributions with positive secret-key rate are exactly the ones

with positive intrinsic mutual information. They put forth the following conjecture, which we call the *positivity conjecture*.

Conjecture 4.11 (Positivity conjecture, [26, Conjecture 1]) *For any probability distribution P_{XYZ}, we have $S(X;Y||Z) > 0$ if and only if $I(X;Y \downarrow Z) > 0$.*

If true, the positivity conjecture would yield an easy way of checking whether a given distribution allows information-theoretic secret-key agreement, especially since verifying whether the intrinsic mutual information is positive is relatively simple whenever Z is a finite random variable, due to Lemma 4.3.

Renner and Wolf [36] proved that the intrinsic mutual information, besides being an upper bound on the secret-key rate, is also a lower bound on the amount of secret bits necessary to create a distribution as secure for Alice and Bob as the one under consideration, which is called the *information of formation* of a distribution. We have the following definition.

Definition 4.12 ([36, Definition 5]) *The* information of formation of X and Y given Z, *denoted by $I_{\text{form}}(X;Y|Z)$, is the infimum of all $R \geq 0$ such that, for all $\varepsilon > 0$, there exists an integer N_0 such that, for all $N \geq N_0$, Alice and Bob, sharing a string of $\lceil RN \rceil$ secret bits, can run an interactive protocol with public communication C after which they end up with random variables X' and Y', respectively, such that there exists a channel $P_{\overline{C}|Z^N}$ for which (X', Y', C) is distributed exactly like $(X^N, Y^N, \overline{C})$ with probability at least $1 - \varepsilon$.*

Renner and Wolf [36] proved the following lemma.

Lemma 4.13 ([36, Theorem 5]) *We have*

$$I(X;Y \downarrow Z) \leq I_{\text{form}}(X;Y|Z).$$

Moreover, neither $I(X;Y)$ nor $I(X;Y|Z)$ are lower bounds on $I_{\text{form}}(X;Y|Z)$.

In fact, we can go further: While the positivity conjecture is still open, it holds that the intrinsic mutual information exactly characterizes the distributions P_{XYZ}, where Z is a finite random variable, which have positive information of formation.

Lemma 4.14 *Fix a discrete probability distribution P_{XYZ} such that Z is a finite random variable. Then, $I_{\text{form}}(X;Y|Z) = 0$ if and only if $I(X;Y \downarrow Z) = 0$.*

Proof The key to the proof is the fact that $I(X;Y \downarrow Z) = 0$ implies that there exists a channel $P_{\overline{Z}|Z}$ such that $I(X;Y|\overline{Z}) = 0$, i.e. such that X and Y are conditionally independent given $\overline{Z}$, due to Lemma 4.3, since Z is a finite random variable.

Fix an integer N. The following protocol works for $R = 0$ and for any $\varepsilon > 0$:

1. Alice samples $\overline{Z}^N$ according to $(P_{\overline{Z}})^N$, and then X' according to $(P_{X|\overline{Z}})^N$;

2. Alice sends $\overline{Z}^N$ to Bob;
3. Bob samples Y' according to $(P_{Y|\overline{Z}})^N$.

Note that the communication of the protocol is $C = \overline{Z}^N$. Then (X', Y', C) is distributed according to $(P_{\overline{Z}} P_{X|\overline{Z}} P_{Y|\overline{Z}})^N$. Since $I(X;Y|\overline{Z}) = 0$, it follows that $P_{X|\overline{Z}} P_{Y|\overline{Z}} = P_{XY|\overline{Z}}$. Therefore, (X', Y', C) is distributed exactly like $(X^N, Y^N, \overline{Z}^N)$. □

This insight allows us to rewrite the positivity conjecture for distributions with finite Z as follows.

Conjecture 4.15 (Positivity conjecture for finite Z, rewritten) *Every distribution P_{XYZ} with Z finite and such that $I_{\text{form}}(X;Y|Z) > 0$ also satisfies $S(X;Y||Z) > 0$.*

In other words, if we restrict our attention to distributions with finite Z, the positivity conjecture is equivalent to the statement that we can extract secret bits from every distribution which requires secret bits to be constructed, i.e. which cannot be constructed by Alice and Bob solely from local operations and public discussion.

Gisin, Renner, and Wolf [12] gave some evidence that the positivity conjecture is false by making a connection with an analogous concept in quantum information theory, called bound entanglement. We give an informal overview of their results, based on their exposition.

A quantum state is said to be bound entangled if it is entangled (informally, if it cannot be prepared solely through local operations and classical communication), and, additionally, it has zero entanglement distillation rate (see Section 2.5 for a definition), and so cannot be used for quantum secret-key agreement. It is known that bound entangled quantum states exist [19].

Gisin, Renner, and Wolf give several examples of distributions P_{XYZ} with finite Z and positive intrinsic mutual information (and thus positive information of formation), but for which there are no clear secret-key agreement strategies. These examples are obtained by having Alice, Bob, and Eve measure bound entangled quantum states in some specific bases. More generally, they prove that entangled states can always yield distributions with positive intrinsic information regardless of the way Eve measures them, and that such does not happen when the state is not entangled (in other words, there is a measurement for Eve that kills the intrinsic information). Therefore, there is a correspondence between entanglement and the positivity of the intrinsic mutual information, or, equivalently, the positivity of the information of formation. Inspired by this, they define a classical analogue of bound entanglement.

Definition 4.16 ([12, Definition 2]) *A distribution P_{XYZ} is said to have* bound information *if it satisfies $I(X;Y \downarrow Z) > 0$ and $S(X;Y||Z) = 0$.*

It is now widely believed that distributions with bound information exist, which would imply that the positivity conjecture is false. Note that, based on what we have seen before, a distribution with bound information requires shared secret bits to be created, but no secret-key can be extracted from it. However, it can happen that the positivity conjecture holds (and thus there do not exist distributions with bound information), but there might still exist a distribution P_{XYZ} with Z unbounded such that $S(X;Y||Z) = 0$, $I(X;Y \downarrow Z) = 0$, and $I_{\text{form}}(X;Y|Z) > 0$, since Lemma 4.14 only applies to finite Z.

It was conjectured in [12] that bound entangled quantum states would always give rise to classical distributions with zero secret-key rate. This was shown to be false by Horodecki et al. [16]. Therefore, bound entanglement does not provide an exact correspondence to bound information, although one can still interpret the fact that a distribution arises from a bound entangled quantum state as positive evidence towards the fact that it has bound information. Nevertheless, most of the candidates for bound information arise from measuring bound entangled quantum states, and Z has finite range in all of them.

Besides the connection to bound entanglement in quantum information theory, there are other results that lead us to believe that bound information does exist even for distributions with finite Z, and thus that the positivity conjecture is false. Acín, Cirac, and Masanes [1] showed that tripartite bound information exists. More specifically, they give an example of a distribution P_{ABCE} that cannot be created solely from local operations and public communication, and such that no subset of two parties of A, B, and C can agree on a secret-key, even with the help of the third party. More recently, Prettico and Bae [34] showed that four-partite bound information exists.

We now show that a natural approach towards settling the positivity conjecture is extremely limited by analyzing the relationship between the intrinsic mutual information and the improved upper bounds on the secret-key rate from Section 4.1. If one wants to prove that a certain probability distribution P_{XYZ} with $I(X;Y \downarrow Z) > 0$ has bound information, then one must show that $S(X;Y||Z) = 0$. A natural approach would be to try to compute the value of one of the other improved upper bounds, and hope that it equals zero. This is reasonable, since these upper bounds are generally tighter than the intrinsic mutual information.

The next theorem shows that the reduced intrinsic mutual information is largely useless for establishing that a distribution has bound information.

Theorem 4.17 *Fix a distribution P_{XYZ} such that at least one of X and Y is a finite random variable. Then, $I(X;Y \downarrow Z) > 0$ if and only if $I(X;Y \downarrow\downarrow Z) > 0$.*

Before delving into the proof, we need an auxiliary result. The intrinsic mutual information satisfies a continuity property on the side information U.

Lemma 4.18 *Fix a probability distribution P_{XYZ} such that either X or Y is a finite random variable. Then, for every $\varepsilon > 0$ there is $\delta > 0$ such that $H(U) < \delta$ implies*

$$|I(X;Y \downarrow Z) - I(X;Y \downarrow ZU)| < \varepsilon.$$

Gohari and Anantharam [13, Lemma A1.2] proved this property for X, Y, Z, and U with finite ranges. We extend the lemma to the case where Z and U are arbitrary discrete random variables, provided one of X or Y is a finite random variable, through a slightly simplified proof.

Proof (Lemma 4.18) We follow a similar path to the proof of [13, Lemma A1.2]. Suppose, without loss of generality, that X is a finite random variable with range of cardinality k. Fix $\delta > 0$ and an arbitrary discrete random variable U such that $H(U) < \delta$. Furthermore, let W be a discrete random variable such that $XY \to ZU \to W$ holds. We will construct a random variable W' such that $XY \to Z \to W'$ holds, and furthermore $I(X;Y|W)$ is close to $I(X;Y|W')$. Consider W', with the same range as W, defined by

$$P_{W'|X=x,Y=y,Z=z}(w) = P_{W|U=u_z,Z=z}(w),$$

where $u_z := \operatorname{argmax}_u P_{U|Z=z}(u)$ (if there are several such maximizers, pick one of them arbitrarily). Note that u_z is well-defined for every z in the range of Z since, if $U|Z=z$ has infinite range, then we must have $P_{U|Z=z}(u) \to 0$ as $u \to \infty$. Suppose that $P_{U|Z=z}(0) > 0$ without loss of generality. Therefore, there exist only finitely many u such that $P_{U|Z=z}(u) \geq P_{U|Z=z}(0)$ holds, and u_z is one of them.

It follows immediately that $XY \to Z \to W'$ holds. Additionally, we can couple W and W' as follows: Given $X = x$, $Y = y$, and $Z = z$, we sample U. If $U = u_z$, then we sample W given $Z = z$ and $U = u_z$ and set $W' = W$. If $U \neq u_z$, then we sample W given $Z = z$ and U, and W' given $Z = z$ and $U = u_z$. Our goal is to bound $|I(X;Y|W) - I(X;Y|W')|$. Let $V := 1_{\{W=W'\}}$. Note that

$$\begin{aligned}|I(X;Y|WW') - I(X;Y|W)| &= |I(X;Y|WW'V) - I(X;Y|W)| \leq \\ &\leq |I(X;Y|WW'V) - I(X;Y|WV)| + H(V) \leq \Pr[V=0]\log(k) + H(V),\end{aligned}$$

where the first inequality follows from Lemma 2.13, and the second inequality holds because $I(X;Y|WW', V=1) = I(X;Y|W, V=1)$ and $I(X;Y|Z) \leq$

$\log|\mathcal{X}|$. It remains to bound $\Pr[V=0]$. We have

$$\Pr[V=0] = \sum_z P_Z(z)\Pr[V=0|Z=z] \leq \sum_z P_Z(z)\Pr[U \neq u_z|Z=z].$$

We also know that

$$\delta > H(U) \geq H(U|Z) = \sum_z P_Z(z) H(U|Z=z) \geq \sum_z P_Z(z) H_\infty(U|Z=z) =$$
$$= -\sum_z P_Z(z)\log P_{U|Z=z}(u_z) \geq -\log\left(\sum_z P_Z(z) P_{U|Z=z}(u_z)\right),$$

where the first inequality follows by hypothesis, the second inequality follows because conditioning reduces entropy, the third inequality follows because Shannon entropy is at least as large as min-entropy, and the fourth inequality follows by Jensen's inequality, since log is concave. Therefore,

$$\sum_z P_Z(z) P_{U|Z=z}(u_z) > 2^{-\delta},$$

which implies that

$$\Pr[V=0] \leq \sum_z P_Z(z)\Pr[U \neq u_z|Z=z] < 1 - 2^{-\delta}.$$

Finally, we can conclude that

$$|I(X;Y|WW') - I(X;Y|W)| < (1-2^{-\delta})\log(k) + h(2^{-\delta}) =: f(\delta).$$

By an analogous reasoning, we have

$$|I(X;Y|WW') - I(X;Y|W')| < f(\delta).$$

Fix now a sequence (W_i) of discrete random variables such that $XY \to ZU \to W_i$ holds, and

$$I(X;Y|W_i) \to I(X;Y \downarrow ZU).$$

Then, there exists a sequence (W_i') of discrete random variables such that $XY \to Z \to W_i'$ holds, and

$$|I(X;Y|W_i) - I(X;Y|W_i')| < 2f(\delta).$$

Therefore,

$$I(X;Y \downarrow ZU) \leq I(X;Y \downarrow Z) \leq I(X;Y \downarrow ZU) + 2f(\delta),$$

because $I(X;Y \downarrow Z) \leq I(X;Y|W_i')$ for all i. The first inequality is due to the fact that, if Eve has access to ZU, she can in particular discard U and consider channels $P_{\overline{Z}|Z}$ only. Since $f(\delta) \to 0$ when $\delta \to 0$, we obtain the desired result. □

We can now proceed to the proof of Theorem 4.17.

Proof (Theorem 4.17) Suppose we have $I(X;Y \downarrow\downarrow Z) = 0$. Then, there exists a sequence of random variables (U_i) such that

$$I(X;Y \downarrow ZU_i) + H(U_i) \to 0.$$

In particular, this implies that $H(U_i) \to 0$. By Lemma 4.18, it follows that

$$I(X;Y \downarrow ZU_i) \to I(X;Y \downarrow Z),$$

which implies that $I(X;Y \downarrow Z) = 0$. □

Theorem 4.17 implies that we can largely ignore the reduced intrinsic mutual information when studying bound information and the positivity conjecture, as it does not provide extra insight beyond the intrinsic mutual information for most tractable distributions, and is usually much harder to compute. Nevertheless, the reduced intrinsic mutual information allows us to show the existence of classes of distributions with *asymptotic bound information*, that is, a sequence of distributions $P_{X_iY_iZ_i}$ such that $I_{\text{form}}(X_i;Y_i|Z_i) > c$ for all i and some constant $c > 0$, but $S(X_i;Y_i||Z_i) \to 0$. Renner and Wolf [36] show this by presenting a sequence of distributions that satisfies $I(X_i;Y_i \downarrow Z_i) > 1/2$ for all i, but $I(X_i;Y_i \downarrow\downarrow Z_i) \to 0$.

We now turn our attention to the bound $GA(X;Y|Z)$. In this case, we are looking for distributions P_{XYZ} such that $I(X;Y \downarrow Z) > 0$, but

$$GA(X;Y|Z) = \inf_J [I(X;Y|J) + I(XY;J|Z)] = 0,$$

where J is an arbitrary discrete random variable. The following theorem states that $GA(X;Y|Z)$ cannot help us settle the positivity conjecture whenever X, Y, and Z are all finite.

Theorem 4.19 *Let X, Y, and Z be random variables with finite ranges. Then $GA(X;Y|Z) = 0$ if and only if $I(X;Y \downarrow Z) = 0$.*

Proof Suppose $G(X;Y|Z) = 0$. Since X, Y, and Z are finite, it follows, by Lemma 4.7, that there exists a discrete random variable J such that

$$I(X;Y|J) + I(XY;J|Z) = 0.$$

In particular, we have $I(XY;J|Z) = 0$, which implies that $XY \to Z \to J$ holds. Therefore,

$$I(X;Y \downarrow Z) \leq I(X;Y|J) = 0,$$

which concludes the proof. □

In the remainder of this section, we extend the connection between the positivity of $GA(X;Y|Z)$ and of the intrinsic mutual information to distributions P_{XYZ} with unbounded ranges, under certain conditions. The associated result can be seen in two ways, either as making the computation of $I(X;Y \downarrow Z)$ slightly simpler when Z is infinite, or as a pessimistic result showing that a separation between $I(X;Y \downarrow Z)$ and $GA(X;Y|Z)$ when Z has unbounded range is hard to prove, if it exists.

If we wish to prove that a distribution satisfies $I(X;Y \downarrow Z) = 0$ where Z has unbounded range, then we cannot use Lemma 4.3 to deduce the existence of a channel $P_{\overline{Z}|Z}$ such that $I(X;Y|\overline{Z}) = 0$. Therefore, the computation of the intrinsic mutual information can be very complex in this case. The following theorem slightly simplifies the process. Intuitively, it states that, for distributions P_{XYZ} with possibly unbounded ranges satisfying some constraints, if we have a sequence (J_i) of finite random variables whose ranges do not grow too fast with i such that J_i is almost conditionally independent of XY given Z for large i, and $I(X;Y|J_i) \to 0$, then this is already proof that $I(X;Y \downarrow Z) = 0$.

On the other hand, the following theorem also provides some evidence that using $GA(X;Y|Z)$ to establish the existence of bound information when Z has unbounded range is likely fruitless. In fact, if $I(X;Y \downarrow Z) > 0$, we would be forced to analyze complex sequences of random variables with fast-growing ranges in order to prove that $GA(X;Y|Z) = 0$.

Theorem 4.20 *Let P_{XYZ} be a discrete distribution such that X and Y are finite random variables. If there exists a sequence (J_i) of finite random variables with corresponding ranges $\mathcal{J}_i$ such that*

$$\varepsilon_i := I(X;Y \downarrow J_i) + I(XY;J_i|Z) \to 0$$

and $\log|\mathcal{J}_i| = o(1/\sqrt{\varepsilon_i})$, then $I(X;Y \downarrow Z) = 0$.

Note that the condition of the theorem implies that $GA(X;Y|Z) = 0$. Before we delve into the proof of Theorem 4.20, we need an auxiliary lemma. The following lemma was proved by Ho and Yeung [15]. Intuitively, it states that if two finite distributions are close in statistical distance, then the difference between their entropies is small.

Lemma 4.21 ([15, Theorem 6]) *Suppose P and Q are probability distributions with finite ranges contained in a set of cardinality M, and assume $\Delta(P,Q) \leq \varepsilon < \frac{2(M-1)}{M}$. Then*

$$|H(P) - H(Q)| \leq h\left(\frac{\varepsilon}{2}\right) + \frac{\varepsilon}{2}\log(M-1).$$

In particular, the entropy $H(\cdot)$ is continuous with respect to the statistical distance metric for finite random variables.

We proceed to the proof of Theorem 4.20.

Proof (Theorem 4.20) Fix random variables X, Y, and Z satisfying the conditions of the theorem statement. Suppose there exists a sequence (J_i) of random variables such that

$$\varepsilon_i := I(X;Y|J_i) + I(XY;J_i|Z) \to 0,$$

with corresponding finite range $\mathcal{J}_i$ such that

$$\log|\mathcal{J}_i| = o(1/\sqrt{\varepsilon_i}).$$

In particular, we have $I(X;Y|J_i) \leq \varepsilon_i$ and $I(XY;J_i|Z) \leq \varepsilon_i$. The idea of the proof is as follows: We will define a random variable J_i' from J_i such that $XY \to Z \to J_i'$ holds, and $I(X;Y|J_i')$ is close to $I(X;Y|J_i)$. We can then conclude that $I(X;Y \downarrow Z) = 0$ by making $i \to \infty$. Intuitively, this is possible because $I(XY;J_i|Z) \leq \varepsilon_i$ implies that J_i is almost conditionally independent of X and Y given Z.

Consider J_i', with the same range as J_i, such that

$$P_{J_i'|X=x,Y=y,Z=z}(j) = P_{J_i|Z=z}(j),$$

for all x, y, z, and j. It follows immediately that $XY \to Z \to J_i'$ holds. We have that $I(XY;J_i|Z) \leq \varepsilon_i$ is equivalent to

$$\sum_z P_Z(z) D(P_{XYJ_i|Z=z} || P_{XY|Z=z} P_{J_i|Z=z}) \leq \varepsilon_i. \tag{4.1}$$

Let $D_z(\varepsilon_i) := D(P_{XYJ_i|Z=z} || P_{XY|Z=z} P_{J_i|Z=z})$. Note that $\mathrm{E}_Z[D_Z(\varepsilon_i)] \leq \varepsilon_i$ by Inequality 4.1. By Pinsker's inequality, it follows that

$$\Delta(P_{XYJ_i|Z=z}, P_{XY|Z=z} P_{J_i|Z=z}) \leq \sqrt{\frac{2D_z(\varepsilon_i)}{\log(e)}} =: \delta_z(\varepsilon_i), \tag{4.2}$$

for all z. Furthermore, we have

$$\begin{aligned}\delta_z(\varepsilon_i) &\geq \sum_{x,y,j} |P_{XYJ_i|Z=z}(x,y,j) - P_{XY|Z=z}(x,y) P_{J_i|Z=z}(j)| = \\ &= \sum_{x,y,j} P_{XY|Z=z}(x,y) |P_{J_i|X=x,Y=y,Z=z}(j) - P_{J_i'|Z=z}(j)|, \end{aligned} \tag{4.3}$$

from expanding the expression for the statistical distance in Inequality 4.2 and noting that $P_{J_i|Z=z} = P_{J_i'|Z=z}$. For simplicity, let

$$d_{xyz}(\varepsilon_i) := \sum_j |P_{J_i|X=x,Y=y,Z=z}(j) - P_{J_i'|Z=z}(j)| = \Delta(P_{J_i|X=x,Y=y,Z=z}, P_{J_i'|Z=z}).$$

Note that $\mathrm{E}_{XY}[d_{XYZ}(\varepsilon_i)|Z=z] \leq \delta_z(\varepsilon_i)$ holds by Inequality 4.3. We are interested in bounding the difference between $I(X;Y|J_i)$ and $I(X;Y|J_i')$. Note that

$$I(X;Y|J_i') = I(X;Y) + H(J_i'|X) + H(J_i'|Y) - H(J_i'|XY) - H(J_i').$$

It holds that $H(J_i') = H(J_i)$. Thus, it suffices to show that $H(J_i'|X)$ is very close to $H(J_i|X)$, and similarly for Y and XY in place of X. We have, for all z, and every x in the range of $X|Z=z$,

$$\begin{aligned}\Delta(P_{J_i|X=x,Z=z}, P_{J_i'|Z=z}) &= \sum_j |P_{J_i|X=x,Z=z}(j) - P_{J_i'|Z=z}(j)| \leq \\ &\leq \sum_y P_{Y|X=x,Z=z}(y) \sum_j |P_{J_i|X=x,Y=y,Z=z}(j) - P_{J_i'|Z=z}(j)| = \\ &= \mathrm{E}_Y[d_{XYZ}(\varepsilon_i)|X=x, Z=z], \quad (4.4)\end{aligned}$$

by introducing Y and applying the triangle inequality. We will now proceed to bound $\Delta(P_{J_i|X=x}, P_{J_i'|X=x})$ for all x, which will allow us to bound the difference between $H(J_i|X)$ and $H(J_i'|X)$. We have

$$\begin{aligned}\Delta(P_{J_i|X=x}, P_{J_i'|X=x}) &\leq \sum_z P_{Z|X=x}(z) \sum_j |P_{J_i|X=x,Z=z}(j) - P_{J_i'|Z=z}(j)| \leq \\ &\leq \mathrm{E}_{YZ}[d_{XYZ}(\varepsilon_i)|X=x] =: d_x(\varepsilon_i),\end{aligned}$$

for all x, where the first inequality follows by introducing Z and applying the triangle inequality, and the second inequality follows from Inequality 4.4. We can apply Lemma 4.21 to obtain

$$|H(J_i|X=x) - H(J_i'|X=x)| \leq h\left(\frac{d_x(\varepsilon_i)}{2}\right) + \frac{d_x(\varepsilon_i)}{2}\log|\mathcal{J}_i|, \quad (4.5)$$

for all x such that $d_x(\varepsilon_i) < \frac{2(|\mathcal{J}_i|-1)}{|\mathcal{J}_i|}$. If x is such that $d_x(\varepsilon_i) \geq \frac{2(|\mathcal{J}_i|-1)}{|\mathcal{J}_i|}$, then the best bound we have is

$$|H(J_i|X=x) - H(J_i'|X=x)| \leq \log|\mathcal{J}_i|.$$

Note that

$$\begin{aligned}\mathrm{E}_X[d_X(\varepsilon_i)] = \mathrm{E}_{XYZ}[d_{XYZ}(\varepsilon_i)] \leq \mathrm{E}_Z[\delta_Z(\varepsilon_i)] &= \mathrm{E}_Z\left[\sqrt{\frac{2D_Z(\varepsilon_i)}{\log(e)}}\right] \leq \\ &\leq \sqrt{\frac{2\mathrm{E}_Z[D_Z(\varepsilon_i)]}{\log(e)}} \leq \sqrt{\frac{2\varepsilon_i}{\log(e)}} =: \delta(\varepsilon_i), \quad (4.6)\end{aligned}$$

where the first inequality follows from the fact that

$$\mathrm{E}_{XY}[d_{XYZ}(\varepsilon_i)|Z=z] \leq \delta_z(\varepsilon_i)$$

by Inequality 4.3, the second inequality follows from Jensen's inequality, since the square root is concave, and the third inequality follows from the fact that $\mathrm{E}_Z[D_Z(\varepsilon_i)] \leq \varepsilon_i$ by Inequality 4.1.

We now need to control how much probability is put into x where we cannot apply the bound of Lemma 4.21. Fix i, and call x *bad* if $d_x(\varepsilon_i) \geq \frac{2(|\mathcal{J}_i|-1)}{|\mathcal{J}_i|}$. Note that we can assume that $|\mathcal{J}_i| > 1$ for all i, since otherwise J_i is constant and we can take $J_i' = J_i$. Then we must have

$$\Pr[X \text{ bad}] \leq \delta(\varepsilon_i), \tag{4.7}$$

since

$$\delta(\varepsilon_i) \geq \mathrm{E}_X[d_X(\varepsilon_i)] \geq \Pr[X \text{ bad}] \cdot \frac{2(|\mathcal{J}_i|-1)}{|\mathcal{J}_i|} \geq \Pr[X \text{ bad}].$$

Therefore, for large enough i,

$$\begin{aligned} |H(J_i|X) - H(J_i'|X)| &\leq \sum_x P_X(x)|H(J_i|X=x) - H(J_i'|X=x)| \leq \\ &\leq \mathrm{E}_X\left[h\left(\frac{d_X(\varepsilon_i)}{2}\right) + \frac{d_X(\varepsilon_i)}{2}\log|\mathcal{J}_i|\right] + \delta(\varepsilon_i)\log|\mathcal{J}_i| \leq \\ &\leq h\left(\frac{\delta(\varepsilon_i)}{2}\right) + \frac{3\delta(\varepsilon_i)}{2}\log|\mathcal{J}_i|, \end{aligned}$$

where the first inequality follows from the triangle inequality, the second inequality follows from Inequalities 4.5 and 4.7, and the third inequality holds because of Jensen's inequality, since h is concave, Inequality 4.6, and the fact that $\delta(\varepsilon_i) < 1$ for large enough i. Note also that $h\left(\frac{d_X(\varepsilon_i)}{2}\right)$ is well-defined since $d_{xyz}(\varepsilon_i) \leq 2$ always.

By an analogous reasoning with Y and XY in place of X, we obtain the same bound for $|H(J_i|Y) - H(J_i'|Y)|$ and $|H(J_i|XY) - H(J_i'|XY)|$. Finally, it follows, by combining the previous inequalities, that

$$|I(X;Y|J_i) - I(X;Y|J_i')| \leq 3\left(h\left(\frac{\delta(\varepsilon_i)}{2}\right) + \frac{3\delta(\varepsilon_i)}{2}\log|\mathcal{J}_i|\right),$$

for large enough i. Note that we have $\delta(\varepsilon_i) = O(\sqrt{\varepsilon_i})$. Furthermore, since $\log|\mathcal{J}_i| = o(1/\sqrt{\varepsilon_i})$, it follows that

$$\lim_{i\to\infty} |I(X;Y|J_i) - I(X;Y|J_i')| = 0.$$

Therefore, we obtain a sequence (J_i') of random variables such that $XY \to Z \to J_i'$ holds for all i, and

$$I(X;Y|J_i') \to 0,$$

which implies that $I(X;Y \downarrow Z) = 0$. □

In the next section, we study a candidate distribution for bound information, proposed by Renner and Wolf [36].

4.3 A candidate for bound information

In this section, we introduce and analyze a class of distributions suggested by Renner and Wolf [36, Section 5] as a potential example of distributions with bound information. Unlike the examples suggested in [12], this class of distributions is not motivated by bound entangled quantum states, but rather by a distribution P_{XYZ} satisfying

$$S(X;Y||Z) < I(X;Y \downarrow Z).$$

Nevertheless, it was shown that such a class of distributions actually gives rise to a (at the time previously unknown) bound entangled quantum state [36]. While we know that in general this is not a guarantee that the classical distribution has bound information (by results of [16] already mentioned), it provides some positive evidence for this fact.

The marginal distributions for X and Y, parameterized by a positive real number a, are defined in Table 4.1. The random variable Z is defined as

		X			
		0	1	2	3
Y	0	1/8	1/8	a	a
	1	1/8	1/8	a	a
	2	a	a	1/4	0
	3	a	a	0	1/4

Table 4.1: A candidate for bound information [36]. Note that the entries must be normalized by $8a + 1$.

follows:

$$Z = \begin{cases} X \oplus Y & \text{if } X, Y \in \{0,1\} \\ X \mod 2 & \text{if } X, Y \in \{2,3\} \\ (X,Y) & \text{otherwise.} \end{cases}$$

We have the following conjecture.

Conjecture 4.22 ([36]) *The distribution P_{XYZ} defined above has bound information for some $a \geq 0$.*

Let us first look at the case $a = 0$. In this case, it holds that $S(X;Y||Z) \geq 1$, since Alice and Bob can extract $U := \lfloor X/2 \rfloor$ from both X and Y, which is completely unknown to Eve. Furthermore, note that Z and U jointly determine X and Y completely, and so $I(X;Y|ZU) = 0$, which implies that

$$I(X;Y \downarrow\downarrow Z) \leq H(U) = 1.$$

We can then conclude that $S(X;Y||Z) = 1$. It can also be shown that $I(X;Y \downarrow Z) = 3/2$. This was the distribution used by Renner and Wolf [36] to prove that $S(X;Y||Z) \neq I(X;Y \downarrow Z)$ in general.

We proceed to show that every distribution in this class has positive intrinsic mutual information. Note that, since Z has finite range, this distribution has positive information of formation if and only if it has positive intrinsic mutual information, due to Lemma 4.14.

Lemma 4.23 *We have $I(X;Y \downarrow Z) > 0$ for all $a \geq 0$. Therefore, $I_{\text{form}}(X;Y|Z) > 0$ for all $a \geq 0$.*

Proof Fix $a \geq 0$. Suppose, in view of a contradiction, that $I(X;Y \downarrow Z) = 0$. This implies, by Lemma 4.3, that there is a channel $P_{\overline{Z}|Z}$ such that $I(X;Y|\overline{Z}) = 0$. Let $\overline{\mathcal{Z}}$ be the range of $\overline{Z}$, and fix $\overline{z} \in \overline{\mathcal{Z}}$. We must have $I(X;Y|\overline{Z} = \overline{z}) = 0$, or equivalently, it must hold that X and Y are conditionally independent given $\overline{Z} = \overline{z}$. We will show that this implies that $\Pr[X = 2, Y = 2|\overline{Z} = \overline{z}] = 0$. Since the choice of $\overline{z}$ was arbitrary, it follows that $\Pr[X = 2, Y = 2] = 0$, which is clearly false.

We know that

$$\Pr[X = 2, Y = 3] = \Pr[X = 3, Y = 2] = 0,$$

which implies that

$$\Pr[X = 2, Y = 3|\overline{Z} = \overline{z}] = \Pr[X = 3, Y = 2|\overline{Z} = \overline{z}] = 0.$$

By the conditional independence hypothesis, we must have

$$\Pr[X = 2|\overline{Z} = \overline{z}] = 0 \quad \text{or} \quad \Pr[Y = 3|\overline{Z} = \overline{z}] = 0,$$

and

$$\Pr[X = 3|\overline{Z} = \overline{z}] = 0 \quad \text{or} \quad \Pr[Y = 2|\overline{Z} = \overline{z}] = 0.$$

We proceed by cases:

1. $\Pr[X = 2|\overline{Z} = \overline{z}] = 0$:

   Then we immediately have

   $$\Pr[X = 2, Y = 2|\overline{Z} = \overline{z}] = \Pr[X = 2|\overline{Z} = \overline{z}] \Pr[Y = 2|\overline{Z} = \overline{z}] = 0$$

   via the conditional independence hypothesis.

2. $\Pr[Y = 2|\overline{Z} = \overline{z}] = 0$:

   Reasoning analogous to (i).

3. $\Pr[X = 3|\overline{Z} = \overline{z}] = 0$ and $\Pr[Y = 3|\overline{Z} = \overline{z}] = 0$:

   We must have

   $$\Pr[X = 3, Y = 3|\overline{Z} = \overline{z}] = 0.$$

   Let $p := P_{\overline{Z}|Z}(\overline{z}, 1)$. Then we can rewrite

   $$\Pr[X = 3, Y = 3|\overline{Z} = \overline{z}] = \frac{p}{4\Pr[\overline{Z} = \overline{z}](8a + 1)},$$

   which implies that $p = 0$. It follows that

   $$\Pr[X = 0, Y = 1|\overline{Z} = \overline{z}] = \frac{p}{8\Pr[\overline{Z} = \overline{z}](8a + 1)} = 0,$$

   and

   $$\Pr[X = 1, Y = 0|\overline{Z} = \overline{z}] = \frac{p}{8\Pr[\overline{Z} = \overline{z}](8a + 1)} = 0.$$

   By the conditional independence hypothesis, we have that either

   $$\Pr[X = 0, Y = 0|\overline{Z} = \overline{z}] = 0,$$

   or

   $$\Pr[X = 1, Y = 1|\overline{Z} = \overline{z}] = 0.$$

   Suppose $\Pr[X = 0, Y = 0|\overline{Z} = \overline{z}] = 0$. The reasoning in the other case is analogous. Let $q := P_{\overline{Z}|Z}(\overline{z}, 0)$. Then we have

   $$\Pr[X = 0, Y = 0|\overline{Z} = \overline{z}] = \frac{q}{8\Pr[\overline{Z} = \overline{z}](8a + 1)} = 0,$$

   which implies that $q = 0$. Therefore

   $$\Pr[X = 2, Y = 2|\overline{Z} = \overline{z}] = \frac{q}{4\Pr[\overline{Z} = \overline{z}](8a + 1)} = 0.$$ □

The reasons why this class of distributions seems like a good example for proving the existence of bound information are, first, that all distributions in the class have positive intrinsic mutual information (and thus positive information of formation), and second, that there seem to be no useful strategies for secret-key agreement besides Alice and Bob computing $U_X := \lfloor X/2 \rfloor$ and $U_Y := \lfloor Y/2 \rfloor$, respectively, discarding everything else, and then running some protocol on the (U_X, U_Y) pairs. It seems intuitive that such a strategy stops working for a large enough value of a. In fact, for large a, Eve knows a considerable fraction of the (X, Y) pairs, and thus the corresponding (U_X, U_Y) pairs. At the same time, Alice and Bob initially have no idea about which pairs (U_X, U_Y) satisfy $U_X \neq U_Y$ (and thus they do not know

which pairs Eve knows), and public discussion may yield too much information to Eve about the remaining pairs (her initial information about them consists only of the fact that $U_X = U_Y$).

We will now study this strategy more carefully. Suppose Alice and Bob receive X and Y, respectively, then compute bits U_X and U_Y, and discard everything else. For a fixed $a \geq 0$, we have

$$\Pr[U_X = 0] = \Pr[U_Y = 0] = \frac{1}{2} \quad \text{and} \quad \Pr[U_X = U_Y] = \frac{1}{8a+1}.$$

If $U_X = U_Y$, then Eve receives $Z = 0$ or $Z = 1$, which gives no information about U_X and U_Y, besides the fact that $U_X = U_Y$. On the other hand, if $U_X \neq U_Y$, then Eve receives (X, Y), which allows her to compute U_X and U_Y. We are then interested in the quantity $S(U_X; U_Y || Z)$.

Consider now random variables $X', Y' \in \{0,1\}$, and $Z' \in \{0\} \cup \{(1,0),(1,1)\}$ with joint probability distribution $P_{X'Y'Z'}$ defined in Table 4.2. Whenever

		X'	
		0	1
Y'	0	1 [0]	$8a$ [(1,1)]
	1	$8a$ [(1,0)]	1 [0]

Table 4.2: The distribution obtained if Alice and Bob follow a strategy based only on U_X and U_Y. Note that each entry in the table must be normalized by $2(8a+1)$. If the entry in row i, column j is p [k], then this means that $\Pr[X = j, Y = i, Z = k] = p$.

$X' = Y'$, we have $Z = 0$, while if $X' \neq Y'$, Eve gets $Z = (1, X')$. This means that, if $X' = Y'$, Eve only learns that $X' = Y'$ and nothing else, while if $X' \neq Y'$, then Eve learns both X' and Y'. From these observations and the previous discussion, it is clear that

$$S(U_X; U_Y || Z) = S(X'; Y' || Z'),$$

and thus we will turn to analyzing $P_{X'Y'Z'}$ instead. We can try applying the lower and upper bounds of Lemma 2.15 first. We have

$$I(X';Y') - I(X';Z') = H(X'|Z') - H(X'|Y') = \frac{1}{8a+1} - h\left(\frac{1}{8a+1}\right) < 0,$$

for large enough a, since $p < h(p)$ for p small enough, and similarly for X' switched with Y'. On the other hand,

$$I(X';Y') = 1 - h\left(\frac{1}{8a+1}\right) > 0,$$

and

$$I(X';Y'|Z') = \frac{1}{8a+1} I(X';Y'|Z' = 0) = \frac{1}{8a+1}.$$

While these bounds do not help us settle whether the secret-key rate is positive or zero, they do show that the secret-key rate approaches zero as a grows. We will turn to the intrinsic mutual information for a sharper analysis of the secret-key rate.

The following theorem implies that, for a large enough, no secret-key agreement strategy for P_{XYZ} that depends only U_X and U_Y can succeed.

Theorem 4.24 *We have*

$$I(X';Y' \downarrow Z') > 0$$

if and only if $a < 1/8$. Therefore,

$$S(U_X;U_Y||Z) = S(X';Y'||Z') = 0$$

holds for $a \geq 1/8$.

Before we proceed to the proof of this theorem, we need the following auxiliary result.

Lemma 4.25 ([12, Part of Example 2]) *Let (b_i) and (c_i) be two finite sequences such that $b_i \geq 0$, $c_i \geq 0$ for all i, $\sum_i b_i = \sum_i c_i = 1$, and such that $c_i = 0$ implies $b_i = 0$. Then*

$$\sum_{i:c_i>0} \frac{b_i^2}{c_i} \geq 1.$$

We are now ready to prove Theorem 4.24.

Proof (Theorem 4.24) The proof of this theorem is inspired by the analysis of Example 2 in [12].

Fix a channel $P_{\overline{Z}|Z'}$, and let $\overline{\mathcal{Z}}$ be the range of $\overline{Z}$. Furthermore, define

$$b_i := P_{\overline{Z}|Z'}(i,0), \quad c_i := P_{\overline{Z}|Z'}(i,(1,0)), \quad d_i := P_{\overline{Z}|Z'}(i,(1,1)), \tag{4.8}$$

for each $i \in \overline{\mathcal{Z}}$. We will derive conditions which the sequences (b_i), (c_i), and (d_i) must satisfy so that

$$I(X';Y'|\overline{Z}=i) = 0$$

holds for all i, or, in other words, such that X' and Y' are conditionally independent given $\overline{Z} = i$ for all i.

For simplicity, let $\alpha := 2(8a+1)$, and let $\alpha_i := \alpha \Pr[\overline{Z} = i]$. The distribution $P_{X'Y'|\overline{Z}=i}$ is shown in Table 4.3.

Our goal will be to re-derive the value of entry (1,0) in the table using the hypothesis that X' and Y' are conditionally independent given $\overline{Z} = i$.

We have

$$\Pr[X'=0|\overline{Z}=i] = \frac{b_i + 8ac_i}{\alpha_i}.$$

		X'	
		0	1
Y'	0	b_i	$8ad_i$
	1	$8ac_i$	b_i

Table 4.3: Distribution $P_{X'Y'|\overline{Z}=i}$. Each entry must be normalized by α_i.

Suppose now that $I(X';Y'|\overline{Z}=i)=0$ and $c_i>0$. We then have

$$\Pr[X'=0|\overline{Z}=i]\cdot\Pr[Y'=0|\overline{Z}=i]=\frac{b_i}{\alpha_i},$$

and thus

$$\Pr[Y'=0|\overline{Z}=i]=\frac{b_i}{b_i+8ac_i}.$$

Furthermore,

$$\Pr[X'=1|\overline{Z}=i]\cdot\Pr[Y'=1|\overline{Z}=i]=\frac{b_i}{\alpha_i},$$

which implies that

$$\Pr[X'=1|\overline{Z}=i]=\frac{b_i(b_i+8ac_i)}{\alpha_i\cdot 8ac_i}.$$

We can then compute

$$\Pr[X'=1|\overline{Z}=i]\cdot\Pr[Y'=0|\overline{Z}=i]=\frac{b_i^2}{\alpha_i\cdot 8ac_i}.$$

By inspecting the table, it finally follows that

$$8ad_i=\frac{b_i^2}{8ac_i},$$

or, equivalently, that

$$d_i=\frac{b_i^2}{64a^2c_i},$$

whenever $c_i>0$. On the other hand, if $c_i=0$, we must have either

$$\Pr[X'=0|\overline{Z}=i]=0 \quad \text{or} \quad \Pr[Y'=1|\overline{Z}=i]=0.$$

Both cases imply that $b_i=0$.

Note that we must have $\sum_i d_i=1$. Therefore, it must hold that

$$\sum_{i:c_i>0}\frac{b_i^2}{64a^2c_i}\leq 1,$$

or, equivalently,

$$\sum_{i:c_i>0} \frac{b_i^2}{c_i} \leq 64a^2. \tag{4.9}$$

We can now make some observations about this. If $b_i \geq 0$, $c_i \geq 0$ for all i, $\sum_i b_i = \sum_i c_i = 1$, and if $c_i = 0$ implies $b_i = 0$, then it holds that

$$\sum_{i:c_i>0} \frac{b_i^2}{c_i} \geq 1$$

by Lemma 4.25. Therefore, in order for Inequality 4.9 to hold, we must have $64a^2 \geq 1$, or equivalently, $a \geq 1/8$. We can then conclude that

$$I(X';Y' \downarrow Z') > 0$$

whenever $a < 1/8$, by Lemma 4.3, since in this case there is no channel $P_{\overline{Z}|Z}$ such that $I(X';Y'|\overline{Z}) = 0$.

If $a \geq 1/8$, it is straightforward to see that the random variable $\overline{Z}$ with range $\{0,1\}$ induced by the distribution $P_{\overline{Z}|Z}$ given by $b_0 = d_0 = 1$, $b_1 = d_1 = 0$, $c_0 = 1/64a^2$, and $c_1 = 1 - 1/64a^2$ satisfies $I(X;Y|\overline{Z}) = 0$. We can then conclude that $I(X';Y' \downarrow Z') = 0$ whenever $a \geq 1/8$. □

The result of Theorem 4.24 provides good evidence that the joint probability distribution in this section has bound information, since it shows that the only clear strategy for secret-key agreement cannot work whenever a is large enough.

In fact, we can easily get a slightly stronger result. Note that, given $U_X = 0$, Alice can perfectly simulate X by sampling uniformly random bits, and the same is true for Bob with U_Y and Y in place of U_X and X. Let X'' be defined as $X'' = (0,R)$, whenever $U_X = 0$, where R is a random bit, and $X'' = 1$ whenever $U_X = 1$. Let Y'' be defined analogously. Since X'' and Y'' can be perfectly simulated from X' and Y' alone, the secret-key rate for (X'',Y'',Z') is zero whenever secret-key agreement is impossible from (X',Y',Z').

Corollary 4.26 *We have*

$$S(X'';Y''||Z') = 0$$

whenever $a \geq 1/8$. Therefore, any secret-key agreement strategy for P_{XYZ} which may depend on X when $U_X = 0$ but only depends on U_X when $U_X = 1$, and analogously for Y and U_Y, cannot possibly work.

In other words, Corollary 4.26 implies that a successful secret-key agreement strategy for P_{XYZ} must use the fact that $X = 2$ or $X = 3$ (respectively $Y = 2$ or $Y = 3$) in some non-trivial way, besides using it to compute U_X (respectively U_Y).

Nevertheless, while it seems intuitive that the distribution studied in this section has bound information, it is still not clear exactly how one could prove that no strategy produces a secret-key. A possible approach to this challenge is to make use of the concept of binarizations, introduced by Gisin, Renner, and Wolf [12].

Definition 4.27 ([12, Part of Proposition 5]) *A distribution P_{XYZ} is said to have a* binarization *if for all $\varepsilon > 0$ there exists an integer N and channels $P_{\overline{X}|X^N}$ and $P_{\overline{Y}|Y^N}$ with $\overline{X}, \overline{Y} \in \{0, 1, \perp\}$ such that, if E denotes the event that $\overline{X} \neq \perp \neq \overline{Y}$ and E' denotes the event that $\overline{X} = \overline{Y} \neq \perp$, then*

- $\Pr[E] > 0$;
- $\Pr[E'|E] > 1 - \varepsilon$;
- $\Pr[\overline{X} = 0|E'] = \Pr[\overline{X} = 1|E'] = 1/2$;
- $H(\overline{X}|Z^N, E) > 1 - \varepsilon$.

The following lemma, proved in [12], showcases why binarizations may be relevant to settle the existence of bound information.

Lemma 4.28 ([12, Proposition 5]) *A distribution P_{XYZ} satisfies $S(X;Y||Z) > 0$ if and only if it has a binarization.*

Therefore, a possible way of proving that our distribution of interest has zero secret-key rate is to prove that it does not have a binarization. A plausible way of achieving this would be to show that any binarization for P_{XYZ} could be transformed into a binarization for $P_{X'Y'Z'}$, which we know does not have a binarization for large a by Theorem 4.24. Intuitively, such a binarization cannot depend much on something other than U_X and U_Y. If it did, then ensuring that $\overline{X}$ and $\overline{Y}$ agree with high probability whenever E happens should cause Eve's uncertainty about $\overline{X}$ to be relatively small.

Another possibility is to attempt to relate the existence of binarizations to the problem of *non-interactive correlation distillation* [28][47]. We will see that known results about non-interactive correlation distillation slightly improve our understanding of binarizations for a wide class of distributions, and point towards possibly deeper connections between the two topics.

Suppose Alice and Bob receive i.i.d. realizations X^N and Y^N of random variables X and Y which are arbitrarily jointly distributed, for some integer N. Intuitively, non-interactive correlation distillation consists in Alice and Bob applying (randomized) functions to X^N and Y^N so that they obtain bits $\overline{X}$ and $\overline{Y}$ that are close to uniform and which coincide with high probability. More precisely, we have the following definition.

Definition 4.29 ([47, Adaptation of Definitions 2, 4, and 5]) *Given random variables X and Y with respective ranges $\mathcal{X}$ and $\mathcal{Y}$, a (c, δ)-*protocol for non-interactive

correlation distillation *consists of a sequence of randomized functions* $\varphi_N^A : \mathcal{X}^N \to \{0,1\}$ *and* $\varphi_N^B : \mathcal{Y}^N \to \{0,1\}$ *such that, if* $\overline{X}_N$ *(resp.* $\overline{Y}_N$*) is the output of* $\varphi_N^A(X^N)$ *(resp.* $\varphi_N^B(Y^N)$*), then*

$$\liminf_N \{2\Pr[\overline{X}_N = \overline{Y}_N] - 1\} = c,$$

and

$$|\Pr[\overline{X}_N = 1] - 1/2| \leq \delta, \quad |\Pr[\overline{Y}_N = 1] - 1/2| \leq \delta$$

for large enough N*.*

The goal is to find protocols where c is arbitrarily close to 1 and δ is arbitrarily close to 0, or prove that such protocols cannot exist. Yang [47] showed a trade-off between c and δ for so-called *regular* distributions P_{XY}. For a given distribution P_{XY} where X and Y are distributed over the same finite range $\{1,\ldots,n\}$, we associate with it the matrix $M(P_{XY})$, where $M_{ij} = P_{XY}(i,j)$. We have the following definition.

Definition 4.30 ([47, Definition 7]) *A finite distribution* P_{XY} *where* X *and* Y *have the same range is said to be* regular *if* $M(P_{XY})$ *is symmetric and its largest absolute eigenvalue is 1 with corresponding unique eigenvector* $(1,\ldots,1)$*.*

Theorem 4.31 ([47, Theorem 1]) *If* P_{XY} *is a regular distribution where* X *and* Y *have range* $\{1,\ldots,n\}$*, and such that the difference between the two largest absolute eigenvalues of* $M(P_{XY})$ *is* $g \cdot n$*, then for any* (c,δ)*-protocol we must have*

$$c \leq 1 - g(1 - 4\delta^2).$$

Suppose P_{XYZ} has a binarization where, for large enough N, we have

$$\Pr[\overline{X} = \perp] = \Pr[\overline{Y} = \perp] = 0,$$

i.e. Alice and Bob never abort. Then, one can see this binarization as a $(1,\delta)$-protocol for non-interactive correlation distillation for any $\delta > 0$. In fact, it is something much stronger, since it has to satisfy an additional secrecy requirement. Theorem 4.31 implies that such a protocol cannot exist when P_{XY} is regular, since if δ is small enough, then c must be bounded away from 1. The corollary below follows immediately from this discussion.

Corollary 4.32 *There exists no binarization for* P_{XYZ} *that satisfies*

$$\Pr[\overline{X} = \perp] = \Pr[\overline{Y} = \perp] = 0$$

for large enough N *if* P_{XY} *is a regular distribution.*

Consider the candidate distribution of this section, whose marginal P_{XY} can be found in Table 4.1. It is easy to see the following.

Lemma 4.33 *The distribution P_{XY} in Table 4.1 is regular for all $a \geq 0$.*

We have the following corollary, which limits the useful binarizations for our candidate distribution.

Corollary 4.34 *The candidate distribution of this section (Table 4.1) has no binarization which satisfies*

$$\Pr[\overline{X} = \perp] = \Pr[\overline{Y} = \perp] = 0$$

for large enough N.

Note that Corollary 4.34 provides only very slight evidence that the candidate distribution has bound information for large enough a. In fact, there are many regular distributions which nevertheless have positive secret-key rate, such as distributions arising from the satellite setting of Chapter 3, or the candidate distribution with $a = 0$. Nevertheless, the regularity of the candidate distribution, together with the other negative results of this section, does seem to point that its structure makes it especially ammenable to an impossibility argument for secret-key agreement, just like it was done for non-interactive correlation distillation.

The main difference between a binarization and a protocol for non-interactive correlation distillation is that Alice and Bob have the freedom to abort the protocol if necessary. Nevertheless, it may be possible to translate some techniques used to prove impossibility results in non-interactive correlation distillation (namely, techniques from the analysis of boolean functions) to the binarization setting.

In Section 4.4, we investigate ways of relaxing the definition of the secret-key rate, their respective notions of bound information, and separations between the various rates.

4.4 Relaxing bound information

Proving the existence of bound information seems to be currently a daunting task. A natural way to proceed is by relaxing the notion of bound information. For example, one can restrict the types of protocols allowed for Alice and Bob, or change the setting slightly, and attempt to prove that the new secret-key rate is zero while the intrinsic mutual information is positive, which should be easier to tackle.

For example, one could restrict protocols to have one-way communication between Alice and Bob. This yields the one-way secret-key rate $S_{\text{ow}}(X;Y||Z)$ of Ahlswede and Csiszár [2] we have seen before. It is easy to see that bound information exists with respect to S_{ow}, i.e. there exist distributions P_{XYZ} such that $S_{\text{ow}}(X;Y||Z) = 0$, but $I(X;Y \downarrow Z) > 0$. In fact, something stronger is

known to hold: There exist distributions P_{XYZ} such that $S_{\text{ow}}(X;Y||Z) = 0$, but $S(X;Y||Z) > 0$, which yields a sharp separation between these two notions of secret-key rate. Examples of such distributions arise, for example, from the satellite setting of Chapter 3 (see Theorem 3.1).

Another separation was recently studied by Ozols, Smith, and Smolin [30]. The authors showcase a separation between the usual secret-key rate and the so-called *secret-key rate by public discussion*. A more precise definition follows below.

Definition 4.35 ([30, Section I], rewritten) *The* secret-key rate by public discussion of X and Y given Z, *denoted by* $S_{\text{PD}}(X;Y||Z)$, *is defined like the secret-key rate, except that Alice and Bob are restricted to protocols where every auxiliary random variable used by Alice and Bob is made public, i.e. Alice's secret-key is a deterministic function of* X^N *and the communication of the protocol* C, *and likewise for Bob with* Y^N *in place of* X^N.

We have the following theorem.

Theorem 4.36 ([30, Theorem 2]) *There exist distributions* P_{XYZ} *such that* $S(X;Y||Z) > 0$ *and* $S_{\text{PD}}(X;Y||Z) = 0$ *hold simultaneously.*

We obtain the existence of bound information with respect to S_{PD} as an immediate corollary.

Corollary 4.37 *There exists bound information with respect to* S_{PD}, *i.e. there exists a distribution* P_{XYZ} *such that* $S_{\text{PD}}(X;Y||Z) = 0$ *but* $I_{\text{form}}(X;Y|Z) > 0$.

In the remainder of this section, we discuss the techniques used to prove Theorem 4.36, and then investigate interesting concepts stemming from their limitations. We finish this section by showing that the techniques developed in [30] provide yet more evidence that the distribution studied in Section 4.3 has bound information for large enough a. The following definition will be important throughout the rest of the section.

Definition 4.38 ([30, Section II]) *A distribution* P_{XYZ} *is said to be* unambiguous *if each of* X, Y, *and* Z *is a deterministic function of the other two random variables.*

Recall the concepts introduced in Section 2.5. The strategy used in [30] is as follows. To each finite distribution P_{XYZ}, we can associate the following pure quantum state, $|\varphi_{XYZ}\rangle$, defined as

$$|\varphi_{XYZ}\rangle := \sum_{x,y,z} \sqrt{P_{XYZ}(x,y,z)}|xyz\rangle.$$

Let ρ_{XY} denote the trace over Eve's environment of $|\varphi_{XYZ}\rangle$.

The main theorem of [30] is the following.

Theorem 4.39 ([30, Theorem 1]) *If P_{XYZ} is an unambiguous distribution, then*

$$D(|\varphi_{XYZ}\rangle) \geq S_{\mathrm{PD}}(X;Y||Z).$$

The bound of Theorem 4.39 can then be used to prove Theorem 4.36 by exhibiting unambiguous distributions P_{XYZ} such that $D(|\varphi_{XYZ}\rangle) = 0$, while $S(X;Y||Z) > 0$.

A natural question follows: How far can this technique take us? In other words, is there a class of distributions for which such a technique cannot possibly succeed in presenting a separation between S and S_{PD}? Obviously, Theorem 4.39 does not tell us anything about distributions which are not unambiguous. However, one can find such a limitation even among unambiguous distributions. The next lemma states that the technique of [30] does not work for unambiguous distributions where X and Y are binary random variables.

Lemma 4.40 *If P_{XYZ} is an unambiguous distribution with $X, Y \in \{0,1\}$ such that $S(X;Y||Z) > 0$, then $D(|\varphi_{XYZ}\rangle) > 0$.*

Proof Fix a distribution P_{XYZ} satisfying the hypotheses of the lemma. By Theorem 2.19, it suffices to show that ρ_{XY} is entangled. Since P_{XYZ} is an unambiguous distribution, we must have $|\mathcal{Z}| \leq 4$. Moreover, since its secret-key rate is positive, it must hold that $|\mathcal{Z}| \leq 3$. Let z_{ij} be the value of Z determined by $X = i$ and $Y = j$. Furthermore, let $p_{ij} := P_{XYZ}(i,j,z(i,j))$. We consider three cases:

1. $\mathcal{Z} = \{0\}$:

   Since P_{XYZ} is unambiguous and $S(X;Y||Z) > 0$, we can assume without loss of generality that $p_{00}, p_{11} > 0$ and $p_{01} = p_{10} = 0$. Let ρ_{XY} be the trace over the environment of $|\varphi_{XYZ}\rangle$. It can be seen that the partial transpose of ρ_{XY} with respect to Bob is

   $$\rho_{XY}^{\top_B} = \begin{bmatrix} p_{00} & 0 & 0 & 0 \\ 0 & 0 & \sqrt{p_{00}p_{11}} & 0 \\ 0 & \sqrt{p_{00}p_{11}} & 0 & 0 \\ 0 & 0 & 0 & p_{11} \end{bmatrix}.$$

   Note that $\rho_{XY}^{\top_B}$ has eigenvalue $-\sqrt{p_{00}p_{11}} < 0$. By Theorem 2.19, it follows that $D(|\varphi_{XYZ}\rangle) > 0$.

2. $\mathcal{Z} = \{0,1\}$:

   Since P_{XYZ} is unambiguous and $S(X;Y||Z) > 0$, we can assume without loss of generality that $p_{00}, p_{01}, p_{10} > 0$, and that $z_{00} = z_{11} = 0$ and

$z_{01} = z_{10} = 1$. Then, the partial transpose of ρ_{XY} with respect to Bob is

$$\rho_{XY}^{\top_B} = \begin{bmatrix} p_{00} & 0 & 0 & \sqrt{p_{01}p_{10}} \\ 0 & p_{01} & \sqrt{p_{00}p_{11}} & 0 \\ 0 & \sqrt{p_{00}p_{11}} & p_{10} & 0 \\ \sqrt{p_{01}p_{10}} & 0 & 0 & p_{11} \end{bmatrix}.$$

We now make use of Theorem 2.19. The partial transpose above can be seen to have eigenvalues

$$\lambda_1 = \frac{1}{2}\left(p_{01} + p_{10} - \sqrt{p_{01}^2 - 2p_{01}p_{10} + 4p_{00}p_{11} + p_{10}^2}\right),$$

and

$$\lambda_2 = \frac{1}{2}\left(p_{00} + p_{11} - \sqrt{p_{00}^2 - 2p_{00}p_{11} + 4p_{01}p_{10} + p_{11}^2}\right).$$

It follows that the partial transpose has a negative eigenvalue, and thus $D(|\varphi_{XYZ}\rangle) > 0$, whenever

$$p_{01}p_{10} \neq p_{00}p_{11}.$$

On the other hand, if $p_{01}p_{10} = p_{00}p_{11}$, then X and Y are independent, and so $S(X;Y||Z) \leq I(X;Y) = 0$.

3. $\mathcal{Z} = \{0,1,2\}$:

In this case we can assume without loss of generality that $p_{ij} > 0$ for all pairs (i,j), and that $z_{00} = z_{11} = 0$, $z_{01} = 1$, $z_{10} = 2$. Then, the partial transpose of ρ_{XY} with respect to Bob is

$$\rho_{XY}^{\top_B} = \begin{bmatrix} p_{00} & 0 & 0 & 0 \\ 0 & p_{01} & \sqrt{p_{00}p_{11}} & 0 \\ 0 & \sqrt{p_{00}p_{11}} & p_{10} & 0 \\ 0 & 0 & 0 & p_{11} \end{bmatrix},$$

which has eigenvalue

$$\lambda = \frac{1}{2}\left(p_{01} + p_{10} - \sqrt{p_{01}^2 - 2p_{01}p_{10} + 4p_{00}p_{11} + p_{10}^2}\right),$$

along with three more non-negative eigenvalues. By Theorem 2.19, it follows that $D(|\varphi_{XYZ}\rangle) > 0$ whenever $p_{00}p_{11} > p_{01}p_{10}$. If $p_{00}p_{11} = p_{01}p_{10}$, then X and Y are independent, and so $S(X;Y||Z) = 0$. It remains to consider the case $p_{00}p_{11} < p_{01}p_{10}$. The reasoning used in the proof of Theorem 4.24 can be used to show that, in this case, one has $I(X;Y \downarrow Z) = 0$, and so $S(X;Y||Z) = 0$. □

Lemma 4.40 implies that a separation between S and S_{PD} is still unknown for unambiguous distributions with binary X and Y, and that the technique previously discussed cannot help find such a separation directly. Therefore, S_{PD} seems more difficult to understand in this case. In view of this, we will introduce another notion of secret-key rate, directly inspired by S_{PD}. Then, we will see that this notion can be further simplified, and we will investigate possible techniques for showcasing a separation with respect to the secret-key rate S. We also present a candidate witness distribution with binary X and Y for this separation.

As we have previously seen, S_{PD} can be thought of as the secret-key rate restricted to protocols where every auxiliary random variable used to compute the secret-key (besides X^N and Y^N) must be made public. Note that this does not mean that the local randomness used by Alice and Bob must be made available to Eve, but rather that the secret-keys only depend on the local randomness through the realizations received and the communication of the protocol. It is thus natural to consider what happens if one forces Alice's and Bob's local randomness to be public. This new setting can be modelled by allowing Alice, Bob, and Eve to access a *common reference string*, which consists of an arbitrarily long sequence of i.i.d. uniformly random bits, and enforcing that all local randomness used in Alice's and Bob's protocol must be extracted from it. More precisely, we have the following definitions.

Definition 4.41 *A* common reference string *(CRS) is a randomized function with input alphabet* $\mathbb{N}$ *and output alphabet* $\{0,1\}$ *which works as follows: If it receives a query* $i \in \mathbb{N}$ *for the first time, it samples a bit uniformly at random and independently of the outputs of other queries, and outputs the bit. If it receives a repeated query* $i \in \mathbb{N}$*, then it outputs the same bit that was originally sampled to answer the first such query.*

Definition 4.42 *The* secret-key rate with a common reference string of X and Y given Z*, denoted by* $S_{\text{CRS}}(X;Y||Z)$*, is defined like the secret-key rate, except that Alice and Bob are restricted to deterministic protocols, i.e. protocols where* $H(R_A) = H(R_B) = 0$*, and furthermore Alice, Bob, and Eve have access to the same common reference string.*

It is clear that $S_{\text{CRS}}(X;Y||Z) \leq S_{\text{PD}}(X;Y||Z)$, since a protocol in the CRS setting can be used in the PD setting by having Alice and Bob make their local randomness public before starting the protocol. In fact, we will show that in the CRS setting we only need to consider a much simpler class of protocols that do not use any local randomness (private or public), which we will see makes S_{CRS} a potentially much more tractable quantity.

Definition 4.43 ([13, Definitions 1 and 2]) *The* deterministic secret-key rate of X and Y given Z*, denoted by* $S_{\text{det}}(X;Y||Z)$*, is defined like the secret-key rate,*

except that Alice and Bob are restricted to deterministic protocols, i.e. protocols where Alice and Bob do not use any local randomness ($H(R_A) = H(R_B) = 0$).

The following theorem shows that the common reference string is useless.

Theorem 4.44 *We have*

$$S_{\text{CRS}}(X;Y||Z) = S_{\text{det}}(X;Y||Z)$$

for all distributions P_{XYZ}.

Proof This proof is very similar to the first part of the proof of [12, Proposition 5].

It is clear that

$$S_{\text{det}}(X;Y||Z) \leq S_{\text{CRS}}(X;Y||Z)$$

holds, so it suffices to prove the opposite inequality.

We can assume without loss of generality that, in a given protocol for fixed N, Alice and Bob access exactly the first ℓ bits of the common reference string, for some ℓ independent of X^N and Y^N, but which depends on N and the protocol in question. The reason for this is the following: Let R denote the bits of the common reference string which were accessed by Alice or Bob. We can assume that if the i-th bit was accessed, then so was the j-th bit, for $j < i$ (the rate is not affected if Alice and Bob access bits they will not use). Fix a protocol for some N and $\varepsilon > 0$. Then, for any $1/2 > \delta > 0$, there is ℓ large enough such that

$$\Pr[|R| > \ell] < \delta.$$

Consider the modified protocol where Alice and Bob both initially access the first ℓ bits of the common reference string, and abort (i.e. they set $S_A = S_B = \perp$) if any of them needs more than the first ℓ bits. Let $F := 1_{\{\text{Alice and Bob abort}\}}$, let S'_A and S'_B denote the resulting keys of the modified protocol, and let C' denote the communication of the modified protocol. Then

$$\Pr[F = 1] < \delta.$$

Note that

$$\begin{aligned} H(S'_A) \geq H(S'_A|F) = \Pr[F=0]H(S_A|F=0) = \\ = H(S_A|F) - \Pr[F=1]H(S_A|F=1) \geq H(S_A) - H(F) - \delta \log|\mathcal{S}| \geq \\ \geq H(S_A) - h(\delta) - \delta \log|\mathcal{S}|, \end{aligned}$$

where the first inequality follows because conditioning reduces entropy, the first equality follows because $H(S'_A|F=0) = H(S_A|F=0)$, the second inequality follows because $H(S_A|F) = H(S_AF) - H(F) \geq H(S_A) - H(F)$

and because $\Pr[F = 1]H(S_A|F = 1) \leq \delta \log |\mathcal{S}|$, and the third inequality follows because $H(F) \leq h(\delta)$. Furthermore,

$$H(S'_A|Z^N C'F) = \Pr[F = 0]H(S_A|Z^N C, F = 0) =$$
$$= H(S_A|Z^N CF) - \Pr[F = 1]H(S_A|Z^N C, F = 1) \geq H(S_A|Z^N C) - h(\delta) - \delta \log |\mathcal{S}|,$$

following the same reasoning as before. Finally,

$$\Pr[S'_A = S'_B] = \Pr[|R| > \ell] + \sum_{n \leq \ell} \Pr[|R| = n] \Pr[S'_A = S'_B||R| = n] \geq$$
$$\geq \sum_n \Pr[|R| = n] \Pr[S_A = S_B||R| = n] = \Pr[S_A = S_B],$$

because $\Pr[S'_A = S'_B||R| = n] = \Pr[S_A = S_B||R| = n]$ for $n \leq \ell$.

Therefore, making ℓ arbitrarily large (and thus δ arbitrarily close to 0) allows us to obtain a secure protocol where $R \in \{0,1\}^\ell$ with rate arbitrarily close to that of the original protocol. It follows that Alice and Bob can be assumed to access only at most a fixed number of bits ℓ in each protocol.

Suppose there is a protocol for N realizations in the CRS setting, where Alice, Bob, and Eve share a string R consisting of ℓ i.i.d. uniformly random bits, with communication C, and after which Alice and Bob hold random variables S_A and S_B with finite range $|\mathcal{S}|$ satisfying

$$\Pr[S_A \neq S_B] < \varepsilon$$

and

$$H(S_A|Z^N CR) > \log |\mathcal{S}| - \varepsilon$$

for some $\varepsilon > 0$.

Consider the functions $e : \{0,1\}^\ell \to [0,1]$ defined by

$$e(r) := \Pr[S_A \neq S_B|R = r],$$

and $k : \{0,1\}^\ell \to [0, \log |\mathcal{S}|]$ defined by

$$k(r) := \log |\mathcal{S}| - H(S_A|Z^N C, R = r).$$

Note that

$$\mathrm{E}_R[e(R)] = \Pr[S_A \neq S_B] < \varepsilon,$$

and

$$\mathrm{E}_R[k(R)] = \log |\mathcal{S}| - H(S_A|Z^N CR) < \varepsilon.$$

It then follows that

$$\Pr[e(R) < 2\varepsilon] > 1/2 \quad \text{and} \quad \Pr[k(R) < 2\varepsilon] > 1/2,$$

since, if $\Pr[e(R) < 2\varepsilon] \leq 1/2$, then $\Pr[e(R) \geq 2\varepsilon] \geq 1/2$, and so

$$\mathrm{E}_R[e(R)] \geq 1/2 \cdot 2\varepsilon = \varepsilon,$$

which is a contradiction, and analogously for $k(R)$. Therefore, there must exist $r^* \in \{0,1\}^\ell$ in the range of R such that $e(r^*) < 2\varepsilon$ and $k(r^*) < 2\varepsilon$, since otherwise we would have

$$\Pr[e(R) < 2\varepsilon \text{ or } k(R) < 2\varepsilon] = \Pr[e(R) < 2\varepsilon] + \Pr[k(R) < 2\varepsilon] > 1.$$

Consider now the deterministic protocol where Alice and Bob run the protocol above with $R = r^*$. Suppose they obtain secret-keys S'_A and S'_B, respectively, and communication C'. Then S'_A and S'_B have range contained in $\mathcal{S}$, and furthermore

$$\Pr[S'_A \neq S'_B] = \Pr[S_A \neq S_B | R = r^*] = e(r^*) < 2\varepsilon,$$

and

$$H(S'_A|Z^N C') = H(S_A|Z^N C, R = r^*) = \log|\mathcal{S}| - k(r^*) > \log|\mathcal{S}| - 2\varepsilon.$$

Thus, we also have

$$H(S'_A) = H(S_A|R = r^*) \geq H(S_A|Z^N C, R = r^*) > \log|\mathcal{S}| - 2\varepsilon.$$

Therefore, since the choices of protocol and $\varepsilon > 0$ were arbitrary, and since the deterministic protocol achieves the same asymptotic rate as the original one in the CRS setting, it follows that

$$S_{\mathrm{det}}(X;Y||Z) \geq S_{\mathrm{CRS}}(X;Y||Z),$$

which concludes the proof. □

A first reason why the equality showcased in Theorem 4.44 is useful for analyzing the corresponding separation is that Gohari and Anantharam [13] proved a connection between the deterministic secret-key rate and the problem of *communication for omniscience by a neutral observer*, which is a generalization of an earlier problem studied by Csiszár and Narayan [7]. Intuitively, the goal of a protocol for communication for omniscience by a neutral observer is for Alice and Bob to agree with high probability on a random variable T_A which allows Eve, who knows Z^N, to have very low uncertainty about $X^N Y^N$.

Definition 4.45 ([13, Definition 3]) *The* communication for omniscience capacity of X and Y by a neutral observer Z, *denoted by* $T(X;Y||Z)$, *is the infimum of all real numbers* $R \geq 0$ *such that for every* $\varepsilon > 0$ *and large enough* N *there exists a deterministic protocol for Alice and Bob, who start with* X^N *and* Y^N, *respectively, with communication* C, *such that* $R = H(C|Z^N)/N$, *and, at the end*

of the protocol, Alice and Bob end up with random variables T_A and T_B, respectively, satisfying

$$\Pr[T_A = T_B] \geq 1 - \varepsilon$$

and

$$H(X^N Y^N | Z^N T_A) \leq \varepsilon N.$$

Gohari and Anantharam proved the following connection between the deterministic secret-key rate and the problem of communication for omniscience by a neutral observer.

Theorem 4.46 ([13, Theorem 3]) *We have*

$$S_{\text{det}}(X;Y||Z) \leq H(XY|Z) - T(X;Y||Z)$$

for all finite distributions P_{XYZ}.

Therefore, a way to prove that $S_{\text{det}}(X;Y||Z) = 0$ is to alternatively prove that $T(X;Y||Z) = H(XY|Z)$. Note that a rate of $H(XY|Z)$ can be easily achieved by having Alice and Bob simply publish X^N and Y^N, and set $T_A = T_B = X^N Y^N$. This alternative characterization of $S_{\text{det}}(X;Y||Z)$ has the advantage of eliminating the secrecy constraint. By following the proof of [13, Theorem 2] with $H(XY|Z)$ in place of the function ψ, we obtain the next lemma, which yields a sufficient condition for $T(X;Y||Z) = H(XY|Z)$ to hold.

Lemma 4.47 *If for small enough $\varepsilon > 0$ and large enough N every protocol for omniscience with communication C satisfies*

$$H(T_A|Z^N C) < N \cdot f(N,\varepsilon),$$

or, equivalently,

$$H(X^N Y^N | Z^N C) < N \cdot f(N,\varepsilon),$$

where $f(N,\varepsilon) \to 0$ whenever $N \to \infty$ and $\varepsilon \to 0$, then $T(X;Y||Z) = H(XY|Z)$.

Intuitively, Lemma 4.47 states that in order to verify that $T(X;Y||Z) = H(XY|Z)$, and hence that $S_{\text{det}}(X;Y||Z) = 0$, it suffices to prove that the communication between Alice and Bob in any successful protocol for omniscience is already enough for Eve to reconstruct most of $X^N Y^N$ with high probability.

Recall that binarizations (Definition 4.27) are thought to be a potentially viable way of proving the existence of bound information, mostly due to the fact that they eliminate the need to consider the communication of protocols for secret-key agreement. Interestingly, binarizations for deterministic protocols have a nice, simpler structure.

Definition 4.48 *Recall the definition of binarization (Definition 4.27). A binarization is said to be* deterministic *if X^N (resp. Y^N) completely determines $\overline{X}$ (resp. $\overline{Y}$). Moreover, a binarization is said to be* semi-deterministic *if X^N (resp. Y^N) completely determines $\overline{X}$ (resp. $\overline{Y}$) when conditioned on $\overline{X} \neq \perp$ (resp. $\overline{Y} \neq \perp$).*

The next lemma is obtained by following the proof of [12, Proposition 5] with our restricted types of protocols.

Lemma 4.49 *It holds that $S_{\text{det}}(X;Y||Z) > 0$ if and only if there exists a deterministic binarization for P_{XYZ}. Furthermore, $S_{\text{PD}}(X;Y||Z) > 0$ holds if and only if there exists a semi-deterministic binarization for P_{XYZ}.*

It follows that, for S_{det}, one may replace the channels $P_{\overline{X}|X^N}$ and $P_{\overline{Y}|Y^N}$ by deterministic functions of X^N and Y^N, respectively. Therefore, the problem of proving that the deterministic secret-key rate is positive is reduced to a problem which is almost purely combinatorial. Another advantage of studying deterministic binarizations is that this equivalence between binarizations and the secret-key rate holds at the most general level, which is not true for the other techniques. Therefore, techniques used to show the non-existence of deterministic binarizations can potentially pave the way for the more general case. It is also interesting to note that the S, S_{PD}, S_{det} hierarchy of secret-key rate definitions implies a natural hierarchy of notions of binarizations.

Recall that Lemma 4.40 states that the techniques of [30] cannot be applied to unambiguous distributions where X and Y are binary. We now present a class of distributions P_{XYZ}, where X and Y are binary, such that $S(X;Y||Z) > 0$ always holds, but for which it is plausible that $S_{\text{det}}(X;Y||Z) = 0$. These distributions are obtained by restricting the range of X and Y of the distribution in [12, Example 3]. The precise definition can be found in Table 4.4.

		X	
		0	1
Y	0	2 [0]	α [2]
	1	$5-\alpha$ [1]	2 [0]

Table 4.4: Class of candidate unambiguous distributions for a separation between S_{det} and S when X and Y are binary, obtained by restricting the range of X and Y in the class of distributions of Example 3 in [12]. We consider $\alpha \in [0,1)$. Note that each entry in the table must be normalized by 9. If the entry in row i, column j is p [k], then this means that $\Pr[X = j, Y = i, Z = k] = p$.

It is clear that the class of distributions defined in Table 4.4 is unambiguous. Furthermore, it can be shown that, for $\alpha \in [0,1)$, it holds that $S(X;Y||Z) > 0$. The secret-key agreement strategy in this case is the following: Alice and

Bob transform each realization of X and Y through appropriate channels $P_{\overline{X}|X}$ and $P_{\overline{Y}|Y}$, and then run an advantage distillation protocol on the outputs of the channels, in order to obtain a positive secret-key rate. Note that, since the strategy involves sending X and Y through channels, it requires local randomness, and therefore it is not a deterministic protocol. For $\alpha \geq 1$, the reasoning in the proof of Theorem 4.24 can be used to prove that $I(X;Y \downarrow Z) = 0$. We put forth the following conjecture.

Conjecture 4.50 *For $\alpha < 1$, the class of distributions defined in Table 4.4 satisfies $S_{\text{det}}(X;Y||Z) = 0$.*

We will focus on the case where $\alpha = 0$, as it entails a separation between S and S_{det} for distributions where all of X, Y, and Z are binary random variables. We provide an informal argument for why we believe that $T(X;Y||Z) = H(XY|Z)$ holds for this distribution. Given Z^N, Eve knows the positions where $X = Y$ and $X \neq Y$. Furthermore, if $X \neq Y$, then she knows that $(X,Y) = (1,0)$. However, for positions such that $X = Y$, she cannot predict whether $(X,Y) = (0,0)$ or $(X,Y) = (1,1)$, since both cases are equally likely. Therefore, in a communication for omniscience protocol, T_A must contain information about the positions of most entries satisfying $(X,Y) = (0,0)$ (at least relative to the entries where $(X,Y) = (1,1)$), or vice-versa. Suppose we have a protocol where T_A contains information about the location of most of the $(0,0)$-entries. Clearly, Alice can create T_A without communicating with Bob, since $X = 0$ implies $Y = 0$. Nevertheless, it must be the case that T_A and T_B agree with high probability. For most realizations $X^N Y^N$, Bob cannot predict the relative location of the $(0,0)$-entries, although he knows the positions of the $(1,1)$-entries. Therefore, in order for T_A and T_B to agree with high probability, the communication C between Alice and Bob must contain, in some way, most positions of the $(0,0)$-entries. Ideally, Z^N and C would then jointly determine most of $X^N Y^N$, and so the result would be a consequence of Lemma 4.47.

We believe that attempting to answer Conjecture 4.50, either through the problem of communication for omniscience or through deterministic binarizations, will give rise to interesting techniques and concepts which might pave the way to a negative answer to the positivity conjecture.

Lastly, we note that Theorem 4.39 provides yet more evidence that the distribution studied in Section 4.3 has bound information. Clearly, the distribution in question is unambiguous. We can compute the trace over Eve's environment ρ_{XY} of its associated quantum state and its partial transpose $\rho_{XY}^{\top_B}$, and check that the only eigenvalue of $\rho_{XY}^{\top_B}$ that can be negative is

$$\lambda = \frac{8a - \sqrt{2}}{8(8a+1)},$$

which is non-negative whenever $a \geq \frac{1}{4\sqrt{2}}$. We have the following theorem, which complements the results obtained about this distribution in Section 4.3.

Theorem 4.51 *The candidate for bound information P_{XYZ} studied in Section 4.3 satisfies $S_{\mathrm{PD}}(X;Y||Z) = 0$ for $a \geq \frac{1}{4\sqrt{2}}$.*

Bibliography

[1] Antonio Acín, Juan Cirac, and Lluis Masanes. Multipartite bound information exists and can be activated. *Physical Review Letters*, 92(10):107903, 2004.

[2] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography. i. secret sharing. *IEEE Transactions on Information Theory*, 39(4):1121–1132, 1993.

[3] Matthieu Bloch and João Barros. *Physical-layer security: from information theory to security engineering*. Cambridge University Press, 2011.

[4] Chris Calabro. *The Exponential Complexity of Satisfiability Problems*. PhD thesis, University of California, San Diego, 2009.

[5] Matthias Christandl, Renato Renner, and Stefan Wolf. A property of the intrinsic mutual information. In *Proceedings of 2003 IEEE International Symposium on Information Theory*, page 258. IEEE, 2003.

[6] Thomas Cover and Joy Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.

[7] Imre Csiszár and Prakash Narayan. Secrecy capacities for multiple terminals. *IEEE Transactions on Information Theory*, 50(12):3047–3061, 2004.

[8] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

[9] Abbas El Gamal and Young-Han Kim. *Network information theory*. Cambridge University Press, 2011.

[10] Willliam Feller. *An introduction to probability theory and its applications*, volume 1. John Wiley & Sons, 2008.

[11] Martin Gander and Ueli Maurer. On the secret-key rate of binary random variables. In *Proceedings of 1994 IEEE International Symposium on Information Theory*, page 351. IEEE, 1994.

[12] Nicolas Gisin, Renato Renner, and Stefan Wolf. Linking classical and quantum key agreement: Is there a classical analog to bound entanglement? *Algorithmica*, 34(4):389–412, 2002.

[13] Amin Gohari and Venkat Anantharam. Information-theoretic key agreement of multiple terminals—part i. *IEEE Transactions on Information Theory*, 56(8):3973–3996, 2010.

[14] Robert Gray. *Entropy and information theory*. Springer Science & Business Media, 2011.

[15] Siu-Wai Ho and Raymond Yeung. The interplay between entropy and variational distance. *IEEE Transactions on Information Theory*, 56(12):5906–5929, 2010.

[16] Karol Horodecki, Michał Horodecki, Paweł Horodecki, and Jonathan Oppenheim. Secure key from bound entanglement. *Physical Review Letters*, 94(16):160502, 2005.

[17] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Separability of mixed states: necessary and sufficient conditions. *Physics Letters A*, 223(1):1–8, 1996.

[18] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Inseparable two spin-1/2 density matrices can be distilled to a singlet form. *Physical Review Letters*, 78(4):574, 1997.

[19] Paweł Horodecki. Separability criterion and inseparable mixed states with positive partial transposition. *Physics Letters A*, 232(5):333–339, 1997.

[20] Shengli Liu, Henk Van Tilborg, and Marten Van Dijk. A practical protocol for advantage distillation and information reconciliation. *Designs, Codes and Cryptography*, 30(1):39–62, 2003.

[21] Ueli Maurer. Protocols for secret key agreement by public discussion based on common information. In *Annual International Cryptology Conference*, pages 461–470. Springer, 1992.

[22] Ueli Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, 1993.

[23] Ueli Maurer. The strong secret key rate of discrete random triples. In *Communication and Cryptography — Two Sides of One Tapestry*, pages 271–285. Kluwer Academic Publishers, 1994.

[24] Ueli Maurer. Information-theoretic cryptography. In *Annual International Cryptology Conference*, pages 47–65. Springer, 1999.

[25] Ueli Maurer and Stefan Wolf. Towards characterizing when information-theoretic secret key agreement is possible. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 196–209. Springer, 1996.

[26] Ueli Maurer and Stefan Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Transactions on Information Theory*, 45(2):499–514, 1999.

[27] Ueli Maurer and Stefan Wolf. Information-theoretic key agreement: From weak to strong secrecy for free. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 351–368. Springer, 2000.

[28] Elchanan Mossel and Ryan O'Donnell. Coin flipping from a cosmic source: On error correction of truly random bits. *Random Structures & Algorithms*, 26(4):418–436, 2005.

[29] Michael Nielsen and Isaac Chuang. *Quantum computation and quantum information*. AAPT, 2002.

[30] Maris Ozols, Graeme Smith, and John Smolin. Bound entangled states with a private key and their classical counterpart. *Physical Review Letters*, 112(11):110502, 2014.

[31] Asher Peres. Separability criterion for density matrices. *Physical Review Letters*, 77(8):1413, 1996.

[32] Mark Pinsker. *Information and information stability of random variables and processes*. Holden-Day, 1964. Translated and edited by Amiel Feinstein.

[33] H Vincent Poor and Rafael Schaefer. Wireless physical layer security. *Proceedings of the National Academy of Sciences*, 114(1):19–26, 2017.

[34] Giuseppe Prettico and Joonwoo Bae. Superactivation, unlockability, and secrecy distribution of bound information. *Physical Review A*, 83(4):042336, 2011.

[35] Renato Renner, Juraj Skripsky, and Stefan Wolf. A new measure for conditional mutual information and its properties. In *Proceedings of 2003 IEEE International Symposium on Information Theory*, page 259. IEEE, 2003.

[36] Renato Renner and Stefan Wolf. New bounds in secret-key agreement: The gap between formation and secrecy extraction. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 562–577. Springer, 2003.

[37] Claude Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[38] Claude Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(4):623–666, 1948.

[39] Claude Shannon. Communication theory of secrecy systems. *Bell Labs Technical Journal*, 28(4):656–715, 1949.

[40] Joel Spencer and Laura Florescu. Asymptopia, volume 71 of student mathematical library. *American Mathematical Society, Providence, RI*, 2014.

[41] Himanshu Tyagi and Shun Watanabe. A bound for multiparty secret key agreement and implications for a problem of secure computing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 369–386. Springer, 2014.

[42] Himanshu Tyagi and Shun Watanabe. Converses for secret key agreement and secure computing. *IEEE Transactions on Information Theory*, 61(9):4809–4827, 2015.

[43] Gilbert Vernam. Cipher printing telegraph systems: For secret wire and radio telegraphic communications. *Journal of the AIEE*, 45(2):109–115, 1926.

[44] Stefan Wolf. *Information-theoretically and computationally secure key agreement in cryptography*. PhD thesis, Diss. Techn. Wiss. ETH Zürich, Nr. 13138, 1999. Ref.: Ueli Maurer; Korref.: Claude Crépeau, 1999.

[45] Stefan Wolf. Unconditional security in cryptography. *Lecture notes in computer science*, pages 217–250, 1999.

[46] Aaron Wyner. The wire-tap channel. *Bell Labs Technical Journal*, 54(8):1355–1387, 1975.

[47] Ke Yang. On the (im) possibility of non-interactive correlation distillation. *Theoretical Computer Science*, 382(2):157–166, 2007.