

The Physics of Quantum Information

Complementarity, Uncertainty, and Entanglement

Joseph M. Renes

12 December 2012

Habilitation

TU Darmstadt, Fachbereich Physik

2011

Contents

Preface	1
1 Introduction: What is Quantum Information?	5
1.1 Understanding Classical Information via the Information Game	6
1.2 Complementarity in the Information Game	7
1.3 Entanglement in the Information Game	8
2 Illustrations and Motivations	11
2.1 Teleportation and Superdense Coding	11
2.2 Quantum Error-Correction	12
2.3 Entanglement Distillation	17
2.4 Quantum Key Distribution	18
3 Characterizing Quantum Information	23
3.1 Amplitude and Phase Predictability & Entanglement	23
3.2 Duality & Decoupling	28
3.3 Entropic Characterizations	31
3.4 Secret Keys & Private States	33
4 Processing Quantum Information	37
4.1 Optimal Entanglement Distillation	37
4.2 Optimal State Merging	43
4.3 Secret Key Distillation and Private Communication	45
5 Duality of Protocols	47
5.1 Duality of Privacy Amplification and Information Reconciliation	47
5.2 Different Approaches to Entanglement Distillation	51
5.3 Classical Channel Coding	54
6 Security of Quantum Key Distribution	59
6.1 Notions of Security	59
6.2 Entanglement in Prepare and Measure QKD	60
6.3 Private States in Quantum Key Distribution	64
7 Summary and Outlook	67
Bibliography	69
Included Papers	79

Preface

Complementarity is one of the central mysteries of quantum mechanics. First put forth by Bohr [1, 2, 3], complementarity holds that the attributes of a physical system familiar from classical mechanics do not all simultaneously exist and are not entirely independent of how they are measured. Famously, if the momentum of a particle is known then its position must be unknown, and *vice versa*, a fact encapsulated in Heisenberg’s uncertainty relation $\Delta x \Delta p \geq \hbar/2$ [4]. Even more dramatic is the wave-particle duality encountered in Young’s double-slit experiment, which illustrates the important role of observation. Light passing through the double slit setup produces an interference pattern on a screen beyond the slits, as would be characteristic of a wave. But a closer examination reveals that light arrives in particle-like “packets” at the screen, and the interference pattern only arises as a statistical average of these particle arrival events. This particle picture tempts us to observe which slit the light went through, which we find destroys the interference pattern! Feynman regarded this bizarre phenomena as characteristic of all the seemingly-paradoxical quantum behavior, claiming that the double-slit experiment is “impossible, *absolutely impossible* to describe classically, [and which] has in it the heart of quantum mechanics”, and that “in reality, it contains the *only* mystery” (emphasis original) [5].

The overarching goal of this thesis is to demonstrate that complementarity is also at the heart of quantum information theory, that it allows us to make (some) sense of just what information "quantum information" refers to, and that it is useful in understanding and constructing quantum information processing protocols. The detailed research results which form the basis of these claims are to be found in the included papers, and the aim here is to present an overview comprehensible to a more general audience.¹

As we shall see in Chapter 1, quantum information can heuristically be thought of as a kind of combination of two types of normal “classical” information, specifically, classical information about the result of measuring one of two complementary observables. Due to the uncertainty principle, we can expect both pieces of information are not simultaneously realizable, and indeed the uncertainty principle will play a central quantitative role throughout this work. Particularly relevant will be the entropic uncertainty relation of [RB09] and its generalization in [BCC⁺10], which state that the more that can be known by one party about one observable, the less can be known by another party about a complementary observable. That complementary observables play an important role in quantum information theory is not new to this thesis, and Chapter 2 discusses several fundamental quantum information processing tasks based on their use, such as teleportation and quantum error-correction. This chapter also provides some relevant formal background for the remainder of this work and establishes the notation used herein.

Chapter 3 begins the overview of the new results obtained in the included papers. Here we show that information about complementary observables not only plays an important role, but indeed a central one, and that possession of both complementary pieces of classical information is strictly equivalent to the existence of entanglement between the physical system the information pertains to and the system in which the information is stored. Moreover, the uncertainty principle provides a dual characterization, saying that entanglement between these two systems exists when the “environment”, i.e. any and all other degrees of freedom, has no information about either complementary observable. Both characterizations can be modified to describe secret keys useful in cryptography instead of entangled states. Because Chapter 3 gathers and mixes results from several of the included papers, it is entirely self-contained, whereas subsequent chapters do not go into as much detail.

¹The included papers are referenced in alphabetical style, while references to other works are numeric.

In Chapter 4 we show that this complementary approach is also useful in constructing quantum information processing protocols and understanding why they work. Especially relevant is the process of entanglement distillation, that is, extracting maximal entanglement from an imperfectly-entangled bipartite resource system. The entanglement distillation process can be built up from two instances, one for each of two complementary observables, of a simpler distillation process for classical information called information reconciliation or data compression with side information. Here partial classical correlation between two systems is refined into maximal correlation, and reconciling classical information about two complementary observables. Protocols for entanglement distillation can then be adapted to a large variety of quantum information processing tasks, such as quantum communication over noisy channels or distillation of secret keys.

Chapter 5 extends the duality in characterizing entanglement afforded by the uncertainty principle to two fundamental information processing tasks, the information reconciliation task of establishing correlations with the first party on the one hand, and the task of *removing* all correlations with the second party on the other. The latter is known as privacy amplification, and it turns out that the ability to perform one protocol implies the ability to perform the other in certain circumstances. This duality also implies alternative methods of entanglement distillation, in particular one which proceeds by destroying all classical correlations with the environment that pertain to two complementary observables. We shall also see that information reconciliation and privacy amplification can be combined to enable classical communication over noisy quantum channels.

Finally, Chapter 6 describes the usefulness of this approach to establishing the security of quantum key distribution (QKD). QKD is perhaps the most natural setting in which the uncertainty principle and corresponding issues of complementarity are immediately relevant, as the goal of this protocol is to establish a secret key between two spatially-separated parties, a shared piece of classical information which no one else should know. Since the uncertainty principle can be understood as a limitation on who can know how much about what sorts of information, we shall see that complementarity-based arguments form the basis for the security of QKD protocols. These allow us to increase the security threshold, the maximum amount of tolerable noise, of several protocols beyond the previously-known values.

The following table summarizes which included papers form the basis for the various sections.

Chapter	Section			
1 Introduction	1.1	1.2	1.3	
		[RB09], [BCC ⁺ 10]		
2 Illustrations & Motivations	—			
3 Characterizing Quantum Information	3.1	3.2	3.3	3.4
	[RB08]	[Ren11]	[RB09]	[RB08]
4 Processing Quantum Information	4.1	4.2	4.3	
	[RB08]	[BR09]	[RB08]	
5 Duality of Protocols	5.1	5.2	5.3	
	[Ren11]	new	[RR11]	
6 Security of QKD	6.1	6.2	6.3	
	—	[RG06]	[RS07]	
			[SRS08]	
			[KR08]	

Introduction: What is Quantum Information?

At a stroke, Shannon’s landmark 1948 publication *A Mathematical Theory of Communication* [6] established the field of information theory, laying out the fundamental lines of inquiry and answering some of the important basic questions. The fundamental problem, according to Shannon, “is that of reproducing at one point either exactly or approximately a message selected at another point.” The different points may be different places, in which case we are interested in transmitting messages from one party to another, such as in a telephone conversation, or they could be different times, and the message should be reliably stored, such as on a sheet of paper. The physical systems used to convey the message carry *information*, which is measured by the entropy in units of bits, short for binary digits.¹

The fact that abstract information must always be instantiated in some physical system and that this results in a connection between physics and information theory was stressed by Landauer. He observed this implies that logically irreversible operations, like erasure of information, are therefore physically irreversible and must be driven by a source of energy [8, 9]. This was later used to resolve the paradox of Maxwell’s Demon in which an intelligent being can apparently violate the second law by sorting the molecules of a gas into hot (fast) and cold (slow) [10]. Building on Szilárd’s simplification of the paradox to a one-atom gas occupying either the left or right side of a divided container [11], Bennett showed that the work gained by the demon is precisely balanced by the work needed to reset the demon’s memory in a cyclic process [12]. It should be noted that Szilárd’s simplification of the problem to a gas occupying one of two nearby anticipates the information-theoretic idea of a bit, also demonstrating the connections between these two fields.

The field of quantum information grew out of this connection by asking the question: What happens to information processing and information theory in general when the information carriers are described by quantum mechanics? One immediate implication is the possibility of quantum superpositions of information states of a bit. Instead of just the usual 0s and 1s, which might be encoded quantum mechanically as $|0\rangle$ and $|1\rangle$, we can also have states of the form $\alpha|0\rangle + \beta|1\rangle$ for $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. This change in structure requires us to reexamine the entirety of Shannon’s information theory, rather than being able to only slightly modify the results to account for quantum effects, as pointed out by Ingarden [13]: “The old theory [Shannon’s theory] cannot be improved only by inserting into it some quantum formulae.”²

By now, a new, explicitly quantum information theory has been constructed by asking many of the same questions as before, but answering them with the tools and methods of quantum mechanics; see for instance the textbook of Nielsen and Chuang [14]. It has also been possible to adapt many of the techniques of usual, *classical* information theory to the quantum setting. For instance, Schumacher’s result that quantum information emitted from a source can be compressed at a rate equal to the von Neumann entropy of the source follows Shannon’s original result quite closely [15]. Nevertheless, in contrast to the classical case, we are still left with the question of what quantum information is information about.

The core theme of this thesis is that quantum information is in a certain sense a combination of two pieces of classical information, information about two physical observables which are *complementary* in the sense first put forth by Bohr [1, 2, 3] and exemplified by the wave-particle duality

¹Interestingly, Vannevar Bush had already used the phrase of “bits of information” in 1936 to describe information encoded into punchcards [7], though his meaning is different from Shannon’s.

²Ingarden also gives a very lucid description of the historical development of quantum information theory for the interested reader.

in the double-slit experiment [5]. Moreover, this point of view is useful in understanding and constructing protocols in quantum information theory. To appreciate this view of quantum information more clearly, the focus of this chapter, it is useful to first make the notions of classical information concrete in the following exceedingly simple game, the information game.

1.1 Understanding Classical Information via the Information Game

The information game has two players, Alice and Bob, and begins with Bob placing a coin, either heads or tails, in a box, and giving the box to Alice. At some point later she asks Bob whether she will see heads or tails when she opens the box. Bob's goal is to win the game by correctly matching Alice's observation.

Is there a strategy with which Bob can always win the game? Of course. For instance, Bob could always place the coin heads up in the box and answer "heads" whenever Alice comes asking. He could also just randomly place the coin heads up or down in the box, as long as he remembers which it was when Alice asks; this task of remembering is precisely Shannon's fundamental problem. To solve it, Bob could just write down "heads" or "tails" on a piece of paper and save it for later. In this sense, the paper carries information about the coin, in particular about what Alice will observe when she opens the box. Because there are two equally-likely possibilities, Bob could just as well use one binary digit, a zero or one, to remember the state of the coin. Therefore the paper carries one bit of information.

Formally, Bob's choice of the state of the coin can be represented as a binary-valued random variable X , taking on the values "heads" and "tails" with whatever probabilities p_{heads} and $p_{\text{tails}} = 1 - p_{\text{heads}}$ he decides. The state of the memory system he uses to remember the state of the coin can likewise be represented by a random variable, M , and a winning strategy simply has $M = X$ for any choice of X .

The amount of information stored the memory can be quantified by the Shannon entropy, defined for an arbitrary random variable Y as

$$H(Y) = - \sum_y p_y \log p_y, \quad (1.1)$$

using $\log = \log_2$ to measure in bits, a choice we shall make henceforth. The entropy of a random variable Y quantifies its uncertainty and is equal to the expected number of binary (yes/no) questions one would need to ask about Y in order to determine its actual value y [16]. A more concentrated distribution is less uncertain and makes guessing easier, and therefore has lower entropy, whereas the uniform distribution has maximum entropy and requires the most questions.

To win the game, the contents of the memory must determine the state of the coin, and thus contain information equal to the entropy of the coin $H(X)$. Thus, for the original winning strategy no information is stored in the memory at all—the memory is not even needed—as the coin *always* shows heads. Correspondingly, the entropy of this distribution is zero. In the second strategy, the memory stores one bit of information, since the coin is placed randomly in the box and $H(X) = 1$. For distributions in between these two limiting cases, we can imagine many playing many rounds of the game and the entropy gives the ratio of number of questions needed to number of rounds. For the distribution $p_{\text{heads}} = \frac{7}{8}$, $p_{\text{tails}} = \frac{1}{8}$, which has entropy $H(X) = 3 - \frac{7}{8} \log_2 7 \approx 0.54$, only 54 questions would be needed to determine the state of the coin in 100 rounds of play. In this case each memory register stores roughly one-half a bit of information.

On the other hand, given the value stored in the memory, the entropy of the coin random variable X is zero for every winning strategy. Formally, we can describe this using the conditional

entropy, defined using the probability of $X = x$ conditional on $M = m$, $p_{x|m} = p_{xm}/p_m$,³

$$H(X|M) = \sum_m p_m H(X|M = m), \quad \text{for} \quad H(X|M = m) = - \sum_x p_{x|m} \log p_{x|m}. \quad (1.2)$$

The conditional entropy can also be shown to satisfy $H(X|M) = H(XM) - H(M)$, and we can interpret it as the uncertainty of X given knowledge of M . Since a winning strategy only requires $M = X$, it is easy to work out that $H(X|M) = 0$ regardless of Bob's choice of X , the probability distribution he uses to decide whether to place the coin heads up or down. If the memory is faulty, then the stored value will not precisely match the state of the coin. For instance, if there is one chance in eight of a memory error and the coin was placed randomly in the box, then $p_{\text{heads}|\text{heads}} = \frac{7}{8}$, $p_{\text{tails}|\text{heads}} = \frac{1}{8}$, and similarly for the probability conditioned on tails. Working out the conditional entropy, we find $H(X|M) \approx 0.54$, meaning roughly half the information about the coin has been corrupted!

1.2 Complementarity in the Information Game

What changes if Alice and Bob play the game with the quantum version of coins, *qubits*, instead of classical bits? Qubits are any quantum system with two levels, which we denote $|0\rangle$ and $|1\rangle$, for instance the polarization degree of freedom of a single photon (horizontal versus vertical polarization) or the angular momentum of a spin- $\frac{1}{2}$ particle (angular momentum aligned or antialigned with a fixed spatial axis). Quantum-mechanical complementarity now comes into play and we can alter the game to illustrate the various effects concretely. Before doing so, let us discuss more precisely what is meant by complementarity, adopting the language of the wave-particle duality simplified to a single photon in a Mach-Zehnder interferometer.

Thinking of light as a particle, we expect to find the photon in one or the other of the two modes. By placing a photodetector in each arm of the interferometer, we can determine where the photon is by looking to see which of the photodetectors is triggered. Let us call this the amplitude measurement. Associating the states $|0\rangle$ and $|1\rangle$ to the two modes, the amplitude measurement corresponds to a projective measurement in this basis. We may also define the amplitude observable by assigning values to the two possible outcomes. The usual choice comes from thinking of a qubit as a spin- $\frac{1}{2}$ particle and using the angular momentum, and we define the amplitude observable as $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$. That is, a photon in the first mode takes the value $+1$ and in the second -1 .

If we instead think of light as a wave, we expect there to be a certain phase relationship between the two arms, and in this case the light can interfere either constructively (in phase, $+$) or destructively (out of phase, $-$). To determine which, we allow the two modes to interfere at a beamsplitter and then check in which mode the photon emerges with a photodetector. Let us call this the phase measurement. Like the amplitude measurement, the phase measurement is a projective measurement, but in the basis $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. Again we can define a corresponding observable, the phase observable, which for later convenience is defined exactly as the amplitude observable, but in the new basis: $X = |+\rangle\langle +| - |-\rangle\langle -|$. In the original basis this works out to be $X = |1\rangle\langle 0| + |0\rangle\langle 1|$.

Amplitude and phase are complementary properties precisely as in the double slit setup, in the sense that if the photon is in a definite mode, then the phase relationship is completely undefined, and *vice versa*. This can be immediately seen from the two sets of basis states, as measurement of either

³We follow physicists' conventions of naming arguments of functions and expressions, so that e.g. $p_{m|x}$ is the probability of $M = m$ given $X = x$, not the probability of $X = m$ given $M = x$.

eigenstate of amplitude produces a completely random outcome. At the level of observables, we can quantify this by an uncertainty relation. The most famous of these is the Heisenberg-Robertson relation relating the variances of the observables to the expectation of their commutator [4, 17],

$$\Delta X \Delta Z \geq \frac{1}{2} \left| \langle [X, Z] \rangle_\psi \right|, \quad (1.3)$$

where $\langle \dots \rangle_\psi$ denotes the expectation value evaluated for the quantum state $|\psi\rangle$ of the system. In this case, however, the bound is trivial. Since the operators X and Z anticommute ($XZ + ZX = 0$), the righthand side reduces to $|\langle XZ \rangle_\psi|$. Choosing $|\psi\rangle = |0\rangle$ immediately yields zero, and a simple calculation shows this conclusion holds for *any* possible choice of amplitude and phase observables.

Fortunately, there exist uncertainty relations for which the bound is state-independent. In particular, a version due to Maassen and Uffink is formulated in terms of entropy [18],⁴

$$H(X)_\psi + H(Z)_\psi \geq \log \frac{1}{c}. \quad (1.4)$$

The quantity c is related to the commutativity of the observables, $c = \max_{j,k} \left| \langle \psi_j | \phi_k \rangle \right|^2$ for $|\psi_j\rangle$ the eigenvectors of X and $|\phi_k\rangle$ those of Z , while the entropies are independently evaluated for the outcomes of the two observables, respectively, given that the system is originally in the quantum state ψ . In addition to the state-independent bound, the values of the observable can take play no role in the measure of uncertainty, only the probabilities of the various values. This makes the entropy a somewhat more natural measure than the variance. In the present case the two observables are complementary, meaning c takes on its maximal value, 1 (for observables on a d -level quantum system $c_{\max} = \log d$). Thus, the amplitude and phase measurements cannot both be certain, and there must be at least one bit of total entropy.

Alice and Bob can still play the classical information game with qubits, provided Alice only ever makes, say, the amplitude measurement. Bob is free to prepare amplitude eigenstates at random, just as before. In this sense the formalism of quantum information theory encapsulates classical information theory, as anything we wish to express in the latter can be done by working in a fixed basis in the former.⁵

Now suppose we alter the game so that Alice is free to make either an amplitude or a phase measurement, but she does not tell Bob which. Bob can prepare arbitrary qubit states, but to win the game he would need to be certain of the outcomes of both possible measurements. According to the Maassen-Uffink relation, Equation (1.4), this is impossible. There is no quantum state $|\psi\rangle$ Bob can send to Alice such that $H(X)_\psi$ and $H(Z)_\psi$ are both zero, and therefore he cannot win the game with certainty. A simple calculation shows that the best chance Bob has to win the game is to send Alice a state like $|\psi\rangle = \cos \frac{\pi}{8} |0\rangle + \sin \frac{\pi}{8} |1\rangle$, which is “in between” the amplitude and a phase eigenstates $|0\rangle$ and $|+\rangle$ in that $|\langle 0 | \psi \rangle| = |\langle + | \psi \rangle|$. Using $|\psi\rangle$, Bob has a roughly 85% chance ($\frac{1}{2} + \frac{1}{2\sqrt{2}}$) of correctly predicting that outcomes of either measurement is $+1$.

⁴Entropic uncertainty relations for position and momentum were first conjectured by Everett [19, 20] and Hirschmann [21] and proven by Becker [22]. Generalizations to arbitrary observables were made by Bialynicki-Birula and Mycielski [23] and Deutsch [24]. Kraus [25] first conjectured the stronger form (1.4).

⁵Here we consider only finite and not continuous alphabets.

1.3 Entanglement in the Information Game

What if, after receiving the qubit from Bob, Alice decides on a measurement at random and only asks for a prediction to this particular measurement? Since Bob does not know in advance which measurement Alice will perform, it would seem that this does not help. After all, he is still faced with the impossible task of preparing a state whose amplitude and phase are both predictable. Surprisingly, however, there does exist a winning strategy! The trick is for Bob to store *quantum information* about the system he sends to Alice. Note that in the game as played in the previous section, Bob really only makes use of classical memory. He may store information such as how he prepared the state for Alice, but this is effectively a recipe for making the state and there is nothing intrinsically quantum about such a recipe.

To win this version of the game, Bob should create an *entangled* state of two qubits A and B ,

$$|\Phi\rangle^{AB} \equiv \frac{1}{\sqrt{2}} (|0\rangle^A |0\rangle^B + |1\rangle^A |1\rangle^B) = \frac{1}{\sqrt{2}} (|+\rangle^A |+\rangle^B + |-\rangle^A |-\rangle^B), \quad (1.5)$$

and send the A system to Alice. Such entangled states were first by Einstein, Podolsky, and Rosen (EPR) [26] and later translated into this 2-level system language by Bohm [27]. EPR pointed out the paradoxical property that identical measurements on the two systems always produce identical results—the amplitude of A always matches that of B and likewise for phase—even though amplitude and phase for the individual systems cannot both be simultaneously well-defined.⁶ In a sense, entangled states display correlations even though there is nothing there to correlate!

However paradoxical, with entanglement Bob can always win the modified game. When Alice asks him to predict a particular measurement, he can simply consult his quantum memory, system B , by performing the same measurement Alice will make. Since the results are correlated, B in some sense contains one bit of classical information about both the amplitude and phase of system A . However, only one of these can ever be accessed because Bob cannot perform both measurements simultaneously; being able to do so would run afoul of the uncertainty principle. This peculiar combination of classical information about complementary physical properties is the essence of *quantum information*. Demonstrating this more concretely will be the topic of Chapter 3.

At first glance it would seem that this behavior violates the entropic uncertainty relation Equation (1.4). Now, however, Bob makes use of system B , so we should consider the entropies of the measurements conditioned on this fact. Thus Equation (1.4) does not apply. Just such a conditional version was conjectured and proven for the particular observables under consideration here in [RB09] and extended to general observables in [BCC⁺10].⁷ It states

$$H(X^A|B)_\psi + H(Z^A|B)_\psi \geq \log \frac{1}{c} + H(A|B)_\psi, \quad (1.6)$$

where now we make use of the quantum conditional entropy, defined using the von Neumann entropy (the Shannon entropy of the eigenvalues of the density matrix) as $H(A|B)_\psi = H(AB)_\psi - H(B)_\psi$. The entropies $H(X^A|B)_\psi$ and $H(Z^A|B)_\psi$ refer to quantum conditional entropies evaluated for the state after the respective observable of system A has been measured. The interpretation of a classical

⁶The EPR-Bohm states were actually states of two spin-1/2 systems with total angular momentum zero, so that identical measurements are always anticorrelated, but the point is the same.

⁷Uncertainty principles involving conditional entropy were first investigated by Hall [28] and extended to the case of separate conditional systems by Cerf *et al.* [29]. Christandl and Winter [30] gave a version for quantum channels which was the inspiration for the work in [RB09]. A much simpler proof of Equation (1.6) using the relative entropy was discovered by Coles *et al.* [31].

entropy conditioned on a quantum system is not as clear as entropy conditioned on a classical system, but Holevo has shown that it provides a lower bound on the classical conditional entropy of the stated measurement on system A given the result of the optimal measurement on system B [32, 33].⁸

Although the additional term on the righthand side might appear to make the bound tighter, the quantum conditional entropy of A given B can in fact be *negative*. For example, entangled states such as $|\Phi\rangle$ have $H(A|B) = -1$ since the AB state is pure (whence $H(AB) = 0$) but the state of B alone is completely random (whence $H(B) = 1$). This reflects another strange nature of the EPR state in that our uncertainty of the whole system AB appears to be *less* than that of one of its parts. In the present context, $H(A|B) = -1$ implies that the righthand side of (1.6) is zero. Thus, the bound is trivial, and conditioned on the quantum information B , the entropy of X and Z *can* both be zero.

An alternate and fully equivalent form of Equation (1.6) ensures that, even if Bob makes use of quantum information in the original version of the game where he has to predict both outcomes, no winning strategy can exist. It now involves three systems: the system to be measured, A , and two memory systems B and C ,

$$H(X^A|B)_\psi + H(Z^A|C)_\psi \geq \log \frac{1}{c}. \quad (1.7)$$

In order to make a prediction of both amplitude and phase, Bob would need two physical systems in which to store this information. Even if he uses systems B and C as quantum memory, Equation (1.7) ensures that amplitude and phase are still not simultaneously predictable. Put differently, although Bob can store classical information about both properties in the EPR state, there is no way to separate the amplitude and phase information without losing some of each in the process.

Note that we were able to define entropy conditioned on a quantum system via the alternate form of the conditional entropy expression, $H(A|B) = H(AB) - H(B)$. In retrospect, it is extremely fortunate that this form exists, because even the very notion of conditioning on quantum information is itself suspect. After all, the very nature of quantum systems is that their physical properties are not well-defined, so it is unclear what one should condition on. For instance, one might also like to define the variance of an observable on system A conditioned on the state of a quantum memory, system B . But how can the presence of B be incorporated into a variance calculation? We could stipulate that B is to be measured, calculate the variance of A for each outcome, and take the average, but this leads to unwieldy expressions. In the case of entropy, the formal structure rescues us and allows us to meaningfully speak of uncertainty conditioned on quantum information.

⁸This result was first proven by Forney [34], who did not make the connection to the conditional entropy.

Illustrations and Motivations

That complementary observables play an important role in quantum information processing is not original to this thesis, though we shall see new, more concrete characterizations of quantum information in terms of classical information pertaining to complementary observables and uses for these characterizations in subsequent chapters. In this chapter we recount several protocols in quantum information theory that anticipated and motivated the work presented herein. Among these are teleportation, where a qubit is sort of transmitted by two classical bits, and superdense coding, where conversely a qubit carries two bits of classical information. An even more concrete prior manifestation comes from quantum error-correction, which is crucial to the possibility of ever constructing a working quantum computer, and its use in protocols for entanglement distillation and quantum key distribution (QKD). This we discuss in more detail, as the structure of error-correcting codes will be useful in later chapters. But first we turn to teleportation and superdense coding.

2.1 Teleportation and Superdense Coding

Teleportation and superdense coding are two simple quantum information processing protocols which rather dramatically demonstrate how different quantum information is from classical information. They also indicate a connection between quantum information and complementary classical information. Both involve two parties, a sender Alice and a receiver Bob, who share an EPR pair as given in Equation (1.5). In superdense coding, Alice would like to transmit classical information to Bob, but using the quantum channel. One method is for both parties to fix a basis, Alice only sending amplitude eigenstates $|0\rangle$ or $|1\rangle$ and Bob only measuring what he receives in the same basis. This allows them to send one bit of classical information per qubit.

However, they can do better by making use of their shared entanglement, and Alice can send Bob two classical bits per qubit [35]. The trick is to use the Bell basis,¹ a basis of two maximally-entangled qubit states, defined as follows,

$$|\beta_{jk}\rangle^{AB} \equiv (X^j Z^k \otimes \mathbb{1})|\Phi\rangle^{AB}, \quad (2.1)$$

using the amplitude and phase operators as defined in the previous chapter. For completeness, we again write them here, in the basis $\{|0\rangle, |1\rangle\}$,

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (2.2)$$

To transmit the two bits j and k , first Alice applies $X^j Z^k$ to her half of the entangled state, A , and then sends it to Bob over the quantum channel. Since the Bell states form a basis, Bob can measure the joint system AB in this basis to determine j and k . In this way, one qubit of quantum information can be made to carry two bits of classical information.

The classical information can heuristically be regarded as one bit of amplitude information and one bit of phase information in the following manner. In the original scheme to transmit one bit per qubit using only the amplitude basis, Alice's actions can be described as *modulating* an initial state $|0\rangle$ by the operator X^j , producing $|1\rangle$ if $j = 1$ and leaving the state as $|0\rangle$ otherwise. The same modulation scheme works in the phase basis using the operator Z^k , starting with $|+\rangle$. In superdense

¹So-named as they figure prominently in the study of whether quantum mechanics permits description as a local hidden variable theory by John S. Bell [14].

coding, Alice apparently performs *both* amplitude and a phase modulation, encoding two bits at once. Due to the entanglement with Bob's system, these two actions can coexist without interfering with each other, allowing two bits to be transmitted.

Teleportation is sort of the inverse of superdense coding; now, preshared entanglement enables Alice to send one qubit to Bob by transmitting two classical bits [36]. Again the trick is to use the Bell basis. If Alice measures her half of the entangled state and the qubit to be sent in the Bell basis, she need only forward Bob the measurement results and he will be able to reconstruct the input state.

Formally, we let C be the qubit input, in an arbitrary state $|\psi\rangle^C$. It is not difficult to verify that ${}^{AC}\langle\beta_{jk}|(\Phi)^{AB}|\psi\rangle^C) = \frac{1}{2}(Z^k X^j)^B |\psi\rangle^B$. This means that after Alice measures her two systems in the Bell basis, each outcome occurring with probability $\frac{1}{4}$, Bob ends up with the state $Z^j X^k |\psi\rangle$. Thus, Alice merely has to send Bob the two bits of information j and k , and he can apply $X^j Z^k$ to recover the original state $|\psi\rangle$ in system B . In this way, the qubit is transmitted by two classical bits, with the help of preshared entanglement.

We can heuristically think of the two classical bits as being the amplitude and phase of the input state $|\psi\rangle$ for the following reason. One way to perform a Bell state measurement is to first perform the controlled-NOT (CNOT) operation and then measure each qubit separately in the appropriate basis. The CNOT gate acts on two qubits, applying X to the second qubit (the target) if the first (the control) is $|1\rangle$ and doing nothing to the target otherwise. It can be thought of as coherently copying the amplitude basis of the control qubit to the target, in that a superposition state $\alpha|0\rangle + \beta|1\rangle$ of the control qubit and a "blank" target state $|0\rangle$ become $\alpha|00\rangle + \beta|11\rangle$. To complete the Bell state measurement after applying CNOT, one measures the amplitude of the target qubit and the phase of the control. Therefore, in the teleportation protocol, we can choose the input qubit to be the control and Alice's half of the entangled state as the target, and it then appears as if the amplitude information is first copied to the second qubit and read out, while the phase is read out from the first qubit, the system itself. Of course, this is not precisely what happens, or else Alice would obtain both amplitude and phase information of $|\psi\rangle$, in violation of Equation (1.7). Nonetheless, teleportation indicates the important role played by amplitude and phase information.

2.2 Quantum Error-Correction

In the uncertainty game of the previous chapter, we assumed that the quantum memory used by Bob was noise-free. Clearly this is an unrealistic assumption, and although not particularly relevant for a *gedankenexperiment*, it nevertheless raises the question of what can be done to combat noise in real quantum information processing protocols. The answer, in the quantum case as in the classical case, is to use error-correcting codes. The fact that quantum error correction exists at all is of tremendous importance both practically and conceptually. On the one hand it shows that construction of quantum computers is not in principle a hopeless task, and on the other that quantum information itself is essentially digital (discrete-valued) in nature, despite its outward analog (continuous-valued) appearance. Even more, the way in which the first quantum error-correcting codes were constructed is related to the complementarity of quantum information: Arbitrary quantum errors are digitized into amplitude and phase errors, each of which is then corrected by essentially classical means. Before delving into the details of how quantum error-correction works, which illustrates the point more clearly and will be of use in later chapters, we give a brief overview of the issue of analog versus digital computation and the important role played by error-correction for both classical and quantum computers.

Whether classical or quantum, both analog and digital devices require error-correction to control the effects of noise inescapably present in an actual device. A simple classical error-correction scheme is simply to repeat the calculation three times and take the majority of the results. However, the error-correction procedure itself is presumably not perfect and can only be performed to some finite accuracy in practice. Nonetheless, following an analysis by von Neumann [37], it is possible that the rate at which errors are decreased by the procedure is greater than that at which they are caused.

For digital computers the finite accuracy of the procedure presents no additional difficulties in principle because the device anyway only requires finite accuracy; in a discrete encoding we choose certain continuous parameter ranges of the underlying physical degree of freedom to correspond to discrete logical values. Thus the nature of the encoding accords very well with the finite accuracy of available operations, and errors in the latter transform into errors in the former. Analog computers, however, require error-correction to arbitrary precision, so the buildup of errors due to finite-accuracy of operations is ultimately unavoidable. That reliable digital computers can be constructed from imperfect components was shown rigorously by Gács [38], though in practice current devices usually require error-correction only in the storage of information, not its manipulation, due to the intrinsically low error-rates of semiconductor-based integrated circuits.

Since the quantum state of the quantum computer is determined by the continuous probability amplitudes appearing in the wavefunction, many of the same difficulties were thought to apply to quantum computers, an issue pointed out by Peres [39] and stressed by Landauer [40, 41, 42]. Noise-induced modifications to these amplitudes leads to errors in the computation, just as in the analog computer, so it would seem that any advantage promised by quantum computation in principle cannot be achieved in practice. Worse still, even the ability to perform error-correction seems suspect in the quantum setting, because the information cannot simply be read out to check for errors, as in the von Neumann repetition scheme, without introducing disturbance [42]. Nonetheless, there was reason for optimism: Zurek observed that owing to the different phase-space structures involved, the kind of exponential blow-up of errors that might be expected for a classical continuous computer would not plague a quantum computer with a discrete spectrum [43].

Happily, the construction of quantum error-correcting codes by Shor [44] and Steane [45] demonstrates convincingly that quantum information is not analog, but digital.² Soon thereafter it was established that, just as with classical digital computers, reliable quantum computers could in principle be constructed using imperfect components, a fact known as the threshold theorem [47, 48, 49, 50, 51]. Unlike the situation for classical electronic computers, no medium has yet been discovered or engineered which offers intrinsically low quantum noise rates, though much effort is devoted to this question and many major experimental achievements have been made. The crux of quantum error-correction is that although continuous errors in the state of the computer are indeed possible, they can be digitized without damaging the encoded quantum information. Instead of accessing the quantum information directly, as one would try in a direct analogue of the repetition scheme, the measurements needed in error-correction are designed only to provide information about the error, not the encoded information. In this way the construction very subtly evades the two objections described above.

2.2.1 The Complementarity of Quantum Error-Correcting Codes

Somewhat amazingly, quantum errors of any type can be corrected if discrete errors of two complementary types, amplitude and phase, can be corrected [52, 53]. These two errors result from the

²Some would still dispute this. See, e.g. Laughlin [46].

2. ILLUSTRATIONS AND MOTIVATIONS

action of the already-defined X and Z operators, respectively, acting exactly as an unwanted modulation of the quantum state. Often these errors are referred to as bit flips and phase flips, for the following reason. One commonly fixes a basis and calls it the amplitude basis, and then for an arbitrary qubit state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ an amplitude error resulting from an unwanted X operator just flips the states $|0\rangle$ and $|1\rangle$, hence the name bit flip. Similarly, phase flips interchange the states $|+\rangle$ and $|-\rangle$, or equivalently, flips the phase of $|1\rangle$, taking (α, β) to $(\alpha, -\beta)$.

Either type of error by itself could be corrected in exactly the way a classical error would be corrected, through repetition. To correct a single bit flip error classically, we can *encode* it into three bits as follows,

$$0 \rightarrow \bar{0} = 000 \quad 1 \rightarrow \bar{1} = 111. \quad (2.3)$$

These two bitstrings are called *codewords*, and the overline denotes a logical value of the encoded bit, as opposed to the values of the individual physical bits. Then, if one error occurs, we can correct it by examining each string and flipping the one bit which is different from the other two. Equivalently, the error may be diagnosed by computing the two parities, generally called *syndromes*, $s_1 = b_1 \oplus b_3$ and $s_2 = b_2 \oplus b_3$, where b_1, b_2 , and b_3 are the three bit values. The syndromes associated to each error position are shown in Table 2.1. Note that the bit is encoded in the value of $\bar{b} = b_1 \oplus b_2 \oplus b_3$.

Bitstring pair $(\bar{0}, \bar{1})$	Error Position	Syndrome (s_1, s_2)
(000, 111)	\emptyset	(0, 0)
(100, 011)	1	(1, 0)
(010, 101)	2	(0, 1)
(001, 110)	3	(1, 1)

Table 2.1: The three-bit repetition code. The first column gives the bitstrings corresponding to the encoded logical zero $\bar{0}$ and logical one $\bar{1}$ after a bitflip error whose position is given in the second column. The third column lists the syndrome information which allows the error position to be diagnosed.

Seen from a different perspective, the reason this works is that the eight possible three-bit strings are grouped into four pairs, as in Table 2.1. One pair is given by the codewords themselves, and the other pairs are the images of the codewords under the three single-bit errors. In each pair one string corresponds to $\bar{0}$ and the other to $\bar{1}$ as defined by this mapping. The syndromes reveal precisely which pair is present, but importantly they do not reveal anything about the logical bit value. Error-correction corresponds to mapping the noisy pair of strings back to the original pair.

To correct qubit bit flip errors we may simply use the same repetition code in the computational basis. Since the syndrome and correction procedure for a given error are independent of the encoded information, superpositions are also maintained by the error-correcting code. Thus, the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is encoded as $|\bar{\psi}\rangle = \alpha|000\rangle + \beta|111\rangle$, a process which can be implemented as a unitary transformation on the input and two auxiliary systems, each in some given state we can take to initially be prepared in the state $|0\rangle$. The necessary syndrome information can be generated by measuring the two *stabilizer* operators $Z\mathbb{1}Z = Z \otimes \mathbb{1} \otimes Z$ and $\mathbb{1}ZZ$, which we can write as Z_1Z_3 and Z_2Z_3 . Each of these has the same action on the two logical states in each subspace, returning the values $(-1)^{s_1}$ and $(-1)^{s_2}$, respectively.

The name stabilizer reflects the fact that the code subspace is stabilized by the two operators, as it is the simultaneous $+1$ eigensubspace of both operators. The encoded subspace supports a single

qubit, and so it must be possible to represent its amplitude and phase operators. One possibility is given by $\bar{Z} = Z_1 Z_2 Z_3$ and $\bar{X} = X_1 X_2 X_3$. These each commute with the stabilizers, but anticommute with each other as intended. Note that \bar{Z} gives the encoded bit, just as in the classical case.

We can also think of the stabilizers and encoded amplitude operator as defining a new complete set of commuting observables for the set of physical qubits. Such a set fixes a basis in the state space of the three qubits, and each of the operators is the amplitude operator for a corresponding “virtual” qubit. Labeling the virtual qubit operators with primes, we can write $Z'_1 = Z_1 Z_2$, $Z'_2 = Z_2 Z_3$, and $Z'_3 = \bar{Z} = Z_1 Z_2 Z_3$. Conjugate to the new amplitude observables are phase observables $X'_1 = X_2 X_3$, $X'_2 = X_1 X_3$, and $X'_3 = \bar{X} = X_1 X_2 X_3$, which are found by ensuring that they anticommute with the amplitude operators of the same qubit but commute with all other operators. The entire collection is shown in Table 2.2. The code subspace is then defined by the first two virtual qubits being in the +1 amplitude state. Bit flip errors change the amplitude of the encoded qubit and at least one of the virtual qubits, and the stabilizer measurement determining the location of the error translates into an amplitude measurement of the first two virtual qubits.

Virtual qubit	Amplitude	Phase
1	ZZ $\mathbb{1}$	$\mathbb{1}XX$
2	$\mathbb{1}ZZ$	$XX\mathbb{1}$
3	ZZZ	XXX

Table 2.2: Virtual qubits associated with the three-qubit amplitude repetition code. Note that amplitude and phase anticommute for each qubit, but commute for different qubits.

Discretization is automatically provided by the measurement of the stabilizer operators, which is anyway necessary for error-correction. Consider an error operator of the form $E = e_0 I + e_1 X_1$, with $e_0, e_1 \in \mathbb{C}$, which is a sort of combination bit flip error and no error on the first qubit. It produces a superposition between two code subspaces,

$$|\bar{\psi}'\rangle = E |\bar{\psi}\rangle = e_0 |\bar{\psi}\rangle + e_1 X_1 |\bar{\psi}\rangle = e_0 (\alpha |000\rangle + \beta |111\rangle) + e_1 (\alpha |100\rangle + \beta |011\rangle). \quad (2.4)$$

Measurement of the stabilizer operators destroys this superposition, forcing the system to the state of either one error or no error, but leaves the logical qubit superposition intact. Here the measurement has two possible syndrome outcomes, either (0,0) or (1,0), with probabilities $|e_0|^2 / (|e_0|^2 + |e_1|^2)$ and $|e_1|^2 / (|e_0|^2 + |e_1|^2)$, respectively. Conditioned on these outcomes, the state becomes $|\bar{\psi}\rangle$ or $X_1 |\bar{\psi}\rangle$, respectively, and can therefore be corrected using the syndrome information.

2.2.2 Correcting Both Kinds of Errors

Since phase flips are just bit flips in the basis $|\pm\rangle$, the above analysis immediately applies to this case upon changing $X \leftrightarrow Z$ and working in the new basis. The insight of Shor and Steane was to realize that a single error of either type can be corrected by appropriately combining these procedures. Shor’s scheme is conceptually somewhat simpler, and is based on concatenating the two error-correcting codes. That is, we take the codewords of the phase flip repetition code and replace each of the three qubits with qubits appropriately encoded in the bit flip repetition code. This pro-

duces codewords of nine qubits, as follows (here ignoring normalization),

$$|+\rangle \longrightarrow |\bar{+}\rangle = |+++ \rangle = (|0\rangle + |1\rangle)(|0\rangle + |1\rangle)(|0\rangle + |1\rangle) \quad (2.5)$$

$$|\bar{+}\rangle \longrightarrow |\tilde{+}\rangle = (|\bar{0}\rangle + |\bar{1}\rangle)(|\bar{0}\rangle + |\bar{1}\rangle)(|\bar{0}\rangle + |\bar{1}\rangle), \quad (2.6)$$

$$|-\rangle \longrightarrow |\bar{-}\rangle = |--\rangle = (|0\rangle - |1\rangle)(|0\rangle - |1\rangle)(|0\rangle - |1\rangle) \quad (2.7)$$

$$|\bar{-}\rangle \longrightarrow |\tilde{-}\rangle = (|\bar{0}\rangle - |\bar{1}\rangle)(|\bar{0}\rangle - |\bar{1}\rangle)(|\bar{0}\rangle - |\bar{1}\rangle). \quad (2.8)$$

The repetition in the amplitude basis in the second step protects the encoded qubit from bit flip errors, since a single bit flip can always be detected and corrected by applying the 3-qubit repetition procedure to each block of three qubits. This corresponds to measuring the Z-parity observables Z_1Z_3 , Z_2Z_3 , Z_4Z_6 , Z_5Z_6 , Z_7Z_9 , and Z_8Z_9 . Phase flips are slightly more involved, but consider what happens when a single phase flip error plagues, say, the fourth qubit. This is the first qubit of the second block, so we can zoom in on this block to determine the effect on the encoded states. Applying the error operator Z_1 to the encoded states we find $Z_1|\bar{0}\rangle = Z_1|000\rangle = |000\rangle = |\bar{0}\rangle$, while $Z_1|\bar{1}\rangle = Z_1|111\rangle = -|111\rangle = -|\bar{1}\rangle$. Thus, the error causes the action

$$|\tilde{+}\rangle \rightarrow (|\bar{0}\rangle + |\bar{1}\rangle)(|\bar{0}\rangle - |\bar{1}\rangle)(|\bar{0}\rangle + |\bar{1}\rangle) \quad (2.9)$$

$$|\tilde{-}\rangle \rightarrow (|\bar{0}\rangle - |\bar{1}\rangle)(|\bar{0}\rangle + |\bar{1}\rangle)(|\bar{0}\rangle - |\bar{1}\rangle), \quad (2.10)$$

which is precisely a phase flip at the “inner” level. We could detect and correct this at the inner level by measuring the X-parities X_1X_3 and X_2X_3 . Translating to the outer level of actual qubits, we replace each of the constituent X operators on the inner level by its encoded \bar{X} operator on the outer level and instead measure $X_1X_2X_3X_7X_8X_9$ and $X_4X_5X_6X_7X_8X_9$. The outcomes for the damaged states are +1 and -1 respectively, for both encoded states, implying that to correct the error we merely need apply Z_4 .³

The six amplitude parities and two phase parities commute pairwise and stabilize the code subspace. As with the repetition code, the error analysis is made simpler by thinking in terms of virtual qubits, in this case nine, as shown in Table 2.3. Observe that the concatenated structure is reflected in the operators: three copies of the repetition code in virtual qubits one through six, followed by the same repetition code on the three blocks. The code subspace is fixed by requiring virtual qubits one through six to be in the +1 amplitude eigenstate and virtual qubits seven and eight in the +1 phase eigenstate, but this structure makes it clear that we could have defined the code the other way around.

Using this framework it is easy to see that the Shor code also enables detection and correction of joint bit and phase errors. A joint bit and phase flip of the fourth qubit, for instance, would reveal itself by the fourth virtual qubit having the wrong amplitude and the seventh having the wrong phase, corresponding to -1 eigenvalues of the stabilizers Z_4Z_6 and $X_4X_5X_6X_7X_8X_9$. From the structure of the virtual amplitude and phase operators it is clear that the code can actually detect and correct one bit and one phase error, irrespective of their locations.

Again error discretization is provided by the stabilizer measurement, and fortunately, being able to correct just these two types of error is sufficient to correct any conceivable single-site error. Just as with the repetition code, we can consider the effect of arbitrary errors which are linear combinations of all the correctable errors. Since the Shor code can correct any single flip of bit and/or phase, errors

³ Z_5 or Z_6 would also work just as well. This flexibility is actually a subtle and important feature of quantum error-correcting codes we shall return to in Section 6.3.

Virtual qubit #	Amplitude	Phase
1	ZZ 1 1 1 1 1 1 1 1	1 XX 1 1 1 1 1 1
2	1 ZZ 1 1 1 1 1 1	XX 1 1 1 1 1 1 1
3	1 1 1 ZZ 1 1 1 1	1 1 1 1 XX 1 1 1
4	1 1 1 1 ZZ 1 1 1	1 1 1 XX 1 1 1 1
5	1 1 1 1 1 1 ZZ 1	1 1 1 1 1 1 1 XX
6	1 1 1 1 1 1 1 ZZ	1 1 1 1 1 1 XX 1
7	ZZZZZZ 1 1 1	1 1 1 XXXXXX
8	1 1 1 ZZZZZZ	XXXXXX 1 1 1
9	ZZZZZZZZ	XXXXXXXXXX

Table 2.3: Virtual qubits associated with the nine-qubit Shor code. Note that amplitude and phase anticommute for each qubit, but commute for different qubits.

of the form $E = e_{00}I + e_{10}X_1 + e_{01}Z_1 + e_{11}X_1Z_1$ with $e_{jk} \in \mathbb{C}$ can also be corrected. But, as can be readily verified, any operator can be expressed in this way as a complex combination of these four operators, meaning arbitrary single-site errors can be digitized to amplitude and/or phase errors and corrected. Despite initial appearances to the contrary, quantum information is therefore in a critical sense digital.

2.3 Entanglement Distillation

Quantum error-correction quickly found use in constructing protocols for *distillation* of entanglement, as well as in proving the cryptographic security of quantum key distribution protocols. We give a brief treatment of these uses here, as they will be generalized in later chapters.

Distilling entanglement refers to transforming imperfect EPR states into approximately perfect ones. For instance, if Alice sends halves of maximally-entangled states through a noisy quantum channel to Bob, then the states which emerge will no longer be maximally-entangled. But it may be possible to repair some fraction of the states by actions undertaken on Alice's and Bob's systems alone, plus classical communication between them to coordinate their actions. To see how this is done, suppose that Alice and Bob share many copies of the state

$$\psi^{AB} = \sum_{jk} p_{jk} |\beta_{jk}\rangle \langle \beta_{jk}|^{AB}, \quad (2.11)$$

with $p_{jk} \geq 0$ and $\sum_{jk} p_{jk} = 1$, which is just a probabilistic mixture of the four Bell states. This state is produced, for instance, by sending the B half of the state $|\Phi\rangle = |\beta_{00}\rangle$ through a channel which applies the operator $X^j Z^k$ with probability p_{jk} . In principle, Bob can repair the actual state ψ^{AB} to the desired state Φ^{AB} by determining which of these operators was applied and subsequently undoing it. Thus, the task is reduced to determining the actual sequence of errors, at least for states of this form.

This sounds like a job for a quantum error-correcting code, even though here Alice is not first encoding the qubits she sends to Bob. Nevertheless, Alice and Bob can determine the error pattern by each measuring the stabilizer operators of an error-correcting code. It is simple to show that, just as in the information game, if Alice and Bob make the same stabilizer measurements on collections of EPR states, then they should always obtain the same outcomes. To the extent that they obtain

different outcomes, this indicates an error. For instance, suppose that Alice and Bob divide their systems into groups of three and use the simple bit-flip repetition code described above. On each group of three, both Alice and Bob measure the stabilizers Z_1Z_3 and Z_2Z_3 . Perhaps the simplest way to work out what outcomes will occur is to note the following relationship,

$$\mathbb{1} \otimes E |\Phi\rangle = E^T \otimes \mathbb{1} |\Phi\rangle, \quad (2.12)$$

valid for any operator E , where E^T is the transpose of the operator when expressed in the amplitude basis $\{|0\rangle, |1\rangle\}$. Now we can calculate the effect of the product of Alice's and Bob's stabilizers on the ideal state. Since $Z^T = Z$ and $Z^2 = \mathbb{1}$,

$$Z^{A_1} Z^{A_3} Z^{B_1} Z^{B_3} (|\Phi\rangle^{A_1 B_1} |\Phi\rangle^{A_2 B_2} |\Phi\rangle^{A_3 B_3}) = |\Phi\rangle^{A_1 B_1} |\Phi\rangle^{A_2 B_2} |\Phi\rangle^{A_3 B_3}, \quad (2.13)$$

which implies that Alice and Bob must indeed obtain identical outcomes for their stabilizer measurements since their product must be $+1$. The same clearly holds for Z_2Z_3 . If there is one X error in the state, say in the first position, then we find using the same method

$$Z^{A_1} Z^{A_3} Z^{B_1} Z^{B_2} (X^{B_1} |\Phi\rangle^{A_1 B_1} |\Phi\rangle^{A_2 B_2} |\Phi\rangle^{A_3 B_3}) = Z^{A_1} X^{A_1} Z^{A_1} (|\Phi\rangle^{A_1 B_1} |\Phi\rangle^{A_2 B_2} |\Phi\rangle^{A_3 B_3}) \quad (2.14)$$

$$= -X^{B_1} |\Phi\rangle^{A_1 B_1} |\Phi\rangle^{A_2 B_2} |\Phi\rangle^{A_3 B_3}. \quad (2.15)$$

Now the state is a -1 eigenstate of the product of stabilizers, meaning the product of syndromes is -1 , and hence that Alice and Bob obtain different outcomes for these stabilizer measurements. A single X error on the first qubit will of course not affect the Z_2Z_3 measurements. But together the two stabilizer measurements suffice to locate a single X error in the three pairs, exactly as in the error-correction scenario.

The story is essentially the same for any quantum error-correcting code, so we may create a protocol for entanglement distillation as follows, following Bennett *et al.* [53]. First, Alice and Bob use a small fraction of their pairs in order to determine the number of each type of error X , Z , and XZ , simply by both measuring in the appropriate basis and recording how often they obtained the same outcome. The bases are just the amplitude basis, the phase basis, and the basis consisting of the eigenstates $\frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle)$ of XZ , respectively. Next, given the expected number of errors, they choose an appropriate error-correcting code, but if no suitable codes exist they must abort the procedure. If a suitable code does exist, they proceed by measuring the stabilizers to determine, with high probability, the actual pattern of errors, which can then be corrected by local operations on Bob's systems.

This does not quite leave them with the desired states $|\Phi\rangle$, however, since they have made the stabilizer measurements. Instead, the $|\Phi\rangle$ reside in the encoded subspaces specified by the error-correcting code, their number corresponding to the number of encoded qubits. To recover these states, they each apply the decoding operation (the inverse of the encoding operation) to their systems. The above protocol is designed to work for states of the form given in Equation (2.11), but actually applies to any input state since the stabilizers used in the protocol will automatically digitize arbitrary errors to amplitude and phase errors.

2.4 Quantum Key Distribution

Quantum key distribution (QKD) provides a means for the two separated parties Alice and Bob to communicate in private using only public communication channels. The security of the scheme is

based only on the laws of physics and not the perceived computational difficulty of some task, like factoring large integers, as commonly used in classical schemes today. Needless to say, the problem of private communication is ancient, but it was first put on a firm mathematical footing by Shannon [54]. There the task is broken into two parts, establishing a *secret key* between the two parties, a random string of bits shared by both parties, and then using it to *encrypt* and *decrypt* the actual messages. One can imagine Alice and Bob creating a secret key together at some point in the past when they could do so secretly, but if they are already separated and can only communicate publicly, the situation seems hopeless. They could communicate privately if they had a key, but they need to communicate privately to create the key.

Quantum information offers a way out of this dilemma in the form of entanglement. Returning to the uncertainty game, recall that Bob can, on demand, predict either the amplitude or phase measurement on Alice's system when they share an EPR pair. Moreover, the uncertainty principle Equation (1.7) implies that any would-be eavesdropper Eve could not predict either measurement using her system C any better than by just blindly guessing, a property of entanglement known as *monogamy*. By measuring each of their systems in identical bases, Alice and Bob can therefore generate one bit of a secret key from each entangled pair.

They can attempt to create such pairs by using a public quantum channel in the manner described in the previous subsection: Alice prepares EPR pairs and sends one system of each to Bob. If the channel is noisy, perhaps due to Eve's interference, Alice and Bob can simply first run an entanglement distillation protocol to extract the required high-quality EPR pairs. Even though this requires them to exchange classical syndrome information over a public channel, it does not help any would-be eavesdropper as the measurements on the EPR pairs are completely independent of this information, a fact again insured by the uncertainty principle. The usefulness of entanglement distillation in this context was first treated by Deutsch et al. [55] and the security of this scheme was first rigorously proven by Lo and Chau [56].

The protocol will require a large quantum memory in which to store the various systems, as well as the ability to perform all the necessary stabilizer measurements. We did not worry about the practicalities of doing so in the previous section, but luckily for QKD all of the required operations can be reduced to just measuring in either the amplitude or phase basis, and subsequent processing of the resulting classical data, as shown by Shor and Preskill [57]. The reason this works is that ultimately we want to distill EPR states but then immediately measure them in some basis to generate the key, and this gives us some flexibility in how we describe the entire process. By picking the right kind of error-correction code this flexibility allows us to get rid of essentially all (difficult) operations on quantum systems apart from measuring them individually and replace them with (easy) operations on classical data.

The necessary codes are called Calderbank-Shor-Steane (CSS) codes and include the original codes found by Shor and Steane as mentioned in Section 2.2. Their defining property, as first described by Calderbank and Shor [58] and Steane [59], is that the stabilizers of the code can be broken into two groups, those composed of products of X operators and those composed of products of Z operators. Similarly, the logical amplitude operators only consist of Z -type operators, while the logical phase operators only consist of X -type operators. The more general formalism of stabilizer codes constructed by Gottesman [60] also includes codes whose stabilizers and logical operators are of mixed type, but importantly, these cannot be used for the present purposes.

Consider the QKD scheme above using a CSS-based entanglement distillation scheme to correct for noise in the quantum channel. The entanglement distillation part proceeds in two steps, the first involving measurement of the Z -type stabilizers, which give Alice and Bob information about the bit errors, and the second involving the X -type stabilizers, which give information about the

phase errors. Now assume that the key is generated by measuring, in the amplitude basis, each half of the pairs output by the decoding step of the distillation protocol. This is equivalent to skipping the decoding step and instead measuring the logical \bar{Z} operators directly. But in a CSS code these operators are composed entirely of products of Z operators on the individual qubits, and one can reconstruct the value of any desired product from the collection of all the individual outcomes. Knowing Z_1 , Z_2 , and Z_3 enables us to calculate $Z_1Z_2Z_3$, for instance.

Thus, Alice and Bob could generate the outcomes of measuring the logical amplitude operators as well as all the Z -type stabilizers by first measuring each of their respective qubits in the amplitude basis and then forming the appropriate products of the outcomes. However, the X -type stabilizers cannot be generated in this way; in fact, all phase information will be destroyed by making amplitude measurements. *The crucial fact is that Alice and Bob do not need the X -type stabilizers at all.* Intuitively this makes sense, as these stabilizers give information about phase errors, but Alice and Bob only care about amplitude information.

The protocol now proceeds as follows. Alice transmits halves of entangled pairs to Bob, and a random subset are used to estimate the rate of bit and phase errors in order to choose an appropriate CSS code, while the rest are immediately measured in the amplitude basis. Just as in the entanglement distillation protocol, if no suitable code exists because the noise rates are too high, they must abort the procedure. If one does exist, Alice proceeds by constructing the Z -type stabilizers according to the chosen code and transmitting them to Bob, who corrects the amplitude errors. They then forget about the phase stabilizers and each constructs the outcomes of measuring the logical amplitude operators for use as the secret key.

From the outside there is no way to tell if Alice and Bob have performed the above procedure or actually measured the X - and Z -type stabilizers directly. Although the phase information has not been exchanged, correction of the phase errors is nevertheless possible in principle. Therefore, the procedure inherits the security of the Lo and Chau protocol in which Alice and Bob actually do create EPR pairs.

In contrast, from Alice and Bob's point of view, the key is created by two classical information processing protocols. First, Alice sends Bob the stabilizer information which enables him to correct his observed amplitude measurements to match hers. This step is referred to as *information reconciliation* since the goal is to reconcile Bob's amplitude information with Alice's. In the next step they use the logical operators to construct a function of the amplitude data, which serves as the key. Due to the entanglement-based picture of the protocol, this has the effect of extracting that part of the amplitude data which is completely uncorrelated with any eavesdropper, and this part of the protocol is termed *privacy amplification*. We can think of the amplitude measurements as a sort of raw key which is then distilled to a truly secret key by running these two protocols in succession.

Remarkably, we can also remove the need for entanglement entirely. Suppose that in the above protocol Alice immediately measures her halves of the EPR states as she sends the other halves to Bob. These measurements essentially prepare amplitude and phase basis states in the systems underway to Bob. For instance, if her amplitude measurement is $|0\rangle$, Bob's system is now in the state $|0\rangle$, and so on. Originally Alice and Bob agree in advance which observable to measure for each qubit, but suppose instead that they each make a random choice. Half the time they choose the same basis, and these outputs are "sifted" out by public announcement of the bases and kept for use as the key and for error estimation.

From the outside there is no way to tell if Alice measures her system after the transmission, so that she is distributing half of an entangled pair, or before, so that she is randomly preparing amplitude or phase eigenstates for Bob to measure. Thus, just as the classical key distillation scheme inherits

security from entanglement distillation, the *prepare and measure* protocol inherits security from the EPR based version. In fact, this prepare and measure scheme is the original QKD protocol proposed by Bennett and Brassard [61] and known as BB84; the connection to the version using entanglement was noted by Bennett, Brassard, and Mermin [62]. Shor and Preskill prove that the BB84 protocol is secure using the reduction of entanglement distillation to information reconciliation and privacy amplification using CSS codes and the reduction of an entanglement-based protocol to a prepare and measure protocol.

Characterizing Quantum Information

In the information game of Chapter 1 we made use of the fact that EPR pairs have the property that measurements of either amplitude or phase on one subsystem are entirely predictable using the other subsystem, and we argued that this property is central to the notion of quantum information itself. Here we make good on this claim by showing the converse is true, amplitude and phase predictability implies entanglement, as well as providing two other characterizations of entanglement based on complementarity and showing how these can be extended to characterizations of secret keys.

The present chapter is divided into four sections. The first presents the converse as stated above. Specifically, following [RB08], we show that if there exist measurements on Bob’s system which predict Alice’s amplitude and phase measurements with low error probability, then Bob can adapt these measurements to create a new system forming an approximate EPR pair with Alice’s system. Essentially this is done by coherently performing both measurements in succession, as depicted in Figure 3.1.

Our approach was inspired by Koashi’s complementary control scenario [63] in which Bob either tries to guess Alice’s amplitude information or somehow help her to prepare a phase eigenstate, and we remark on the connections below. Furthermore, the entanglement recovery procedure is useful in several other scenarios, such as approximate quantum error correction and the quantum information processing protocol known as state merging.

In the second section we give two other sufficient conditions for entanglement recovery using the uncertainty principle, recounting the results of [Ren11]. Again amplitude and phase information play the decisive role, but now the conditions involve a third system. In the first of these, entanglement is implicitly present in the systems shared by Alice and Bob if the amplitude measurement is predictable with low error probability using Bob’s system, but high error probability using *any* other system. In the second, entanglement is present if both amplitude and phase are unpredictable in this sense using any other system. These conditions are not as constructive as the first, and instead rely on a powerful method often used in quantum information theory called *decoupling*.

The third section modifies a result of [RB09] and details how the three characterizations above can be formulated in terms of conditional entropy. Finally, by appealing to the uncertainty principle, we can make a slight modification to the entanglement recovery procedure to instead create *private states*, which are the most general quantum-mechanical description of secret keys. Indeed this was actually the original motivation of [RB08].

3.1 Amplitude and Phase Predictability & Entanglement

Let us now specify the setup under consideration more formally. Our two parties Alice and Bob are located some distance apart and each have a technologically-advanced laboratory in which they can manipulate quantum systems. Suppose now that Alice and Bob share a generic bipartite quantum state ψ^{AB} . Without loss of generality this state is the AB subsystem of a pure state $|\psi\rangle^{ABE}$ for E the “environment”.¹ We can express this pure state in two ways by expanding Alice’s system in the amplitude or phase basis,

$$|\psi\rangle^{ABE} = \sum_{z=0}^1 \sqrt{p_z} |z\rangle^A |\varphi_z\rangle^{BE} \quad \text{and} \quad |\psi\rangle^{ABE} = \sum_{x=0}^1 \sqrt{q_x} |\tilde{x}\rangle^A |\vartheta_x\rangle^{BE}. \quad (3.1)$$

¹In the context of the pure state $|\psi\rangle^{ABE}$, ψ^{AB} denotes the marginal state of the AB system.




Figure 3.1: The quantum circuit enabling entanglement recovery from a bipartite state ψ^{AB} by Bob, when he can approximately predict measurement of either conjugate observable X or Z by Alice. It proceeds in three steps. First, Bob coherently performs the measurement \mathcal{M}_Z allowing him to predict Z , storing the result in auxiliary system C_Z (unitary $U_{\mathcal{M}_Z}^{BC_Z}$). Next, he coherently performs the measurement \mathcal{M}_X allowing him to predict X , storing the result in auxiliary system C_X (unitary $U_{\mathcal{M}_X}^{BC_X}$). Finally, to recover a maximally entangled state in system C_Z , he applies a controlled-NOT gate, with control C_Z and target C_X (unitary $U_{\text{CNOT}}^{C_Z C_X}$). This procedure also leaves Bob holding the original input state ψ^{AB} in systems C_X and B .

Here $|z\rangle$ denote amplitude eigenstates according to $Z|z\rangle = (-1)^z|z\rangle$ and similarly $|\tilde{x}\rangle$ denote phase eigenstates according to $X|\tilde{x}\rangle = (-1)^x|\tilde{x}\rangle$. The states $|\varphi_z\rangle^{BE}$ and $|\vartheta_x\rangle^{BE}$ are normalized pure states, but otherwise arbitrary; p_z and q_x are the probabilities that Alice obtains the outcome z and x for amplitude and phase measurements, respectively.

If Alice makes the amplitude measurement corresponding to the observable Z on her system, Bob can attempt to match her outcome by performing some generalized measurement on his system. This measurement is described most generally by a positive operator valued-measure (POVM) \mathcal{M}_Z , which consists of elements positive semidefinite operators Λ_z such that $\sum_z \Lambda_z = \mathbb{1}$. The probability that he can correctly guess her outcome is given by

$$p_{\text{guess}}(Z^A | \mathcal{M}_Z^B)_\psi \equiv \sum_{z=0}^1 \text{Tr} \left[(P_z^A \otimes \Lambda_z^B) \psi^{AB} \right] = \sum_{z=0}^1 p_z \text{Tr} \left[\Lambda_z^B \varphi_z^B \right], \quad (3.2)$$

where P_z is the projector onto the amplitude state $|z\rangle$, i.e. $P_z = |z\rangle\langle z|$. The subscript on the guessing probability denotes which state we should use to evaluate it. To predict Alice's phase measurement, Bob would use a different POVM \mathcal{M}_X with POVM elements Γ_x . His guessing probability is

$$p_{\text{guess}}(X^A | \mathcal{M}_X^B)_\psi \equiv \sum_{x=0}^1 \text{Tr} \left[(\tilde{P}_x^A \otimes \Gamma_x^B) \psi^{AB} \right] = \sum_{x=0}^1 q_x \text{Tr} \left[\Gamma_x^B \vartheta_x^B \right], \quad (3.3)$$

where now \tilde{P}_x is the projector onto the phase state $|\tilde{x}\rangle$.

3.1.1 Approximate Entanglement Implies Approximate Predictability

Given the EPR state $|\Phi\rangle^{AB}$, we saw in Chapter 1 that Bob can perfectly predict Alice's amplitude and phase measurements. Before moving on to the converse, we can strengthen this to an approximate

condition, that if Alice and Bob share a good approximation to $|\Phi\rangle^{AB}$, then the amplitude and phase measurements are approximately predictable in that Bob's error probabilities are small. The trick is to choose the correct notion of "good approximation". Such approximate conditions are important in considering practical scenarios, since Bob's prediction of Alice's measurement outcome will never be perfect. For the same reason, they ensure that the idea that predictability of complementary information is what counts for quantum information is truly a physical statement, not merely a mathematical curiosity of the theory.

If we want the resulting error probabilities to be small, a natural choice is to demand the *trace distance* between the state Alice and Bob actually share, ψ^{AB} , and the ideal state Φ^{AB} be less than some prescribed approximation parameter ϵ . The trace distance between any two states ρ and σ is defined as $\frac{1}{2}\|\rho - \sigma\|_1$, where $\|M\|_1 \equiv \text{Tr}\sqrt{M^\dagger M}$ for any operator M . The reason this is an appropriate choice stems from the fact that the trace distance cannot increase under quantum operations such as measurement, and that the guessing probability is directly related to the trace distance of the measured state. More generally, the trace distance between two states is also related to the maximum probability that the two states give different outcomes under any possible measurement.² Thus, for small trace distance, the two states behave essentially identically under any possible measurement.

Using the trace distance we can show that for approximate EPR pairs amplitude and phase are approximately predictable. Suppose that $\frac{1}{2}\|\psi^{AB} - \Phi^{AB}\|_1 \leq \epsilon$ and imagine both Alice and Bob perform the amplitude measurement on their respective systems. If Φ^{AB} were the actual state, the result would be $\Phi^{Z^A Z^B} = \frac{1}{2}\sum_z P_z^A \otimes P_z^B$, where we use the observable Z^A , respectively Z^B , to denote that the state has been measured and to specify which measurement has been made. For ψ^{AB} they obtain $\psi^{Z^A Z^B} = \sum_{z,z'} p_z \text{Tr}[P_{z'}^B \varphi_z^B] P_z^A \otimes P_{z'}^B$. Computing the trace distance, we find

$$\frac{1}{2}\|\Phi^{Z^A Z^B} - \psi^{Z^A Z^B}\|_1 = \frac{1}{2}\sum_{z,z'} \left| \frac{1}{2}\delta_{z,z'} - p_z \text{Tr}[P_{z'}^B \varphi_z^B] \right| \|P_z^A \otimes P_{z'}^B\|_1 \quad (3.4)$$

$$= \frac{1}{2}\sum_{z,z'} \left| \frac{1}{2}\delta_{z,z'} - p_z \text{Tr}[P_{z'}^B \varphi_z^B] \right|, \quad (3.5)$$

which is just the *variational distance* between the ideal distribution $\frac{1}{2}\delta_{z,z'}$ and the actual distribution $p_z \text{Tr}[P_{z'}^B \varphi_z^B]$. But the variational distance can also be expressed as follows

$$\frac{1}{2}\sum_{z,z'} \left| \frac{1}{2}\delta_{z,z'} - p_z \text{Tr}[P_{z'}^B \varphi_z^B] \right| = \max_S \sum_{(z,z') \in S} \left| \frac{1}{2}\delta_{z,z'} - p_z \text{Tr}[P_{z'}^B \varphi_z^B] \right|, \quad (3.6)$$

where S is any subset of the pairs (z, z') . Choosing $S = (z, z')$ for $z \neq z'$ gives

$$\frac{1}{2}\|\Phi^{Z^A Z^B} - \psi^{Z^A Z^B}\|_1 \geq \sum_{z \neq z'} p_z \text{Tr}[P_{z'}^B \varphi_z^B] \quad (3.7)$$

$$= 1 - p_{\text{guess}}(Z^A|Z^B)_\psi. \quad (3.8)$$

The latter expression is a slight abuse of notation, using the observable Z^B to denote Bob's measurement. Because the trace distance cannot increase under the measurement, $p_{\text{guess}}(Z^A|Z^B)_\psi \geq 1 - \epsilon$. The same conclusion holds for the phase measurement, $p_{\text{guess}}(X^A|X^B)_\psi \geq 1 - \epsilon$.

²See, e.g. [14] for an excellent introduction to and explication of the basic results in quantum information theory.

3.1.2 Approximate Amplitude and Phase Predictability Implies Entanglement

Now we can state and prove the converse, that approximate amplitude and phase predictability implies entanglement. We first establish a lemma which will also be used in the subsequent entanglement characterizations.

Lemma 1. *Given a state $|\psi\rangle^{ABE} = \sum_z \sqrt{p_z} |z\rangle^A |\varphi_z\rangle^{BE}$, let $|\psi\rangle^{C_X BE}$ be the identical state with system C_X replacing A , and define $|\psi_Z\rangle^{AC_Z BE} = \sum_z \sqrt{p_z} |z\rangle^A |z\rangle^{C_Z} |\varphi_z\rangle^{BE}$. If there exist partial isometries $U_1^{B \rightarrow C_Z B}$ and $U_2^{C_Z B \rightarrow C_Z C_X B}$ such that*

$${}_{AC_Z BE} \langle \psi_Z | U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE} \geq 1 - \epsilon_1, \quad \text{and} \quad (3.9)$$

$$({}_{AC_Z} \langle \Phi | {}_{C_X BE} \langle \psi |) U_2^{C_Z B \rightarrow C_Z C_X B} |\psi_Z\rangle^{AC_Z BE} \geq 1 - \epsilon_2, \quad (3.10)$$

then for $U^{B \rightarrow C_Z C_X B} = U_2^{C_Z B \rightarrow C_Z C_X B} U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE}$,

$$\frac{1}{2} \left\| | \Phi \rangle^{AC_Z} |\psi\rangle^{C_X BE} - U^{B \rightarrow C_Z C_X B} |\psi\rangle^{ABE} \right\|_1 \leq \sqrt{2\epsilon_1} + \sqrt{2\epsilon_2}. \quad (3.11)$$

Proof. The fidelity $F(\psi, \phi) = \langle \psi | \phi \rangle$ between two pure states gives an upper bound on their trace distance, $\frac{1}{2} \|\psi - \phi\|_1 \leq \sqrt{1 - F(\psi, \phi)^2}$, so that fidelity greater than $1 - \epsilon$ translates into trace distance less than $\sqrt{2\epsilon}$. Since the trace distance is invariant under unitaries and partial isometries, the lemma follows from the triangle inequality. \square

Theorem 1. *If $p_{\text{guess}}(Z^A | \mathcal{M}_Z^B)_\psi \geq 1 - \epsilon_1$ and $p_{\text{guess}}(X^A | \mathcal{M}_Z^B)_\psi \geq 1 - \epsilon_2$ for some measurements \mathcal{M}_Z^B and \mathcal{M}_X^B on a state ψ^{AB} , then there exists a partial isometry $U^{B \rightarrow C_Z C_X B}$ such that*

$$\frac{1}{2} \left\| | \Phi \rangle^{AC_Z} |\psi\rangle^{C_X BE} - U^{B \rightarrow C_Z C_X B} |\psi\rangle^{ABE} \right\|_1 \leq \sqrt{2\epsilon_1} + \sqrt{2\epsilon_2}. \quad (3.12)$$

Proof. We use the measurements to define the two isometries required for Lemma 1. For the first isometry $U_1^{B \rightarrow C_Z B}$ we may use the coherent implementation of the measurement \mathcal{M}_Z^B , which stores the measurement result in system C_Z . Performing the measurement coherently produces the state $U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE}$, which without loss of generality takes the form

$$U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE} = \sum_{z, z'} \sqrt{p_z} |z\rangle^A |z'\rangle^{C_Z} \sqrt{\Lambda_{z'}^B} |\varphi_z\rangle^{BE}. \quad (3.13)$$

Now compute the overlap of this state with the state $|\psi_Z\rangle^{AC_Z BE} = \sum_z \sqrt{p_z} |z\rangle^A |z\rangle^{C_Z} |\varphi_z\rangle^{BE}$, which would be the ideal output of the coherent measurement process.

$$\langle \psi_Z | U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE} = \sum_z p_z \langle \varphi_z | \sqrt{\Lambda_z^B} |\varphi_z\rangle^{BE} \quad (3.14)$$

$$\geq \sum_z p_z \langle \varphi_z | \Lambda_z^B |\varphi_z\rangle^{BE} \quad (3.15)$$

$$= p_{\text{guess}}(Z^A | \mathcal{M}_Z^B)_\psi, \quad (3.16)$$

using the fact that $\sqrt{\Lambda} \geq \Lambda$ for $0 \leq \Lambda \leq \mathbb{1}$. Hence, we have the first condition of Lemma 1.

For the second, let $V^{B \rightarrow C_X B}$ be the partial isometry which coherently implements the measurement \mathcal{M}_X^B , storing the result in system C_X . Coherently measuring $|\psi_Z\rangle^{AC_Z BE}$ gives

$$V^{B \rightarrow C_X B} |\psi_Z\rangle^{AC_Z BE} = \sum_z \sqrt{p_z} |z\rangle^A |z\rangle^{C_Z} V^{B \rightarrow C_X B} |\varphi_z\rangle^{BE} \quad (3.17)$$

$$= \frac{1}{\sqrt{2}} \sum_z |z\rangle^A |z\rangle^{C_Z} V^{B \rightarrow C_X B} \sum_x (-1)^{xz} \sqrt{q_x} |\vartheta_x\rangle^{BE} \quad (3.18)$$

$$= \frac{1}{\sqrt{2}} \sum_z |z\rangle^A |z\rangle^{C_Z} \sum_{x, x'} (-1)^{xz} \sqrt{q_x} |\tilde{x}'\rangle \sqrt{\Gamma_{x'}^B} |\vartheta_x\rangle^{BE}. \quad (3.19)$$

Here we have made use of the algebraic relationship between the two bases. Ideally the output would be the state

$$|\psi'_Z\rangle^{AC_Z C_X BE} = \frac{1}{\sqrt{2}} \sum_z |z\rangle^A |z\rangle^{C_Z} \sum_x \sqrt{q_x} (-1)^{xz} |\tilde{x}\rangle^{C_X} |\vartheta_x\rangle^{BE}, \quad (3.20)$$

and computing the fidelity between the ideal and actual outputs gives

$${}^{AC_Z C_X BE} \langle \psi'_Z | V^{B \rightarrow C_X B} |\psi_Z\rangle^{AC_Z BE} = \frac{1}{2} \sum_{z, x, x', x''} \sqrt{q_x q_{x''}} (-1)^{z(x-x'')} \langle \tilde{x}'' | \tilde{x}' \rangle \langle \vartheta_{x''} | \sqrt{\Gamma_{x'}^B} |\vartheta_x\rangle^{BE} \quad (3.21)$$

$$= \sum_x q_x \langle \vartheta_x | \sqrt{\Gamma_x^B} |\vartheta_x\rangle^{BE} \quad (3.22)$$

$$\geq \sum_x q_x \langle \vartheta_x | \Gamma_x^B |\vartheta_x\rangle^{BE} \quad (3.23)$$

$$= p_{\text{guess}}(X^A | \mathcal{M}_X^B) \psi. \quad (3.24)$$

We may also express $|\psi'_Z\rangle^{AC_Z C_X BE}$ as $\frac{1}{\sqrt{2}} \sum_z |z\rangle^A |z\rangle^{C_Z} (X^z)^{C_X} |\psi\rangle^{C_X BE}$, and therefore applying a control-NOT $W_{\text{CNOT}}^{C_Z C_X}$ with C_Z as the control and C_X as the target to the ideal output gives $|\Phi\rangle^{AC_Z} |\psi\rangle^{C_X BE}$. Since the fidelity is invariant under partial isometries, the second condition of Lemma 1 holds for $U_2^{C_Z B \rightarrow C_Z C_X B} = W_{\text{CNOT}}^{C_Z C_X} V^{B \rightarrow C_X B}$, completing the proof. \square

3.1.3 Further Uses of the Entanglement Recovery Operation

We originally introduced the environment system E as the purification of the joint state held by Alice and Bob, but of course we can look at it the other way around; generically Alice and Bob jointly hold the purification of system E . However, our entanglement recovery operation has done more than just recover entanglement, as it reveals that when Alice's amplitude and phase measurements are predictable by Bob, he implicitly holds the purification of E by himself. This follows because the recovery operation also produces (a good approximation to) the state $|\psi\rangle^{C_X BE}$, which is identical to the initial state except system A is replaced by C_X , held by Bob. In the following chapter we shall use this property to construct protocols for *state merging*, in which Alice attempts to merge her state with Bob by using classical or quantum communication.

Theorem 1 may be regarded as giving necessary and sufficient conditions on the existence of an approximate quantum error-correction scheme: Approximate error correction is possible when amplitude and phase information can each approximately be recovered. The schemes discussed in Chapter 2 based on quantum error-correcting codes were *perfect* in the sense that the input quantum

state can be perfectly recovered if the error is of the correctable type. Approximate error-correction sets the more modest goal of only recovering the input approximately.

Entanglement recovery is relevant to this goal because we can always mimic the initial single-system input to the error-correction problem as half of an EPR pair, the other half of which is then measured in an appropriate basis. The basis is given by complex conjugating the coefficients of the original basis, which follows because the EPR state can be written as $|\Phi\rangle^{AB} = \frac{1}{\sqrt{2}} \sum_j |\xi_j^*\rangle^A |\xi_j\rangle^B$ for any basis $\{|\xi_j\rangle = \xi_{j0}|0\rangle + \xi_{j1}|1\rangle\}_{j=0}^1$, where $|\xi_j^*\rangle = \xi_{j0}^*|0\rangle + \xi_{j1}^*|1\rangle$. This was precisely the method used in reducing the QKD scheme based on EPR pairs to one involving only preparation and measurement of single systems. From this line of reasoning it follows from a result of Schumacher [64] that if entanglement can be approximately recovered by the scheme, then the approximation parameter sets a lower bound on the average fidelity with which single systems can be recovered by the same procedure.

In the preceding analysis, we have assumed that Alice's system is a qubit, whereas Bob's system is arbitrary. But the result may be easily extended to the case that Alice holds a d -level system by using the more general amplitude and phase operators defined by

$$X = \sum_{k=0}^{d-1} |k \oplus 1\rangle \langle k| \quad \text{and} \quad Z = \sum_{k=0}^{d-1} e^{2\pi i k/d} |k\rangle \langle k|. \quad (3.25)$$

Often these are called the Weyl-Heisenberg operators, as they have similar properties to the position and momentum operators of continuous-variable systems. Here the crucial point is that the algebraic properties of the amplitude and phase operators used in Theorem 1 hold for higher-dimensional systems as well. In the sequel, we shall continue to specialize to the qubit case.

3.2 Duality & Decoupling

The uncertainty principle Equation (1.7) establishes a tradeoff in how well Alice's amplitude measurement can be predicted using system B and how well her phase measurement can be predicted using system E . In the previous section the sufficient conditions for entanglement were of the former type, but the tradeoff suggests that we might to be able to find sufficient conditions of the latter type and focus instead on what information system E does *not* have, rather than what information system B does have. Concentrating on lack of information and building protocols by destroying correlations is the essence of the decoupling approach to quantum information processing, which goes back to work on approximate error-correction by Schumacher and Westmoreland [65] and has found wide application to constructing information processing protocols such as state merging [66, 67] and noisy channel coding [68, 69, 70, 71] that we shall encounter in Chapter 4.³

In the decoupling approach one tries to show that Alice's system is completely uncorrelated with system E in order to infer that Bob's system is entangled with Alice's. Here, however, we shall be able to show that it suffices for this purpose to ensure that E has no information about Alice's amplitude or phase. This reflects our main theme that what really counts in quantum information is classical information about complementary observables. Part of the appeal of decoupling is that it allows us to avoid the problem of constructing the isometries needed for Lemma 1. Instead, the isometries are automatically constructed by appealing to Uhlmann's theorem on the relationship between fidelity of mixed states and that of their possible purifications.

³The decoupling approach has also been extended to quantum channels, instead of quantum states as described here, in [72, 73, 74, 75, 76].

However, from the uncertainty principle we are only entitled to expect that if Z^A is predictable from B , then X^A is unpredictable from E , but not the converse. Were the converse true in general, we could immediately establish that lack of information in E is sufficient to imply the presence of entanglement because it would imply the conditions we already have in Lemma 1 and Theorem 1. Absent the converse, it is not immediately clear that this approach will work.

Note that the converse does hold if it happens the state $|\psi\rangle^{ABE}$ saturates Equation (1.7), so that $H(Z^A|B)_\psi + H(X^A|E)_\psi = 1$. Two separate sufficient conditions for equality in the uncertainty principle are derived in [RB09], and these take the simple form $p_{\text{guess}}(X^A|B)_\psi = 1$ and $p_{\text{guess}}(Z^A|E)_\psi = 1$. Equivalently, each of these conditions implies the complementary form $H(X^A|B) + H(Z^A|E) = 1$ is trivially saturated, since either $H(X^A|B) = 0$ and thus $H(Z^A|E) = 1$ or $H(Z^A|E) = 0$ and thus $H(X^A|E) = 1$. Luckily, it turns out that due to the structure of Lemma 1, either of these two equality conditions can be satisfied without loss of generality to the entanglement criteria. We shall make use of both in the two results presented next.

First, we need to formally characterize the unpredictability of measurements on Alice's system when making use of the purification system E . The most straightforward approach would be to say that the associated guessing probabilities are small, even for the optimal measurement. However, optimal measurements are quite often difficult to specify in quantum information theory. To sidestep this problem, we may instead use the following quantity,

$$p_{\text{secure}}(Z^A|E)_\psi = 1 - \frac{1}{2} \left\| \psi^{Z^A E} - \frac{1}{2} \mathbb{1}^A \otimes \psi^E \right\|_1, \quad (3.26)$$

and say that E has no information about Z^A when $p_{\text{secure}}(Z^A|E)_\psi$ is nearly one. Another possibility would be to phrase matters in terms of the conditional entropy, stating that $H(Z^A|E)$ is large. It turns out that such an entropic condition implies that of Equation (3.26) and we shall return to this point in the next section.

A bound on $p_{\text{secure}}(Z^A|E)_\psi$ also implies a bound on the guessing probability, as follows. Suppose system E is measured with some POVM $\mathcal{M}_Z^E = \{\Lambda_z^E\}$. The probability distributions of measurement outcomes on the real and ideal states are $p_{z,z'} = p_z \text{Tr}[\Lambda_z^E \varphi_z^E]$ and $p'_{z,z'} = \frac{1}{2} \text{Tr}[\Lambda_z^E \varphi^E]$, respectively. For the variational distance we find

$$\frac{1}{2} \sum_{z,z'} |p_{z,z'} - p'_{z,z'}| \geq \frac{1}{2} \sum_z p_{z,z} - p'_{z,z} \quad (3.27)$$

$$= \frac{1}{2} \sum_z \left(p_z \text{Tr}[\Lambda_z^E \varphi_z^E] - \frac{1}{2} \text{Tr}[\Lambda_z^E \varphi^E] \right) \quad (3.28)$$

$$= \frac{1}{2} p_{\text{guess}}(Z^A|\mathcal{M}_Z^E)_\psi - \frac{1}{4}, \quad (3.29)$$

and therefore $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon$ implies $p_{\text{guess}}(Z^A|\mathcal{M}_Z^E)_\psi \leq \frac{1}{2} + 2\epsilon$ for any measurement \mathcal{M}_Z^E . Observe that the quantity p_{secure} also implies the outcome of the Z^A measurement is nearly random. This accounts for the name 'secure' since effectively this means Alice can generate a secure secret key bit by this measurement.

Now we are ready to state the new entanglement conditions. The first says that Alice and Bob implicitly share entanglement if Alice's amplitude measurement can be predicted using B but not E . This is almost the same as saying that Alice and Bob can generate a shared secret key from their state, a point we return to in Section 3.4.

3. CHARACTERIZING QUANTUM INFORMATION

Theorem 2. *If $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon_2$ and $p_{\text{guess}}(Z^A|\mathcal{M}_Z^B)_\psi \geq 1 - \epsilon_1$ for some measurement \mathcal{M}_Z^B , then there exists a partial isometry $U^{B \rightarrow C_Z C_X B}$ such that*

$$\frac{1}{2} \left\| |\Phi\rangle^{AC_Z} |\psi\rangle^{C_X BE} - U^{B \rightarrow C_Z C_X B} |\psi\rangle^{ABE} \right\|_1 \leq \sqrt{2\epsilon_1} + \sqrt{2\epsilon_2}. \quad (3.30)$$

Proof. From the proof of Theorem 1, the first condition of Lemma 1 is fulfilled for $U_1^{B \rightarrow C_Z B}$ the coherent implementation of the measurement \mathcal{M}_Z^B . For the second condition, consider the implications of $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon_2$ for the state $|\psi_Z\rangle^{AC_Z BE}$. Since tracing out $C_Z B$ from $|\psi_Z\rangle^{AC_Z BE}$ gives the same result as measuring the amplitude Z^A of the state $|\psi\rangle^{ABE}$, we have

$$\frac{1}{2} \left\| \psi_Z^{AE} - \frac{1}{2} \mathbb{1}^A \otimes \psi^E \right\|_1 \leq \epsilon_2. \quad (3.31)$$

The trace distance gives a lower bound to the fidelity, so that

$$F\left(\psi_Z^{AE}, \frac{1}{2} \mathbb{1}^A \otimes \psi^E\right) \geq 1 - \epsilon_2. \quad (3.32)$$

By Uhlmann's theorem, the fidelity of two mixed states is identical to the largest fidelity of their possible purifications. Two possible purifications of the two states in question are $|\psi_Z\rangle^{AC_Z BE}$ and $|\Phi\rangle^{AC_Z} |\psi\rangle^{C_X BE}$, and since all other purifications are related by isometries involving the purifying system, we have

$$\left({}^{AC_Z} \langle \Phi | {}^{C_X BE} \langle \psi | \right) U_2^{C_Z B \rightarrow C_Z C_X B} |\psi_Z\rangle^{AC_Z BE} \geq 1 - \epsilon_2 \quad (3.33)$$

for some $U_2^{C_Z B \rightarrow C_Z C_X B}$. This is the sought-after second condition and completes the proof. \square

Observe that although Theorem 2 calls for $p_{\text{secure}}(Z^A|E)_\psi$ to be large, we actually apply this condition to the state $|\psi_Z\rangle^{AC_Z BE}$ for which $p_{\text{guess}}(Z^A|C_Z B) = 1$. Therefore the uncertainty relation $H(X^A|C_Z B)_\psi + H(Z^A|E)_\psi = 1$ holds, and we are essentially able to trade large $p_{\text{secure}}(Z^A|E)_\psi$ for large $p_{\text{guess}}(X^A|C_Z B)_\psi$ as is needed for Lemma 1. Indeed, it follows from the discussion prior to Lemma 1 that the following is an immediate corollary to Theorem 2. Using the isometry $U^{B \rightarrow C_Z C_X B}$ to define the measurement $\mathcal{M}_X^B = \{U^{\dagger B \rightarrow C_Z C_X B} \tilde{P}_x^{C_Z} U^{B \rightarrow C_Z C_X B}\}$, we have

Corollary 2. *If $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon$ and $p_{\text{guess}}(Z^A|\mathcal{M}_Z^B)_\psi = 1$ for some measurement \mathcal{M}_Z^B , then there exists a measurement \mathcal{M}_X^B such that $p_{\text{guess}}(X^A|\mathcal{M}_X^B)_\psi \geq 1 - \sqrt{2\epsilon}$.*

This is Theorem 4.2(a) of [Ren11]. Note that we have dropped the explicit use of $|\psi_Z\rangle$ by stipulating that $p_{\text{guess}}(Z^A|\mathcal{M}_Z^B)_\psi = 1$.

Continuing the trend of denying information to E , we might hope that Alice and Bob share entanglement when the amplitude and phase are unpredictable using E . However, a simple counterexample shows that this cannot be true in general. Define $|\psi\rangle^{ABE} = \frac{1}{\sqrt{2}} (|0\rangle + i|1\rangle)^A \otimes |\varphi\rangle^{BE}$; here Alice's system is an eigenstate of the observable XZ . Due to the product structure, both Z^A and X^A are unpredictable using either E or B . This is to be expected in light of the preceding discussion on the need to saturate the uncertainty principle. One way to avoid this problem is to require that not only is the amplitude measurement unpredictable using E , but the phase measurement is unpredictable using E *even assuming E could predict the amplitude*. Formally, we require $p_{\text{secure}}(X^A|C_Z E)_{\psi_Z}$ to be large, which again involves the state $|\psi_Z\rangle$ that saturates the uncertainty principle, though this time C_Z is joined with E , not B .

Theorem 3. *If $p_{\text{secure}}(X^A|C_Z E)_{\psi_Z} \geq 1 - \epsilon_1$ and $p_{\text{secure}}(Z^A|E)_{\psi} \geq 1 - \epsilon_2$, then there exists a partial isometry $U^{B \rightarrow C_Z C_X B}$ such that*

$$\frac{1}{2} \left\| |\Phi\rangle^{AC_Z} |\psi\rangle^{C_X BE} - U^{B \rightarrow BC_Z C_X} |\psi\rangle^{ABE} \right\|_1 \leq \sqrt{2\epsilon_1} + \sqrt{2\epsilon_2}. \quad (3.34)$$

Proof. From the proof of Theorem 2, the second condition here implies the second condition of Lemma 1 is satisfied. To show the first, consider the implications of $p_{\text{secure}}(X^A|C_Z E)_{\psi_Z} \geq 1 - \epsilon_1$ for the state $|\psi_Z\rangle^{AC_Z BE}$. We may express the state as follows

$$|\psi_Z\rangle^{AC_Z BE} = \frac{1}{\sqrt{2}} \sum_x |\tilde{x}\rangle^A \sum_z \sqrt{p_z} (-1)^{xz} |z\rangle^{C_Z} |\varphi_z\rangle^{BE} \quad (3.35)$$

$$= \frac{1}{\sqrt{2}} \sum_x |\tilde{x}\rangle^A (Z^x)^{C_Z} \sum_z \sqrt{p_z} |z\rangle^{C_Z} |\varphi_z\rangle^{BE} \quad (3.36)$$

$$= \frac{1}{\sqrt{2}} \sum_x |\tilde{x}\rangle^A (Z^x)^{C_Z} |\psi\rangle^{C_Z BE}. \quad (3.37)$$

Defining $|\eta_x\rangle^{C_Z BE} \equiv (Z^x)^{C_Z} |\psi\rangle^{C_Z BE}$ and $\eta \equiv \frac{1}{2} \sum_x \eta_x^{C_Z E}$, we have

$$1 - p_{\text{secure}}(X^A|C_Z E)_{\psi_Z} = \frac{1}{2} \left\| \frac{1}{2} \sum_x |\tilde{x}\rangle \langle \tilde{x}|^A \otimes \eta_x^{C_Z E} - \frac{1}{2} \mathbb{1}^A \otimes \eta^{C_Z E} \right\|_1 \quad (3.38)$$

$$= \frac{1}{4} \sum_x \left\| \eta_x^{C_Z E} - \eta^{C_Z E} \right\|_1. \quad (3.39)$$

But since the trace distance is invariant under unitary operations, in particular $(Z^x)^{C_Z}$, $\left\| \eta_x^{C_Z E} - \eta^{C_Z E} \right\|_1 = \left\| \eta_{x'}^{C_Z E} - \eta^{C_Z E} \right\|_1$ for all x, x' . Observe that $\eta^{C_Z E} = \psi_Z^{C_Z E}$ since the random phase flip has the effect of “measuring” the amplitude of C_Z . As $|\eta_0\rangle^{C_Z BE} = |\psi\rangle^{C_Z BE}$, we can therefore infer that $\frac{1}{2} \left\| \psi^{C_Z E} - \psi_Z^{C_Z E} \right\|_1 \leq \epsilon_1$, or equivalently $\frac{1}{2} \left\| \psi^{AE} - \psi_Z^{AE} \right\|_1 \leq \epsilon_1$. Converting trace distance to fidelity and applying Uhlmann’s theorem, we find there exists an isometry $U_1^{B \rightarrow C_Z B}$ such that

$${}^{AC_Z BE} \langle \psi_Z | U_1^{B \rightarrow C_Z B} |\psi\rangle^{ABE} \geq 1 - \epsilon_1. \quad (3.40)$$

Thus, the first condition of Lemma 1 is satisfied, completing the proof. \square

Again, the resulting isometry can be used to define the measurement \mathcal{M}_Z^B in the following, which is Theorem 4.2(b) of [Ren11].

Corollary 3. *If $p_{\text{secure}}(X^A|E)_{\psi} \geq 1 - \epsilon$ and $p_{\text{guess}}(Z^A|\mathcal{M}_Z^E)_{\psi} = 1$ for some measurement \mathcal{M}_Z^E , then there exists a measurement \mathcal{M}_Z^B such that $p_{\text{guess}}(Z^A|\mathcal{M}_Z^B)_{\psi} \geq 1 - \sqrt{2\epsilon}$.*

Figure 3.2 illustrates the contents of Theorems 1, 2, and 3 by indicating which system must have what kind of information, or lack thereof, in order to infer the presence of entanglement between Alice and Bob.

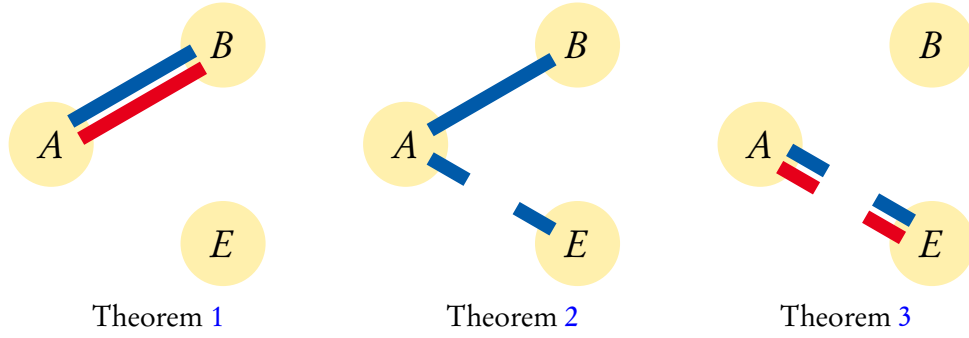


Figure 3.2: Graphical depiction of the contents of Theorems 1, 2, and 3. All three theorems specify conditions under which Alice and Bob can transform their shared state ψ^{AB} into a collection of EPR pairs using only local operations. Theorem 1 shows this can be done when both the amplitude and phase of Alice’s system are correlated with Bob’s system in that he could reliably predict either, depicted by the blue (amplitude) and red (phase) lines joining Alice and Bob. Theorem 2 shows that amplitude correlation with Bob and *uncorrelation* with the environment leads to the same conclusion. Finally, Theorem 3 establishes that appropriate uncorrelation of both amplitude and phase suffices to infer that the Alice-Bob system to be entangled.

3.3 Entropic Characterizations

As advertised after Equation (3.26), another possible formalization of “unpredictability” is using the conditional entropy: The amplitude measurement outcome Z^A is unpredictable using E when $H(Z^A|E)$ is large. Owing to the connection between conditional entropy and the quantity p_{secure} given by the following lemma, we can establish entropic conditions on entanglement from the results of the previous section. This was partially investigated in [RB09].

Lemma 2. *If $H(Z^A|E)_\psi \geq 1 - \epsilon^2$, then $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon$.*

Proof. The proof relies on the connection between relative entropy $D(\rho||\sigma)$ of two states ρ and σ and trace distance between them, in particular the bound $\frac{1}{2\ln 2} \|\rho - \sigma\|_1^2 \leq D(\rho||\sigma)$ [65]. This may be more conveniently expressed as $\frac{1}{2} \|\rho - \sigma\|_1 \leq \sqrt{D(\rho||\sigma)}$. By direct calculation it is easy to show

$$D\left(\psi_Z^{AE} \parallel \frac{1}{2} \mathbb{1}^A \otimes \psi^E\right) = 1 - H(Z^A|E)_\psi. \quad (3.41)$$

Using the bound on $H(Z^A|E)_\psi$ and the definition of $p_{\text{secure}}(Z^A|E)_\psi$ completes the proof. \square

Theorem 4. *Given any of the following pairs of conditions,*

$$\begin{array}{lll} (1) & H(Z^A|B)_\psi \leq \epsilon_1^2, & (2) & H(Z^A|B)_\psi \leq \epsilon_1^2, & (3) & H(X^A|C_Z E)_\psi \geq 1 - \epsilon_1^2, \\ & H(X^A|B)_\psi \leq \epsilon_2^2 & & H(Z^A|E)_\psi \geq 1 - \epsilon_2^2 & & H(Z^A|E)_\psi \geq 1 - \epsilon_2^2 \end{array}$$

there exists a partial isometry $U^{B \rightarrow C_Z C_X B}$ such that

$$\frac{1}{2} \left\| |\Phi\rangle^{AC_Z} |\psi\rangle^{C_X B E} - U^{B \rightarrow BC_Z C_X} |\psi\rangle^{ABE} \right\|_1 \leq \sqrt{2\epsilon_1} + \sqrt{2\epsilon_2}. \quad (3.42)$$

Proof. Using Lemma 2 for the last pair, we can apply Theorem 3. But since $H(Z^A|B)_\psi = H(Z^A|B)_{\psi_Z}$, (1) and (2) each separately imply (3) by the uncertainty principle Equation (1.7). \square

That the first pair of entropic conditions implies that Alice and Bob share entanglement is a variation of a related result found by Christandl and Winter that quantum channels are useful for transmitting entanglement if they could be used to reliably transmit classical amplitude and phase information [30]. The fact that the first pair of entropic conditions in Theorem 4 are sufficient for systems A of arbitrary dimension is actually somewhat surprising, as it is known that just because the conditional entropy $H(Z^A|B)$ is small does not imply that there exists a measurement \mathcal{M}_Z^B such that $H(Z^A|\mathcal{M}_Z^B) = H(Z^A|B)$, let alone that the guessing probability $p_{\text{guess}}(Z^A|\mathcal{M}_Z^B)$ is large. In fact, Ruskai has shown that Bob's conditional marginal states must all commute pairwise for the conditional entropy to be achievable [77].

The gap can be simply illustrated by the following example implicitly given by Holevo [78], in which Bob's state conditioned on Alice's amplitude basis measurement is a randomly-selected amplitude or phase eigenstate,

$$|\psi\rangle^{AB} = \frac{1}{2} \sum_{t=0}^3 |t\rangle^A |\varphi_t\rangle^B, \quad (3.43)$$

for $|\varphi_0\rangle = |0\rangle$, $|\varphi_1\rangle = |1\rangle$, $|\varphi_2\rangle = |+\rangle$, and $|\varphi_3\rangle = |-\rangle$. By direct calculation we find $H(Z^A|B) = 1$, where now Z^A is any non-degenerate observable diagonal in the $|z\rangle$ basis. On the other hand, a derivation by DiVincenzo *et al.* [79] using the Maassen-Uffink uncertainty relation Equation (1.4) shows that the optimal measurement \mathcal{M}_Z^B is such that $H(Z^A|\mathcal{M}_Z^B) = \frac{1}{2}$. The optimal measurement can also be found by exploiting the group covariance of Bob's states and appealing to a theorem of Davies [80]. Nonetheless, fulfilling *both* entropy conditions evidently circumvents this issue, as the necessary measurements are defined by Corollaries 2 and 3.

3.4 Secret Keys & Private States

With a very slight modification, we can extend the results above to give necessary and sufficient conditions on the ability to extract a secret key instead of an EPR pair from the state ψ^{AB} . In Section 2.4 we discussed the fact that EPR pairs can be used to create secret keys, but entanglement of this form is not actually necessary, a fact first observed by Aschauer and Briegel [81]. Instead, bipartite quantum states which are capable of producing secret keys are called *private states* and their general form was established by Horodecki *et al.* [82]. In this section we show that just like entanglement, knowledge of complementary observables plays a decisive role in characterizing private states.

Private states have two defining features, as alluded to prior to Theorem 2. First, the key measurements by Alice and Bob clearly must produce identical results. Second, the key should be completely random and uncorrelated with any third party, i.e. a would-be eavesdropper Eve. Without loss of generality we can assume that Alice and Bob have two systems each, A, A' and B, B' , respectively, and the key bit is generated by amplitude measurements of A and B . If they start with any other state having only systems A' and B' , they can coherently perform the key generation measurements and store the result in the amplitude of systems A and B , respectively. Horodecki *et al.* [82] give the following characterization of *ideal* private states, whose proof we include here for completeness.

Theorem 5 (Horodecki *et al.* [82]). $\psi^{AA'BB'}$ is a private state iff there exists a twisting operator $U^{AA'B'}$ of the form $U^{AA'B'} = \sum_{z=0}^1 |z\rangle\langle z|^A \otimes V_z^{A'B'}$ with V_z unitary such that for some $\xi^{A'B'}$,

$$\psi^{AA'BB'} = U^{AA'B'} \left(\Phi^{AB} \otimes \xi^{A'B'} \right) U^{\dagger AA'B'}. \quad (3.44)$$

3. CHARACTERIZING QUANTUM INFORMATION

Proof. Consider a purification of a private state. By the first requirement, it must have the form

$$|\psi\rangle^{AA'BB'E} = \frac{1}{\sqrt{2}} \sum_{z=0}^1 |z, z\rangle^{AB} |\varphi_z\rangle^{A'B'E}. \quad (3.45)$$

The second requirement implies that the states φ_z^E are all identical, so that the key bit z is secret from any eavesdropper. All possible purifications of a state are related by unitaries on the purifying system, meaning $|\varphi_z\rangle^{A'B'E} = V_z^{A'B'} |\varphi_0\rangle^{A'B'E}$ for some unitaries V_z . Using these to define the twisting operator and letting $\xi^{A'B'} = \varphi_0^{A'B'}$ completes the proof. \square

Thus, private states are “twisted” versions of entangled states in which the $A'B'$ system is transformed in some way conditioned on the value of the key. Since the function of the $A'B'$ system is to block correlations of the key with E , it is called the *shield*. Here we have defined the twisting operator as conditioning on Alice’s key system A , but since her key is always equal to Bob’s, the twisting operator can just as well be conditioned on B . Private states are conceptually distinct from entangled states because the distributed nature of the $A'B'$ system prevents Alice and Bob from undoing the twisting operator on a general private state. Indeed, there exist private states from which no entanglement can be locally extracted [82].

As with entanglement, we are more interested in characterizations of approximate secret keys, since perfection will be impossible to achieve in practice. The following lemma shows that the above definition of secret keys can be extended to a sensible approximate version. Here we denote by $\psi^{Z^A Z^B E}$ the state ψ^{ABE} after measurement of the observables Z^A and Z^B , and we say that an approximate secret key is ϵ -good when its trace distance to a perfect key is less than ϵ .

Lemma 3. *If $p_{\text{guess}}(Z^A|Z^B)_\psi \geq 1 - \epsilon_1$ and $p_{\text{secure}}(Z^A|E)_\psi \geq 1 - \epsilon_2$, then $\psi^{Z^A Z^B E}$ is an $(\epsilon_1 + \epsilon_2)$ -good secret key.*

Proof. Start with $p_{\text{guess}}(Z^A|Z^B)_\psi \geq 1 - \epsilon_1$. By the triangle inequality we have

$$\frac{1}{2} \left\| \sum_{z,z'} p_{zz'} P_z^A \otimes P_{z'}^B \otimes \varphi_{zz'}^E - \sum_{z,z'} p_{zz'} P_z^A \otimes P_z^B \otimes \varphi_{zz'}^E \right\|_1 = \frac{1}{2} \sum_{z,z' \neq z} p_{zz'} + \frac{1}{2} \sum_z \left\| \sum_{z' \neq z} p_{zz'} \varphi_{zz'}^E \right\|_1 \quad (3.46)$$

$$\leq \sum_{z,z' \neq z} p_{zz'} \quad (3.47)$$

$$\leq \epsilon_1. \quad (3.48)$$

But the state $\sum_{z,z'} p_{zz'} P_z^A \otimes P_z^B \otimes \varphi_{zz'}^E$ can be thought of as $U_{\text{CNOT}}^{AB} \left(\sum_{z,z'} p_{zz'} P_z^A \otimes P_0^B \otimes \varphi_{zz'}^E \right) U_{\text{CNOT}}^{\dagger AB}$. From the second condition it follows, for $\varphi^E = \sum_{zz'} \varphi_{zz'}^E$, that

$$\frac{1}{2} \left\| \sum_{z,z'} p_{zz'} P_z^A \otimes P_0^B \otimes \varphi_{zz'}^E - \frac{1}{2} \mathbb{1}^A \otimes P_0^B \otimes \varphi^E \right\|_1 \leq \epsilon_1, \quad (3.49)$$

since the presence of P_0^B doesn’t change the trace distance. Using unitary invariance of the trace distance and the triangle inequality once more completes the proof. \square

To give an approximate characterization of private states based on knowledge of complementary information, we merely need to show that a converse of Corollaries 2 and 3 holds, namely that if Bob can accurately guess the amplitude of Alice’s system, then the phase is unpredictable using the purification E . We formalize this in the following lemma, which is Theorem 4.1 of [Ren11].

Lemma 4. *If there exists a measurement \mathcal{M}_Z^B such that $p_{\text{guess}}(Z^A | \mathcal{M}_Z^B)_\psi \geq 1 - \epsilon_2$ for pure state $|\psi\rangle^{ABE}$, then $p_{\text{secure}}(X^A | E)_\psi \geq 1 - \sqrt{2\epsilon}$.*

Proof. Following the proof of Theorem 1, we know that $\langle \psi_Z | U_{\mathcal{M}_Z}^{B \rightarrow C_Z B} | \psi \rangle^{ABE} \geq 1 - \epsilon$. Using the form of $|\psi_Z\rangle^{AC_Z BE}$ in Equation (3.37), it follows that E is completely decoupled from the phase measurement of A , i.e. the post-measurement state is $\frac{1}{2} \mathbb{1}^A \otimes \psi^E$. Since $U_{\mathcal{M}_Z}^{B \rightarrow C_Z B}$ does not involve A or E , the post-measurement state of AE is the same for $|\psi\rangle^{ABE}$ as for $U_{\mathcal{M}_Z}^{B \rightarrow C_Z B} |\psi\rangle^{ABE}$. Converting fidelity to trace distance, we find that $\frac{1}{2} \left\| \sum_x q_x \tilde{P}_x^A \otimes \vartheta_x^E - \frac{1}{2} \mathbb{1}^A \otimes \psi^E \right\|_1 \leq \sqrt{2\epsilon}$. \square




Figure 3.3: The quantum circuit implementing the (un)twisting operator on the state $\psi^{AA'BB'}$, when Bob can approximately predict Alice's key (amplitude) measurement of system A and there exists a measurement $\mathcal{M}_X^{A'BB'}$ approximately predicting her phase measurement. It proceeds in three steps. First, Bob coherently copies his key (amplitude) to an auxiliary system C_Z using a controlled-NOT gate (unitary $U_{\text{CNOT}}^{BC_Z}$). Next, he coherently performs the measurement \mathcal{M}_X allowing him to predict X , storing the result in auxiliary system C_X (unitary $U_{\mathcal{M}_X}^{C_X A' C_Z B'}$). Finally, to recover a maximally entangled state in system B , he applies another controlled-NOT gate, with control B and target C_X (unitary $U_{\text{CNOT}}^{BC_X}$). Observe that the overall action is a controlled operation with Bob's key as the control and the shield and auxiliary systems the target, i.e. a twisting operator.

With this lemma the following theorem, first shown in [RB08], is immediate.

Theorem 6. *Suppose $p_{\text{guess}}(Z^A | Z^B)_\psi \geq 1 - \epsilon_1$ and there exists a measurement $\mathcal{M}_X^{A'BB'}$ for which $p_{\text{guess}}(X^A | \mathcal{M}_X^{A'BB'})_\psi \geq 1 - \epsilon_2$. Then $\psi^{Z^A Z^B E}$ is an $(\epsilon_1 + \sqrt{2\epsilon_2})$ -good secret key.*

It is also interesting to see how the untwisting operator can be directly constructed using the measurement $\mathcal{M}_X^{A'BB'}$. First write the initial state as $|\psi\rangle^{AA'BB'E} = \sum_{z,z'} \sqrt{p_{z,z'}} |z\rangle^A |z'\rangle^B |\varphi_{z,z'}\rangle^{A'B'E}$ consider the action of a CNOT operation from B to an ancilla system C_Z prepared in the state $|0\rangle^{C_Z}$. This copies the value of z' and gives

$$U_{\text{CNOT}}^{BC_Z} |\psi\rangle^{AA'BB'E} = \sum_{z,z'} \sqrt{p_{z,z'}} |z\rangle^A |z'\rangle^{C_Z} |z'\rangle^B |\varphi_{z,z'}\rangle^{A'B'E}. \quad (3.50)$$

3. CHARACTERIZING QUANTUM INFORMATION

From the first condition it follows that $\langle \psi_Z | U_{\text{CNOT}}^{BC_Z} | \psi \rangle^{AA'BB'E} \geq 1 - \epsilon_1$, where

$$|\psi_Z\rangle^{AC_ZA'BB'E} = \sum_{z,z'} \sqrt{p_{z,z'}} |z\rangle^A |z\rangle^{C_Z} |z'\rangle^B |\varphi_{z,z'}\rangle^{A'B'E}. \quad (3.51)$$

Now make the replacement $A'BB' \rightarrow B$ in this state and apply the latter half of the proof of Theorem 1, from which it follows that

$${}^{AC_Z} \langle \Phi | {}^{C_X A' B B' E} \langle \psi | U_{\text{CNOT}}^{C_Z C_X} V^{A' B B' \rightarrow C_X A' B B'} |\psi_Z\rangle^{AC_Z A' B B' E} \geq 1 - \epsilon_2. \quad (3.52)$$

The fidelity is unchanged by inserting the identity operator in the form $U_{\text{SWAP}}^{BC_Z} U_{\text{SWAP}}^{\dagger BC_Z}$, yielding

$${}^{AB} \langle \Phi | {}^{C_X A' C_Z B' E} \langle \psi | U_{\text{CNOT}}^{BC_X} V^{A' C_Z B' \rightarrow C_X A' C_Z B'} U_{\text{SWAP}}^{BC_Z} |\psi_Z\rangle^{AC_Z A' B B' E} \geq 1 - \epsilon_2. \quad (3.53)$$

Since $U_{\text{SWAP}} U_{\text{CNOT}}^{BC_Z} |\psi\rangle^{AA'BB'E} = U_{\text{CNOT}}^{BC_Z} |\psi\rangle^{AA'BB'E}$, the same method applied to the first fidelity condition gives

$$\langle \psi_Z | U_{\text{SWAP}}^{\dagger BC_Z} U_{\text{CNOT}}^{BC_Z} |\psi\rangle^{AA'BB'E} \geq 1 - \epsilon_1. \quad (3.54)$$

Lemma 1 now implies that the operator $U_{\text{CNOT}}^{BC_X} V^{A' C_Z B' \rightarrow C_X A' C_Z B'} U_{\text{CNOT}}^{BC_Z}$ produces a high-fidelity entangled state in systems AB . But owing to its form, this is a twisting operator, as depicted in Figure 3.3.

Processing Quantum Information

Having concretely developed the relationship between quantum information in the form of entanglement and classical information about complementary amplitude and phase observables in Chapter 3, we may now apply it to the problem of constructing various quantum information processing protocols and understanding why they work. Being able to do so is the second stated goal of this thesis, and will be the subject of this and the remaining chapters. This chapter considers the particular tasks of entanglement distillation, quantum state merging, and secret key distillation in three respective sections. The complementarity approach to the first and last was developed in detail in [RB08], while state merging was treated from this approach in [BR09].

Entanglement distillation is one of the fundamental protocols in quantum information processing and can be used as a building block in a variety of other protocols. In particular, one-way protocols for entanglement distillation can be repurposed for use in reliable communication of quantum information over noisy channels. This allows us to apply our results to that problem and show that the quantum capacity of a channel can be achieved when the sender uses CSS codes. Meanwhile, the secret key distillation results imply that the capacity of a quantum channel to send classical information privately can likewise be achieved when the sender uses CSS codes.

4.1 Optimal Entanglement Distillation

We begin by returning to the problem of entanglement distillation, introduced in Section 2.3. In this setting, Alice and Bob share a supply of identical, somewhat-entangled bipartite resource states which they would like to use to create EPR pairs. An entanglement distillation protocol is a sequence of local operations they should perform on their respective systems, supplemented by classical communication to coordinate their actions and exchange information. The protocol produces approximate EPR pairs at a given rate r , converting n resource states to nr pairs. For instance, the rate of the protocol described in Chapter 2 using the Shor 9-qubit code is given by the rate of the error-correcting code, namely $1/9$, since the output was taken from the encoded subspace of the code. The asymptotically optimal rate is the largest r one can find among protocols for $n \rightarrow \infty$ such that the approximation parameter vanishes in this limit.

As we saw in Chapter 3, Alice and Bob implicitly share an entangled state if Alice's amplitude and phase measurements are predictable by Bob. But a generic bipartite state does not share this property; at best Bob has only partial information about either observable. Heuristically, one way to manufacture entangled states would therefore be to increase Bob's information about these measurements somehow. And since such information is classical, we may be able to arrange for Alice to send it over the classical communication channel. In the following we shall develop this heuristic notion into a concrete protocol. To do so we must first overcome two immediate hurdles. First, what sort of information can she send which will be sufficient for this purpose? And second, how do we make sure Alice does not violate the uncertainty principle when sending information about complementary observables? We take up these two questions in turn.

4.1.1 Information Reconciliation

If we consider either observable alone, the present task is a more general version of the information reconciliation task mentioned in conjunction with QKD in Section 2.4. If we only care about, say, amplitude, then we can imagine Alice measures the amplitude of all of her systems, and these out-

comes are described as a classical random variable. Formally, we can describe the state of all their systems after the measurement by

$$\Psi^{Z^A B} = \sum_{\mathbf{z}} p_{\mathbf{z}} |\mathbf{z}\rangle \langle \mathbf{z}|^A \otimes \varphi_{\mathbf{z}}^B = \sum_{\mathbf{z}} P_{\mathbf{z}}^A \otimes \text{Tr}_A [P_{\mathbf{z}}^A (\psi^{AB})^{\otimes n}] \quad (4.1)$$

using the state $|\psi\rangle^{ABE} = \sum_{\mathbf{z}} p_{\mathbf{z}} |\mathbf{z}\rangle^A |\varphi_{\mathbf{z}}\rangle^{BE}$ and defining $p_{\mathbf{z}} = p_{z_1} \cdots p_{z_n}$ and $\varphi_{\mathbf{z}} = \varphi_{z_1} \otimes \cdots \otimes \varphi_{z_n}$. We use boldface to denote sequences or strings of indices. For each sequence of outcomes Bob is left with the quantum state $\varphi_{\mathbf{z}}^B$, but generally there is no measurement which will indicate which one he has with any accuracy.

However, if Alice gives him some extra information about her outcome \mathbf{z} , then the set of states he is attempting to distinguish between gets smaller, and the task gets easier. For instance, if Alice simply tells him that the sum of the first two outcomes (thought of as binary outcomes) is 0 modulo 2, then he excludes from consideration all the $\varphi_{\mathbf{z}}^B$ for which this is not true and attempts to distinguish between the remaining states with a new measurement. Of course, she could just send Bob her entire measurement record \mathbf{z} , but the goal of information reconciliation is for Alice to transmit as few bits as necessary to enable Bob to reconstruct \mathbf{z} with high probability.

It turns out that in the asymptotic limit $n \rightarrow \infty$, Alice only needs to send information at rate $H(Z^A|B)_{\psi}$. This expression accords with the interpretation of conditional entropy as the uncertainty about Z^A given B : Bob is missing this much information about Z^A and in the protocol Alice simply provides it. Importantly, the information in question can be generated by the technique of *universal hashing* and Alice does not need to know anything about Bob's system except the value of $H(Z^A|B)_{\psi}$. In universal hashing, Alice randomly picks a so-called hash function f from a universal family of hash functions and sends Bob a description of f along with the output $f(\mathbf{z})$.

First defined by Carter and Wegman [83, 84], universal hashing is meant to mimic certain behavior of random functions: A family of functions is universal when the probability that two different inputs to a randomly-chosen family member have the same output is the same as if the function had been chosen at random from all possible functions. This latter probability is simply the inverse of the number of possible function outputs, so formally we say a set \mathcal{F} of functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is universal when

$$\Pr_f [f(x) = f(y)] \leq \frac{1}{2^m} \quad \forall x, y \in \{0, 1\}^n. \quad (4.2)$$

Above we illustrated the information Alice might send to Bob by a linear function, and in fact the set of all linear functions forms a universal family [83, 84]. We shall make extensive use of linear functions for hashing in the next section.

As shown in [RB08], when the size of the hash is roughly $nH(Z^A|B)_{\psi}$ bits, Bob can reliably predict Z^A . More concretely, for each hash value $\hat{\mathbf{z}} = f(\mathbf{z})$ there exists a measurement $\mathcal{M}_{Z;\hat{\mathbf{z}}}^B$ with elements $\Lambda_{z;\hat{\mathbf{z}}}^B$ such that the guessing probability averaged over $\hat{\mathbf{z}}$ is nearly one, $\sum_{\mathbf{z}} p_{\mathbf{z}} \text{Tr}[\Lambda_{z;\hat{\mathbf{z}}}^B \varphi_{\mathbf{z}}^B] \approx 1$. The proof, following ideas from Holevo [85] and Schumacher and Westmoreland [86] in the study of transmission of classical information over quantum channels, explicitly constructs $\mathcal{M}_Z^{BB'}$ as a variant of the *pretty-good measurement* first used by Holevo for pure states [87] and later extended to mixed states (and so-named) by Hausladen and Wootters [88]. Essentially, Bob's measurement is given by

$$\Lambda_{z;\hat{\mathbf{z}}} \approx p_{\mathbf{z}} \varphi_{\hat{\mathbf{z}}}^{-1/2} \varphi_{\mathbf{z}} \varphi_{\hat{\mathbf{z}}}^{-1/2}, \quad \varphi_{\hat{\mathbf{z}}} = \sum_{\mathbf{z}:f(\mathbf{z})=\hat{\mathbf{z}}} p_{\mathbf{z}} \varphi_{\mathbf{z}}, \quad (4.3)$$

with some small modifications. We can simplify the formalism somewhat by imagining that Bob stores the hash value in an auxiliary system B' and uses the measurement $\mathcal{M}_Z^{BB'}$ with elements $\Gamma_z^{BB'} = \Lambda_{z;\hat{z}}^B \otimes P_{\hat{z}}^{B'}$, for which $p_{\text{guess}}(Z^A | \mathcal{M}_Z^{BB'}) \approx 1$.

In the case Bob holds classical information, i.e. the states φ_z are all simultaneously diagonalizable, information reconciliation is closely related to the famous Slepian-Wolf problem of coding of correlated sources [89]. The case of Bob having quantum information was studied and solved in the present i.i.d. scenario by Winter [90] and Devetak and Winter [91] using random coding techniques. In Section 5.1.4 we very briefly describe how the result can be generalized to the case of arbitrary resources.

4.1.2 Reconciling Complementary Information

Having seen that the output of a suitably-sized random hash function enables Bob to reconstruct the outcome of Alice's amplitude or phase measurement, we now turn to the problem of how Alice can generate *both* pieces of information without violating the uncertainty principle. Calling the hash function used for the amplitude measurement f and the phase measurement g , Bob separately requires both $f(\mathbf{z})$ and $g(\mathbf{x})$ so that he can predict the amplitude outcome \mathbf{z} and the phase outcome \mathbf{x} . Naively, it seems impossible to generate both $f(\mathbf{z})$ and $g(\mathbf{x})$, since this would apparently require Alice to measure both the amplitude and phase of her systems.

Crucially, however, the *input* \mathbf{x} (\mathbf{z}) is not required to fix the *output* $g(\mathbf{x})$ ($f(\mathbf{z})$). Instead, Alice need only measure appropriate observables which generate the output directly, *and the necessary observables for $f(\mathbf{z})$ and $g(\mathbf{x})$ can commute*. Such a structure is in fact provided by CSS codes. Recall again the very simple example above, in which Alice transmitted the output of the linear function $z_1 \oplus z_2$ to Bob. As we saw in Section 2.2.1, this can equally-well be thought of as the outcome of measuring the operator $Z_1 Z_2$, since its eigenvalues are $(-1)^{z_1 \oplus z_2}$. But every linear function is a sequence of one-bit functions, each of which is just a sum of particular amplitude outcomes z_k , so to each linear function corresponds a sequence of products of amplitude operators. In other words, every linear function of the amplitude measurement outcome is associated with a collection of Z -type stabilizers, and similarly for X -type stabilizers and functions of the phase measurement outcomes.

If the two functions f and g are chosen so that the corresponding Z - and X -type stabilizers commute, together they define a CSS code, and Alice can then generate both $f(\mathbf{z})$ and $g(\mathbf{x})$ by measuring the stabilizers of the code. The commutation condition on the stabilizers can be succinctly expressed in the following way. For an n -qubit stabilizer, the corresponding linear function can be specified by the n -dimensional binary \mathbb{F}_2 vector with entries 1 at position k if z_k appears in the sum, and zero otherwise. For instance, in the 9-qubit Shor code, the stabilizer $Z_1 Z_2$ corresponds to the vector $(1, 1, 0, 0, 0, 0, 0, 0, 0)$ while the stabilizer $X_1 X_2 X_3 X_7 X_8 X_9$ corresponds to $(1, 1, 1, 0, 0, 0, 1, 1, 1)$. In this representation, two stabilizers commute if the corresponding vectors are orthogonal over \mathbb{F}_2 .

The only requirement on the functions f and g is that they come from universal families \mathcal{F} and \mathcal{G} of hash functions, respectively. Suppose n_Z and n_X are the required number of amplitude and phase type stabilizers, respectively, as determined by the rate requirements of the respective information reconciliation tasks. Then it is easy to show that one simple universal family encompassing both hash functions is the set of $(n_Z + n_X) \times n$ matrices over \mathbb{F}_2 consisting of pairwise orthogonal rows. The first n_Z rows give the Z -type stabilizers and the remaining n_X rows the X -type stabilizers.

Given these stabilizers, it is convenient to think of the code as partitioning Alice's qubits in system A into three different sets of virtual qubits, the encoded qubits in subsystem \bar{A} , the n_Z qubits whose amplitude measurement gives $f(\mathbf{z})$ in \hat{A} , and the n_X qubits whose phase measurement gives

$g(\mathbf{x})$ in \tilde{A} . Then by the properties of the stabilizer operators, $\hat{\mathbf{z}} = f(\mathbf{z})$ and $\tilde{\mathbf{x}} = g(\mathbf{x})$, where $\hat{\mathbf{z}}$ denotes a particular sequence of amplitude measurement outcomes for system \hat{A} .

Now we have all the pieces needed to construct an entanglement distillation protocol. Starting from n copies of the resource state, Alice will measure n_Z Z -type stabilizers and n_X X -type stabilizers and communicate the resulting syndromes to Bob. Then, for large enough n , he will be able to predict Alice's measurement of encoded amplitude and phase operators using the appropriate pretty-good measurements, and thus create an approximate EPR state following Theorem 1.

Formally, they begin with the state

$$|\Psi\rangle^{ABE} = \sum_{\mathbf{z}} \sqrt{p_{\mathbf{z}}} |\mathbf{z}\rangle^A |\varphi_{\mathbf{z}}\rangle^{BE} \quad (4.4)$$

$$= \sum_{\bar{\mathbf{z}}, \hat{\mathbf{z}}, \tilde{\mathbf{z}}} \sqrt{p_{\mathbf{z}}} |\bar{\mathbf{z}}\rangle^{\bar{A}} |\hat{\mathbf{z}}\rangle^{\hat{A}} |\tilde{\mathbf{z}}\rangle^{\tilde{A}} |\varphi_{\mathbf{z}}\rangle^{BE}, \quad (4.5)$$

where in the second line we use the decomposition of Alice's system into the three sets of virtual qubits and consider \mathbf{z} to be a function of $(\bar{\mathbf{z}}, \hat{\mathbf{z}}, \tilde{\mathbf{z}})$. The number n_Z is chosen so that $p_{\text{guess}}(Z^A | \mathcal{M}_Z^{BB'})_{\psi} \approx 1$, where again B' is the system in which Bob stores $\hat{\mathbf{z}}$. Since $\bar{\mathbf{z}}$ is a (linear) function of \mathbf{z} , this implies that $p_{\text{guess}}(\bar{Z} | \mathcal{M}_Z^{BB'}) \approx 1$. As much is true for the phase in that given the value of $\tilde{\mathbf{x}}$ stored in B'' , there exists a measurement $\mathcal{M}_X^{BB''}$ for which $p_{\text{guess}}(\bar{X} | \mathcal{M}_X^{BB''})_{\psi} \approx 1$. Therefore Bob can recover approximate EPR pairs by performing these measurements coherently, as shown in Theorem 1. In this way they can distill $n - n_X - n_Z$ approximate EPR pairs, provided this quantity is positive. Note that here we have only utilized communication from Alice to Bob, making this a *one-way* protocol. Using back and forth communication Alice and Bob could in principle increase the distillation rate, as pointed out by Bennett *et al.* [53] for protocols where *both* parties use quantum error correction, as described in Section 2.3.

4.1.3 Constructing an Optimal Protocol

The final question is how small n_Z and n_X can be made, and here there arises an additional subtlety. From the above discussion, we would expect that $n_Z \approx nH(Z^A|B)_{\psi}$ and $n_X \approx nH(X^A|B)_{\psi}$. However, Alice and Bob can do better. Initially, the purification of their shared state is

$$|\Psi\rangle^{ABE} = \sum_{\mathbf{z}} \sqrt{p_{\mathbf{z}}} |\mathbf{z}\rangle^A |\varphi_{\mathbf{z}}\rangle^{BE}. \quad (4.6)$$

After receiving the amplitude information, Bob has full information about \mathbf{z} , which he could store in system C_Z . Then, for the purposes of predicting Alice's hypothetical phase measurement, it is as if they originally shared (a close approximation to) the following state,

$$|\Psi_Z\rangle^{AC_ZBE} = \sum_{\mathbf{z}} \sqrt{p_{\mathbf{z}}} |\mathbf{z}\rangle^A |\mathbf{z}\rangle^{C_Z} |\varphi_{\mathbf{z}}\rangle^{BE}, \quad (4.7)$$

and this may simplify Bob's phase-prediction task in general.

One might worry that Alice's phase measurement is no longer possible even hypothetically, due to the amplitude stabilizer measurement. However, Bob can still use the conditional marginal states $\vartheta_{\mathbf{x}}^{C_Z B}$ for $\sqrt{q_{\mathbf{x}}} |\vartheta_{\mathbf{x}}\rangle^{C_Z BE} = \frac{1}{\sqrt{2^n}} \sum_{\mathbf{z}} \sqrt{p_{\mathbf{z}}} (-1)^{\mathbf{x} \cdot \mathbf{z}} |\mathbf{z}\rangle^{C_Z} |\varphi_{\mathbf{z}}\rangle^{BE} = \frac{1}{\sqrt{2^n}} (Z^{\mathbf{x}})^{C_Z} |\Psi\rangle^{C_Z BE}$ to build the unitary operator $U_{\mathcal{M}_X}$. This gives him what would have been the phase measurement outcome, and therefore the outcome of the encoded phase measurement. The existence of the former has indeed been destroyed by the amplitude stabilizer measurement, but the latter has not.

A concrete example in which amplitude information is relevant to phase is provided by the following. Suppose the state Alice and Bob share is a maximally-entangled state Φ^{AB} afflicted only with errors of the form XZ . Then amplitude Z errors are completely correlated with phase X errors. Thus, Bob need only know the positions of amplitude errors in order to infer the positions of phase errors. In other words, $H(X^A|BC_Z)_\psi = 0$.

Taking the above consideration into account, the rate of entanglement distillation becomes $r(\psi) = 1 - H(Z^A|B)_\psi - H(X^A|BC)_\psi$, as all the approximation parameters can be taken to arbitrarily small values by choosing a large enough n . It turns out that $r(\psi) = -H(A|B)_\psi$, which we can see by direct computation. First evaluate the latter entropy $H(X^A|C_ZB)_\psi$, using the form of $|\psi_Z\rangle$ derived in (3.37). We find

$$H(X^A|C_ZB)_\psi \equiv H(X^A C_Z B)_\psi - H(C_Z B)_\psi \quad (4.8)$$

$$= H(X^A)_\psi + H(C_Z B|X^A)_\psi - H(C_Z B)_\psi \quad (4.9)$$

$$= 1 - H(C_Z B)_\psi - H(C_Z B)_\psi \quad (4.10)$$

$$= 1 - H(E)_\psi - H(AE)_\psi \quad (4.11)$$

$$= 1 - H(E)_\psi - H(Z^A E)_\psi. \quad (4.12)$$

The first step follows from the general relation between conditional and unconditional von Neumann entropies, while the second follows because the state of $C_Z B E$ conditioned on outcome $X^A = x$ is $(Z^x)^{C_Z} |\psi\rangle^{C_Z B E}$. As these are all unitarily equivalent, each term $H(C_Z B|X^A = x)_\psi$ has the same value $H(C_Z B)_\psi$. In the third step we have used the fact that $H(S_1) = H(S_2)$ for a bipartite pure state on systems S_1 and S_2 . The last step follows because ψ_Z^{AE} is identical to the result of measuring the amplitude of A for the initial state ψ^{AE} . Hence $r(\psi) = H(Z^A|E)_\psi - H(Z^A|B)_\psi$. But since the BE system given the measurement outcome $Z^A = z$ is pure, $H(B|Z^A = z)_\psi = H(E|Z^A = z)_\psi$. Therefore $H(Z^A E)_\psi = H(Z^A B)_\psi$ and $r(\psi) = -H(A|B)_\psi = H(A|E)_\psi$. This rate is sometimes called the *hashing bound*.

Two further modifications lead to the optimal entanglement distillation rate. First, Alice is free to first apply any quantum operation \mathcal{Q}^A to her system before the protocol begins, and this increases the rate to

$$D_1(\psi) = \max_{\mathcal{Q}} \left(-H(A|B)_{\psi_{\mathcal{Q}}} \right). \quad (4.13)$$

Second, the rate can be further improved by *regularization*. Although we have described the protocol above for system A a qubit, it works almost precisely the same for any dimension d which is a prime power.¹ Given a state ψ^{AB} , we could then imagine considering $\Psi^{AB} = (\psi^{AB})^{\otimes m}$ to be the fundamental input to the protocol, and Alice and Bob starting with n copies thereof. The difference is that now Alice and Bob can ignore the product structure of Ψ^{AB} , which leads to the possibly-higher rate

$$D(\psi) = \lim_{n \rightarrow \infty} \frac{1}{n} D_1(\psi^{\otimes n}). \quad (4.14)$$

Devetak and Winter show this rate, the *distillable entanglement*, is in fact optimal in [92].

¹The restriction to prime powers comes from the structure of the stabilizer operators. These require the vector-representation described in Section 4.1.2 which only exists when the symbols come from a finite field.

4.1.4 Quantum Noisy Channel Coding

With a small modification, this entanglement distillation protocol can be used for reliable transmission of quantum information over a noisy channel \mathcal{N} . As mentioned in the discussion of approximate error-correction in Section 3.1, we can always mimic the quantum communication task by sending half of an EPR pair through the channel and measuring the half remaining with Alice in the appropriate basis. Thus, by deferring the measurement indefinitely, we only need to consider reliably transmitting halves of EPR states.

Now consider the entanglement distillation protocol applied to $\psi^{AB} = [\text{id}^A \otimes \mathcal{N}^B](\theta^{AB})$, for an arbitrary pure state $|\theta\rangle^{AB}$. The protocol is constructed so that, averaged over all values that the syndromes could take on, the distilled state closely approximates the ideal of nr EPR pairs. Pick the syndrome with the best approximation parameter, which is surely better than the average. Since in the communication scenario Alice can choose the input, she can always do so in a way which ensures her stabilizer measurement yields precisely this syndrome. Therefore, Alice and Bob could agree on the syndrome value in advance.

But this defines an encoder and decoder in an error-correction scheme! Alice directly creates the bipartite state resulting from measuring the code stabilizers on many instances of θ^{AB} and obtaining the specified syndrome. She then sends Bob's halves through the channel, and he is able to decode the result by applying the entanglement distillation procedure. Since entanglement can be faithfully transmitted, so could any particular single-system state.

Applied to single inputs, this implies that reliable quantum communication must be possible over the channel at rate (here we dispense with the operation \mathcal{Q})

$$Q_1(\mathcal{N}) = \max_{\theta} \left(-H(A|B)_{\psi} \right). \quad (4.15)$$

Despite its nonstandard appearance, this is equal to a maximization over the *coherent information* I_c introduced by Schumacher and Nielsen [93] and more frequently used in this context. To see this, write $|\theta\rangle^{AB} = \sum_k \sqrt{p_k} |k\rangle^A |\vartheta_k\rangle^B$ for some probabilities p_k and normalized states $|\vartheta_k\rangle^B$. The action of the channel on B can be thought of as an isometry $U_{\mathcal{N}}^{B \rightarrow BE}$ and $|\psi\rangle^{ABE} = \sum_k \sqrt{p_k} |k\rangle^A U_{\mathcal{N}}^{B \rightarrow BE} |\vartheta_k\rangle^B$ is the output. Computing the conditional entropy $H(A|B)_{\psi}$ we find

$$Q_1(\mathcal{N}) = \max_{\theta} \left(H(B)_{\psi} - H(R)_{\psi} \right) \quad (4.16)$$

$$= \max_{\vartheta} \left(H(\mathcal{N}(\vartheta)) - H(\mathcal{N}^*(\vartheta)) \right) \quad (4.17)$$

$$\equiv \max_{\vartheta} I_c(\vartheta, \mathcal{N}), \quad (4.18)$$

where $\vartheta = \sum_k p_k \vartheta_k^B$ and \mathcal{N}^* is the channel complementary to \mathcal{N} obtained by applying $U_{\mathcal{N}}^{B \rightarrow BE}$ and keeping R instead of B . In the first line $H(AB)_{\psi} = H(R)_{\psi}$ since ψ is pure, and maximization over θ is equivalent to maximization over ϑ in the second line.

Regularization could improve the result, and we have therefore we have constructed a noisy-channel coding scheme which achieves a rate $Q(\mathcal{N})$, where

$$Q(\mathcal{N}) = \lim_{n \rightarrow \infty} \frac{1}{n} Q_1(\mathcal{N}^{\otimes n}). \quad (4.19)$$

In fact, this is the ultimate capacity of the channel. In a sequence of papers [64, 93, 94, 95], Barnum, Knill, Nielsen, and Schumacher established Q as an upper bound on the capacity, while Lloyd [96], Shor [97], and Devetak [98] used random-coding arguments to show that Q can be attained.

Here we have shown that CSS codes can achieve the capacity, since the resulting code inherits this structure from Alice's use of CSS-type stabilizers in the entanglement distillation protocol. Previously, CSS codes were only known to achieve a lower rate, as implicitly shown by Shor and Preskill [57] and explicitly by Hamada [99]. The more-general stabilizer codes were shown to achieve the capacity by Hayden *et al.* [70].

Devetak's coding scheme has some CSS-like properties in that it essentially consists of an amplitude error-correction step followed by a privacy amplification step. From the discussion of the previous chapter, particularly Lemma 4 but with amplitude and phase trading places, we are tempted to view the latter step as error correction of a phase observable, and indeed we shall examine this in more detail in Section 5.2.1, but the amplitude and phase observables implicitly used in [98] are functions of the coding scheme itself and not identical to the (code-independent) amplitude and phase as we have used here.

One appealing aspect of the use of CSS structure is the possibility of constructing efficiently encodable and decodable codes which approach or even achieve the capacity. For classical communication over classical channels, Forney exhibited such a construction by concatenating random codes with structured codes known as polynomial or Reed-Solomon codes [100]. Hamada has extended this to the quantum case in a sequence of papers [101, 102, 103], but only up to the suboptimal rate mentioned above. It would be interesting to see if the methods presented here can be combined with those of Hamada to reach the capacity efficiently.

4.2 Optimal State Merging

Since the unitary Bob eventually uses to distill the entangled states also transfers the state of Alice's system A to his laboratory, the above protocol can be used for state merging, a process first studied by Horodecki *et al.* [66]. Here the goal is to merge Alice's part ψ^{AB} of the joint state ψ^{AB} with Bob's so that he ends up with ψ^{AB} , using as little quantum or classical communication as possible. Additionally, if we consider the purification $|\psi\rangle^{ABE}$, all correlations with the purifying system R should be transferred to Bob as well. Not only should Bob end up with a good approximation to ψ^{AB} , but together with R the final state should closely approximate $|\psi\rangle^{ABE}$.

When ψ^{AR} is itself pure, state merging reduces to quantum data compression. Since Bob has no initial information about Alice's state, whatever she sends must be sufficient to reconstruct her state and can be regarded as the compressed version of it. Schumacher has shown that a state ψ^A can be compressed at rate no greater than $H(A)$ [15], meaning Alice and Bob will need to use a quantum channel at this rate.

However, when Bob's system is correlated with Alice's, they can take advantage of these correlations to reduce the amount of communication needed. Indeed, if Alice and Bob share the EPR state $|\Phi\rangle^{AB}$, then no communication is required at all! This follows because a maximally-entangled state is not correlated with any third system, and so Bob can simply recreate the state at his end. For example, applying Theorem 1 to the input state $|\Phi\rangle^{AB}$ yields output $|\Phi\rangle^{AD}|\Phi\rangle^{BC}$ upon application of the partial isometry $U^{B \rightarrow BCD}$.

Moreover, sometimes sending only classical information is sufficient for transferring a quantum state. This is precisely the case when using the entanglement distillation protocol, which works for all ψ^{AB} such that $H(A|B) < 0$. That classical communication is sometimes sufficient is somewhat surprising, but with entanglement Alice could teleport her system to Bob using only classical information, and this is effectively what happens as a byproduct of the entanglement distillation protocol. By expressly using teleportation, we can also apply the distillation protocol to cases when

$H(A|B) > 0$, as described in [66]. For n resource states ψ^{AB} Alice and Bob can create $nH(A|B)_\psi$ EPR pairs to go with their n resource states, and the overall conditional entropy of the entire collection of systems is now roughly zero. Running the entanglement distillation protocol produces no new EPR pairs, but does transfer Alice's part of the resource state to Bob.

Horodecki *et al.* have shown that state merging requires quantum communication at the rate $H(A|B)_\psi$ when this quantity is positive, but only classical communication at rate $I(A : E)_\psi$ when $H(A|B)_\psi$ is negative [66, 67]. Using the entanglement distillation procedure above is therefore optimal in the first setting but not always in the second, as the rate of classical communication needed is $(n_Z + n_X)/n = 1 - H(A|E)_\psi \geq I(A : E)_\psi$.

However, we can make a small alteration to the protocol to make it optimal, as shown in [BR09]. Observe that when $H(A)_\psi = 1$, the procedure is in fact optimal. This suggests that we ought to first compress system A and then perform entanglement distillation. The difficulty in making this work is to ensure that the compression step does not interfere with the amplitude and phase information reconciliation steps. Since compression of quantum systems can be thought of as essentially just classical compression in the eigenbasis, it simplifies matters to choose the amplitude basis to be the eigenbasis of Alice's state ψ^A .

Formally, the tripartite system ABE starts in the pure state given in Equation (4.6). The compressor projects the system onto a subspace spanned by a set of eigenvectors $|\mathbf{z}\rangle$ whose total probability is nearly equal to one, a so-called *typical set*. Even though the typical set contains almost all of the probability, it only contains roughly $2^{nH(Z)}$ of the 2^n total eigenvectors. Thus, with probability nearly one the projection operation succeeds and the subspace needed to support the state drastically shrinks. Rarely, the projection operation fails, and the state must be written off as a total loss.

When the compressor succeeds, the state can be expressed as

$$|\psi'\rangle^{ABE} = \frac{1}{\mathcal{N}} \sum_{\mathbf{z} \in \text{Typ}} \sqrt{p_{\mathbf{z}}} |\mathbf{z}\rangle^A |\varphi_{\mathbf{z}}\rangle^{BE}, \quad (4.20)$$

where Typ is the typical set and \mathcal{N} is the required normalization factor. On the *typical subspace* we can order the basis elements lexicographically and define a new amplitude observable Z' as in Equation (3.25), as well as the phase observable corresponding to the shift operator of said basis. After the compression step, the idea is for Alice and Bob to run the entanglement distillation procedure for the new observables Z' and X' . However, the distribution of measurement results for these two operators is no longer i.i.d., and thus the results of information reconciliation we used previously no longer apply. We have no direct way of knowing how many stabilizers Alice should measure, nor how Bob should construct his measurement.

This poses no serious problem for the new amplitude observable, since it is essentially the same as the old one, just missing the non-typical values. Indeed, the information reconciliation protocol also makes use of typicality in that Bob's measurement does not bother to look for non-typical \mathbf{z} in the first place. Thus, explicitly rejecting these possibilities in the compression step will only serve to reduce the error probability for information reconciliation of Z' . Alice can perform precisely the same Z -type stabilizer measurements as before, and Bob's original measurement will accurately reconstruct Z and therefore Z' .

However, this sort of argument does not work for the new phase observable X' . Since X and X' are not so simply related, Bob's knowledge of X generally does not pertain at all to his knowledge of X' . Luckily, the extra system C_Z which was used to achieve the optimal entanglement distillation rate comes the rescue. In the entanglement distillation protocol it gave Bob's marginal states conditioned on Alice's phase measurement a group-covariant structure, and it does so in the present scheme as

well. In turn, this makes it possible to transform the information reconciliation protocol in the original i.i.d. setting to one appropriate for the new non-i.i.d. setting.

After the amplitude information reconciliation step, phase information reconciliation proceeds as if Alice and Bob shared the state

$$|\psi'_z\rangle^{ABE} = \frac{1}{\mathcal{N}} \sum_{z \in \text{Typ}} \sqrt{p_z} |z\rangle^A |z\rangle^C |\varphi_z\rangle^{BE}. \quad (4.21)$$

The group covariance arises just as before, due to the “copy” of \mathbf{z} in system C_Z . The required number of X^A -stabilizer outcomes must be computed in the construction of the reconciliation protocol, and it turns out to be $n_{X'} = H(A)_{\psi_z} + H(C_Z B | X^A)_{\psi_z} - H(C_Z B)_{\psi_z}$. Following the calculation in the previous section, this is just $n_{X'} = H(A)_{\psi} - H(Z^A | E)_{\psi}$. For the first term we have used the fact that $H(A)_{\psi_z} = H(Z^A)_{\psi} = H(A)_{\psi}$ since the amplitude basis is the eigenbasis. The communication cost of the protocol is now $n_Z + n_{X'} = H(A)_{\psi} + H(Z^A | B)_{\psi} - H(Z^A | E)_{\psi} = I(A : E)_{\psi}$.

Since they are working in the typical subspace, Alice and Bob can expect to extract roughly $n \log |\text{Typ}| - n_Z - n_{X'}$ entangled pairs, where $|\text{Typ}| \approx 2^{nH(A)}$ is the size of the typical set. This works out to an entanglement distillation rate of $H(A)_{\psi} - I(A : E)_{\psi} = -H(A | B)_{\psi}$, just as before. Therefore, by adding a compression step and choosing the amplitude basis to be the eigenbasis of Alice’s system, we have managed to convert the optimal entanglement distillation protocol into an optimal state merging protocol.

4.3 Secret Key Distillation and Private Communication

Section 3.4 detailed the close connection between private and entangled states, and in this section we show that the same methods used in Section 4.1 to construct entanglement distillation protocols can be used to construct protocols for creating a shared secret key from a supply of bipartite quantum states. Due to the CSS nature of this approach, we really only need to construct a private state distillation scheme, and it will work for secret key distillation as well. As explained in Section 2.4, Alice and Bob ultimately only need to ensure that the information about Alice’s hypothetical phase measurement is somewhere to be found in the systems under their control.

The private state distillation protocol works almost exactly as the entanglement distillation protocol. Given n copies of the resource state ψ^{AB} , Alice is free to decide how to define the prospective key and first performs a quantum operation $\mathcal{Q}^{A \rightarrow AA'}$ which maps her system A into two systems AA' . The first is used as the key and the second as a shield. This operation may additionally involve a measurement whose outcome T is publicly transmitted to Bob, and the resulting state is $\psi_{\mathcal{Q}}^{AA'BT} = \sum_t \mathcal{Q}_t^{A \rightarrow AA'}(\psi^{AB}) \otimes |t\rangle \langle t|^T$.

In the second step Alice measures enough amplitude and phase stabilizers on A so that the amplitude Z^A can be reconstructed from system B and the phase X^A from the compound system $A'B$. The number of stabilizers needed is set by the requirements for information reconciliation of each task separately, and again the amplitude information may be useful in recovering the phase information. Therefore the number of stabilizers needed amounts to $n_Z \approx nH(Z^A | BT)_{\psi_{\mathcal{Q}}}$ and $n_X \approx nH(X^A | C_Z A' BT)_{\psi_{\mathcal{Q}, Z}}$, where $\psi_{\mathcal{Q}, Z}^{AC_Z A' B}$ is the state defined by coherently copying the ampli-

tude in A to system C_Z . Alice may choose the optimal operation \mathcal{Q} , yielding the distillation rate

$$\kappa_1(\psi) = \max_{\mathcal{Q}} \left(1 - H(Z^A|BT)_{\psi_{\mathcal{Q}}} - H(X^A|C_Z A' B B' T)_{\psi_{\mathcal{Q}, Z}} \right) \quad (4.22)$$

$$= \max_{\mathcal{Q}} \left(H(Z^A|RT)_{\psi_{\mathcal{Q}}} - H(Z^A|BT)_{\psi_{\mathcal{Q}}} \right), \quad (4.23)$$

where the second line follows by the same calculations which led to Equation (4.12).

Alice need only transmit the amplitude syndromes to Bob since they use the encoded amplitude \bar{Z} as the final key. The phase syndromes need not be transmitted, since according to Theorem 6 the mere existence of a phase-predicting measurement $\mathcal{M}_X^{C_Z A' B B'}$ ensures the secrecy of the key. This means the protocol can be immediately converted into a secret-key distillation scheme in which Alice and Bob make their amplitude measurements first, Alice then transmits the amplitude syndromes, and finally Alice and Bob compute the final key from the encoded amplitude operator \bar{Z} . From the outside, they *could* have actually run the private state distillation protocol, phase stabilizer measurement and all, and so the secret key distillation protocol inherits security from the private state distillation protocol.

Regularization can again in principle increase the rate further, and the resulting rate is identical to the upper bound found by Devetak and Winter [92]. Thus we have constructed a secret key distillation protocol which achieves the optimal rate

$$\mathcal{K}(\psi) = \lim_{n \rightarrow \infty} \frac{1}{n} \kappa_1(\psi^{\otimes n}). \quad (4.24)$$

Given a shared, secret key Alice can transmit secret messages to Bob over a public communication channel simply by encrypting the message with the key. For absolute security, Shannon showed that one requires a key exactly as long as the message [54], and the message may be encrypted by simply computing the exclusive-OR of the key, a scheme known as a *one-time pad* or *Vernam cipher* after its inventor [104].

Therefore Alice and Bob may use the secret key distillation scheme above for private communication over public channels. As Alice can choose the input to the channel, she may simply select that input which gives the output with the largest distillable key. Then they proceed with secret key distillation and the one-time pad. This gives a private communication rate of at least $P_1(\mathcal{N})$ using κ_1 above, at least when assisted with public communication. This quantity is sometimes referred to as the *private information* and we shall encounter it again in Section 6.3. Once more, regularization may improve the rate, and the resulting expression $P(\mathcal{N})$ was shown to be an upper bound in [98]. The the protocol for private communication constructed in this way achieves the capacity. Here we have not attempted to remove the public communication from Alice to Bob as we did in the case of quantum communication, but it is also shown in [98] that the private capacity can be achieved even without such assistance.

Duality of Protocols

In Chapter 4 we saw that reconciling Bob’s quantum information in system B with Alice’s amplitude observable Z^A requires her to send Bob extra information about Z^A at rate $H(Z^A|B)_\psi$. This quantity trades off with $H(X^A|E)_\psi$ in the uncertainty principle Equation (1.7), $H(Z^A|B)_\psi + H(X^A|E)_\psi \geq 1$. As it happens, $H(X^A|E)_\psi$ is also the rate at which Alice can perform privacy amplification of X^A , extracting uniformly-distributed bits from X^A which are completely uncorrelated with E . Thus, the less information Alice has to send to Bob about Z^A , the more randomness she can extract from X^A unknown to E . There exists a duality between these two protocols due to the uncertainty principle. The fact that the rates of the two protocols are connected invites us to think that the protocols themselves may be connected as well—that it may be able to transform one protocol into the other.

Here we show that this is indeed the case, recounting results from [Ren11, RR11] and presenting some new material. This chapter is divided into four main sections. In the first, we recount how information reconciliation and privacy amplification protocols based on linear hash functions can be transformed into each other, following [Ren11]. The duality extends to non-i.i.d. resources where the notion of asymptotic rates is no longer valid, and we remark that this implies a more general form of the uncertainty principle in terms of generalized entropies suitable for such unstructured resources. In the second section, we explore the implications of this duality for constructing entanglement distillation protocols, and by extension, the other related protocols discussed in Chapter 4. This material has not been previously published. The third section is devoted to the result of [RR11] which shows that coding schemes for communication of either public or private classical information over noisy channels can be constructed by combining privacy amplification and information reconciliation. Thus, the two dual protocols occupy a very fundamental place in the study of information theory, as they can be combined to generate a variety of protocols for other tasks.

5.1 Duality of Privacy Amplification and Information Reconciliation

The duality of information reconciliation and privacy amplification protocols both based on linear universal hashing essentially comes down to complementarity, specifically the fact that amplitude measurements destroy phase information and vice versa. Roughly speaking, if Alice measures amplitude stabilizers to perform information reconciliation of Z^A with Bob, this can also be seen as randomizing the conjugate phase X^A stabilizers, as would be useful in privacy amplification. With Lemma 4 in mind, we expect that if information reconciliation succeeds and Bob can reliably recover the encoded amplitude \bar{Z} , then the encoded phase \bar{X} must be uncorrelated with system E . Making this work in reverse is slightly more complicated, and there are two versions, corresponding to Corollaries 2 and 3.

5.1.1 Privacy Amplification

Before delving into the duality of these protocols, we first describe the process of privacy amplification and the known results in more detail. Imagine that Alice has an n -bit classical random variable X^A which is correlated with an external system E in some way. Letting X^A be the phase observable, we can describe this state of affairs as

$$\Psi^{AE} = \sum_{\mathbf{x}} q_{\mathbf{x}} |\mathbf{x}\rangle \langle \mathbf{x}|^A \otimes \vartheta_{\mathbf{x}}^E. \quad (5.1)$$

If the $\vartheta_{\mathbf{x}}$ were identical for all \mathbf{x} , then E would have no information about the value of X^A . Conversely, if the $\vartheta_{\mathbf{x}}$ have disjoint supports, then a measurement of E projecting onto these supports can determine \mathbf{x} without error.

First introduced by Bennett, Brassard, and Robert [105, 106], the goal of privacy amplification is twofold, to compute some function $\bar{X}^A = f(X^A)$ of X^A which is both uniformly distributed and independent of E . Keeping only the function output $f(\mathbf{x})$ means that, for a given output $\bar{\mathbf{x}}$, the state in E is averaged over all the \mathbf{x} for which $f(\mathbf{x}) = \bar{\mathbf{x}}$. The goal is then to average over enough values of \mathbf{x} so that the conditional states $\vartheta_{\mathbf{x}} = \sum_{\mathbf{x}:f(\mathbf{x})=\bar{\mathbf{x}}} q_{\mathbf{x}} \vartheta_{\mathbf{x}}^E$ are identical for all \mathbf{x} . Of course, it is unrealistic to expect such an ideal output, so we settle for $p_{\text{secure}}(\bar{X}^A | E) \geq 1 - \epsilon$. When the state Ψ^{AE} is n instances of a state ψ^{AE} pertaining to a single bit in A , the asymptotically-optimal rate at which private random bits can be extracted is defined by the largest rate achievable in the simultaneous limits $n \rightarrow \infty$, $\epsilon \rightarrow 0$.

In the case that the states $\vartheta_{\mathbf{x}}^E$ are classical, i.e. simultaneously diagonalizable, Bennett *et al.* have shown that universal hashing can be used for privacy amplification [105, 106, 107]. Using random coding techniques in the i.i.d. setting, Devetak and Winter proved that the rate $H(Z^A | E)_{\psi}$ is achievable in the asymptotic limit for quantum $\vartheta_{\mathbf{x}}^E$ [92], while Renner and König show that universal hashing is also effective against quantum adversaries even for unstructured, non-i.i.d. resources [108].

One drawback of approaches based on universal hashing is the need for a large amount of randomness to select the hash function from the family, $\Theta(n)$ seed bits for n input bits. Smaller function families would naturally be preferable. If we are unconcerned with privacy, the task reduces to extracting the maximum amount of randomness inherent in the distribution of Z^A , and constructing efficient *extractors* has been the subject of much research in theoretical computer science (see e.g. Shaltiel [109] for a review).

In particular, Trevisan's breakthrough construction showed that essentially all the randomness may be extracted from the input using extractors with seeds of size $O(\text{polylog}(n))$ [110, 111]. Recently De *et al.* showed that Trevisan's construction can be extended to privacy amplification against quantum adversaries [112].

5.1.2 Privacy Amplification from Information Reconciliation

Now we examine how an information reconciliation protocol using linear functions for universal hashing can be used for privacy amplification. Use of CSS codes makes this simple. Consider, as usual, a tripartite pure state $|\psi\rangle^{ABE}$. Instead of taking system A to be a qubit, we now assume that it has dimension 2^n for some n . This can be done without loss of generality by embedding A into a state space larger than the support of ψ^A , and allows us to think of system A as a collection of n qubits.

Suppose that there exists a protocol for information reconciliation of Bob's information with the Alice's amplitude Z^A which calls for Alice to compute a linear function of Z^A and send it to Bob. This computation can be thought of as measuring the stabilizers of a CSS code which contains only Z -type stabilizers. In terms of virtual qubits as described in Section 2.2.1, the entire collection of qubits can be grouped into two subsets, the encoded qubits and the stabilizer qubits. Denoting the amplitude values of the encoded qubits by $\bar{\mathbf{z}}$ and those of the stabilizer qubits by $\hat{\mathbf{z}}$, we can express the initial state as (abusing notation slightly)

$$|\psi\rangle^{ABE} = \sum_{\mathbf{z}} \sqrt{p_{\mathbf{z}}} |\mathbf{z}\rangle^A |\varphi_{\mathbf{z}}\rangle^{BE} = \sum_{\bar{\mathbf{z}}, \hat{\mathbf{z}}} \sqrt{p_{\bar{\mathbf{z}}, \hat{\mathbf{z}}}} |\bar{\mathbf{z}}\rangle^{\bar{A}} |\hat{\mathbf{z}}\rangle^{\hat{A}} |\varphi_{\bar{\mathbf{z}}, \hat{\mathbf{z}}}\rangle^{BE}, \quad (5.2)$$

where \bar{A} (\hat{A}) denotes the virtual subsystem of the encoded (stabilizer) qubits.

The information reconciliation protocol assures us that given the value of \hat{z} , Bob can determine the value of z and therefore \bar{z} . That is, there exists a measurement on $\hat{A}B$ which can reliably predict the amplitude of \bar{A} with guessing probability greater than $1 - \epsilon$ for some small ϵ . Then by Lemma 4, $p_{\text{secure}}(X^{\bar{A}}|E)_\psi \geq 1 - \sqrt{2\epsilon}$. Therefore, to generate a random secret string from the phase observable X^A , Alice can simply compute the encoded phase \bar{X} .

5.1.3 Information Reconciliation from Privacy Amplification

Showing that a privacy amplification protocol can be repurposed for information reconciliation is somewhat more involved. Here we encounter the same complications as in Section 3.2: Just because E has no knowledge of X^A does not imply that B can predict Z^A . But the same technique used there of imposing extra conditions so that the uncertainty principle is saturated works here as well. There are two separate cases to consider.

In the first of these we require $p_{\text{guess}}(Z^A|E)_\psi = 1$, meaning we might as well write the state as

$$|\psi\rangle^{ABE} = \sum_z \sqrt{p_z} |z\rangle^A |z\rangle^{E_1} |\varphi_z\rangle^{BE_2}, \quad (5.3)$$

for $E = E_1E_2$. This is somewhat more natural for the goal of amplitude information reconciliation, as it ensures that the AB state describes a classical variable in A and a quantum state in B : $\psi^{AB} = \sum_z p_z P_z^A \otimes \varphi_z^B$.

Now suppose that there exists an encoded \bar{X} such that $p_{\text{secure}}(\bar{X}^A|E)_\psi \geq 1 - \epsilon$. Again using the encoded and stabilizer qubits for system A , it follows from Corollary 3 that there exists a measurement $\mathcal{M}_{\bar{Z}}$ on $\hat{A}B$ which can recover \bar{Z} with error probability less than $\sqrt{2\epsilon}$. However, Bob does not have access to \hat{A} , and Alice must take care in what information she sends to Bob, lest it leak any information about the phase to E . Intuitively, however, measuring amplitude stabilizers on \hat{A} destroys any phase information that might be present, so it should be safe to transmit the resulting syndromes to Bob.

Indeed, the formal nature of the state shared by Alice and Bob makes this clear, since \hat{A} is effectively already measured. Tracing out E , we obtain

$$\psi^{AB} = \sum_{\bar{z}, \hat{z}} p_{\bar{z}, \hat{z}} P_{\bar{z}}^{\hat{A}} \otimes P_{\hat{z}}^{\hat{A}} \otimes \varphi_{\bar{z}, \hat{z}}^B. \quad (5.4)$$

Due to the classical structure of system \hat{A} , we can assume without loss of generality that the measurement $\mathcal{M}_{\bar{Z}}^{\hat{A}B}$ has this structure, too. For let $\Lambda_{\bar{z}}^{\hat{A}B}$ be the POVM elements of the $\mathcal{M}_{\bar{Z}}^{\hat{A}B}$ and consider the joint probability of obtaining the outcome $\mathcal{M}_{\bar{Z}}^{\hat{A}B} = \bar{z}'$ and $\bar{Z} = \bar{z}$,

$$\Pr[\mathcal{M}_{\bar{Z}}^{\hat{A}B} = \bar{z}', \bar{Z} = \bar{z}] = \sum_{\hat{z}} p_{\bar{z}, \hat{z}} \text{Tr}[\Lambda_{\bar{z}'}^{\hat{A}B} P_{\hat{z}}^{\hat{A}} \otimes \varphi_{\bar{z}, \hat{z}}^B]. \quad (5.5)$$

Clearly the same probability results if we first determine the value of \hat{z} and then use a POVM on B having elements $\Pi_{\bar{z}, \hat{z}}^B = \text{Tr}[P_{\hat{z}}^{\hat{A}} \Lambda_{\bar{z}}^{\hat{A}B}]$. But this is precisely how we expected the information reconciliation process to work: after learning \hat{z} , Bob can measure B and recover \bar{z} .

In the second case we require $p_{\text{guess}}(X^A|B)_\psi = 1$, so that Bob already has information about the phase. Should he also learn the amplitude, Alice and Bob would have created an entangled state, so this scenario is essentially the latter half of an entanglement distillation scheme. In fact, the protocol of Devetak and Winter in [92] is constructed along these lines. Just as in the previous scenario, if a privacy amplification protocol can construct an encoded phase \bar{X} uncorrelated with E , then the conjugate encoded amplitude \bar{Z} must be reliably recoverable by measurement on $\hat{A}B$, though now the implication follows from Corollary 2. However, we cannot use the same argument to show that measurement of the amplitude of \hat{A} is sufficient to enable information reconciliation using system B .

Instead, we can proceed as follows. From the requirement $p_{\text{guess}}(X^A|B)_\psi = 1$, the marginal state of the AE subsystems takes the form $\psi^{AE} = \sum_{\mathbf{x}} q_{\mathbf{x}} P_{\mathbf{x}}^A \otimes \vartheta_{\mathbf{x}}^E$ for some probabilities $q_{\mathbf{x}}$ and normalized states $\vartheta_{\mathbf{x}}^E$. Decomposing Alice's qubits into virtual encoded and stabilizer qubits, the state is, in a slight abuse of notation, just $\psi^{AE} = \sum_{\bar{\mathbf{x}}, \hat{\mathbf{x}}} q_{\bar{\mathbf{x}}, \hat{\mathbf{x}}} P_{\bar{\mathbf{x}}}^{\bar{A}} \otimes P_{\hat{\mathbf{x}}}^{\hat{A}} \otimes \vartheta_{\bar{\mathbf{x}}, \hat{\mathbf{x}}}^E$. Since the \hat{A} system is in a phase eigenstate, measuring its amplitude delivers a completely random outcome and results in precisely the same state as if \hat{A} were traced out. But the encoded phase is chosen by the privacy amplification protocol so that disposing of the stabilizer qubits leaves a nearly ideal key, and the amplitude measurement of the stabilizer qubits does not change this. Thus, for every measurement result we can conclude by Corollary 2 that there exists a measurement on B which gives \bar{z} with high probability.

In both of these situations the desired measurement is only shown to exist, but is not directly constructed. However, due to a result by Barnum and Knill, this presents no real difficulty, as the pretty-good measurement has an error probability which is at most a factor of two worse than the optimal case [95]. Thus, if privacy amplification is possible so that $p_{\text{secure}}(\bar{X}^A|E)_\psi \geq 1 - \epsilon$, then using the amplitude stabilizer measurement and the pretty good measurement for Bob's conditional marginal states results in information reconciliation protocols with error probability less than $2\sqrt{2\epsilon}$.

5.1.4 One-Shot Protocols and a Generalized Uncertainty Principle

In the preceding sections we have treated Alice's system as a collection of n qubits, but it is important to note that the duality holds for arbitrary resource states, not just i.i.d. states. The i.i.d. setting is only necessary to define the asymptotically-achievable rates of the various protocols. Recently, a new framework has been constructed which makes it possible to characterize protocols operating on arbitrary, structureless resource states in terms of *smooth entropies*. A proper treatment of smooth entropies and their calculus is beyond the scope of this thesis, but we remark that they can be thought of as generalizations of Rényi entropies which are somewhat more familiar in standard information theory and obey many of the same chain rules as the usual Shannon or von Neumann entropies. Here we wish to point out that the duality above, in particular the former duality of Section 5.1.3, implies a new entropic uncertainty principle formulated in terms of smooth entropies.

There are two different smooth entropies, the smooth min-entropy and the smooth max-entropy, and each comes in both conditional and unconditional varieties. It turns out that the number $\ell_{\text{ext}}^\epsilon(X^A|E)_\psi$ of ϵ -good random bits one can extract from Z^A which are secret from E is characterized by the smooth min-entropy, $\ell_{\text{ext}}^\epsilon(X^A|E)_\psi \approx H_{\text{min}}^\epsilon(Z^A|E)_\psi$ [108, 113, 114, 115]. More precisely, $\ell_{\text{ext}}^\epsilon(X^A|E)_\psi$ equals $H_{\text{min}}^\epsilon(Z^A|E)_\psi$ up to small deviations involving the smoothing parameter ϵ . Much the same holds for information reconciliation, except using the smooth max-entropy. As shown by the present author and Renner [116], the number of bits Alice needs to send to Bob, generated by universal hashing, is given by $\ell_{\text{rec}}^\epsilon(Z^A|B)_\psi \approx H_{\text{max}}^\epsilon(Z^A|B)_\psi$. Though it might not appear so, the definitions of the smooth entropies are logically distinct from the operational quantities $\ell_{\text{ext}}^\epsilon$ and $\ell_{\text{rec}}^\epsilon$.

It should be noted, however, that the smooth entropies are themselves related to the operational quantities p_{guess} and p_{secure} , a fact discovered by König *et al.* [117].

Now consider a quantum state of the form given in Equation (5.2). Information reconciliation of the amplitude requires that Alice send $\ell_{\text{rec}}^{\epsilon}(Z^A|B)_{\psi}$ bits obtained via universal hashing of Z^A to Bob. But this implies Alice can equally-well use the encoded phase to generate random bits uncorrelated with E . In all she can create $n - \ell_{\text{rec}}^{\epsilon}(Z^A|B)_{\psi}$ random bits this way, which must of course be less than the bound on privacy amplification established by the smooth min-entropy. Similarly, $\ell_{\text{rec}}^{\epsilon}(Z^A|B)_{\psi}$ is bounded by the smooth max-entropy, so we anticipate from this heuristic argument that

$$H_{\min}^{\epsilon}(X^A|E)_{\psi} + H_{\max}^{\epsilon}(Z^A|E)_{\psi} \gtrsim n. \quad (5.6)$$

Indeed, the full analysis performed in [Ren11] shows that the above expression is correct, up to terms of order $\log(1/\epsilon)$. The state in Equation (5.2) is arbitrary, so this generalized uncertainty principle holds for conjugate observables and any tripartite quantum state. Recently, Tomamichel and Renner have found a simple proof which extends the above uncertainty relation to arbitrary observables in the manner of Equation (1.7) [118].

5.2 Different Approaches to Entanglement Distillation

The entanglement distillation protocol presented in Section 4.1 was built by combining information reconciliation protocols for both Alice’s amplitude and phase observables. By the duality of information reconciliation and privacy amplification, we expect to be able to trade one task for the other, and base the construction of the protocol on either Theorem 2 or Theorem 3 rather than Theorem 1. In the following we present these two alternate approaches. It should be stressed that ultimately the alternate approaches followed here yield the same protocol as in Section 4.1, but they have completely independent justifications.

In the first approach, we may think of the phase information reconciliation in the original protocol as amplitude privacy amplification, which makes the goal of entanglement distillation to simultaneously give Bob full information about Alice’s amplitude while ensuring that E has none. Formally, the goal in constructing the protocol is to fulfill the conditions of Theorem 2. Clearly this approach is quite closely related to secret-key distillation, which has nearly the same goals, and indeed was the original approach followed by Devetak and Winter [92] for entanglement distillation and Devetak in establishing the quantum capacity of a quantum channel [98]. Here we construct an entanglement distillation protocol having the same aims but a somewhat different structure, namely the use of CSS codes by Alice.

In the second approach, we can give up on Bob altogether and focus entirely on removing amplitude and phase correlations from E , with the aim of fulfilling the conditions of Theorem 3. To our knowledge, this approach is new. It shows that the commonly used quantum decoupling method can be broken down into two classical decoupling steps, further reinforcing the claim that quantum information processing can be understood as a combination of classical information processing of amplitude and phase information. Figure 5.1 depicts the relationship between the three approaches.

5.2.1 Amplitude Information Reconciliation & Privacy Amplification

Although the approach based on Theorem 2 is substantially similar to that pursued in [92], we include it here for completeness. Again we consider the case in which Alice and Bob share asymptotically many copies of a resource state ψ^{AB} which may be purified to $|\psi\rangle^{ABE}$. We will construct the protocol




Figure 5.1: Breakdown of Alice’s n physical qubits into three subsets of virtual qubits in subsystems \bar{A} , \hat{A} , and \tilde{A} and what the different subsets are used for in the various approaches to entanglement distillation presented here. Theorem 1 is the goal of the construction in Section 4.1.3, where \hat{A} is used to reconcile the amplitude information with Bob and \tilde{A} the phase information. The construction in Section 5.2.1 takes Theorem 2 as its goal, and \hat{A} is again used for amplitude information reconciliation with Bob, but \tilde{A} is used for privacy amplification of the same amplitude information against the environment. Finally, Theorem 3 is the aim of construction in Section 5.2.2, where \hat{A} is used to decouple Alice’s phase information from the environment and \tilde{A} her amplitude information.

by choosing two sets of amplitude stabilizers, first a number n_Z large enough to enable information reconciliation with Bob and the second n_X to achieve privacy amplification against E . Thinking in terms of virtual qubits and their associated amplitude and phase operators, let us call the encoded amplitude operators \bar{Z} , of which there are $n - n_Z - n_X$, the n_Z stabilizers associated with information reconciliation \hat{Z} , and those n_X associated with privacy amplification \tilde{Z} .

To ensure that Bob can reconstruct the original amplitude, and therefore the encoded \bar{Z} , Alice measures the \hat{Z} stabilizers and sends the resulting syndromes to Bob. This could give additional information about \bar{Z} to E , but if the \tilde{Z} stabilizers are numerous enough, averaging over their syndromes destroys whatever information E had about the original amplitude Z . By itself, \tilde{Z} is independent of \bar{Z} , since they belong to different sets of virtual qubits, so Alice can be certain that no information leaks to E in this process.

We are not ready to apply Theorem 2, however. The shared state at this step in the protocol is

$$|\Psi'\rangle_{\bar{A}\tilde{A}BB'E'} = \sum_{\bar{z}, \hat{z}, \tilde{z}} \sqrt{p_{\bar{z}, \hat{z}, \tilde{z}}} |\bar{z}\rangle_{\bar{A}} |\tilde{z}\rangle_{\tilde{A}} |\hat{z}\rangle_{B'} |\hat{z}\rangle_{E'} |\varphi_{\bar{z}, \hat{z}, \tilde{z}}\rangle_{BE}, \quad (5.7)$$

where the amplitude of \hat{A} has been transferred and copied to new systems B' and E' , which mimics the classical measurement of \hat{Z} and broadcast of the result \hat{z} . From information reconciliation there is a measurement $\mathcal{M}_{\bar{z}}$ on BB' such that $p_{\text{guess}}(\bar{Z}|\mathcal{M}_{\bar{z}}^{BB'})$ is close to one, and via the above discussion of privacy amplification $p_{\text{secure}}(\bar{Z}|EE')$ is likewise nearly one. To apply Theorem 2 we still need to discard \tilde{A} without changing either of these conditions.

This situation is precisely that of the second case of the previous section, from which it follows that measuring the phase \tilde{X} will not decrease Bob’s guessing probability and will also not leak any information about Z to E . Formally, we can see this by examining the state after the phase stabilizer

measurement,

$$|\Psi''\rangle^{\bar{A}BB'B''EE'E''} = \frac{1}{\sqrt{2^{n_{\text{PA}}}}} \sum_{\tilde{x}, \tilde{z}, \tilde{z}} \sqrt{p_{\tilde{z}, \tilde{z}, \tilde{z}}} (-1)^{\tilde{x}\tilde{z}} |\tilde{z}\rangle^{\bar{A}} |\tilde{z}\rangle^{B'} |\tilde{x}\rangle^{B''} |\tilde{z}\rangle^{E'} |\tilde{x}\rangle^{E''} |\varphi_{\tilde{z}, \tilde{z}, \tilde{z}}\rangle^{BE} \quad (5.8)$$

$$= \frac{1}{\sqrt{2^{n_{\text{PA}}}}} \sum_{\tilde{x}, \tilde{z}, \tilde{z}} \sqrt{p_{\tilde{z}, \tilde{z}, \tilde{z}}} |\tilde{z}\rangle^{\bar{A}} |\tilde{z}\rangle^{B'} (X^{\tilde{z}})^{B''} |\tilde{x}\rangle^{B''} |\tilde{z}\rangle^{E'} |\tilde{x}\rangle^{E''} |\varphi_{\tilde{z}, \tilde{z}, \tilde{z}}\rangle^{BE}. \quad (5.9)$$

Because \tilde{z} only shows up as part of a unitary operator on B'' , tracing out all of Bob's systems means the state in E is averaged over these values, which was precisely the goal of privacy amplification. Moreover, \tilde{z} by itself is uncorrelated with \tilde{z} . Thus, in transferring the phase of \tilde{A} to systems B'' and E'' , we have $p_{\text{guess}}(\tilde{Z} | \mathcal{M}_{\tilde{Z}}^{BB'B''}), p_{\text{secure}}(\tilde{Z} | EE'E'') \approx 1$. Hence we can apply Theorem 2 to infer that Alice and Bob can recover a high-quality entangled state from their systems. By the known results on information reconciliation and privacy amplification, we can pick $n_Z \approx nH(Z^A|B)_\psi$ and $n_X \approx n - nH(Z^A|E)_\psi$ so that the rate achievable by this protocol is $H(Z^A|E)_\psi - H(Z^A|B)_\psi = -H(A|B)_\psi$, the hashing bound.

5.2.2 Privacy Amplification of Both Amplitude and Phase

The method of the previous section can serve as a stepping stone towards a protocol which is based entirely on decoupling both amplitude and phase from E . All we have to do is turn the amplitude information reconciliation into privacy amplification of phase. From the discussion prior to Theorem 3, we know that it will be insufficient to decouple E from X and Z , rather we must aim to simultaneously decouple E from Z on the one hand, and $C_Z E$ from X on the other. Note that in the latter case the state $|\psi_Z\rangle^{AC_ZBE}$ is only a device used in the proof; it does not need to show up in the protocol directly.

To achieve this simultaneous decoupling, we again begin by specifying two sets of stabilizers, n_Z Z -type stabilizers to decouple the amplitude and n_X X -type stabilizers to decouple the phase. As before, Alice's n qubits can be grouped into three sets of virtual qubits, the $n - n_Z - n_X$ encoded qubits in \bar{A} , n_X qubits in \hat{A} , and n_Z qubits in \tilde{A} . If n_X and n_Z are chosen appropriately, we can be sure that both $p_{\text{secure}}(\bar{X}^A | C_Z E)_{\psi_Z}$ and $p_{\text{secure}}(\tilde{Z}^A | E)_\psi$ are nearly one. Therefore system \bar{A} is implicitly in a maximally-entangled state with the joint system $\hat{A}\tilde{A}B$, and the remaining task is to classically transfer $\hat{A}\tilde{A}$ to Bob without violating the privacy conditions.

Following the method of the previous construction, suppose Alice makes *amplitude* measurements on \hat{A} and *phase* measurements on \tilde{A} , which she then broadcasts this information publicly. While E now receives extra information about the original amplitude and phase, no information about the encoded amplitude and phase has been leaked for the same reason as in the previous construction. The marginal states in E conditioned on the encoded amplitude (phase) value are still averaged over enough \mathbf{z} (\mathbf{x}) values to make them essentially identical.

Formally, the situation is very similar to the previous case as well. In fact, for the observable \tilde{Z} , the state of $|\psi\rangle^{ABE}$ after the measurements described above is precisely that of Equation (5.8), and so we can immediately conclude that $p_{\text{secure}}(\tilde{Z} | EE'E'') \approx 1$. The state relevant to privacy amplification of the phase can be expressed as, following Equation (3.37),

$$|\psi_Z\rangle^{AC_ZBE} = \frac{1}{\sqrt{2^n}} \sum_{\tilde{x}, \tilde{x}, \tilde{x}} |\tilde{x}\rangle^{\bar{A}} |\hat{x}\rangle^{\hat{A}} |\tilde{x}\rangle^{\tilde{A}} (Z^{\tilde{x}})^{C_Z} |\psi\rangle^{C_ZBE}, \quad (5.10)$$

and after the measurement it becomes

$$|\Psi_Z''\rangle^{\overline{ABB'B''EE'E''}} = \frac{1}{\sqrt{2^{n+n_X}}} \sum_{\hat{z}, \bar{x}, \tilde{x}} (-1)^{\hat{x}\hat{z}} |\bar{x}\rangle^{\overline{A}} |\hat{z}\rangle^{B'} |\hat{z}\rangle^{E'} |\bar{x}\rangle^{B''} |\tilde{x}\rangle^{E''} (Z^{\hat{x}})^{C_Z} |\psi\rangle^{C_Z BE} \quad (5.11)$$

$$= \frac{1}{\sqrt{2^{n+n_X}}} \sum_{\hat{z}, \bar{x}, \tilde{x}} |\bar{x}\rangle^{\overline{A}} |\hat{z}\rangle^{B'} |\hat{z}\rangle^{E'} |\tilde{x}\rangle^{B''} |\tilde{x}\rangle^{E''} (Z^{\bar{x}})^{\overline{C_Z}} (Z^{\tilde{x}})^{\tilde{C_Z}} |\psi\rangle^{C_Z BE}. \quad (5.12)$$

Now the phase $(-1)^{\hat{x}\hat{z}}$ cancels the similar phase inherent in the operator $(Z^{\hat{x}})^C$. Again this enforces an average over \hat{x} for the states in system E , ensuring that they are completely uncorrelated with \mathbf{x} and therefore \bar{x} . Just as before, \tilde{x} does not add any additional information about \bar{x} , so we can conclude that $p_{\text{secure}}(\overline{X}|EE'E'')_{\psi_Z} \approx 1$ and therefore Theorem 3 is applicable. For n_X and n_Z we can pick $n - nH(X^A|C_Z E)_{\psi_Z}$ and $n - nH(Z^A|E)_{\psi}$, respectively, yielding an overall rate of $H(X^A|C_Z E)_{\psi_Z} + H(Z^A|E)_{\psi} - 1$. This works out to be $H(A|E) = -H(A|B)$, which is the hashing bound once again.

5.3 Classical Channel Coding

In Section 4.1.4 we described how a protocol for entanglement distillation using one-way communication can be used to reliably send quantum information over a noisy channel, and that protocols achieving the optimal rate of entanglement distillation lead to optimal channel coding. A similar result holds for classical information, as demonstrated in [RR11], albeit using information reconciliation and randomness extraction or privacy amplification. This leads not only to a new proof of Shannon's original noisy channel coding theorem in the case the channel is classical, but also to one-shot results for both public and private communication of classical information over noisy quantum channels. Moreover, using the results of Section 5.1, we can exchange the use of information reconciliation with privacy amplification of a complementary observable, and thereby construct a channel coding scheme which is entirely based on decoupling-type arguments. That is, we can construct a means for noisy channel communication not by directly ensuring that the receiver can properly decode the transmissions, but rather by ensuring that complementary information does not leak to the environment.

On a heuristic level, the approach itself is quite similar to that of Section 4.1.4, not just the result. We can make the same sort of modification to an appropriate information reconciliation protocol as we did to entanglement distillation in order to create a coding scheme for the channel scenario. Suppose that Alice can send classical messages $z \in \{0, 1\}$ to Bob over a quantum channel such that he receives the corresponding state φ_z . If they are in possession of an information reconciliation protocol for the state $\psi^{AB} = \frac{1}{2} \sum_z |z\rangle \langle z|^A \otimes \varphi_z^B$, then they can use this to communicate reliably over the channel. In the information reconciliation scheme Alice would compute a hash function of n instances of the random variable Z^A , and with this information $f(\mathbf{z})$ Bob could determine the actual \mathbf{z} from his state φ_z^B .

In the channel scenario this can be used to specify a code by the set of all possible inputs \mathbf{z} (codewords) which hash to a specified value, say $\hat{\mathbf{z}}$. Ordering the elements of this set in some way, Alice can then map her actual message to the corresponding codeword. This defines an encoder. Presumably they have chosen an $\hat{\mathbf{z}}$ for which the information reconciliation decoder has a small probability of error, and thus Bob can use that decoder to determine \mathbf{z} and therefore Alice's intended message.

In fact, when the original inputs \mathbf{z} are uniformly distributed as above, one can easily show that not only will Bob have a small average probability of decoding error, but also a low error probability

for every message. To determine the number of messages Alice can send, it is simplest to consider the case of linear hash functions, where every output has the same number of preimages, namely the ratio of input to output size.¹ Thus, if information reconciliation requires an m -bit hash for an n -bit string \mathbf{z} , the resulting code can be used to transmit $n - m$ bits, remembering that an n -bit input corresponds to 2^n possible input strings. As we are working in the asymptotic i.i.d. scenario, we can apply the result mentioned in Section 4.1.1 that information reconciliation is possible at the rate $r = H(Z^A|B)_\psi$, so that $m \approx nr$. Therefore Alice can reliably send messages at rate $1 - H(Z^A|B)_\psi$.

There is still room for improvement, however, as the Holevo-Schumacher-Westmoreland (HSW) theorem (the quantum version of Shannon’s noisy channel coding theorem) assures us that rates up to at least the Holevo quantity $\chi = \max_{P_Z} I(Z^A:B)_\psi = \max_{P_Z} H(Z^A)_\psi - H(Z^A|B)_\psi$ are possible [85, 86].² Clearly something is missing in the above, unless it happens that the optimal distribution is uniform. We have restricted attention to the uniform distribution for convenience in the proof, in particular to easily determine the number of messages which can be send with small worst-case probability of error. If we only cared about average probability of error, any distribution could be used for the purposes of converting an information reconciliation protocol to a channel code. The difficulty is then to exploit this freedom without requiring a substantially new proof.

Fortunately, there is a simple way to deal with this problem by making use of the randomness extractors described in Section 5.1.1, though here the privacy properties will only be relevant to the case of private channel communication. Alice can use the extractor in reverse as a *distribution shaper* to simulate a random variable Z with arbitrary distribution P_Z using a uniformly-distributed random variable U . To do so, Alice chooses an extractor output u at random and then maps it to a possible preimage \mathbf{z} using the conditional distribution $P_{Z|U=u}$. This requires an additional source of randomness, as the extractor function is not one-to-one.

When Z is destined to be the input to the communication channel, we can instead think of U as the input to the “superchannel” composed of the shaper and the original channel. This is depicted in Figure 5.2, taken from [RR11]. Note that for this step we must rely on the recently-established one-shot results on information reconciliation, as mentioned in Section 5.1.4, because the joint state shared by Alice and Bob which describes the input and output is generally not i.i.d. However, in the one-shot framework all the previous results linking information reconciliation to channel coding can be applied to the superchannel. Alice encodes messages into the outputs \mathbf{u} of the extractor and then sends these first through the shaper and then through the communication channel to Bob. Information reconciliation of U with B enables Bob to recover the original message.

Using the smooth entropy results on structureless resources we can determine the (logarithm of the) raw number of messages Alice can reliably send to Bob, instead of the rate as appropriate to the i.i.d. setting. The details of the derivation are given in [RR11], and the result is that Alice can reliably transmit N bits to Bob, for

$$N_{\text{class}} \approx \max_{P_Z} \left[H_{\min}^\epsilon(Z)_\psi - H_{\max}^\epsilon(Z|B)_\psi - O(\log \frac{1}{\epsilon}) \right]. \quad (5.13)$$

Here ϵ characterizes the worst-case error probability of the coding scheme, and this expression agrees with a result found for classical channels found by Renner *et al.* [119].³ This result applies to completely arbitrary channels, but when Alice and Bob would like to communicate using n

¹The general case can be handled by probabilistic arguments [RR11].

²As with quantum communication and private classical communication over quantum channels, regularization can increase the rate further. Indeed, as discussed at the end of Chapter 6, regularization is necessary to reach the capacity.

³Wang and Renner have recently derived a one-shot result for classical communication over quantum channels via a different method [120].

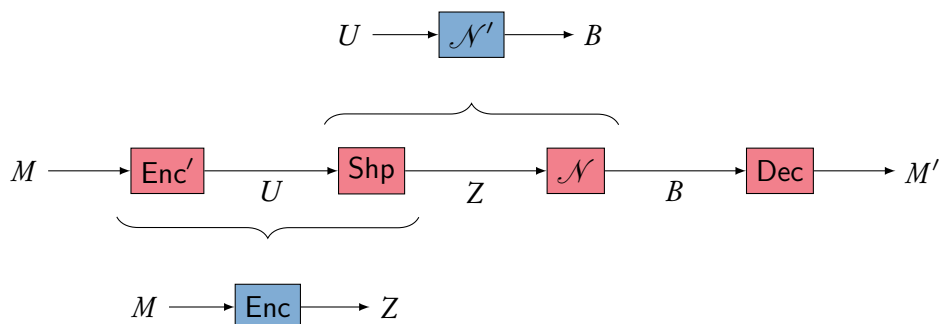


Figure 5.2: Schematic of using randomness extraction and information reconciliation to perform noisy channel communication. Messages $m \in M$ are input to the encoder Enc' and subsequently to the shaper Shp , which is a randomness extractor run in reverse. Then they are then transmitted over the channel \mathcal{N} to the receiver, who uses the decoder Dec to construct a guess $m' \in M'$ of the original input. Concatenating the shaper and channel gives a new effective channel \mathcal{N}' , for which an encoder/decoder pair $(\text{Enc}', \text{Dec})$ can be constructed by repurposing an information reconciliation scheme that operates on the joint input-output UB of the channel. Ultimately, the shaper can instead be regarded as part of the encoder Enc , which is formed by concatenating Enc' and Shp .

uses of a memoryless channel we can appeal to the asymptotic equipartition property (AEP) of the smooth min- and max-entropies, proven by Tomamichel *et al.* [121]. Roughly speaking, it states that $H_{\min}^{\epsilon}(Z|B)_{\psi^{\otimes n}} \approx nH(Z|B)_{\psi}$ and similarly for the max-entropy. We then recover the rate given by the HSW theorem; for channels with purely classical outputs, i.e. quantum states which all pairwise commute, we recover Shannon’s noisy channel coding theorem [6].

Besides an appealing modular proof of the noisy channel coding problem based on the simpler primitives of randomness extraction and information reconciliation, another appeal of this approach is that by using privacy amplification instead of just randomness extraction for the distribution shaper, we automatically obtain a construction suitable for private communication of classical information over a noisy quantum channel. In that case we find that the (logarithm of the) number of private messages which can be reliably sent is given by

$$N_{\text{priv}} \approx \max_{p_z} \left[H_{\min}^{\epsilon}(Z|E)_{\psi} - H_{\max}^{\epsilon}(Z|B)_{\psi} - O(\log \frac{1}{\epsilon}) \right], \quad (5.14)$$

where system E is the “other half” of the channel output. That is, upon input of z the channel produces the pure state $|\varphi_z\rangle^{BE}$ shared between Bob and the environment or eavesdropper. As before, an application of the AEP recovers the rate relevant in the asymptotic i.i.d. setting, namely $\max_{p_z} [H(Z|E)_{\psi} - H(Z|B)_{\psi}]$. This agrees with the findings of Devetak [98] for quantum channels, and those of Wyner [122], Ahlswede and Csiszar [123], and Maurer and Wolf [124] for classical channels.

Finally, we note that combining this proof technique with the duality between information reconciliation and privacy amplification it is possible to prove that reliable communication is possible by ensuring that not too much information leaks to the environment. This decoupling approach was heretofore unknown to work for channel coding of classical information, and in fact this was the one major protocol not known to be amenable to a decoupling analysis. The encoding and decoding procedure is precisely the same as before, using a distribution shaper and information reconciliation to

create an encoder and decoder. But instead of relying on constructions of information reconciliation protocols, we use privacy amplification and duality. Thus, the size of the code is fixed by how much privacy amplification is needed for the observable conjugate to the uniform input U , and is therefore given by a smooth min-entropy. Using the uncertainty principle for smooth entropies formulated in [118] we can relate this to the smooth max entropy of U conditioned on B , and obtain again Equations (5.13) and (5.14).

Security of Quantum Key Distribution

Quantum key distribution is one of the major current applications of quantum information processing, requiring only minimal ability to coherently manipulate quantum information. Devices implementing QKD protocols such as BB84 are even currently available commercially. But where does the security of QKD come from? That is to say, how can we prove that a given protocol is truly secure and no would-be eavesdropper has any information about the key?

There have been three main approaches to answering this question, each with its own advantages and disadvantages, which we briefly describe in the first of three sections in this chapter. In the second section we follow one of these methods, treating QKD as a means for virtual creation of entanglement as described in Section 4.3, and recount the results of [RG06] showing that it applies to a wide class of protocols, not just the original BB84 scheme.

From Section 3.4 we know that entanglement is not strictly necessary for generating secret keys, and that in general private states suffice. In the third section of this chapter we describe how alterations to the BB84 protocol which improve the maximum tolerable error rates can be understood as part of a virtual private state distillation scheme, and that combining this additional step with similar enhancements to quantum error-correction lead to still better tolerable error rates. This work was first reported in [RS07, SRS08, KR08].

6.1 Notions of Security

The first proofs of *unconditional* security of the BB84 protocol—that is, security of the protocol under arbitrary attacks on the public quantum channel by the eavesdropper Eve—were given by Biham *et al.* [125, 126] and Mayers [127, 128]. Their methods are similar, and essentially rest on an implicit use of the uncertainty principle to bound Eve’s information about Alice’s key by Bob’s information about the conjugate basis to the key.¹ Biham *et al.* characterized the security as due to an information-disturbance tradeoff, the fact that eavesdropper cannot acquire information about Alice’s signals without disturbing them. Such a tradeoff follows immediately from Equation (1.7), as to be able to gain information about e.g. the phase without disturbing the amplitude information would imply a violation of the entropic bound.

At the same time, efforts to base the security of QKD on virtual entanglement distribution as described in Section 2.4 were underway, culminating in Shor and Preskill’s proof for BB84 shortly after the two mentioned above. Their proof was a good deal simpler than the earlier versions, and achieved a higher error threshold, the maximum error rate at which the protocol can still safely generate secret keys (albeit at vanishingly small rates). The new proof established a threshold of 11%, the previous proofs 7.56%. The simplicity also enabled the method to be extended to other protocols. Lo [130] established the unconditional security of the six-state protocol proposed by Brass [131] which uses the eigenstates of the XZ operator as signals in addition to those of X and Z . Tamaki, Koashi, and Imoto [132] extended the method to a proof of Bennett’s two-state protocol (B92) [133], while Gottesman and Lo showed that it could also treat information reconciliation steps involving two-way communication [134], greatly increasing the error rate tolerable by BB84 to 18.9%. Boileau *et al.* (including the present author) [135] proved the security of a B92-like protocol involving three states which was originally proposed by Phoenix *et al.* [136].

The original approach of Biham *et al.* and Mayers has its own advantages within the realm of the BB84 protocol, however, as it is not actually concerned with the details of Bob’s measurement appara-

¹Both of their formal statements make use of a related result by Yao [129].

tus, only Alice's preparation device. This can be anticipated from the implicit use of the uncertainty principle: From Equation (1.7) it is clear that to bound Eve's knowledge of the key it suffices to have a bound on Bob's knowledge of the conjugate observable to the key. It is not necessary to have an accurate physical description of how he comes by such knowledge, which greatly extends the practicality of the proof. Koashi and Preskill combined techniques from both methods to treat the problem of an uncharacterized source [137] (but characterized detector), and later Koashi gave an even simpler proof which was the first to quantitatively appeal to the uncertainty principle [138, 139]. Although the proof itself is constructed via other means, Koashi used the Maassen and Uffink relation, Equation (1.4), as a guide to determine the size of the secret key. Very recently, Tomamichel *et al.* [140] have directly used the smooth entropy uncertainty relation of [118] to give a simple security proof of BB84 with uncharacterized detectors.

Meanwhile, a third general approach focused on showing that privacy amplification produces secure keys even when the adversary holds quantum instead of classical information. To make use of privacy amplification one then needs to characterize the quantum states held by the eavesdropper, or at least give a bound on the size of their overall support. Ben-Or showed that a result from quantum communication complexity implies the efficacy of privacy amplification and that the knowledge gained by Alice and Bob in the BB84 protocol can be used to bound the effective size of Eve's system [141]. König *et al.* demonstrated that privacy amplification works against quantum adversaries generally [142], and Christandl *et al.* developed this into a generic security proof which replicated the one-way results above, even improving the threshold for the B92 protocol [143]. Kraus, Gisin, and Renner [144, 145] extended this to establish that many protocols are not only unconditional secure, but also safely *composable* with other cryptographic primitives to create larger cryptographic schemes which are themselves secure, following composability results by Renner and König [108] and Ben-Or *et al.* [146]. Renner provided another method also suitable for two-way protocols in his thesis [113].

It should be noted that the task of key *distribution* is considerably more involved than the task of key *distillation* as discussed in Section 4.3, and the security issue all the more complex. There the input state shared by Alice and Bob is known in advance, and moreover it is assumed to consist of n copies of some state ψ . Neither of these statements hold in general in the present context, for although Alice sends n quantum systems to Bob, these travel over an insecure communication channel which could in principle be under the control of the would-be eavesdropper Eve. The difficulty lies in the fact that the eavesdropper could in principle attack all the signals jointly, what is termed a *coherent attack*. If Eve attacked each signal separately, a *collective attack*, then Alice's and Bob's state would have the aforementioned i.i.d. form, and could be handled by those methods.

Unsurprisingly, then, one widely-used method of handling coherent attacks is to reduce them in some way to collective attacks. Originally this was done on a more *ad hoc* basis for particular protocols, but has been made more systematic by Renner [113, 147], culminating in a very general statement by Christandl *et al.* [148]. This states that as long as the key distribution protocol is unconcerned with the order in which Alice transmits the signals, which can be enforced by arbitrarily permuting them, then security against collective attacks implies security against coherent attacks.

6.2 Entanglement in Prepare and Measure QKD

Quantum key distribution can be formulated as a virtual entanglement distribution scheme for a wide class of protocols and the Shor-Prekill approach used to prove their security. In this section we briefly sketch out how this can be done, following [RG06] and simplifying some issues in light of

intervening research advances. The main conceptual difficulty in considering protocols other than BB84 in the Shor-Preskill framework is that it appears as if the CSS structure of information reconciliation and privacy amplification are directly related to the use of amplitude and phase eigenstates as the signals and for measurement. However, this is not actually the case, and in fact these two parts of the protocol have nothing to do with each other. This was already noted in the proofs by Tamaki *et al.* [132] and Boileau *et al.* [135], but [RG06] show how it can be made to work more generally.

First let us settle on the general framework of prepare and measure protocol. A generic protocol consists of five main stages. First Alice prepares quantum states and transmits them over the insecure quantum channel to Bob, who measures them; this is the only step in which quantum operations are actually needed. Second, they transform their classical transmission and measurement records to a prospective *raw key*. This step is usually called sifting, after the specific mapping used in BB84, and usually the transformation is chosen so that the raw key would be a truly secret key if the quantum channel were noiseless.

As real channels are inevitably noisy, Alice and Bob need to distill a truly shared, secret key from the raw key. In stage three, parameter estimation, they compare some random subset of the raw key to determine the likely number of errors. This serves two purposes. In the fourth stage, information reconciliation, they use the knowledge from parameter estimation to agree on an identical refined key. Usually this involves Bob reconciling his raw key to Alice's, hence the name. Finally, they also use this knowledge to perform privacy amplification and thereby generate the final secure key.

The trick to applying the Shor-Preskill framework more generally is to first formulate the prepare and measure process coherently, i.e. in quantum-mechanical language, and then regard Alice's and Bob's systems in this setting as being composed of two virtual subsystems. One subsystem (quantumly) records the key value, while the other (quantumly) records the sifting information. The sifting stage can then be seen as a measurement of the latter subsystems, plus postselection by public communication to select appropriately matching sifting outcomes. The virtual key subsystems remain, and it is their entanglement which is at issue in the Shor-Preskill framework. The amount of entanglement, and thus secret key, which can be distilled may be estimated by making use of the symmetries of the signal states and measurement.

We can illustrate this most easily using the BB84 protocol itself and then describe how it can be made to work more generally. As discussed in Section 2.4, the BB84 protocol can be described coherently by pretending that Alice first creates EPR pairs and then sends one subsystem of each pair to Bob. Here, however, it is more appropriate to describe each signal sent by Alice as her preparation of the state

$$|\psi_0\rangle = \frac{1}{2} \sum_{j,k} |j\rangle^{A_K} |k\rangle^{A_S} |\xi_{jk}\rangle^B, \quad (6.1)$$

and transmission of the B subsystem to Bob. The indices j and k specify the eigenvalue and observable, respectively, of the state $|\xi_{jk}\rangle$ transmitted by Alice; $k = 0$ denotes amplitude Z and $k = 1$ phase, while the eigenvalue is given by $(-1)^j$. Bob makes a random measurement of the two observables, which can be described by the isometry $U_{\mathcal{M}}^{B \rightarrow B_K B_S} = \frac{1}{\sqrt{2}} \sum_{j,k} |j\rangle^{B_K} |k\rangle^{B_S} \langle \eta_{jk}|^B$, where here $|\eta_{jk}\rangle = |\xi_{jk}\rangle$ but the distinction will be useful later. For a noiseless channel, his measurement process results in the state

$$|\psi_1\rangle^{A_K A_S B_K B_S} = \frac{1}{\sqrt{8}} \sum_{j j' k k'} |j\rangle^{A_K} |k\rangle^{A_S} |j'\rangle^{B_K} |k'\rangle^{B_S} \langle \eta_{j'k'} | \xi_{jk} \rangle. \quad (6.2)$$

From the form of the inner products $\langle \eta_{j'k'} | \xi_{jk} \rangle$ one can easily work out that if Alice and Bob each measure their S -labeled subsystems and obtain the same result, the remaining K -labeled systems are in the state $|\Phi\rangle^{A_K B_K}$ and thus measurement produces a secret key. This mimics the sifting process of the actual protocol, as Alice and Bob perform the measurements separately and compare their results by public discussion. Also crucial is the fact that the overall probability distribution for signals and measurement outcomes found here is precisely the same as in the prepare and measure scheme. Thus, this state has the form claimed above: It provides a coherent description of the real protocol in which Alice and Bob each have key K and sifting S subsystems, and sifting is accomplished by local measurement of the latter subsystems and postselection.

Noisy channels require the additional steps of parameter estimation, information reconciliation, and privacy amplification, but change the above picture only slightly. Describing the channel resulting from Eve's attack by its decomposition into Kraus operators, and assuming the attack is collective, the state $|\psi_1\rangle$ is altered by the noise to

$$|\psi_1'\rangle^{A_K A_S B_K B_S E} = \frac{1}{\sqrt{8}} \sum_{jj'kk'\ell} |j\rangle^{A_K} |k\rangle^{A_S} |j'\rangle^{B_K} |k'\rangle^{B_S} |\ell\rangle^E \langle \eta_{j'k'} | E_\ell | \xi_{jk} \rangle. \quad (6.3)$$

In the sifting stage, Alice and Bob keep only the cases in which $k = k'$ and subsequently discard the information specifying which value of k they observed. We can model this process as keeping only the $k = k'$ terms in (6.3) and then giving the A_S and B_S systems to Eve. Alice and Bob keep only the raw key, and the state becomes (slightly redefining E)

$$|\psi_2'\rangle^{A_K B_K E} \propto \sum_{jj'k\ell} c_{jj'}^{k\ell} |j\rangle^{A_K} |j'\rangle^{B_K} |k, \ell\rangle^E, \quad c_{jj'}^{k\ell} = \langle \eta_{j'k} | E_\ell | \xi_{jk} \rangle. \quad (6.4)$$

Following the Shor-Preiskill idea, as generalized in Section 4.3, Alice and Bob can construct the information reconciliation and privacy amplification protocols necessary to turn the raw key into a secret key once they are able to estimate $p_{\text{guess}}(Z^{A_K} | B_K)_{\psi_2'}$ and $p_{\text{guess}}(X^{A_K} | C_Z B_K)_{\psi_2'}$. A bound on the former is given directly by parameter estimation, but the latter is not so straightforward. The joint state of the key systems is determined via the coefficients $c_{jj'}^{k\ell}$, creating a connection between the two guessing probabilities, albeit in general a not at all straightforward one. The structure of the sifting and of the signals and measurements greatly simplifies the connection, and makes it possible to find useful bounds on the latter guessing probability as a function of the former. This enables Alice and Bob to construct the remainder of the protocol to be provably secure.

For BB84, one finds by direct calculation that $p_{\text{guess}}(Z^{A_K} | Z^{B_K}) = p_{\text{guess}}(X^{A_K} | X^{B_K})$ regardless of the value of ℓ . That is, the correlation in the amplitude basis (which gives the key itself) is precisely the same as the correlation in the phase basis (conjugate to the key). This was to be expected from the original coherent description of BB84 which explicitly uses EPR pairs from the beginning, since half the time the key comes from the original amplitude basis, and half the time from the phase basis, so the correlations ought to be the same. Using this relationship in the formula for the rate of secret key distillation, Equation (4.22) (ignoring \mathcal{Q} and T), we recover the rate $r_{\text{BB84}} = 1 - 2h_2(\delta)$, for δ the observed error rate in the raw key and $h_2(\delta) = -\delta \log_2 \delta - (1 - \delta) \log_2 (1 - \delta)$ the binary entropy, which leads to the threshold of 11%. Security against general coherent attacks is then ensured by the result of Christandl *et al.* [148].

A great advantage of the above approach is the modularity of the security proof. The details of the signals, measurements, and sifting are logically completely separate from the details of information reconciliation and privacy amplification. The former enter only into the coefficients $c_{jj'}^{k\ell}$, which are

used to select a CSS code for the latter. This approach is developed in [RG06] as a generalization of that used by Tamaki *et al.* [132] and Boileau *et al.* [135], and it is shown that it applies to a wide class of protocols, particularly those based on so-called *equiangular spherical codes*. These are constellations of pure states $|\xi_j\rangle$ whose pairwise overlaps are all identical, as in the three-state protocol of Phoenix [136] mentioned above, and were adapted for use in QKD generally by the author [149, 150, 151]. The other main contribution of [RG06] is the development of a method of exploiting the symmetries of the sifting scheme and the signal and measurement states to simplify this task, relying on results from group representation theory.

To see how this works, consider the protocol in which Alice's signals are four qubit states for which $|\langle \xi_j | \xi_k \rangle|^2 = \frac{1}{3}$, as described in [150]. These form a regular tetrahedron in the Bloch-sphere representation of a qubit, and Bob's measurement is comprised of appropriately-normalized projectors onto the states $|\eta_k\rangle$ for which $\langle \eta_k | \xi_k \rangle = 0$, i.e. the inverse tetrahedron in the Bloch-sphere. Due to symmetry, Bob's measurement would randomly reveal one state which Alice did not send if the channel were noiseless, and so the information exchanged by Alice in the sifting stage consists of a random choice of two states she did not send.

In one-third of cases these two pieces of information specify which state she did send, and Bob publicly announces that he has successfully decoded the transmission. From this they generate one secret bit corresponding to which of the two signals Alice did send, given the public exclusion of two of the initial possibilities. There are 12 possible announcements by Alice, since she must also specify how the two possible signal states are to be decoded into the raw key, and we may label the signal states by the combination of sifting announcement and raw key value. In this way each signal is counted six times, but this presents no difficulty as each is counted the same number of times. Much the same holds for Bob, and so the state in Equation (6.1) can be used to describe the protocol coherently.

The remaining task is to use the $c_{jj'}^{k\ell}$ to bound $p_{\text{guess}}(X^{A_K} | C_Z B_K)$ in terms of $p_{\text{guess}}(Z^{A_K} | Z^{B_K})$. By exploiting symmetries of the QKD protocol as in [RG06], we can greatly simplify this task. Suppose that the sifting step of the protocol is such that there exist unitaries U_k and V_k for which $|\xi_{jk}\rangle = U_k |\xi_{j0}\rangle$ and $|\eta_{jk}\rangle = V_k |\eta_{j0}\rangle$. Then the $c_{jj'}^{k\ell}$ become $c_{jj'}^{k\ell} = \langle \eta_{j'0} | V_k^\dagger E_\ell U_k |\xi_{j0}\rangle$. Now let us focus on a particular Kraus operator E_ℓ by fixing the value of ℓ , but average over the value of k , which corresponds to Alice and Bob throwing away the information specifying which particular sifting map they applied. Their shared state given the value of ℓ has the form

$$\psi_\ell^{A_K B_K} \propto \sum_{ii'jj'} |ii'\rangle \langle jj'|^{A_K B_K} \sum_k \langle \eta_{i'0} | V_k^\dagger E_\ell U_k |\xi_{i0}\rangle \langle \xi_{j'0} | U_k^\dagger E_\ell^\dagger V_k |\eta_{j0}\rangle. \quad (6.5)$$

Examining the form of the matrix elements, we see that the sifting symmetries U_k and V_k create an effective channel having Kraus operators $V_k^\dagger E_\ell U_k$. Moreover, the group nature of these operators enables us to compute the action of the channel by appealing to representation theory. In the particular case of the tetrahedral protocol, one finds that the effective channel is just a depolarizing channel, irrespective of the value ℓ . The depolarizing rate can be determined by the noise rate observed in the parameter estimation phase. Computing the state after the sifting step reveals that Alice and Bob can describe their shared key state by a Bell-diagonal state $\psi^{A_K B_K} = \sum_{jk} p_{jk} |\beta_{jk}\rangle \langle \beta_{jk}|^{A_K B_K}$, as in Equation (2.11), with the p_{jk} satisfying $p_{01} = p_{11} = 2p_{10}$.

This implies $p_{\text{guess}}(Z^{A_K} | Z^{B_K}) = \delta$ and $p_{\text{guess}}(X^{A_K} | X^{B_K}, Z^{A_K} = Z^{B_K}) = \frac{1}{3}$ while $p_{\text{guess}}(X^{A_K} | X^{B_K}, Z^{A_K} \neq Z^{B_K}) = 1 - 2\delta/3(1 - \delta)$. The latter guessing probabilities are directly related to $p_{\text{guess}}(X^{A_K} | C_Z B)$ since Bob's knowledge of Z^{A_K} stored in C_Z can be equivalently thought of as the information as to whether

or not an amplitude error $Z^{A_K} \neq Z^{B_K}$ occurred or not. Using these guessing probabilities in Equation (4.22), we obtain the rate $r_{\text{tetra}} = 1 - h_2(\delta) - \delta h_2(\frac{1}{3}) - (1 - \delta)h_2(2\delta/3(1 - \delta))$, which has a threshold of 11.56%. In [RG06] the method is applied to several other spherical code protocols with signal and measurement states having Hilbert space dimension three.

6.3 Private States in Quantum Key Distribution

By clever if perhaps unintuitive choice of *preprocessing* operations \mathcal{Q} in Equations (4.22) and (4.24) the error thresholds of QKD can be pushed higher than those found by the Shor-Preskill method alone. Understanding how this can be the case requires interpreting QKD as a virtual means of private state distillation rather than just entanglement distillation, as first shown in [RS07]. Furthermore, the private state distillation approach suggests that it would be beneficial to combine two types of preprocessing operations previously studied, and this was indeed shown to be the case for the BB84 protocol in [SRS08]. Further improvements and an extension of the method to the six-state protocol were reported in [KR08], and we describe both of these results here.

That private state distillation is actually needed to give a fully quantum-mechanical description of QKD was necessitated by the work of Kraus, Gisin, and Renner [144, 145]. They established the seemingly-paradoxical result that the noise threshold of BB84 can be improved if Alice randomly flips some of her raw key bits before performing the final three steps of the protocol, and they reported a threshold improvement from 11% to 12.4%. From the viewpoint of QKD as a virtual scheme for entanglement distillation this additional step would seem to be counterproductive, as noise inflicted by Alice behaves the same as noise inflicted by Eve. However, we saw in Chapter 3 that entanglement is not actually necessary for secret key creation, private states are. This raises the question of whether or not one can view the noisy preprocessing step as part of a virtual scheme for private state distillation, which [RS07] answers in the affirmative.

The crux of understanding such noisy preprocessing in a private state picture is to include the system Alice uses to impart the noise to her raw key and observe that it functions as a shield system. The overhead in the protocol of additional information reconciliation needed due to the noisy preprocessing is then more than made up for by a reduction in the required amount of privacy amplification. The particular guessing probabilities found in the previous section imply that the state of Alice's and Bob's raw keys immediately after the sifting stage takes a Bell-diagonal form in which the probabilities of amplitude and phase error are independent and equal. That is, in the state $\psi^{A_K B_K} = \sum_{jk} p_{jk} |\beta_{jk}\rangle \langle \beta_{jk}|^{A_K B_K}$, one has $p_{00} = (1 - \delta)^2$, $p_{10} = p_{01} = \delta(1 - \delta)$, and $p_{11} = \delta^2$ for δ the probability of amplitude (or phase) error.

Now suppose that Alice randomly flips each raw key bit independently with some probability q . This process may be modelled as a CNOT gate whose control is an ancillary system A' prepared by Alice in the state $|\varphi\rangle = \sqrt{1 - q}|0\rangle + \sqrt{q}|1\rangle$ and whose target is her raw key A_K . The error rate in Alice's and Bob's keys has jumped to $\delta' = \delta(1 - q) + q(1 - \delta)$, but the crucial difference from the entanglement distillation scenario is that for security it is relevant how well A' and B_K together can predict X^{A_K} , not merely how well B_K could alone.

The resulting state of $A_K A' B_K$ can be used to compute $H(X^{A_K} | A' B)$ for use in Equation (4.22); observe that we do not need to make use of the C_Z system here because knowing if there is an amplitude error tells Bob nothing about the likelihood of a phase error. Using Equation (4.22) and optimizing over the choice of q we recover the threshold of 12.4%. A similar calculation (now requiring the use of C_Z) recovers the six-state threshold of 14.1%. Actually, [RS07] follows a different approach than what we have outlined here, directly constructing the twisting operator, but this can

be seen as a particular case of the general results on private states and secret key distillation presented in Sections 3.4 and 4.3.

In his security proof of the six-state protocol, Lo observed [130] that the threshold can be improved from the nominal 12.6% one would find following the Shor-Prekill method to 12.7% by employing so-called *degenerate* quantum error-correcting codes first discussed by DiVincenzo, Shor, and Smolin [152]. This code consists of a concatenation of an amplitude repetition code with a random CSS code and in the present context corresponds to a preprocessing operation \mathcal{Q} on blocks of inputs, as in Equation (4.24).

The degeneracy of the code refers to the fact that several different errors can share the same recovery operation and the syndrome need only reveal which recovery operation is required, a phenomenon which is not possible for classical error-correcting codes. For example, in the amplitude repetition code of Section 2.2.1, the three possible phase errors acting on single qubits all have the same effect on the encoded quantum information, namely as a phase flip. Thus, if we concatenate the repetition code with another code, we need not determine the precise location of a phase error on the physical qubits. Reducing the number of stabilizers needed to enable correction of phase errors implies a reduction in the necessary amount of privacy amplification in the context of QKD, and thus the threshold increases.

Shor's nine-qubit code described in Section 2.2.2 provides a simple example. There we considered the effect of a single phase flip error on the fourth qubit and found that it would be detected by measuring certain stabilizer operators. But it is clear from the argument there that the same result is obtained for a phase error on either qubit five or six. This is reflected in the fact that associated with the code are six amplitude stabilizers and only two phase stabilizers. The former determine the precise location of an amplitude error, but the latter only fix the location of the phase error up to the position in the block. This is all that is necessary.

It is possible to combine the noisy preprocessing discussed above with degenerate codes to improve the threshold of BB84 still further, as described in [SRS08]. The original protocol is modified as follows. After the raw key is created in the sifting phase, Alice performs a noisy preprocessing step in which she independently flips each raw key bit with some probability q . Then she computes the syndromes of an amplitude repetition code encoding one qubit into m qubits, i.e. $z_1 \oplus z_2, z_1 \oplus z_3, \dots, z_1 \oplus z_m$ and transmits these publicly to Bob. The first bit of each block she saves for further use as the key. Bob then computes the syndromes of his block, and attempts to correct his key bit so that the syndromes match Alice's, exactly as was done in the entanglement distillation protocol discussed in Section 2.3. Alice and Bob then repeat this process for many blocks, collecting one key bit per block. On these refined key bits they then perform information reconciliation and privacy amplification as needed.

To determine the threshold, for which the main difficulty is, as usual, to determine the amount of privacy amplification needed, it is simpler to focus on Eve's states and compute $H(\bar{A}_K|ES)$, where S denotes the syndrome information and \bar{A}_K the key bit encoded in the repetition code. Again the symmetries of the problem enable the use of group representation theory to make the calculation numerically tractable, allowing thresholds for blocklengths in the hundreds to be determined. The best threshold found in [SRS08] is 12.9%, corresponding to $q \approx 0.32$ and $m = 400$. A more elaborate analysis is required for the six-state protocol, and this is carried out in [KR08], with the result that the threshold is at least 14.59%. Additionally, the effects of iterating the entire noisy preprocessing plus repetition code procedure are investigated therein, and this is found to offer substantial increases in the key distribution rate of the protocol at high error rates, though the overall threshold is not as large.

As mentioned previously, the use of repetition codes is a type of blockwise preprocessing, in contrast to the noisy preprocessing which is applied to single key bits. As blockwise preprocessing is more complicated, and the expression for the optimal rate for secret key distillation, Equation (4.24), essentially impossible to evaluate, the question arises whether blockwise preprocessing, i.e. regularization are truly necessary. Unfortunately, the answer is yes, as observed in [SRS08]. One can show that the threshold found by noisy preprocessing, 12.4%, is the optimal threshold using single-bit, or single-letter preprocessing. Since the combination of noisy processing and repetition codes leads to a higher threshold, regularization must in general be necessary. This result then applies to the private capacity of a channel as well, since one way to communicate privately is to first generate secret keys and then encrypt the actual message to be sent.

Thus, neither the secret key distillation rate nor the private capacity are *single-letterizeable* quantities. This reveals a large distinction between classical and quantum information theory, as single-letter quantities are usual in the former, reflecting the fact that the random coding arguments of Shannon are optimal in a wide variety of situations. In quantum information theory this is no longer true. The degenerate codes found by DiVincenzo, Shor, and Smolin [152] show that the quantum capacity is also not single-letterizeable, while Hastings has recently established that the classical capacity of a quantum channel is not single-letterizeable either [153]. Despite the apparent similarities with classical information theory, a full understanding of quantum information theory will require the development of tools beyond the usual random coding methods.

Summary and Outlook

The preceding six chapters demonstrate that, far from being just an abstract mathematical study, the study of quantum information theory is quite closely connected to core physical concepts, namely complementarity and the uncertainty principle. Indeed, although we have presented the topics of this thesis in a logical order, it was actually research into secret key distillation in [RB08] that led to the conjecture of the entropic uncertainty principle of Equation (1.7) in [RB09] and its eventual proof in [BCC⁺10].

The results described in this thesis spring from trying to make sense of what it means to have “quantum” information, working within the formalism of quantum theory itself. Using the conditional entropy $H(Z^A|B)_\psi$ we can describe the information held by B about the amplitude measurement Z on system A , when A and B are jointly in the quantum state ψ^{AB} . Having quantum information then refers to the situation in which B implicitly contains information about two complementary observables X and Z , and the uncertainty principle in the form $H(X^A|B)+H(Z^A|C) \geq \frac{1}{c}$ constrains the extent to which information about both can be simultaneously explicitly realized. Quantum information processing protocols can then be constructed by mimicking related classical information processing protocols for the two complementary pieces of classical information, taking care not to violate the uncertainty principle.

Although complementarity is at the heart of the results presented herein, to complete the proofs we have also relied heavily on certain algebraic properties both of the observables X and Z as defined in Equation (3.25) and of the attendant CSS stabilizer codes. In particular, the algebraic properties of the amplitude and phase observables play important roles in Theorems 1, 3, 4, and 6, while the algebraic structure of CSS codes is used extensively throughout Chapters 4, 5, and 6. Removing the algebraic requirement on the observables is precisely the difference between the uncertainty principle results of [RB09] and [BCC⁺10], and a major goal of future work is to remove this requirement from the aforementioned results as well. The situation is akin to difference between the heuristic use of the uncertainty principle in the early proofs of QKD, where the uncertainty principle provided guidance for the actual algebraic arguments, and the recently formulated BB84 security proof of Tomamichel *et al.* [118, 140] based directly on the uncertainty principle formulated in terms of smooth-entropy.

This goal is likely to be fairly straightforward for the results of Chapter 3, but the use of CSS codes in the protocols of the subsequent chapters appears much more central to those results. The difficulty lies in the need to combine classical protocols for complementary observables in such a way that all the important quantities can actually simultaneously exist, i.e. the corresponding operators all commute. In the entanglement distillation scheme of Chapter 4 for instance, the use of CSS codes ensures that the syndrome information needed to establish strong phase correlations does not interfere with either the amplitude syndromes nor the final encoded amplitude.

Another goal of future work will be to extend all the results beyond the realm of asymptotic i.i.d. resources and into the one-shot domain of structureless resources briefly described in Section 5.1.4. Here we have presented optimal protocols in the former scenario, but it is not clear whether this will be possible in the more general setting. One cause for hope is that the uncertainty principle already plays a fundamental role in the one-shot setting. Tomamichel *et al.* [154] have shown that the smooth min- and max-entropies are not independent: One may be defined in terms of the other using a purification system. The smooth entropy uncertainty relation then follows from this duality [118].

Finally, a much more ambitious goal is to extend the notion of quantum information as complementary classical information past the simple two-party communication scenarios studied here. Can this point of view shed some light into how quantum computers work?

Bibliography

- [1] N. Bohr, “Das Quantenpostulat und die neuere Entwicklung der Atomistik,” *Naturwissenschaften* **16**(15), 245–257 (1928).
- [2] N. Bohr, “The Quantum Postulate and the Recent Development of Atomic Theory,” *Nature* **121**(3050), 580–590 (1928).
- [3] Niels Bohr, “Discussion with Einstein on Epistemological Problems in Atomic Physics,” in *Albert Einstein Philosopher-Scientist*, volume 7 of *The Library of Living Philosophers*, 201–241 (MJP Books, New York, 1949).
- [4] W. Heisenberg, “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik,” *Zeitschrift für Physik* **43**(3), 172–198 (1927).
- [5] Richard P. Feynman, *Feynman Lectures On Physics* (Addison Wesley Longman, Boston, 1970).
- [6] C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal* **27**(3), 379–423 (1948).
- [7] Vannevar Bush, “Instrumental Analysis,” *Bulletin of the American Mathematical Society* **42**(10), 649–669 (1936).
- [8] R Landauer, “Irreversibility and Heat Generation in the Computing Process,” *IBM Journal of Research and Development* **5**(3), 183 (1961).
- [9] Rolf Landauer, “Information is Physical,” *Physics Today* **44**(5), 23 (1991).
- [10] Harvey Leff and Andrew F. Rex, *Maxwell’s Demon 2: Entropy, Classical and Quantum Information, Computing* (Institute of Physics Publishing, Bristol, 2002).
- [11] L. Szilárd, “Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen,” *Zeitschrift für Physik* **53**(11-12), 840–856 (1929).
- [12] Charles H. Bennett, “The thermodynamics of computation—a review,” *International Journal of Theoretical Physics* **21**(12), 905–940 (1982).
- [13] Roman S. Ingarden, “Quantum information theory,” *Reports on Mathematical Physics* **10**(1), 43–72 (1976).
- [14] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [15] Benjamin Schumacher, “Quantum coding,” *Physical Review A* **51**(4), 2738 (1995).
- [16] David J. C. MacKay, *Information Theory, Inference & Learning Algorithms* (Cambridge University Press, 2002).
- [17] H. P. Robertson, “The Uncertainty Principle,” *Physical Review* **34**(1), 163 (1929).
- [18] Hans Maassen and J. B. M. Uffink, “Generalized entropic uncertainty relations,” *Physical Review Letters* **60**(12), 1103 (1988).

- [19] Hugh Everett, *The Theory of the Universal Wavefunction*, Ph.D. thesis, Princeton University (1957).
- [20] Bryce S DeWitt and Neill Graham, *Many-worlds Interpretation of Quantum Mechanics* (Princeton University Press, 1973).
- [21] I. I. Hirschman, "A Note on Entropy," *American Journal of Mathematics* **79**(1), 152–156 (1957).
- [22] William Beckner, "Inequalities in Fourier Analysis," *The Annals of Mathematics* **102**(1), 159–182 (1975).
- [23] Iwo Białynicki-Birula and Jerzy Mycielski, "Uncertainty relations for information entropy in wave mechanics," *Communications in Mathematical Physics* **44**(2), 129–132 (1975).
- [24] David Deutsch, "Uncertainty in Quantum Measurements," *Physical Review Letters* **50**(9), 631 (1983).
- [25] K. Kraus, "Complementary observables and uncertainty relations," *Physical Review D* **35**(10), 3070 (1987).
- [26] A. Einstein, B. Podolsky, and N. Rosen, "Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?" *Physical Review* **47**(10), 777 (1935).
- [27] David Bohm, *Quantum Theory* (Dover Publications, New York, 1989).
- [28] Michael J. W. Hall, "Information Exclusion Principle for Complementary Observables," *Physical Review Letters* **74**(17), 3307 (1995).
- [29] Nicolas J. Cerf, Mohamed Bourennane, Anders Karlsson, and Nicolas Gisin, "Security of Quantum Key Distribution Using d-Level Systems," *Physical Review Letters* **88**(12), 127902 (2002).
- [30] M. Christandl and A. Winter, "Uncertainty, Monogamy, and Locking of Quantum Correlations," *IEEE Transactions on Information Theory* **51**(9), 3159–3165 (2005).
- [31] Patrick J. Coles, Li Yu, Vlad Gheorghiu, and Robert B. Griffiths, "Information-theoretic treatment of tripartite systems and quantum channels," *Physical Review A* **83**(6), 062338 (2011).
- [32] A. S. Holevo, "Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel," *Problemy Peredachi Informatsii* **9**(3), 3–11 (1973).
- [33] A. S. Holevo, "Statistical problems in quantum physics," in *Proceedings of the Second Japan-USSR Symposium on Probability Theory*, volume 330 of *Lecture Notes in Mathematics*, 104–119 (Springer, Berlin, 1973).
- [34] G. David Forney, *The concepts of state and entropy in quantum mechanics*, M.S. thesis, Massachusetts Institute of Technology (1963).
- [35] Charles H. Bennett and Stephen J. Wiesner, "Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states," *Physical Review Letters* **69**(20), 2881 (1992).

-
- [36] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters, “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels,” *Physical Review Letters* **70**(13), 1895 (1993).
- [37] John von Neumann, “Probabilistic logics and the synthesis of reliable organisms from unreliable components,” in *Automata Studies*, volume 34 of *Annals of Mathematical Studies*, 43–98 (Princeton University Press, 1956).
- [38] Peter Gács, “Reliable computation with cellular automata,” *Journal of Computer and System Sciences* **32**(1), 15–78 (1986).
- [39] Asher Peres, “Reversible logic and quantum computers,” *Physical Review A* **32**(6), 3266 (1985).
- [40] Rolf Landauer, “Computation and physics: Wheeler’s meaning circuit?” *Foundations of Physics* **16**(6), 551–564 (1986).
- [41] Rolf Landauer, “Dissipation and noise immunity in computation and communication,” *Nature* **335**(6193), 779–784 (1988).
- [42] Rolf Landauer, “The Physical Nature of Information,” *Physics Letters A* **217**(4-5), 188–193 (1996).
- [43] W. H. Zurek, “Reversibility and Stability of Information Processing Systems,” *Physical Review Letters* **53**(4), 391 (1984).
- [44] Peter W. Shor, “Scheme for reducing decoherence in quantum computer memory,” *Physical Review A* **52**(4), R2493 (1995).
- [45] A. M. Steane, “Error Correcting Codes in Quantum Theory,” *Physical Review Letters* **77**(5), 793 (1996).
- [46] Robert B. Laughlin, *A different universe: reinventing physics from the bottom down* (Basic Books, New York, 2006).
- [47] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error,” in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing, STOC ’97*, 176–188 (ACM, New York, 1997).
- [48] Dorit Aharonov and Michael Ben-Or, “Fault-Tolerant Quantum Computation with Constant Error Rate,” *SIAM Journal on Computing* **38**(4), 1207 (2008).
- [49] A Yu Kitaev, “Quantum computations: algorithms and error correction,” *Russian Mathematical Surveys* **52**(6), 1191–1249 (1997).
- [50] Emanuel Knill, Raymond Laflamme, and Wojciech H. Zurek, “Resilient Quantum Computation,” *Science* **279**(5349), 342–345 (1998).
- [51] Emanuel Knill, Raymond Laflamme, and Wojciech H. Zurek, “Resilient Quantum Computation: Error Models and Thresholds,” *Proceedings: Mathematical, Physical and Engineering Sciences* **454**(1969), 365–384 (1998).
- [52] Emanuel Knill and Raymond Laflamme, “Theory of quantum error-correcting codes,” *Physical Review A* **55**(2), 900 (1997).

- [53] Charles H. Bennett, David P. DiVincenzo, John A. Smolin, and William K. Wootters, “Mixed-state entanglement and quantum error correction,” *Physical Review A* **54**(5), 3824 (1996).
- [54] C. E. Shannon, “Communication Theory of Secrecy Systems,” *Bell System Technical Journal* **28**(4), 656 (1949).
- [55] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera, “Quantum Privacy Amplification and the Security of Quantum Cryptography over Noisy Channels,” *Physical Review Letters* **77**(13), 2818 (1996).
- [56] Hoi-Kwong Lo and H. F. Chau, “Unconditional Security of Quantum Key Distribution over Arbitrarily Long Distances,” *Science* **283**(5410), 2050–2056 (1999).
- [57] Peter W. Shor and John Preskill, “Simple Proof of Security of the BB84 Quantum Key Distribution Protocol,” *Physical Review Letters* **85**(2), 441 (2000).
- [58] A. R. Calderbank and Peter W. Shor, “Good quantum error-correcting codes exist,” *Physical Review A* **54**(2), 1098 (1996).
- [59] Andrew Steane, “Multiple-Particle Interference and Quantum Error Correction,” *Proceedings of the Royal Society A* **452**(1954), 2551–2577 (1996).
- [60] Daniel Gottesman, *Stabilizer Codes and Quantum Error Correction*, Ph.D. thesis, California Institute of Technology (1997).
- [61] Charles H. Bennett and Gilles Brassard, “Quantum Cryptography: Public Key Distribution and Coin Tossing,” in *Proceedings of IEEE International Conference on Computers Systems and Signal Processing*, 175–179 (IEEE, New York, 1984).
- [62] Charles H. Bennett, Gilles Brassard, and N. David Mermin, “Quantum cryptography without Bell’s theorem,” *Physical Review Letters* **68**(5), 557 (1992).
- [63] Masato Koashi, “Complementarity, distillable secret key, and distillable entanglement,” arXiv:0704.3661v1 [quant-ph] (2007).
- [64] Benjamin Schumacher, “Sending entanglement through noisy quantum channels,” *Physical Review A* **54**(4), 2614 (1996).
- [65] Benjamin Schumacher and Michael D. Westmoreland, “Approximate Quantum Error Correction,” *Quantum Information Processing* **1**(1), 5–12 (2002).
- [66] Michał Horodecki, Jonathan Oppenheim, and Andreas Winter, “Partial quantum information,” *Nature* **436**(7051), 673–676 (2005).
- [67] Michał Horodecki, Jonathan Oppenheim, and Andreas Winter, “Quantum State Merging and Negative Information,” *Communications in Mathematical Physics* **269**(1), 107–136 (2007).
- [68] Michał Horodecki, Seth Lloyd, and Andreas Winter, “Quantum coding theorem from privacy and distinguishability,” *Open Systems & Information Dynamics* **15**(1), 47–69 (2008).
- [69] Patrick Hayden, Peter W. Shor, and Andreas Winter, “Random quantum codes from Gaussian ensembles and an uncertainty relation,” *Open Systems & Information Dynamics* **15**(1), 71–89 (2008).

-
- [70] Patrick Hayden, Michał Horodecki, Andreas Winter, and Jon Yard, “A decoupling approach to the quantum capacity,” *Open Systems & Information Dynamics* **15**(1), 7–19 (2008).
- [71] Rochus Klesse, “A Random Coding Based Proof for the Quantum Coding Theorem,” *Open Systems & Information Dynamics* **15**(1), 24–45 (2008).
- [72] Rochus Klesse, “Approximate quantum error correction, random codes, and quantum channel capacity,” *Physical Review A* **75**(6), 062315–13 (2007).
- [73] D. Kretschmann, D. Schlingemann, and R.F. Werner, “The Information-Disturbance Trade-off and the Continuity of Stinespring’s Representation,” *IEEE Transactions on Information Theory* **54**(4), 1708–1717 (2008).
- [74] Dennis Kretschmann, David W. Kribs, and Robert W. Spekkens, “Complementarity of private and correctable subsystems in quantum cryptography and error correction,” *Physical Review A* **78**(3), 032330 (2008).
- [75] Cédric Bény and Ognian Oreshkov, “General Conditions for Approximate Quantum Error Correction and Near-Optimal Recovery Channels,” *Physical Review Letters* **104**(12), 120501 (2010).
- [76] Hui Khoon Ng and Prabha Mandayam, “Simple approach to approximate quantum error correction based on the transpose channel,” *Physical Review A* **81**(6), 062342 (2010).
- [77] Mary Beth Ruskai, “Inequalities for quantum entropy: A review with conditions for equality,” *Journal of Mathematical Physics* **43**(9), 4358–4375 (2002).
- [78] A. S. Holevo, “Statistical decision theory for quantum systems,” *Journal of Multivariate Analysis* **3**(4), 337–394 (1973).
- [79] David P. DiVincenzo, Michał Horodecki, Debbie W. Leung, John A. Smolin, and Barbara M. Terhal, “Locking Classical Correlations in Quantum States,” *Physical Review Letters* **92**(6), 067902 (2004).
- [80] E. Davies, “Information and quantum measurement,” *IEEE Transactions on Information Theory* **24**(5), 596–599 (1978).
- [81] Hans Aschauer and Hans J. Briegel, “Security proof of quantum cryptography based entirely on entanglement purification,” *Physical Review A* **66**(3), 032302 (2002).
- [82] Karol Horodecki, Michał Horodecki, Paweł Horodecki, and Jonathan Oppenheim, “Secure Key from Bound Entanglement,” *Physical Review Letters* **94**(16), 160502–4 (2005).
- [83] J. Lawrence Carter and Mark N. Wegman, “Universal classes of hash functions (Extended Abstract),” in *Proceedings of the ninth annual ACM symposium on Theory of computing, STOC ’77*, 106–112 (ACM, New York, 1977).
- [84] J. Lawrence Carter and Mark N. Wegman, “Universal classes of hash functions,” *Journal of Computer and System Sciences* **18**(2), 143–154 (1979).
- [85] A.S. Holevo, “The capacity of the quantum channel with general signal states,” *IEEE Transactions on Information Theory* **44**(1), 269–273 (1998).

- [86] Benjamin Schumacher and Michael D. Westmoreland, “Sending classical information via noisy quantum channels,” *Physical Review A* **56**(1), 131 (1997).
- [87] A. S. Holevo, “On asymptotically optimal hypotheses testing in quantum statistics,” *Teoriya Veroyatnostei i ee Primeneniya* **23**(2), 429–432 (1978).
- [88] Paul Hausladen and William K. Wootters, “A ‘Pretty Good Measurement for Distinguishing Quantum States,” *Journal of Modern Optics* **41**(12), 2385 (1994).
- [89] David Slepian and Jack K. Wolf, “Noiseless coding of correlated information sources,” *IEEE Transactions on Information Theory* **19**(4), 471–480 (1973).
- [90] Andreas Winter, “Coding theorem and strong converse for quantum channels,” *IEEE Transactions on Information Theory* **45**(7), 2481–2485 (1999).
- [91] I. Devetak and A. Winter, “Classical data compression with quantum side information,” *Physical Review A* **68**(4), 042301 (2003).
- [92] Igor Devetak and Andreas Winter, “Distillation of secret key and entanglement from quantum states,” *Proceedings of the Royal Society A* **461**(2053), 207–235 (2005).
- [93] Benjamin Schumacher and M. A. Nielsen, “Quantum data processing and error correction,” *Physical Review A* **54**(4), 2629 (1996).
- [94] Howard Barnum, M. A. Nielsen, and Benjamin Schumacher, “Information transmission through a noisy quantum channel,” *Physical Review A* **57**(6), 4153 (1998).
- [95] Howard Barnum, Emanuel Knill, and Michael A. Nielsen, “On quantum fidelities and channel capacities,” *IEEE Transactions on Information Theory* **46**(4), 1317–1329 (2000).
- [96] Seth Lloyd, “Capacity of the noisy quantum channel,” *Physical Review A* **55**(3), 1613 (1997).
- [97] Peter W. Shor, “The quantum channel capacity and coherent information,” Presented at the MSRI Workshop on Quantum Computation (2002), URL <http://www.msri.org/publications/ln/msri/2002/quantumcrypto/shor/1/>.
- [98] Igor Devetak, “The private classical capacity and quantum capacity of a quantum channel,” *IEEE Transactions on Information Theory* **51**(1), 44– 55 (2005).
- [99] Mitsuru Hamada, “Reliability of Calderbank–Shor–Steane codes and security of quantum key distribution,” *Journal of Physics A: Mathematical and General* **37**(34), 8303–8328 (2004).
- [100] G. David Forney, *Concatenated Codes*, Ph.D. thesis, Massachusetts Institute of Technology (1966).
- [101] Mitsuru Hamada, “Conjugate Codes for Secure and Reliable Information Transmission,” in *Proceedings of the 2006 IEEE Information Theory Workshop*, 149–153 (IEEE, New York, 2006).
- [102] Mitsuru Hamada, “Efficient quotient codes decodable in polynomial time for quantum error correction and cryptography,” in *Proceedings of the 2008 International Symposium on Information Theory and Its Applications*, 1–6 (IEEE, New York, 2008).

-
- [103] Mitsuru Hamada, “Concatenated Quantum Codes Constructible in Polynomial Time: Efficient Decoding and Error Correction,” *IEEE Transactions on Information Theory* **54**(12), 5689–5704 (2008).
- [104] G. S. Vernam, “Cipher Printing Telegraph Systems For Secret Wire and Radio Telegraphic Communications,” *Transactions of the American Institute of Electrical Engineers* **45**, 295–301 (1926).
- [105] Charles Bennett, Gilles Brassard, and Jean-Marc Robert, “How to Reduce your Enemys Information,” in *Advances in Cryptology – CRYPTO 85*, volume 218 of *Lecture Notes in Computer Science*, 468–476 (Springer, Berlin, 1986).
- [106] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert, “Privacy Amplification by Public Discussion,” *SIAM Journal on Computing* **17**(2), 210–229 (1988).
- [107] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Ueli M. Maurer, “Generalized privacy amplification,” *IEEE Transactions on Information Theory* **41**(6), 1915–1923 (1995).
- [108] Renato Renner and Robert König, “Universally Composable Privacy Amplification Against Quantum Adversaries,” in *Proceedings of the Second Theory of Cryptography Conference*, volume 3378 of *Lecture Notes in Computer Science*, 407–425 (Springer, Berlin, 2005).
- [109] Ronen Shaltiel, “Recent Developments in Extractors,” in *Current trends in theoretical computer science. The Challenge of the New Century. Vol 1: Algorithms and Complexity*, 189–228 (World Scientific Publishing, Singapore, 2004).
- [110] Luca Trevisan, “Construction of extractors using pseudo-random generators (extended abstract),” in *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, STOC ’99, 141–148 (ACM, New York, 1999).
- [111] Luca Trevisan, “Extractors and pseudorandom generators,” *Journal of the ACM* **48**(4), 860–879 (2001).
- [112] Anindya De, Christopher Portmann, Thomas Vidick, and Renato Renner, “Trevisan’s extractor in the presence of quantum side information,” *SIAM Journal on Computing* **41**(4), 915–940 (2012).
- [113] Renato Renner, *Security of Quantum Key Distribution*, Ph.D. thesis, ETH Zurich (2005).
- [114] Robert König and Renato Renner, “Sampling of min-entropy relative to quantum knowledge,” *IEEE Transactions on Information Theory* **57**(7), 4760–4787 (2011).
- [115] Marco Tomamichel, Christian Schaffner, Adam Smith, and Renato Renner, “Leftover Hashing Against Quantum Side Information,” *IEEE Transactions on Information Theory* **57**(8), 5524–5535 (2011).
- [116] Joseph M. Renes and Renato Renner, “One-Shot Classical Data Compression with Quantum Side Information and the Distillation of Common Randomness or Secret Keys,” *IEEE Transactions on Information Theory* **58**(3), 1985–1991 (2012).
- [117] Robert König, Renato Renner, and Christian Schaffner, “The Operational Meaning of Min- and Max-Entropy,” *IEEE Transactions on Information Theory* **55**(9), 4337–4347 (2009).

- [118] Marco Tomamichel and Renato Renner, “Uncertainty Relation for Smooth Entropies,” *Physical Review Letters* **106**(11), 110506 (2011).
- [119] Renato Renner, Stefan Wolf, and Jürg Wullschleger, “The Single-Serving Channel Capacity,” in *Proceedings of the 2006 IEEE International Symposium on Information Theory*, 1424–1427 (IEEE, New York, 2006).
- [120] Ligong Wang and Renato Renner, “One-Shot Classical-Quantum Capacity and Hypothesis Testing,” *Physical Review Letters* **108**(20), 200501 (2012).
- [121] Marco Tomamichel, Roger Colbeck, and Renato Renner, “A Fully Quantum Asymptotic Equipartition Property,” *IEEE Transactions on Information Theory* **55**(12), 5840–5847 (2009).
- [122] A. D. Wyner, “The Wire-Tap Channel,” *Bell System Technical Journal* **54**(8), 1355–1387 (1975).
- [123] Rudolph Ahlswede and Imre Csiszàr, “Common randomness in information theory and cryptography. I. Secret sharing,” *IEEE Transactions on Information Theory* **39**(4), 1121–1132 (1993).
- [124] Ueli Maurer and Stefan Wolf, “Information-Theoretic Key Agreement: From Weak to Strong Secrecy for Free,” in *Advances in Cryptology – Eurocrypt 2000*, volume 1807 of *Lecture Notes on Computer Science* (Springer, Berlin, 2000).
- [125] Eli Biham, Michel Boyer, P. Oscar Boykin, Tal Mor, and Vwani Roychowdhury, “A proof of the security of quantum key distribution (extended abstract),” in *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, STOC ’00, 715–724 (ACM, New York, 2000).
- [126] Eli Biham, Michel Boyer, P. Oscar Boykin, Tal Mor, and Vwani Roychowdhury, “A Proof of the Security of Quantum Key Distribution,” *Journal of Cryptology* **19**(4), 381–439 (2006).
- [127] Dominic Mayers, “Quantum Key Distribution and String Oblivious Transfer in Noisy Channels,” in *Advances in Cryptology – CRYPTO 96*, volume 1109 of *Lecture Notes in Computer Science*, 343–357 (Springer, Berlin, 1996).
- [128] Dominic Mayers, “Unconditional security in quantum cryptography,” *Journal of the ACM* **48**(3), 351–406 (2001).
- [129] Andrew Chi-Chih Yao, “Security of quantum protocols against coherent measurements,” in *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, STOC ’95, 67–75 (ACM, New York, 1995).
- [130] Hoi-Kwong Lo, “Proof of unconditional security of six-state quantum key distribution scheme,” *Quantum Information and Computation* **1**(2), 81–94 (2001).
- [131] Dagmar Bruß, “Optimal Eavesdropping in Quantum Cryptography with Six States,” *Physical Review Letters* **81**(14), 3018 (1998).
- [132] Kiyoshi Tamaki, Masato Koashi, and Nobuyuki Imoto, “Unconditionally Secure Key Distribution Based on Two Nonorthogonal States,” *Physical Review Letters* **90**(16), 167904 (2003).

-
- [133] Charles H. Bennett, “Quantum cryptography using any two nonorthogonal states,” *Physical Review Letters* **68**(21), 3121 (1992).
- [134] Daniel Gottesman and Hoi-Kwong Lo, “Proof of security of quantum key distribution with two-way classical communications,” *IEEE Transactions on Information Theory* **49**(2), 457–475 (2003).
- [135] J.-C. Boileau, K. Tamaki, J. Batuwantudawe, R. Laflamme, and J. M. Renes, “Unconditional Security of a Three State Quantum Key Distribution Protocol,” *Physical Review Letters* **94**(4), 040503–4 (2005).
- [136] Simon J. D. Phoenix, Stephen M. Barnett, and Anthony Chefles, “Three-state quantum cryptography,” *Journal of Modern Optics* **47**(2), 507 (2000).
- [137] Masato Koashi and John Preskill, “Secure Quantum Key Distribution with an Uncharacterized Source,” *Physical Review Letters* **90**(5), 057902 (2003).
- [138] Masato Koashi, “Unconditional security of quantum key distribution and the uncertainty principle,” *Journal of Physics: Conference Series* **36**(1), 98–102 (2006).
- [139] M Koashi, “Simple security proof of quantum key distribution based on complementarity,” *New Journal of Physics* **11**(4), 045018 (2009).
- [140] Marco Tomamichel, Charles Ci Wen Lim, Nicolas Gisin, and Renato Renner, “Tight Finite-Key Analysis for Quantum Cryptography,” arXiv:1103.4130v1 [quant-ph] (2011).
- [141] Michael Ben-Or, “Security of quantum key distribution protocols II,” Presented at the MSRI Introductory Workshop in Quantum Computation (2002), URL <http://www.msri.org/realvideo/ln/msri/2002/quantumintro/ben-or/2/>.
- [142] Robert König, Ueli Maurer, and Renato Renner, “On the Power of Quantum Memory,” *IEEE Transactions on Information Theory* **51**(7), 2391–2401 (2005).
- [143] Matthias Christandl, Renato Renner, and Artur Ekert, “A Generic Security Proof for Quantum Key Distribution,” arXiv:quant-ph/0402131v2 (2004).
- [144] B. Kraus, N. Gisin, and R. Renner, “Lower and Upper Bounds on the Secret-Key Rate for Quantum Key Distribution Protocols Using One-Way Classical Communication,” *Physical Review Letters* **95**(8), 080501–4 (2005).
- [145] Renato Renner, Nicolas Gisin, and Barbara Kraus, “Information-theoretic security proof for quantum-key-distribution protocols,” *Physical Review A* **72**(1), 012332–17 (2005).
- [146] Michael Ben-Or, Michał Horodecki, Debbie W. Leung, Dominic Mayers, and Jonathan Oppenheim, “The Universal Composable Security of Quantum Key Distribution,” in *Proceedings of the Second Theory of Cryptography Conference*, volume 3378 of *Lecture Notes in Computer Science*, 386–406 (Springer, Berlin, 2005).
- [147] Renato Renner, “Symmetry of large physical systems implies independence of subsystems,” *Nature Physics* **3**(9), 645–649 (2007).

- [148] Matthias Christandl, Robert König, and Renato Renner, “Postselection Technique for Quantum Channels with Applications to Quantum Cryptography,” *Physical Review Letters* **102**(2), 020504–4 (2009).
- [149] Joseph M. Renes, *Frames, Designs, and Spherical Codes in Quantum Information Theory*, Ph.D. thesis, University of New Mexico (2004).
- [150] Joseph M. Renes, “Spherical-code key-distribution protocols for qubits,” *Physical Review A* **70**(5), 052314 (2004).
- [151] Joseph M. Renes, “Equiangular spherical codes in quantum cryptography,” *Quantum Information and Computation* **5**(1), 081–092 (2005).
- [152] David P. DiVincenzo, Peter W. Shor, and John A. Smolin, “Quantum-channel capacity of very noisy channels,” *Physical Review A* **57**(2), 830 (1998).
- [153] M. B. Hastings, “Superadditivity of communication capacity using entangled inputs,” *Nature Physics* **5**(4), 255–257 (2009).
- [154] M. Tomamichel, R. Colbeck, and R. Renner, “Duality Between Smooth Min- and Max-Entropies,” *IEEE Transactions on Information Theory* **56**(9), 4674–4681 (2010).

Included Papers

- [BCC⁺10] Mario Berta, Matthias Christandl, Roger Colbeck, Joseph M. Renes, and Renato Renner, “The uncertainty principle in the presence of quantum memory,” *Nature Physics* **6**(9), 659–662 (2010).
- [BR09] Jean-Christian Boileau and Joseph M. Renes, “Optimal State Merging without Decoupling,” in *Proceedings of the Fourth Workshop on Theory of Quantum Computation, Communication, and Cryptography*, volume 5906 of *Lecture Notes in Computer Science*, 76–84 (Springer, Berlin, 2009).
- [KR08] Oliver Kern and Joseph M. Renes, “Improved one-way rates for BB84 and 6-state protocols,” *Quantum Information and Computation* **8**(8/9), 756–772 (2008).
- [RB08] Joseph M. Renes and Jean-Christian Boileau, “Physical underpinnings of privacy,” *Physical Review A* **78**(3), 032335–12 (2008).
- [RB09] Joseph M. Renes and Jean-Christian Boileau, “Conjectured Strong Complementary Information Tradeoff,” *Physical Review Letters* **103**(2), 020402–4 (2009).
- [Ren11] Joseph M. Renes, “Duality of privacy amplification against quantum adversaries and data compression with quantum side information,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* **467**(2130), 1604–1623 (2011).
- [RG06] Joseph M. Renes and Markus Grassl, “Generalized decoding, effective channels, and simplified security proofs in quantum key distribution,” *Physical Review A* **74**(2), 022317–9 (2006).
- [RR11] Joseph M. Renes and Renato Renner, “Noisy channel coding via privacy amplification and information reconciliation,” *IEEE Transactions on Information Theory* **57**(11), 7377–7385 (2011).
- [RS07] Joseph M. Renes and Graeme Smith, “Noisy Processing and Distillation of Private Quantum States,” *Physical Review Letters* **98**(2), 020502–4 (2007).
- [SRS08] Graeme Smith, Joseph M. Renes, and John A. Smolin, “Structured Codes Improve the Bennett-Brassard-84 Quantum Key Rate,” *Physical Review Letters* **100**(17), 170502–4 (2008).

