

Improved DIQKD protocols with finite-size analysis

Ernest Y.-Z. Tan¹, Pavel Sekatski^{2,3}, Jean-Daniel Bancal⁴, René Schwonnek⁵, Renato Renner¹, Nicolas Sangouard⁴, and Charles C.-W. Lim^{6,7}

¹*Institute for Theoretical Physics, ETH Zürich, Switzerland*

²*Department of Physics, University of Basel, Klingelbergstrasse 82, 4056 Basel, Switzerland*

³*Department of Applied Physics, University of Geneva, Chemin de Pinchat 22, 1211 Geneva, Switzerland*

⁴*Université Paris-Saclay, CEA, CNRS, Institut de physique théorique, 91191, Gif-sur-Yvette, France*

⁵*Naturwissenschaftlich-Technische Fakultät, Universität Siegen, Germany*

⁶*Department of Electrical & Computer Engineering, National University of Singapore, Singapore*

⁷*Centre for Quantum Technologies, National University of Singapore, Singapore*

Abstract

The security of finite-length keys is essential for the implementation of device-independent quantum key distribution (DIQKD). Presently, there are several finite-size DIQKD security proofs, but they are mostly focused on standard DIQKD protocols and do not directly apply to the recent improved DIQKD protocols based on noisy preprocessing, random key measurements, and modified CHSH inequalities. Here, we provide a general finite-size security proof that can simultaneously encompass these approaches, using tighter finite-size bounds than previous analyses. In doing so, we develop a method to compute tight lower bounds on the asymptotic keyrate for any such DIQKD protocol with binary inputs and outputs. With this, we show that positive asymptotic keyrates are achievable up to depolarizing noise values of 9.33%, exceeding all previously known noise thresholds. We also develop a modification to random-key-measurement protocols, using a pre-shared seed followed by a “seed recovery” step, which yields substantially higher net key generation rates by essentially removing the sifting factor. Some of our results may also improve the keyrates of device-independent randomness expansion.

1 Introduction

Device-independent quantum key distribution (DIQKD) is a cryptographic concept based on the observation that if some quantum devices violate a Bell inequality, then it is possible to distill a secret key from the devices’ outputs, even when the devices are not fully characterized [BHK05, PAB+09, Sca13]. This is a stronger form of security than that offered by standard QKD protocols, which assume that the devices are performing measurements within specified tolerances [SBPC+09]. In recent years, experimental and theoretical developments have brought the possibility of a physical DIQKD demonstration closer to fruition. In particular, a series of experiments have achieved Bell inequality violations while closing the fair-sampling loophole, using NV-centre [HBD+15], photonic [GMR+13, CMA+13, SMSC+15, GVW+15, LZL+18, SLT+18], and cold-atom [RBG+17] implementations. On the theoretical front, several protocol modifications have been explored that improve the asymptotic keyrates and noise tolerance of DIQKD. We focus on those studied in [HST+20] (*noisy preprocessing*, in which a small amount of trusted noise is added to the device outputs), [SGP+21] (*random key measurements*, in which more than one measurement basis

is used to generate the key, preventing an adversary from optimally attacking both bases simultaneously), and [WAP21, SBV+21] (*modified CHSH inequalities*, which certify more entropy than the standard CHSH inequality).

Furthermore, earlier device-independent security proofs [PAB+09] required the assumption that the device behaviour across multiple rounds is independent and identically distributed (IID), sometimes referred to as the assumption of *collective attacks*. With the development of a result known as the *entropy accumulation theorem* (EAT) [DFR20, AFRV19, DF19], this assumption can now¹ be removed for some DIQKD protocols — security proofs based on the EAT are valid against general attacks (sometimes referred to as *coherent attacks*), as long as the device behaviour can be modelled in a sequential manner. The EAT also provides explicit bounds on the finite-size behaviour, hence concretely addressing the question of what sample size is required for a secure demonstration of DIQKD. (Another recent result that can be applied to obtain such bounds against general attacks is the *quantum probability estimation* technique [ZKB18], but in this work we focus on the EAT.)

However, thus far there has been no comprehensive analysis which simultaneously encompasses the various approaches mentioned above. In this work, we present a finite-size security proof that can be applied to protocols which incorporate all of these proposed improvements. In addition, we show that such protocols can achieve substantially higher depolarizing noise tolerance than all previous proven results. We now highlight the main contributions of our work.

1.1 Summary of key results

We consider a protocol based on one-way error correction (Protocol 1) that combines noisy preprocessing and random key measurements, and we perform a finite-size analysis against general attacks (Theorem 1) as well as collective attacks (Theorem 4). Our approach is similar to that used in [AFRV19]; however, we slightly modify the analysis in order to relax the theoretical requirements for the error-correction step, and our bounds are tighter since we use an updated version [DF19, LLR+21] of the EAT. Roughly speaking, the proofs rely on lower bounds on the asymptotic keyrates, and hence in Sec. 5 we also describe an algorithm to compute such bounds, which can be applied independently of our finite-size analysis. This algorithm improves over previous results in [HST+20, SGP+21, WAP21, SBV+21] by having all of the following properties simultaneously: it applies to arbitrary 2-input 2-output protocols, it accounts for noisy preprocessing and random key measurements, and it provably converges to a tight bound² (for protocols of this form).

For simplicity, the protocol we present only uses the CHSH inequality rather than the modified CHSH inequalities of [WAP21, SBV+21], because our results suggest (see Sec. 5.6) that within the 2-input 2-output scenario, there may not be much prospect for improvement by using the latter, at least for the depolarizing-noise model. However, the finite-size analysis we perform can (like [BRC20]) be applied to protocols based on arbitrary Bell inequalities — in Sec. 4.4, we briefly explain the relevant adjustments in that case. Similarly, while we only performed explicit computations of the asymptotic keyrates for CHSH-based protocols, our approach as described in Sec. 5 can be applied to all 2-input 2-output Bell inequalities.

With our formulas for the finite-size keyrates, we computed the keyrates for several scenar-

¹To be precise, other proof techniques [NSPS14, VV14, JMS20] are available that do not require the IID assumption, but the asymptotic keyrates given by those techniques are lower than that of the EAT.

²In fact, similar to [WAP21], we find numerical evidence that the bound in [SGP+21] is not entirely tight, and also that a conjecture proposed there regarding Eve’s optimal attack may not be true after all; see Sec. 5.4.

ios. In particular, we studied the keyrates that could be achieved if the honest devices had performance described by the estimated parameters in [MDR+19] for the NV-centre [HBD+15] and cold-atom [RBG+17] loophole-free Bell tests (detailed keyrate plots are shown in Fig. 1 of Sec. 3.4). For photonic experiments, it is currently somewhat unclear whether the protocol here provides improvements over the results in [HST+20, WAP21, SBV+21] — there are some complications which we explain in Sec. 5.5. We also plot some finite-size keyrates for honest devices subject to a simple depolarizing-noise model (see Fig. 2 of Sec. 3.4).

From those computations, we found that our security proof would require the [HBD+15] and [RBG+17] experiments to run for approximately $n \sim 10^8$ and 10^{10} rounds respectively in order to certify a positive finite-size keyrate against general attacks. While this is a marked improvement over the basic [PAB+09] protocol (which yields zero asymptotic keyrate for those experiments), these requirements still appear to be quite far outside the reach of those implementations — for reference, the [HBD+15] experiment had sample size $n = 245$ (over a 220-hour period), while the [RBG+17] experiment collected a data set of size $n = 10,000$ (over 2 measurement runs in a 10-day period) and also a larger data set of size $n = 55,568$ (over multiple measurement runs in a 7-month period). Our keyrate plots (Fig. 1) also showed that changing the various security parameters (discussed in detail in Sec. 2.1) by several orders of magnitude only results in fairly small changes to the keyrate, so it appears unlikely that the keyrates could be substantially improved by relaxing these security requirements.

To improve on these results, in Sec. 6 we describe two modifications of Protocol 1. Firstly, Protocol 2 is a modification based on a pre-shared key, which achieves a *net* key generation rate approximately double that of Protocol 1, by overcoming a crucial disadvantage of random-key-measurement protocols (namely, the sifting factor). Secondly, Protocol 3 is a modification that is optimized for the collective-attacks assumption — by changing the protocol itself for this scenario rather than just the security proof, we were able to further improve the keyrate. However, we find that even with these modifications and relaxed security requirements, it appears that the required number of rounds for a positive finite-size keyrate is still impractically large, at about $n \sim 10^6$ to 10^7 for the [RBG+17] parameters (see Fig. 6 in Sec. 6).

As another immediate consequence of our work, we computed a lower bound on the asymptotic noise tolerance of our protocols against depolarizing noise, by using our algorithm for evaluating the asymptotic keyrates. Our results (shown in Fig. 4 of Sec. 5.5) certify that positive asymptotic keyrates are possible for depolarizing-noise values of up to 9.33% at least. In comparison, the previous best thresholds were 8.34% [WAP21] (using noisy preprocessing and modified CHSH inequalities) or 8.2% [SGP+21] (using random key measurements). Our improvement over these results is hence of similar magnitude to their improvements over the basic [PAB+09] protocol, which has a threshold of 7.15%. Furthermore, our bounds are in fact close to the highest possible bounds allowed by simple convexity arguments — we prove in Sec. 5.6 that all protocols of this general form cannot achieve a depolarizing-noise threshold beyond 9.57%, so the threshold we obtained is not far from the absolute highest possible value in this setting.

Finally, we remark that several results from this work can also be used for DI randomness expansion (DIRE) [Col06, PAM+10, BRC20, LLR+21]. Specifically, the algorithm we use for computing the asymptotic DIQKD keyrates can also be used to obtain tighter bounds on the asymptotic keyrates for DIRE, after which a finite-size analysis could be performed using the EAT [BRC20, LLR+21]. Moreover, the “key recovery” process in the modified protocol (Protocol 2) can be applied in the context of DIRE, hence allowing one to improve the keyrates by using the random-key-measurement approach of [SGP+21]. (This idea was also independently proposed in a

separate work [BRC21] at a similar time.) We discuss this in detail in Sec. 6.2.

1.2 Paper structure

In Sec. 2, we state the definitions and notation we use in this work. In Sec. 3, we describe the main protocol we consider, and state the main theorem which bounds the finite-size keyrates (Theorem 1), followed by presenting some plots of the resulting values. We give the security proof for this finite-size bound in Sec. 4. In Sec. 5 we present the algorithm to compute the asymptotic keyrates (i.e. the leading-order terms in Theorem 1), and describe the resulting depolarizing-noise thresholds as well as an upper bound on these thresholds. Finally, in Sec. 6 we discuss several variations, such as a modified random-key-measurement protocol which bypasses the sifting factor, and security proofs against collective attacks instead of general attacks.

2 Preliminaries

We define some basic notation in Table 1, and state some further definitions below. We take all systems to be finite-dimensional, but we will not impose any bounds on the system dimensions unless otherwise specified.

Table 1: List of notation

<i>Symbol</i>	<i>Definition</i>
\log	Base-2 logarithm
$H(\cdot)$	Base-2 von Neumann entropy
$D(\cdot\ \cdot)$	Base-2 quantum relative entropy
$\ \cdot\ _p$	Schatten p -norm
$\lfloor \cdot \rfloor$ (resp. $\lceil \cdot \rceil$)	Floor (resp. ceiling) function
$X \geq Y$ (resp. $X > Y$)	$X - Y$ is positive semidefinite (resp. positive definite)
$S_=(A)$ (resp. $S_{\leq}(A)$)	Set of normalized (resp. subnormalized) states on register A
\mathbb{U}_A	Maximally mixed state on register A
$[n]$	Indices from 1 to n , i.e. $\{1, 2, \dots, n\}$
$A_{[n]}$	Registers $A_1 \dots A_n$
Ω^c	Complement (i.e. negation) of an event Ω
freq	See Definition 1
$B_{n,p}(k)$	See Definition 2
$\rho_{\wedge\Omega}$ and $\rho_{ \Omega}$	See Definition 3

Definition 1. (Frequency distributions) For a string $\mathbf{z} \in \mathcal{Z}^n$ on some alphabet \mathcal{Z} , $\text{freq}_{\mathbf{z}}$ denotes the following probability distribution on \mathcal{Z} :

$$\text{freq}_{\mathbf{z}}(z) := \frac{1}{n} \sum_{j=1}^n \delta_{z, z_j}. \quad (1)$$

Definition 2. (Binomial distribution) Let $X \sim \text{Binom}(n, p)$ denote a random variable X following a binomial distribution with parameters (n, p) , i.e. X is the sum of n IID Bernoulli random variables X_j with $\Pr[X_j = 1] = p$. We denote the corresponding cumulative distribution function as

$$B_{n,p}(k) := \Pr_{X \sim \text{Binom}(n,p)}[X \leq k]. \quad (2)$$

Definition 3. (Conditioning on classical events) For a classical-quantum state $\rho \in \mathcal{S}_{\leq}(CQ)$ in the form $\rho_{CQ} = \sum_c |c\rangle\langle c| \otimes \omega_c$ for some $\omega_c \in \mathcal{S}_{\leq}(Q)$, and an event Ω defined on the register C , we define the following “conditional states”:

$$\rho_{\wedge\Omega} := \sum_{c \in \Omega} |c\rangle\langle c| \otimes \omega_c, \quad \rho_{|\Omega} := \frac{\text{Tr}[\rho]}{\text{Tr}[\rho_{\wedge\Omega}]} \rho_{\wedge\Omega} = \frac{\sum_c \text{Tr}[\omega_c]}{\sum_{c \in \Omega} \text{Tr}[\omega_c]} \rho_{\wedge\Omega}. \quad (3)$$

We informally refer to these states as the subnormalized and normalized conditional states respectively (the latter is perhaps a slight misnomer if $\text{Tr}[\rho] < 1$, but this situation does not arise in our proofs). The process of taking subnormalized conditional states is commutative and “associative”, in the sense that for any events Ω, Ω' we have $(\rho_{\wedge\Omega})_{\wedge\Omega'} = (\rho_{\wedge\Omega'})_{\wedge\Omega} = \rho_{\wedge(\Omega \wedge \Omega')}$; hence for brevity we will denote all of these expressions as

$$\rho_{\wedge\Omega \wedge \Omega'} := (\rho_{\wedge\Omega})_{\wedge\Omega'} = (\rho_{\wedge\Omega'})_{\wedge\Omega} = \rho_{\wedge(\Omega \wedge \Omega')}. \quad (4)$$

On the other hand, some disambiguating parentheses are needed when combined with taking normalized conditional states.

Definition 4. (2-universal hashing) A *2-universal family of hash functions* is a set \mathcal{H} of functions from a set \mathcal{X} to a set \mathcal{Y} , such that if h is drawn uniformly at random from \mathcal{H} , then

$$\Pr[h(x) = h(x')] \leq \frac{1}{|\mathcal{Y}|} \quad \forall x \neq x'. \quad (5)$$

To ensure consistency of definitions, we now state the definitions of smoothed entropies relevant for this work, though we will not need to use the explicit expressions. We follow the presentation in [DFR20, DF19], which can be shown to be equivalent to the definitions in [Tom16, TL17].

Definition 5. For $\rho, \sigma \in \mathcal{S}_{\leq}(A)$, the *generalized fidelity* is

$$F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1 + \sqrt{(1 - \text{Tr}[\rho])(1 - \text{Tr}[\sigma])}, \quad (6)$$

and the *purified distance* is $P(\rho, \sigma) := \sqrt{1 - F(\rho, \sigma)^2}$.

Definition 6. For $\rho \in \mathcal{S}_{\leq}(AB)$, the *min- and max-entropies of A conditioned on B* are

$$H_{\min}(A|B)_{\rho} := -\log \min_{\substack{\sigma \in \mathcal{S}_{\leq}(B) \text{ s.t.} \\ \ker(\rho_B) \subseteq \ker(\sigma_B)}} \left\| \rho_{AB}^{\frac{1}{2}} (\mathbb{I}_A \otimes \sigma_B)^{-\frac{1}{2}} \right\|_{\infty}^2, \quad (7)$$

$$H_{\max}(A|B)_{\rho} := \log \max_{\sigma \in \mathcal{S}_{\leq}(B)} \left\| \rho_{AB}^{\frac{1}{2}} (\mathbb{I}_A \otimes \sigma_B)^{\frac{1}{2}} \right\|_1^2, \quad (8)$$

where in the first equation the $(\mathbb{I}_A \otimes \sigma_B)^{-\frac{1}{2}}$ term should be understood in terms of the Moore-Penrose generalized inverse. (Note that the optimum is indeed attained in both equations [Tom16], and it can be attained by a normalized state, so $S_{\leq}(B)$ can be replaced by $S_{=}(B)$ without loss of generality.)

Definition 7. For $\rho \in S_{\leq}(AB)$ and $\varepsilon \in \left[0, \sqrt{\text{Tr}[\rho_{AB}]}\right)$, the ε -smoothed min- and max-entropies of A conditioned on B are

$$H_{\min}^{\varepsilon}(A|B)_{\rho} := \max_{\substack{\tilde{\rho} \in S_{\leq}(AB) \text{ s.t.} \\ P(\tilde{\rho}, \rho) \leq \varepsilon}} H_{\min}(A|B)_{\tilde{\rho}}, \quad H_{\max}^{\varepsilon}(A|B)_{\rho} := \min_{\substack{\tilde{\rho} \in S_{\leq}(AB) \text{ s.t.} \\ P(\tilde{\rho}, \rho) \leq \varepsilon}} H_{\max}(A|B)_{\tilde{\rho}}. \quad (9)$$

2.1 Security definitions

The question of formalizing an appropriate security definition for the device-independent setting has not been definitively resolved yet [AFRV19], due to considerations regarding the device-reuse attacks described in [BCK13]. However, we shall proceed by following the security definitions used in [AFRV19], which were based on strong³ security definitions [PR14] for standard QKD, and should be sufficient under suitable constraints on the nature of the device memories (we shall briefly discuss these in the next section, when listing the assumptions).

Qualitatively, the concepts involved in the security definition we use here are: *completeness*, meaning that the honest devices will accept with high probability, and *soundness*, meaning that the devices remain “secure” (possibly by aborting) even in the presence of dishonest behaviour. Note that the completeness concept relies on having some description of the honest devices, which should be understood to mean the device behaviour in the situation where they perform “according to specifications”, without manipulation or eavesdropping attempts from Eve (we give examples of such descriptions in Sec. 3.3). The following definition formalizes these concepts:

Definition 8. Consider a DIQKD protocol such that at the end, the honest parties either *accept* (producing keys K_A and K_B of length ℓ_{key} for Alice and Bob respectively) or *abort* (producing an abort symbol \perp for all parties). It is said to be ε^{com} -complete and ε^{sou} -sound if the following properties hold:

- (Completeness) The honest protocol implementation aborts with probability at most ε^{com} .
- (Soundness) For any implementation of the protocol, we have

$$\Pr[\text{accept}] \frac{1}{2} \left\| \sigma_{K_A K_B E'} - \left(\frac{1}{2^{\ell_{\text{key}}}} \sum_k |kk\rangle \langle kk|_{K_A K_B} \right) \otimes \sigma_{E'} \right\|_1 \leq \varepsilon^{\text{sou}}, \quad (10)$$

where σ denotes the normalized state conditioned on the protocol accepting, and E' denotes all side-information registers available to the adversary at the end of the protocol.

In the security proof, it is convenient to use the fact that the soundness property is implied by a pair of slightly simpler conditions, as shown in [PR14]. Specifically, to prove a DIQKD protocol is ε^{sou} -sound, it suffices to find $\varepsilon_{\text{QKD}}^{\text{cor}}$, $\varepsilon_{\text{QKD}}^{\text{sec}}$ such that $\varepsilon^{\text{sou}} \geq \varepsilon_{\text{QKD}}^{\text{cor}} + \varepsilon_{\text{QKD}}^{\text{sec}}$ and the protocol is both $\varepsilon_{\text{QKD}}^{\text{cor}}$ -correct and $\varepsilon_{\text{QKD}}^{\text{sec}}$ -secret, defined as follows:

³In the sense that the definitions imply *composable security*, which roughly speaking means that we can safely use this protocol in place of a simpler, more idealized functionality that generates an ideal secret key whenever it does not abort — see [PR14] for details.

Definition 9. A DIQKD protocol as described above is said to be $\varepsilon_{\text{QKD}}^{\text{cor}}$ -correct and $\varepsilon_{\text{QKD}}^{\text{sec}}$ -secret if the following properties hold:

- (Correctness) For any implementation of the protocol, we have⁴

$$\Pr[K_A \neq K_B \wedge \text{accept}] \leq \varepsilon_{\text{QKD}}^{\text{cor}}. \quad (11)$$

- (Secrecy) For any implementation of the protocol, we have

$$\Pr[\text{accept}] \frac{1}{2} \|\sigma_{K_A E'} - \mathbb{U}_{K_A} \otimes \sigma_{E'}\|_1 \leq \varepsilon_{\text{QKD}}^{\text{sec}}, \quad (12)$$

where σ is as described in Definition 8, and \mathbb{U}_{K_A} denotes the maximally mixed state (i.e. a uniformly random key for Alice).

3 Main protocol

The overall structure of our main protocol is stated as Protocol 1 below, with the details of some steps to be specified in the following subsections. We make the following fairly standard assumptions, following the presentation in [AFRV19]:

- Alice and Bob can prevent unwanted information from leaking outside of their respective locations (for instance, the inputs and outputs to the devices remain private for each party until/unless they are revealed in the public communication steps).
- Alice and Bob can generate trusted (local) randomness.
- Alice and Bob have trusted post-processing units to perform classical computations.
- Alice and Bob perform all classical communication using an authenticated public channel.
- All systems can be modelled as finite-dimensional quantum registers (though we do not impose any bounds on the dimensions unless otherwise specified).

Regarding the first point in particular, we remark that while the loophole-free Bell tests in [HBD+15, SMSC+15, GVW+15, RBG+17] used spacelike separation to motivate the assumption that the devices do not reveal their inputs to each other, this may not be strictly necessary for a DIQKD implementation. It could be possible, and perhaps more reasonable, to instead justify this assumption by implementing some “shielding” measures on the devices, to prevent them from leaking unwanted information. In any case, “shielding” measures of this nature are likely necessary to prevent the raw *outputs* of the devices from leaking to the adversary, so it may be expedient to use these measures to prevent the inputs from leaking as well.

The above assumptions will be sufficient for us to show that the protocol satisfies the completeness and soundness definitions stated above. However, when considering whether these formal definitions yield security in an intuitive (or composable) sense, some additional assumptions are needed to address the memory attack of [BCK13]. One approach would be to impose the condition

⁴In [AFRV19], this definition is stated slightly differently since the accept condition is implicitly absorbed into the K_A, K_B condition by defining $K_A = K_B = \perp$ when the protocol aborts.

that the devices do not access any registers storing “private” data from previous protocols, and that this condition continues to hold if the devices are reused in the future. (Note that this condition is always implicit in standard QKD, because the register being measured is inherently specified when describing the trusted measurements the devices perform.) It remains to fully formalize this condition in a suitable framework for composable security, but this would be beyond the scope of this work.

Protocol 1

The protocol is defined in terms of the following parameters (chosen before the protocol begins), which we qualitatively describe:

- n : Total number of rounds
- γ : Probability of a test round
- p : Noisy-preprocessing bias
- EC_{\max} : Bound on number of bits used for error correction
- w_{exp} : Expected winning probability for the (IID) honest devices
- δ_{tol} : Tolerated deviation from expected winning probability w_{exp}
- ℓ_{key} : Length of final key

The honest behaviour consists of n IID copies of a device characterized by w_{exp} and an error-correction parameter h_{hon} (details in Sec. 3.3).

1. **Measurement:** For each $j \in [n]$, perform the following steps:
 - 1.1. Alice and Bob’s devices each receive some share of a quantum state.
 - 1.2. Alice chooses a uniform input $X_j \in \{0, 1\}$. With probability γ , Bob chooses a uniform input $Y_j \in \{2, 3\}$, otherwise Bob chooses a uniform input $Y_j \in \{0, 1\}$.
 - 1.3. Alice and Bob supply their inputs to their devices, and record the outputs as A_j and B_j respectively.
2. Alice and Bob publicly announce their input strings \mathbf{X} and \mathbf{Y} .
3. **Sifting:** For all rounds such that $Y_j \in \{0, 1\}$ and $X_j \neq Y_j$, Alice and Bob overwrite their outputs with $A_j = B_j = 0$.
4. **Noisy preprocessing:** For all rounds such that $Y_j \in \{0, 1\}$ and $X_j = Y_j$, Alice generates a biased random bit F_j with $\Pr[F_j = 1] = p$, and overwrites her output A_j with $A_j \oplus F_j$.
5. **Error correction:** Alice and Bob publicly communicate some bits $\mathbf{L} = (\mathbf{L}_{\text{EC}}, \mathbf{L}_{\text{h}})$ for error correction as follows (see Sec. 3.1):
 - 5.1. Alice and Bob communicate some bits \mathbf{L}_{EC} to allow Bob to produce a guess $\tilde{\mathbf{A}}$ for \mathbf{A} . If at some point the number of communicated bits reaches EC_{\max} , Alice and Bob immediately cease communication and proceed to the next step, with Bob producing $\tilde{\mathbf{A}}$ using only the information he has at that point.
 - 5.2. Alice computes a 2-universal hash $\mathbf{L}_{\text{h}} = \text{hash}(\mathbf{A})$ of length $\lceil \log(1/\varepsilon_{\text{h}}) \rceil$. She sends \mathbf{L}_{h} (and the choice of hash function) to Bob.
6. **Parameter estimation:** For all $j \in [n]$, Bob sets $C_j = \perp$ if $Y_j \in \{0, 1\}$; otherwise he sets $C_j = 0$ if $\tilde{A}_j \oplus B_j \neq X_j \cdot (Y_j - 2)$ and $C_j = 1$ if $\tilde{A}_j \oplus B_j = X_j \cdot (Y_j - 2)$.
7. Bob checks whether $\mathbf{L}_{\text{h}} = \text{hash}(\tilde{\mathbf{A}})$, as well as whether the value \mathbf{c} on registers \mathbf{C} satisfies $\text{freq}_{\mathbf{c}}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma$ and $\text{freq}_{\mathbf{c}}(0) \leq (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma$. If all those conditions hold, Alice

and Bob proceed to the next step. Otherwise, the protocol aborts.

8. **Privacy amplification:** Alice and Bob apply privacy amplification (see Sec. 3.2) on \mathbf{A} and $\tilde{\mathbf{A}}$ respectively to obtain final keys K_A and K_B of length ℓ_{key} .
-

The rounds in which $Y_j \in \{2, 3\}$ will be referred to as test rounds, and the rounds in which $Y_j \in \{0, 1\}$ will be referred to as generation rounds (though strictly speaking, the final key in this protocol is obtained from all the rounds, not merely the generation rounds alone). In each round, Eve is allowed to hold some extension of the state distributed to the devices. We will use E to denote the collection of all such quantum side-information she retains over the entire protocol. (We do not denote this using separate registers E_j for individual rounds, because in a general scenario, Eve’s side-information may not necessarily “factorize” into a tensor product across the individual rounds.)

We briefly highlight some aspects of this protocol that may differ slightly as compared to more commonly used QKD protocols. Firstly, we do not choose a random subset of fixed size as test rounds, but rather, each round is independently chosen to be a test or generation round, following [AFRV19]. This was in order to apply the entropy accumulation theorem, which holds for processes that can be described using a sequence of maps. Furthermore, the parameter-estimation check is performed on *both* $\text{freq}_{\mathbf{c}}(1)$ and $\text{freq}_{\mathbf{c}}(0)$. This was required in order to derive a critical inequality in the security proof (following [BRC20]), though in some cases it is possible to omit the $\text{freq}_{\mathbf{c}}(1)$ check (see Eq. (70) and the subsequent discussion).

We now describe some of the individual steps in more detail.

3.1 Error correction

We first discuss Step 5.2, because it will have an impact on our discussion of Step 5.1. Given any $\varepsilon_h \in (0, 1]$, if we consider a 2-universal family of hash functions where the output is a bitstring of length $\lceil \log(1/\varepsilon_h) \rceil$, then the defining property of 2-universal hashing guarantees that

$$\Pr \left[\text{hash}(\mathbf{A}) = \text{hash}(\tilde{\mathbf{A}}) \mid \mathbf{A} \neq \tilde{\mathbf{A}} \right] \leq \varepsilon_h. \quad (13)$$

In other words, the probability of getting matching hashes from different strings can be made arbitrarily small, by using sufficiently long hashes. Informally speaking, this gives us some laxity in Step 5.1, because regardless of how much the devices deviate from the honest behaviour, the guarantee (13) will still hold, providing a final “check” on how bad the guess $\tilde{\mathbf{A}}$ could be. Importantly, our later proof of the *soundness* of the protocol will not rely on any guarantees regarding the procedure in Step 5.1 — only the *completeness* of the protocol (i.e. the probability that the honest devices mistakenly abort) requires such guarantees.

We now study Step 5.1. EC_{\max} is defined as follows: it is the length of \mathbf{L}_{EC} required such that given the honest devices, Bob can use \mathbf{L}_{EC} and \mathbf{B} to produce a guess $\tilde{\mathbf{A}}$ satisfying

$$\Pr \left[\mathbf{A} \neq \tilde{\mathbf{A}} \right]_{\text{hon}} \leq \varepsilon_{\text{EC}}^{\text{com}}. \quad (14)$$

(In this section, we will use the subscript $_{\text{hon}}$ to emphasize quantities computed with respect to an honest behaviour.) We stress that while some preliminary characterization of the devices can be performed beforehand to choose a suitable EC_{\max} , this parameter *must not* be changed once the protocol begins.

This immediately raises the question of what value should be chosen for EC_{\max} in order to achieve a desired $\varepsilon_{\text{EC}}^{\text{com}}$. Theoretically, there exists a protocol [RR12] with one-way communication that achieves (14) as long as we choose EC_{\max} such that

$$\text{EC}_{\max} \geq H_{\max}^{\tilde{\varepsilon}_s}(\mathbf{A}|\mathbf{BXY})_{\text{hon}} + 2 \log \frac{1}{\varepsilon_{\text{EC}}^{\text{com}} - \tilde{\varepsilon}_s} + 4, \quad (15)$$

where $\tilde{\varepsilon}_s \in [0, \varepsilon_{\text{EC}}^{\text{com}})$ is a parameter that can be optimized over. (This bound is essentially tight for one-way protocols.) Since the honest behaviour is IID, the max-entropy can be bounded by using the asymptotic equipartition property (AEP) in the form stated as Corollary 4.10 of [DFR20], which yields⁵

$$H_{\max}^{\tilde{\varepsilon}_s}(\mathbf{A}|\mathbf{BXY})_{\text{hon}} \leq nh_{\text{hon}} + \sqrt{n} (2 \log 5) \sqrt{\log \frac{2}{\tilde{\varepsilon}_s^2}}, \quad (16)$$

where (using the decomposition $H(Q|Q'C) = \sum_c \Pr[c] H(Q|Q'; C=c)$ for classical C)

$$\begin{aligned} h_{\text{hon}} := H(A_j|B_j X_j Y_j)_{\text{hon}} &= \frac{1-\gamma}{4} \sum_{z \in \{0,1\}} H(A_j|B_j; X_j = Y_j = z)_{\text{hon}} \\ &\quad + \frac{\gamma}{4} \sum_{x \in \{0,1\}, y \in \{2,3\}} H(A_j|B_j; X_j = x, Y_j = y)_{\text{hon}}, \end{aligned} \quad (17)$$

where the terms in the summation over z should be understood to refer to the $A_j B_j$ values *after* the noisy-preprocessing step. (Any value of j can be used in the above equation since the honest behaviour is IID.)

However, the protocol achieving the bound (15) may not be easy to implement. In practice, error-correction protocols typically achieve performance described by

$$\text{EC}_{\max} \approx \xi(n, \varepsilon_{\text{EC}}^{\text{com}}) n h_{\text{hon}}, \quad (18)$$

where $\xi(n, \varepsilon_{\text{EC}}^{\text{com}})$ lies between 1.05 and 1.2 for “typical” values of n and $\varepsilon_{\text{EC}}^{\text{com}}$. (More precise characterizations can be found in [TMMP+17], which gives for instance an estimate

$$\text{EC}_{\max} \approx \xi_1 n h_{\text{hon}} + \tilde{\xi}(\varepsilon_{\text{EC}}^{\text{com}}, h_{\text{hon}}) \sqrt{n}, \quad (19)$$

for a constant ξ_1 and a specific function $\tilde{\xi}$.) Furthermore, some protocols used in practice do not have a theoretical bound on $\varepsilon_{\text{EC}}^{\text{com}}$ (for a given EC_{\max}), only heuristic estimates.

Fortunately, as mentioned earlier, the choice of error-correction procedure in Step 5.1 will have no effect in our proof of the soundness of Protocol 1 (as long as EC_{\max} is a fixed parameter), only its completeness. This means that as long as we are willing to accept heuristic values for ε^{com} , we can use the heuristic values of $\varepsilon_{\text{EC}}^{\text{com}}$ provided by using some “practical” error-correction procedure in Step 5.1, and the value of ε^{sou} (i.e. how “secure” the protocol is) will be completely unaffected. The critical point to remember is that EC_{\max} is a value to be fixed before the protocol begins, and Alice and Bob *must* stop Step 5.1 once they have reached that number of communicated bits. With this in mind, we remark that while we mainly focus on protocols using one-way error correction,

⁵In this analysis, we deviated slightly from [AFRV19] by using the error-correction protocol from [RR12] instead of [RW05], and the AEP stated in [DFR20] rather than [TCR09]. Both of these yield slight improvements in the bounds (the former at large n , and the latter when $\dim(A_j)$ is not too large).

this is not quite a strict requirement — in principle, one could use a procedure involving two-way communication (such as Cascade), as long as EC_{\max} includes all the communicated bits, not just those sent from Alice to Bob. Another possibility worth considering might be adaptive procedures that adjust to the noise level encountered during execution of the protocol, rather than the expected noise level (again, making sure to halt once EC_{\max} bits are communicated, where EC_{\max} is defined beforehand based on the expected behaviour).

We remark that in our situation, the registers **ABXY** have some substructure in the sense that they can be naturally divided into the substrings where $X_j \neq Y_j \in \{0, 1\}$, $X_j = Y_j \in \{0, 1\}$, and $Y_j \in \{2, 3\}$, so the error-correction procedure should ideally take advantage of this substructure (for instance, no error-correction data needs to be sent regarding the rounds where $X_j \neq Y_j \in \{0, 1\}$). Also, if we assume that $\Pr[A_j = B_j | X_j = x, Y_j = y]_{\text{hon}}$ is the same for all $x \in \{0, 1\}, y \in \{2, 3\}$ (in which case it must equal w_{exp}), then for the $y \in \{2, 3\}$ terms in Eq. (17) we have

$$H(A_j | B_j; X_j = x, Y_j = y)_{\text{hon}} = h_2\left(\Pr[A_j \neq B_j | X_j = x, Y_j = y]_{\text{hon}}\right) = h_2(w_{\text{exp}}), \quad (20)$$

which lies in approximately $[0.600, 0.811]$ for $w_{\text{exp}} \in [3/4, (2 + \sqrt{2})/4]$. If the protocol parameters are such that $\xi(n, \varepsilon_{\text{EC}}^{\text{com}})h_2(w_{\text{exp}})$ turns out to be fairly close to 1, there is not much loss incurred by simply sending the outputs of the test rounds directly rather than expending the effort to compute appropriate error-correction data.

3.2 Privacy amplification

Privacy amplification is essentially centred around the *Leftover Hashing Lemma*, which we state below in the form described in [TL17] (obtained via a small modification of the proof in [Ren05]):

Proposition 1. (*Leftover Hashing Lemma*) *Consider any $\sigma \in \mathcal{S}_{\leq}(CQ)$ where C is a classical n -bit register. Let \mathcal{H} be a 2-universal family of hash functions from \mathbb{Z}_2^n to \mathbb{Z}_2^ℓ , and let H be a register of dimension $|\mathcal{H}|$. Define the state*

$$\omega_{KCH} := \mathcal{E}(\sigma_{CQ} \otimes \mathbb{U}_H), \quad (21)$$

where the map \mathcal{E} represents the (classical) process of applying the hash function specified in the register H to the register C , and recording the output in register K . Then for any $\varepsilon \in [0, \sqrt{\text{Tr}[\sigma_{CQ}]}]$, we have

$$\frac{1}{2} \|\omega_{KCH} - \mathbb{U}_K \otimes \omega_{QH}\|_1 \leq 2^{-\frac{1}{2}(H_{\min}^\varepsilon(C|Q)_{\sigma} - \ell + 2)} + 2\varepsilon. \quad (22)$$

Practically speaking, the privacy amplification step simply consists of Alice choosing a random function from the 2-universal family and publicly communicating it to Bob, followed by Alice and Bob applying that function to **A** and **Ā** respectively. The Leftover Hashing Lemma then ensures that the output of this process is close to an ideal key, as long as the conditional min-entropy of the original state was sufficiently large. (Notice that the register H is included in the “side-information” term in Eq. (22), so it can be publicly communicated.)

3.3 Honest behaviour

For this protocol, the honest implementation consists of n IID copies of a device characterized by 2 parameters, w_{exp} and h_{hon} . The first parameter w_{exp} is the probability with which the device wins

the CHSH game when supplied with uniformly random inputs $X_j \in \{0, 1\}, Y_j \in \{2, 3\}$, while the second parameter h_{hon} is defined in Eq. (17). While h_{hon} does not explicitly appear in the protocol description, it is implicitly used to define the parameter EC_{max} , as was described in Sec. 3.1. (Since h_{hon} has a dependence on γ , strictly speaking it may be more precise to instead view the honest device behaviour as being parametrized by a tuple specifying all the individual entropies in Eq. (17), but for brevity we shall summarize this as the honest behaviour being parametrized by h_{hon} . In the more specific models of honest devices described below, these entropies are expressed in terms of some simpler parameters.)

When computing the keyrates shown in Figs. 1 and 6 later for honest devices corresponding to the Bell tests in [HBD+15] (resp. [RBG+17]), we used the following model of the honest devices: following the estimates given in [MDR+19], we characterize them via the parameters $w_{\text{exp}} = 0.797$ (resp. 0.777) and $p_{\text{err}} = 0.06$ (resp. 0.035), where p_{err} is a parameter such that the probabilities *before* noisy preprocessing satisfy⁶

$$\Pr[A_j B_j | X_j = Y_j = z]_{\text{hon}} = (1 - 2p_{\text{err}}) \frac{\delta_{A_j, B_j}}{2} + 2p_{\text{err}} \frac{1}{4} \quad \text{for all } z \in \{0, 1\}, \quad (23)$$

where $\delta_{j,k}$ is the Kronecker delta. Furthermore, we take $\Pr[A_j = B_j | X_j = x, Y_j = y]_{\text{hon}}$ to be the same for all $x \in \{0, 1\}, y \in \{2, 3\}$ (in which case it must equal w_{exp}). Under this model, the expression (17) for h_{hon} can be simplified:

$$h_{\text{hon}} = \frac{1 - \gamma}{2} h_2(p + (1 - 2p)p_{\text{err}}) + \gamma h_2(w_{\text{exp}}), \quad (24)$$

where the $p + (1 - 2p)p_{\text{err}}$ term is obtained by an explicit computation [WAP21].

When computing the keyrates shown in Figs. 2–3 later for a depolarizing-noise scenario, we modelled the honest devices as being described by a parameter $q \in [0, 1/2]$, such that the devices hold the two-qubit Werner state $(1 - 2q) |\Phi^+\rangle\langle\Phi^+| + 2q \mathbb{I}/4$ (where $|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$). The measurements corresponding to Alice and Bob's inputs $X_j \in \{0, 1\}, Y_j \in \{2, 3\}$ are the ideal CHSH measurements (see e.g. [PAB+09]), and the measurements corresponding to Bob's inputs $Y_j \in \{0, 1\}$ are measurements in the same bases as Alice's measurements. In terms of w_{exp} and h_{hon} , this yields

$$w_{\text{exp}} = (1 - 2q) \frac{2 + \sqrt{2}}{4} + 2q \frac{1}{2}, \quad (25)$$

$$h_{\text{hon}} = \frac{1 - \gamma}{2} h_2(p + (1 - 2p)q) + \gamma h_2(w_{\text{exp}}), \quad (26)$$

where the second expression follows by the same computation as above, noting that essentially we have $p_{\text{err}} = q$ in this case.

Finally, for modelling photonic experiments in Fig. 5 later, we follow [Ebe93] and use a highly simplified model with a single parameter $\eta \in [0, 1]$, which is intended to represent the overall detection efficiency (grouping together various effects such as fibre-optic losses and photodetector efficiency into this single parameter). In this model, we take the honest devices to be able to implement arbitrary two-qubit states and measurements perfectly well, but then with probability $1 - \eta$ the outcome is replaced with a no-detection symbol ϕ . In order to apply our security proof,

⁶It can be shown that this is equivalent to taking $\Pr[A_j B_j | X_j = Y_j = z]_{\text{hon}}$ to be independent of z and taking the marginal distributions $\Pr[A_j | X_j = Y_j = z]_{\text{hon}}, \Pr[B_j | X_j = Y_j = z]_{\text{hon}}$ to be uniform, then defining $p_{\text{err}} := \Pr[A_j \neq B_j | X_j = Y_j = z]_{\text{hon}}$.

which requires the test-round measurements to have binary outcomes, we impose that for inputs $X_j \in \{0, 1\}$ by Alice and inputs $Y_j \in \{2, 3\}$ by Bob, the no-detection outcome ϕ is deterministically mapped to the output value 0 (this is a common approach for Bell tests and/or DIQKD using this model [Ebe93, PAB+09]). However, for inputs $Y_j \in \{0, 1\}$ by Bob, we preserve the no-detection outcome, as it slightly improves the keyrates by reducing the error-correction term h_{hon} [ML12]. Unlike the depolarizing-noise model, in this case we do not stick to a fixed choice of states and measurements for all η , but rather we (heuristically) optimize the states and measurements to maximize the keyrate for each η — as is typical in this model, this makes a significant difference in the threshold η value required to achieve e.g. Bell violation [Ebe93] or nonzero keyrates [BFF21].

3.4 Finite-size keyrates

We now present our main theorem, giving the length of the final key as a function of the number of rounds and the desired security parameters. To do so, we need to introduce some notation and ancillary functions. First, we require an function r_p satisfying some properties we now describe. Consider any tripartite quantum state $\rho_{\bar{A}\bar{B}\bar{E}}$, and suppose there are two possible binary-outcome measurements (indexed by x) on register \bar{A} , and similarly two binary-outcome measurements (indexed by y) on register \bar{B} . Let w be the probability of winning the CHSH game (with uniform inputs) for these measurements on the state $\rho_{\bar{A}\bar{B}}$. Let \hat{A}_x be a register that stores the result if measurement x is performed on register \bar{A} and noisy preprocessing (Step 4) with bias p is then applied to the outcome. We require r_p to be an affine function such that for any choice of $\rho_{\bar{A}\bar{B}\bar{E}}$ and measurements, we have

$$\sum_{x \in \{0,1\}} \frac{1}{2} H(\hat{A}_x | \bar{E}) \geq r_p(w). \quad (27)$$

(The factor of $1/2$ arises from the input distributions in the protocol.) We give details on how to obtain such a bound in Sec. 5 (in principle, the $p = 0$ case could also be obtained from the results of [SGP+21]). Given such a function r_p , we then define the affine function

$$g(w) := \frac{1-\gamma}{2} r_p(w) + \gamma r_0(w). \quad (28)$$

Informally, g can be interpreted as a lower bound on the entropy “accumulated” in one round of the protocol.

Finally, for an affine function f defined on all probability distributions on some register C , and any subset \mathcal{S} of its domain, we will define

$$\begin{aligned} \text{Max}(f) &:= \max_q f(q), & \text{Min}_{\mathcal{S}}(f) &:= \inf_{q \in \mathcal{S}} f(q), \\ \text{Var}_{\mathcal{S}}(f) &:= \sup_{q \in \mathcal{S}} \left(\sum_c q(c) f(\delta_c)^2 - \left(\sum_c q(c) f(\delta_c) \right)^2 \right), \end{aligned} \quad (29)$$

where \max_q is taken over all distributions on C , and δ_c denotes the distribution with all its weight on the symbol c .

With these definitions, we can state the security guarantees of the protocol. The following theorem involves various parameters in addition to those listed at the start of Protocol 1, which we first qualitatively describe (note that several descriptions are somewhat informal, not meant to be entirely precise on their own):

- $\varepsilon_{\text{EC}}^{\text{com}}$: Bound on the probability for the honest implementation that Bob's guess $\tilde{\mathbf{A}}$ for \mathbf{A} is wrong
- $\varepsilon_{\text{PE}}^{\text{com}}$: Bound on the probability for the honest implementation that Bob's guess $\tilde{\mathbf{A}}$ for \mathbf{A} is correct but $\text{freq}_{\mathbf{c}}$ does not satisfy the parameter-estimation checks
- ε_{EA} : Informally, a bound on the probability that a "virtual" parameter estimation step (in a related virtual protocol) accepts when given devices that produce insufficient min-entropy
- ε_{PA} : Informally, a bound on the secrecy parameter (12) of the keys after privacy amplification (in a related virtual protocol) when given devices that produce sufficient min-entropy, up to some smoothing corrections
- ε_{h} : Bound on the probability that $\text{hash}(\tilde{\mathbf{A}})$ matches $\text{hash}(\mathbf{A})$ even when Bob's guess $\tilde{\mathbf{A}}$ for \mathbf{A} is wrong
- $\varepsilon_s, \varepsilon'_s, \varepsilon''_s$: Smoothing parameters for several entropic terms
- α, α' : Rényi entropy parameters arising in the EAT
- β : Free parameter in constructing min-tradeoff function (see Eq. (68))

These parameters, together with the protocol parameters $\gamma, p, \delta_{\text{tol}}$ described earlier, can be considered to be variational parameters that should be chosen to optimize the keyrate as much as possible. The formal theorem statement is as follows:

Theorem 1. *Take any $\varepsilon_{\text{EC}}^{\text{com}}, \varepsilon_{\text{PE}}^{\text{com}}, \varepsilon_{\text{EA}}, \varepsilon_{\text{PA}}, \varepsilon_{\text{h}}, \varepsilon_s, \varepsilon'_s, \varepsilon''_s \in (0, 1]$ such that $\varepsilon_s > \varepsilon'_s + 2\varepsilon''_s$, and any $\alpha \in (1, 2)$, $\alpha' \in (1, 1 + 2/V')$, $\beta \in [g(0), g(1)]$, $\gamma \in (0, 1)$, $p \in [0, 1/2]$, where $V' := 2 \log 5$. Protocol 1 is $(\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}})$ -complete and $(\max\{\varepsilon_{\text{EA}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + 2\varepsilon_{\text{h}})$ -sound when performed with any choice of EC_{max} such that Eq. (14) holds, and $\delta_{\text{tol}}, \ell_{\text{key}}$ satisfying*

$$\varepsilon_{\text{PE}}^{\text{com}} \geq B_{n, \gamma w_{\text{exp}}}(\lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor) + B_{n, 1 - \gamma + \gamma w_{\text{exp}}}(\lfloor (1 - \gamma + w_{\text{exp}}\gamma - \delta_{\text{tol}}\gamma)n \rfloor), \quad (30)$$

$$\begin{aligned} \ell_{\text{key}} \leq & ng(w_{\text{exp}} - \delta_{\text{tol}}) - n \frac{(\alpha - 1) \ln 2}{2} V^2 - n(\alpha - 1)^2 K_{\alpha} - n\gamma - n \left(\frac{\alpha' - 1}{4} \right) V'^2 \\ & - \frac{\vartheta_{\varepsilon'_s}}{\alpha - 1} - \frac{\vartheta_{\varepsilon''_s}}{\alpha' - 1} - \left(\frac{\alpha}{\alpha - 1} + \frac{\alpha'}{\alpha' - 1} - 2 \right) \log \frac{1}{\varepsilon_{\text{EA}}} - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s} \\ & - \text{EC}_{\text{max}} - \left\lceil \log \left(\frac{1}{\varepsilon_{\text{h}}} \right) \right\rceil - 2 \log \frac{1}{\varepsilon_{\text{PA}}} + 2, \end{aligned} \quad (31)$$

where $B_{n,p}(k)$ is the cumulative distribution function of a binomial distribution (Definition 2), and

$$\begin{aligned} \vartheta_{\varepsilon} &:= \log \frac{1}{1 - \sqrt{1 - \varepsilon^2}} \leq \log \frac{2}{\varepsilon^2}, \\ V &:= \sqrt{\text{Var}_{\mathcal{Q}_f}(f_{\text{min}})} + 2 + \log 33, \\ K_{\alpha} &:= \frac{2^{(\alpha-1)(2 \log 4 + \text{Max}(f_{\text{min}}) - \text{Min}_{\mathcal{Q}_f}(f_{\text{min}}))}}{6(2 - \alpha)^3 \ln 2} \ln^3 \left(2^{2 \log 4 + \text{Max}(f_{\text{min}}) - \text{Min}_{\mathcal{Q}_f}(f_{\text{min}})} + e^2 \right), \end{aligned} \quad (32)$$

with f_{min} and \mathcal{Q}_f being a function and a set (defined explicitly in Sec. 4.3.1) that satisfy

$$\text{Max}(f_{\text{min}}) = \frac{1}{\gamma} g(1) + \left(1 - \frac{1}{\gamma} \right) \beta, \quad \text{Min}_{\mathcal{Q}_f}(f_{\text{min}}) = g \left(\frac{2 - \sqrt{2}}{4} \right), \quad (33)$$

$$\text{Var}_{\mathcal{Q}_f}(f_{\text{min}}) \leq \frac{2 - \sqrt{2}}{4\gamma} \min \{ \Delta_0^2, \Delta_1^2 \} + \frac{2 + \sqrt{2}}{4\gamma} \max \{ \Delta_0^2, \Delta_1^2 \}, \quad \text{where } \Delta_w := \beta - g(w).$$

Qualitatively, the ϑ_ε terms in the above theorem arise from various properties of smoothed entropies, while the V, V' terms are measures of “variance” of some functions considered in the the EAT, and the K_α term is an additional correction that essentially depends on the range of these functions.

To compute the finite-size keyrates presented in this work, for instance in Figs. 1 and 2 later, we optimized the parameter choices by using the inbuilt (heuristic) constrained-optimization functions in Mathematica or MATLAB to maximize ℓ_{key} while imposing the constraint (30) on the parameters $\varepsilon_{\text{PE}}^{\text{com}}, \gamma, \delta_{\text{tol}}$.⁷ We highlight that the exact expression for ϑ_ε is numerically unstable, and hence we replaced it with the upper bound of $\log(2/\varepsilon^2)$ (this bound is basically tight at small ε , so it makes little difference). Furthermore, the optimization over β also appears to be somewhat unstable. We observed heuristically that the optimal value of β appears to typically be very close to $g(1)$, and hence for simplicity in some cases we did not optimize over it but instead simply fixed $\beta = g(1)$ (or slightly below it, to avoid some instabilities at $\beta - g(1) = 0$). Finally, we found that direct computation of $B_{n,p}(k)$ (e.g. via the regularized beta function) could sometimes be slow or unstable, and in such cases we followed [LLR+21] and replaced it with the upper bound in the following theorem:

Proposition 2. [ZS13] *Let*

$$C_{n,p}(k) := \Phi \left(\text{sign} \left(\frac{k}{n} - p \right) \sqrt{2n D_e \left(\frac{k}{n} \parallel p \right)} \right), \quad (34)$$

where $D_e(q\|p) := q \ln \frac{q}{p} + (1-q) \ln \frac{1-q}{1-p}$ is the base- e relative entropy between the distributions $\{q, 1-q\}$ and $\{p, 1-p\}$, while $\Phi(z) := \int_{-\infty}^z (2\pi)^{-1/2} e^{-u^2/2} du = (1/2) \text{erfc}(-z/\sqrt{2})$ is the cumulative distribution function of the standard normal distribution. Then for any $k \in [n-1]$,

$$C_{n,p}(k) \leq B_{n,p}(k) \leq C_{n,p}(k+1). \quad (35)$$

Replacing $B_{n,p}(k)$ with $C_{n,p}(k+1)$ and computing the latter (which is a Gaussian integral) appeared to be faster and more stable than computing $B_{n,p}(k)$ directly. There is little loss incurred by performing this replacement — the inequalities (35) imply $C_{n,p}(k+1) \leq B_{n,p}(k+1)$, so the effect is no larger than replacing $B_{n,p}(k)$ by $B_{n,p}(k+1)$, which is basically negligible in the parameter regimes studied in this work.

Note that in general, the optimal parameter values (especially for α and α') would depend heavily on n . To get an estimate for the asymptotic scaling of ℓ_{key} , we can choose all the ε parameters to take some constant values satisfying the desired completeness and soundness bounds, then choose [DFR20, DF19]

$$\begin{aligned} \alpha - 1 &= \frac{1}{\sqrt{n}} \sqrt{\frac{2}{V^2 \ln 2} \left(\vartheta_{\varepsilon'_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right)}, \\ \alpha' - 1 &= \frac{1}{\sqrt{n}} \sqrt{\frac{4}{V'^2} \left(\vartheta_{\varepsilon''_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right)}, \end{aligned} \quad (36)$$

⁷To enforce that for instance the parameters $\varepsilon_{\text{EC}}^{\text{com}}, \varepsilon_{\text{PE}}^{\text{com}}$ satisfy the condition $\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}} = \varepsilon^{\text{com}}$ for the desired completeness value ε^{com} , rather than imposing this condition as a constraint in the optimization, we instead introduced a reparametrization $\varepsilon_{\text{EC}}^{\text{com}} = \varepsilon^{\text{com}} \sin^2 \theta$, $\varepsilon_{\text{PE}}^{\text{com}} = \varepsilon^{\text{com}} \cos^2 \theta$ and optimized over $\theta \in (0, \pi/2)$. This ensures that $\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}} = \varepsilon^{\text{com}}$ automatically holds without imposing it as an explicit constraint. A similar approach (with more parameters) was taken for conditions such as $\max\{\varepsilon_{\text{EA}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + 2\varepsilon_h = \varepsilon^{\text{son}}$ and so on.

taking n to be large enough such that $\alpha \in (1, 3/2)$ and $\alpha' \in (1, 1 + 2/V')$. Furthermore, introduce a constant

$$\hat{K} := \frac{2^{(2 \log 4 + \text{Max}(f_{\min}) - \text{Min}_{\mathcal{Q}_f}(f_{\min}))}}{(3/4) \ln 2} \ln^3 \left(2^{2 \log 4 + \text{Max}(f_{\min}) - \text{Min}_{\mathcal{Q}_f}(f_{\min})} + e^2 \right), \quad (37)$$

which satisfies $\hat{K} \geq K_\alpha$ for $\alpha \in (1, 3/2)$. With these choices, Eq. (31) can be satisfied by taking

$$\begin{aligned} \ell_{\text{key}} &= \left[ng(w_{\text{exp}} - \delta_{\text{tol}}) - n \frac{(\alpha - 1) \ln 2}{2} V^2 - n(\alpha - 1)^2 \hat{K} - n\gamma - n \left(\frac{\alpha' - 1}{4} \right) V'^2 \right. \\ &\quad - \frac{1}{\alpha - 1} \left(\vartheta_{\varepsilon'_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right) - \frac{1}{\alpha' - 1} \left(\vartheta_{\varepsilon''_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right) + 2 \log \frac{1}{\varepsilon_{\text{EA}}} - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s} \\ &\quad \left. - \text{EC}_{\text{max}} - \left[\log \left(\frac{1}{\varepsilon_{\text{h}}} \right) \right] - 2 \log \frac{1}{\varepsilon_{\text{PA}}} + 2 \right] \\ &= \left[ng(w_{\text{exp}} - \delta_{\text{tol}}) - \sqrt{n} 2 \sqrt{\frac{V^2 \ln 2}{2}} \left(\vartheta_{\varepsilon'_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right) - \frac{2}{V^2 \ln 2} \left(\vartheta_{\varepsilon'_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right) \hat{K} - n\gamma \right. \\ &\quad - \sqrt{n} 2 \sqrt{\frac{V'^2}{4}} \left(\vartheta_{\varepsilon''_s} + 2 \log \frac{1}{\varepsilon_{\text{EA}}} \right) + 2 \log \frac{1}{\varepsilon_{\text{EA}}} - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s} \\ &\quad \left. - \text{EC}_{\text{max}} - \left[\log \left(\frac{1}{\varepsilon_{\text{h}}} \right) \right] - 2 \log \frac{1}{\varepsilon_{\text{PA}}} + 2 \right] \\ &= \left[ng(w_{\text{exp}} - \delta_{\text{tol}}) - n\gamma - O(\sqrt{n}) - O(1) - \text{EC}_{\text{max}} \right]. \end{aligned} \quad (38)$$

Furthermore, Eq. (30) can be satisfied by choosing $\gamma = 3w_{\text{exp}} n^{-1} \delta_{\text{tol}}^{-2} \log(2/\varepsilon_{\text{PE}}^{\text{com}})$ (see Eq. (44)), in which case by taking $\delta_{\text{tol}} \propto n^{-1/3}$ we have $\delta_{\text{tol}}, \gamma \rightarrow 0$ as $n \rightarrow \infty$, and Eq. (38) then yields

$$\lim_{n \rightarrow \infty} \frac{\ell_{\text{key}}}{n} = \frac{1}{2} \left(r_p(w_{\text{exp}}) - \sum_{z \in \{0,1\}} \frac{1}{2} H(A_j | B_j; X_j = Y_j = z)_{\text{hon}} \right), \quad (39)$$

taking EC_{max} according to Eqs. (15)–(17). This is the expected asymptotic result, e.g. according to the Devetak-Winter bound [DW05] (with the prefactor of $1/2$ being due to the sifting).

Given the scaling behaviour shown in Eq. (38), it can be seen that the optimal values of the various ε parameters (given some desired values of ε^{com} and ε^{sou}) may be of rather different orders of magnitude. This is because some of them appear in $O(1/\sqrt{n})$ corrections to the finite-size keyrate while others appear in $O(1/n)$ corrections. Intuitively speaking, the ε parameters in the latter can be chosen to be substantially smaller than the former, since the $O(1/n)$ scaling reduces their contribution to the finite-size effects.

With Theorem 1, we compute the achievable finite-size keyrates for the various types of honest devices described in Sec. 3.3, showing the results in Figs. 1 and 2. As mentioned previously, the former shows the results regarding the Bell tests in [HBD+15] and [RBG+17], while the latter shows the results for a depolarizing-noise model. We remark that for the computations in Fig. 1, noisy preprocessing was not applied because it appears to only slightly improve the keyrates for those experimental parameters; see Fig. 6 later. For reference, in Fig. 1 we also include plots of the finite-size keyrates against collective attacks as derived in Theorem 4 later — comparing them to the keyrates against general attacks given by Theorem 1, we see that there is indeed some difference, but the threshold value of n to achieve positive keyrates only differs by about one order

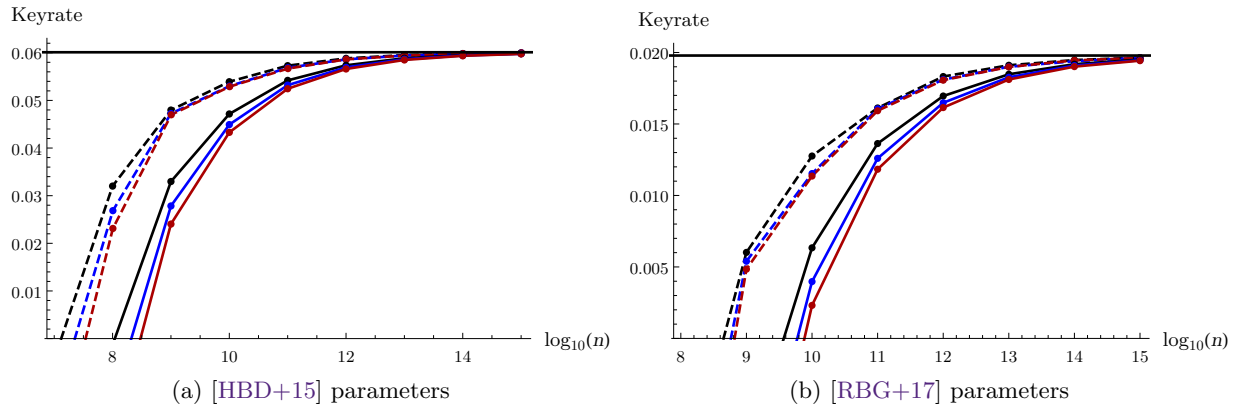


Figure 1: Finite-size keyrates as a function of number of rounds in Protocol 1 without noisy preprocessing ($p = 0$), for honest devices following the estimated parameters in [MDR+19] for the Bell tests in [HBD+15] and [RBG+17]. The solid curves show the results for general attacks (Theorem 1), while the dashed curves show the results under the assumption of collective attacks (Theorem 4), with the error-correction protocol taken to satisfy Eqs. (15)–(17). The colours correspond to soundness parameters (informally, a measure of how “insecure” the key is) of $\varepsilon^{\text{sou}} = 10^{-3}$, 10^{-6} , and 10^{-9} for black, blue, and red respectively, while the completeness parameter (the probability that the honest devices abort) is $\varepsilon^{\text{com}} = 10^{-2}$ in all cases. The horizontal line denotes the asymptotic keyrate. All other parameters in Theorems 1 and 4 were numerically optimized, except β .




Figure 2: Finite-size keyrates secure against general attacks (Theorem 1) as a function of number of rounds in Protocol 1 without noisy preprocessing ($p = 0$), where the honest devices are described by depolarizing noise q (see Sec. 3.3). The solid, dashed and dotted curves denote $q = 5\%$, 6% and 7% respectively, with the error-correction protocol taken to satisfy Eqs. (15)–(17). The soundness parameter is $\varepsilon^{\text{sou}} = 10^{-6}$ and the completeness parameter is $\varepsilon^{\text{com}} = 10^{-2}$. The horizontal lines denote the asymptotic keyrates. All other parameters were numerically optimized.

of magnitude. Furthermore, the various choices we considered for the soundness parameter all yielded fairly similar keyrate curves, indicating that changing the soundness requirements would not significantly change the n threshold.

4 Finite-size security proof

We now prove that Protocol 1 indeed satisfies the security properties claimed in Theorem 1. To do so, we first introduce a virtual protocol that is more convenient to analyze. For the purposes of understanding this construction, it may be helpful to think of it as being based on a *specific* set of states and measurements that could be occurring in a run of Protocol 1 (as opposed to simultaneously considering all possible states and measurements that could be occurring). In particular, this virtual protocol (and the channels \mathcal{M}_j in Sec. 4.3 later) should be understood as being constructed in terms of this specific set of states and measurements. Since we will not impose any additional assumptions on these states and measurements beyond those specified by the protocol, this will still yield a valid way for us to prove the desired security properties (in particular the soundness property, which has to be proven for all possible states and measurements that could occur in a run of the protocol).

Consider the state at the end of Step 6 in Protocol 1. We now describe a virtual protocol⁸ that produces exactly the same state (when it is implemented using the same input state and measurements as those used in a run of Protocol 1), apart from the introduction of two additional registers $\mathbf{B}'\mathbf{C}'$.

Protocol 1' A virtual protocol

1. Alice and Bob's devices each receive and store all quantum states that they will subsequently measure.
 2. For each $j \in [n]$, perform the following steps:
 - 2.1. Alice chooses a uniform input $X_j \in \{0, 1\}$. With probability γ , Bob chooses a uniform input $Y_j \in \{2, 3\}$, otherwise Bob chooses a uniform input $Y_j \in \{0, 1\}$.
 - 2.2. Alice and Bob supply their inputs to their devices, and record the outputs as A_j and B_j respectively.
 - 2.3. Alice and Bob publicly announce their inputs.
 - 2.4. **Sifting:** If $Y_j \in \{0, 1\}$ and $X_j \neq Y_j$, Alice and Bob overwrite their outputs with $A_j = B_j = 0$.
 - 2.5. **Noisy preprocessing:** If $Y_j \in \{0, 1\}$ and $X_j = Y_j$, Alice generates a biased random bit F_j with $\Pr[F_j = 1] = p$, and overwrites her output A_j with $A_j \oplus F_j$.
 - 2.6. If $Y_j \in \{0, 1\}$, Bob sets $B'_j = 0$, otherwise Bob sets $B'_j = B_j$.
 - 2.7. **Virtual parameter estimation:** Set $C'_j = \perp$ if $Y_j \in \{0, 1\}$; otherwise set $C'_j = 0$ if $A_j \oplus B'_j \neq X_j \cdot (Y_j - 2)$ and $C'_j = 1$ if $A_j \oplus B'_j = X_j \cdot (Y_j - 2)$.
 3. **Error correction:** Alice and Bob publicly communicate some bits \mathbf{L} for error correction as previously described, allowing Bob to construct a guess $\tilde{\mathbf{A}}$ for \mathbf{A} .
 4. **Parameter estimation:** For all $j \in [n]$, Bob sets $C_j = \perp$ if $Y_j \in \{0, 1\}$; otherwise he sets $C_j = 0$ if $\tilde{A}_j \oplus B'_j \neq X_j \cdot (Y_j - 2)$ and $C_j = 1$ if $\tilde{A}_j \oplus B'_j = X_j \cdot (Y_j - 2)$.
-

⁸We stress that this “protocol” is not performed in practice (and in fact cannot be, since the “virtual parameter estimation” step cannot be performed locally by either party). However, it produces exactly the same state as Protocol 1 on all relevant registers, and can hence be used for the security analysis.

The key changes as compared to Protocol 1 are as follows:

- All the states that the devices will measure are distributed immediately at the start (note that this is possible because in Protocol 1, the measurement choices \mathbf{X}, \mathbf{Y} are not disclosed until all measurements have been performed, and hence the distributed states cannot behave adaptively with respect to the inputs).
- The sifting and noisy preprocessing steps are now performed immediately after each measurement, instead of after all measurements are performed. This is to allow us to subsequently apply the EAT.
- Two additional registers were introduced: \mathbf{B}' , which is equal to \mathbf{B} on all the test rounds but is otherwise set to 0, and \mathbf{C}' , which is analogous to \mathbf{C} but computed using \mathbf{A} in place of $\tilde{\mathbf{A}}$. These registers were used in a *virtual parameter estimation* step.
- All parameter estimation is performed with \mathbf{B}' instead of \mathbf{B} (this substitution has no physical effect since $B'_j = B_j$ in all rounds used for parameter estimation).

Let ρ denote the state on registers $\mathbf{A}\tilde{\mathbf{A}}\mathbf{B}\mathbf{B}'\mathbf{X}\mathbf{Y}\mathbf{L}\mathbf{C}\mathbf{C}'E$ (as well as the choice of hash function in the error-correction step) at the end of Protocol 1'. As mentioned above, the reduced state after tracing out $\mathbf{B}'\mathbf{C}'$ is exactly the same as that at the end of Step 6 in Protocol 1. Since all subsequent steps in Protocol 1 (i.e. simply the accept/abort check and the privacy amplification) only involve this reduced state, to analyze Protocol 1 it suffices to consider the equivalent (apart from $\mathbf{B}'\mathbf{C}'$) process where all the steps up to Step 6 are replaced by this virtual protocol, and then the remaining steps in Protocol 1 are performed. With this in mind, let us define the following events on the state ρ :

$$\Omega_g: \mathbf{A} = \tilde{\mathbf{A}} \text{ (i.e. Bob correctly guesses } \mathbf{A}\text{)}$$

$$\Omega_h: \text{hash}(\mathbf{A}) = \text{hash}(\tilde{\mathbf{A}})$$

$$\Omega_{\text{PE}}: \text{freq}_{\mathbf{c}}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma \text{ and } \text{freq}_{\mathbf{c}}(0) \leq (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma$$

$$\Omega'_{\text{PE}}: \text{freq}_{\mathbf{c}'}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma \text{ and } \text{freq}_{\mathbf{c}'}(0) \leq (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma$$

Note that in terms of these events, the accept condition of the protocol is $\Omega_h \wedge \Omega_{\text{PE}}$. With the virtual protocol and the above events in mind, we now turn to proving completeness and soundness of Protocol 1.

4.1 Completeness

Completeness is defined entirely with respect to the honest behaviour of the devices, hence all discussion in this section is with respect to the situation where the state ρ described above is the one produced by the honest states and measurements. To prove completeness, we simply need to obtain an upper bound on the probability that this honest behaviour yields an abort, i.e. $\Pr[\Omega_h^c \vee \Omega_{\text{PE}}^c]_{\text{hon}}$ (recall we use Ω^c to denote the complement of an event). However, we encounter a slight inconvenience here because the event Ω_{PE}^c involves the register \mathbf{C} produced using Bob's guess $\tilde{\mathbf{A}}$ rather than Alice's actual string \mathbf{A} , and there is some small probability that his guess was wrong. To cope with this, we shall break down $\Pr[\Omega_h^c \vee \Omega_{\text{PE}}^c]_{\text{hon}}$ into simpler terms that can be bounded in terms

of probabilities involving only the “virtual” string \mathbf{C}' rather than \mathbf{C} , where the former is easier to handle since it is produced from the actual value of \mathbf{A} .

We begin by noting that the hashes of \mathbf{A} and $\tilde{\mathbf{A}}$ can only differ if $\mathbf{A} \neq \tilde{\mathbf{A}}$, which is to say that the event Ω_h^c implies the event Ω_g^c . With this, we write

$$\begin{aligned} \Pr[\Omega_h^c \vee \Omega_{\text{PE}}^c]_{\text{hon}} &\leq \Pr[\Omega_g^c \vee \Omega_{\text{PE}}^c]_{\text{hon}} \\ &= \Pr[\Omega_g^c]_{\text{hon}} + \Pr[\Omega_g \wedge \Omega_{\text{PE}}^c]_{\text{hon}} \end{aligned} \quad (40)$$

where in the second line we have partitioned the event $\Omega_g^c \vee \Omega_{\text{PE}}^c$ into the disjoint events Ω_g^c and $\Omega_g \wedge \Omega_{\text{PE}}^c$. We shall now upper bound the probabilities of each of these events.

The $\Pr[\Omega_g^c]_{\text{hon}}$ term is straightforward to handle, since by construction the error-correction step ensures that this probability is at most $\varepsilon_{\text{EC}}^{\text{com}}$. As for the $\Pr[\Omega_g \wedge \Omega_{\text{PE}}^c]_{\text{hon}}$ term, we now make the critical observation that $\Omega_g \wedge \Omega_{\text{PE}}^c = \Omega_g \wedge \Omega_{\text{PE}}^c$ (because the event Ω_g implies that $\mathbf{C} = \mathbf{C}'$). Therefore, we have

$$\Pr[\Omega_g \wedge \Omega_{\text{PE}}^c]_{\text{hon}} = \Pr[\Omega_g \wedge \Omega_{\text{PE}}^c]_{\text{hon}} \leq \Pr[\Omega_{\text{PE}}^c]_{\text{hon}}, \quad (41)$$

which is the desired reduction to a term involving \mathbf{C}' rather than \mathbf{C} . To explicitly upper bound $\Pr[\Omega_{\text{PE}}^c]_{\text{hon}}$, observe that under the honest behaviour, the string \mathbf{C}' consists of n IID rounds such that $\Pr[C'_j = 1]_{\text{hon}} = \gamma w_{\text{exp}}$ and $\Pr[C'_j = 0]_{\text{hon}} = \gamma(1 - w_{\text{exp}})$ in each round. Therefore, we have

$$\begin{aligned} \Pr[\text{freq}_{\mathbf{C}'}(1) < (w_{\text{exp}} - \delta_{\text{tol}})\gamma]_{\text{hon}} &\leq \Pr[\text{freq}_{\mathbf{C}'}(1)n \leq \lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor]_{\text{hon}} \\ &= B_{n, \gamma w_{\text{exp}}}(\lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor), \end{aligned} \quad (42)$$

$$\begin{aligned} \Pr[\text{freq}_{\mathbf{C}'}(0) > (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma]_{\text{hon}} &= \Pr[\text{freq}_{\mathbf{C}'}(-0) \leq 1 - (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma]_{\text{hon}} \\ &= \Pr[\text{freq}_{\mathbf{C}'}(-0)n \leq \lfloor (1 - \gamma + w_{\text{exp}}\gamma - \delta_{\text{tol}}\gamma)n \rfloor]_{\text{hon}} \\ &= B_{n, 1 - \gamma + \gamma w_{\text{exp}}}(\lfloor (1 - \gamma + w_{\text{exp}}\gamma - \delta_{\text{tol}}\gamma)n \rfloor), \end{aligned} \quad (43)$$

where -0 represents all symbols other than 0. Hence by the union bound, $\Pr[\Omega_{\text{PE}}^c]_{\text{hon}}$ is upper bounded by $\varepsilon_{\text{PE}}^{\text{com}}$ is as specified in Eq. (30) (i.e. the sum of the expressions (42) and (43)). This yields a final upper bound of $\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}}$ on the probability of the honest behaviour aborting, as desired.

In principle, somewhat simpler expressions could be obtained using the (multiplicative) Chernoff bound:

$$\begin{aligned} \Pr[\text{freq}_{\mathbf{C}'}(1) < (w_{\text{exp}} - \delta_{\text{tol}})\gamma]_{\text{hon}} &\leq \left(\frac{e^{-t}}{(1-t)^{1-t}} \right)^{n\gamma w_{\text{exp}}} \leq e^{-\frac{n\gamma\delta_{\text{tol}}^2}{2w_{\text{exp}}}}, \\ \Pr[\text{freq}_{\mathbf{C}'}(0) > (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma]_{\text{hon}} &\leq \left(\frac{e^t}{(1+t)^{1+t}} \right)^{n\gamma w_{\text{exp}}} \leq e^{-\frac{n\gamma\delta_{\text{tol}}^2}{3w_{\text{exp}}}}, \end{aligned} \quad (44)$$

where $t = \delta_{\text{tol}}/w_{\text{exp}} \leq 1$. However, it was observed in [LLR+21] that these bounds are weaker than Prop. 2 by a significant amount. As yet another alternative, Hoeffding’s inequality yields a bound of $e^{-2n\gamma^2\delta_{\text{tol}}^2}$, but this is worse than the Chernoff bound whenever $6\gamma \leq 1/w_{\text{exp}}$; furthermore, it does not yield nontrivial bounds if we choose $\gamma \propto 1/n$ (for constant $\varepsilon_{\text{EC}}^{\text{com}}, \delta_{\text{tol}}$), which ought to be possible in principle according to the scaling analysis in [DF19].

4.2 Soundness

Soundness has to be proven against all possible dishonest behaviours (subject to the protocol assumptions as listed in Sec. 3). In this section, we shall consider any particular state ρ (as defined after the Protocol 1' description) that would be produced by a particular choice of such dishonest behaviour, and prove that it satisfies the soundness condition (10) (for a specific value of ε^{sou}) regardless of which dishonest behaviour was considered. All probabilities are to be understood as being defined with respect to that state ρ .

We first note that it is straightforward to show that Protocol 1 is ε_{h} -correct — recalling that the accept condition can be written as $\Omega_{\text{h}} \wedge \Omega_{\text{PE}}$, we obtain the desired upper bound on the probability that $K_A \neq K_B$ and the protocol accepts:

$$\Pr[K_A \neq K_B \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}] \leq \Pr[K_A \neq K_B \wedge \Omega_{\text{h}}] \leq \Pr[\Omega_{\text{g}}^c \wedge \Omega_{\text{h}}] \leq \Pr[\Omega_{\text{h}} | \Omega_{\text{g}}^c] \leq \varepsilon_{\text{h}}, \quad (45)$$

where the last inequality holds by the defining property of 2-universal hashing.

It remains to prove secrecy. Denote the privacy amplification step in Protocol 1 as the map \mathcal{M}_{PA} , so the subnormalized state conditioned on the event of Protocol 1 accepting can be written as $\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})$. Let H be the register storing the choice of hash function used for privacy amplification, and denote $E' := \mathbf{XYLEH}$ for brevity. We can now rewrite the secrecy condition as the requirement

$$\frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{E'}\|_1 \leq \varepsilon_{\text{QKD}}^{\text{sec}}. \quad (46)$$

Now, somewhat similarly to the completeness analysis, we shall find a way to upper bound the left-hand-side of the above expression in terms of Ω'_{PE} rather than Ω_{PE} , as the former is easier to handle. Specifically, by noting that $\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}} = \rho_{\wedge \Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}} + \rho_{\wedge \Omega_{\text{g}}^c \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}}$, we find the following bound using norm subadditivity:

$$\begin{aligned} & \frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{E'}\|_1 \\ & \leq \frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{E'}\|_1 \\ & \quad + \frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}}^c \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}}^c \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}})_{E'}\|_1 \\ & \leq \frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega'_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega'_{\text{PE}}})_{E'}\|_1 + \varepsilon_{\text{h}}, \end{aligned} \quad (47)$$

where in the last line we have performed the substitution $\Omega_{\text{g}} \wedge \Omega_{\text{PE}} = \Omega_{\text{g}} \wedge \Omega'_{\text{PE}}$ in the first term (again, because the event Ω_{g} implies that $\mathbf{C} = \mathbf{C}'$), and bounded the second term using

$$\Pr[\Omega_{\text{g}}^c \wedge \Omega_{\text{h}} \wedge \Omega_{\text{PE}}] \leq \Pr[\Omega_{\text{g}}^c \wedge \Omega_{\text{h}}] \leq \varepsilon_{\text{h}}, \quad (48)$$

as noted in Eq. (45).

We now aim to bound the first term in Eq. (47). To do so, we study three exhaustive possibilities for the state ρ (the first two are not mutually exclusive, but this does not matter):

Case 1: $\Pr[\Omega_{\text{g}} \wedge \Omega_{\text{h}} | \Omega'_{\text{PE}}] \leq \varepsilon_s^2$.

Case 2: $\Pr[\Omega'_{\text{PE}}] \leq \varepsilon_{\text{EA}}$.

Case 3: Neither of the above are true.

In case 1, that term is bounded by

$$\Pr[\Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}] = \Pr[\Omega_g \wedge \Omega_h | \Omega'_{\text{PE}}] \Pr[\Omega'_{\text{PE}}] \leq \varepsilon_s^2. \quad (49)$$

In case 2, that term is bounded by

$$\Pr[\Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}] \leq \Pr[\Omega'_{\text{PE}}] \leq \varepsilon_{\text{EA}}. \quad (50)$$

The main challenge is case 3. To study this case, we first focus on the conditional state $\rho_{|\Omega'_{\text{PE}}}$. Importantly, the relevant smoothed min-entropy of this state can be bounded by using the following theorem, which we prove in Sec. 4.3 using entropy accumulation:

Theorem 2. *For all parameter values as specified in Theorem 1, the state at the end of Protocol 1' satisfies*

$$\begin{aligned} H_{\min}^{\varepsilon_s}(\mathbf{A} | \mathbf{X} \mathbf{Y} E)_{\rho_{|\Omega'_{\text{PE}}}} &> ng(w_{\text{exp}} - \delta_{\text{tol}}) - n \frac{(\alpha - 1) \ln 2}{2} V^2 - n(\alpha - 1)^2 K_\alpha - n\gamma - n \left(\frac{\alpha' - 1}{4} \right) V'^2 \\ &\quad - \frac{\vartheta_{\varepsilon'_s}}{\alpha - 1} - \frac{\vartheta_{\varepsilon''_s}}{\alpha' - 1} - \left(\frac{\alpha}{\alpha - 1} + \frac{\alpha'}{\alpha' - 1} \right) \log \frac{1}{\Pr[\Omega'_{\text{PE}}]} - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s}. \end{aligned} \quad (51)$$

where $V', \vartheta_\varepsilon, V, K_\alpha$ are as defined in Theorem 1.

To relate this to our state of interest, we use [TL17] Lemma 10, which states that for a state $\sigma \in \mathcal{S}_{\leq}(ZZ'QQ')$ that is classical on registers ZZ' , an event Ω on the registers ZZ' , and any $\varepsilon \in [0, \sqrt{\text{Tr}[\sigma_{\wedge \Omega}]}]$, we have

$$H_{\min}^\varepsilon(ZQ|Q')_{\sigma_{\wedge \Omega}} \geq H_{\min}^\varepsilon(ZQ|Q')_\sigma. \quad (52)$$

In our context, we observe that the probability of the event $\Omega_g \wedge \Omega_h$ on the normalized state $\rho_{|\Omega'_{\text{PE}}}$ is $\Pr[\Omega_g \wedge \Omega_h | \Omega'_{\text{PE}}]$, which is greater than ε_s^2 since we are in case 3. Hence the conditions of the lemma are satisfied (identifying \mathbf{A} with Z , $\tilde{\mathbf{A}}$ (and the error-correction hash choice) with Z' , $\mathbf{X} \mathbf{Y} \mathbf{L} E$ with Q' , and leaving Q empty), allowing us to obtain the bound

$$\begin{aligned} H_{\min}^{\varepsilon_s}(\mathbf{A} | \mathbf{X} \mathbf{Y} \mathbf{L} E)_{(\rho_{|\Omega'_{\text{PE}}})_{\wedge \Omega_g \wedge \Omega_h}} &\geq H_{\min}^{\varepsilon_s}(\mathbf{A} | \mathbf{X} \mathbf{Y} \mathbf{L} E)_{\rho_{|\Omega'_{\text{PE}}}} \\ &\geq H_{\min}^{\varepsilon_s}(\mathbf{A} | \mathbf{X} \mathbf{Y} E)_{\rho_{|\Omega'_{\text{PE}}}} - \text{len}(\mathbf{L}) \\ &\geq H_{\min}^{\varepsilon_s}(\mathbf{A} | \mathbf{X} \mathbf{Y} E)_{\rho_{|\Omega'_{\text{PE}}}} - \text{EC}_{\max} - \left\lceil \log \left(\frac{1}{\varepsilon_h} \right) \right\rceil, \end{aligned} \quad (53)$$

where in the second line we have applied a chain rule for smoothed min-entropy (see e.g. [WTH+11] Lemma 11 or [Tom16] Lemma 6.8).

Putting together Eq. (53) and Theorem 2, we find that for a key of length ℓ_{key} satisfying Eq. (31),

we have⁹

$$\begin{aligned}
& \frac{1}{2} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{(\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h} - \ell_{\text{key}} + 2 \right) + \log \frac{1}{\Pr[\Omega'_{\text{PE}}]} \\
& \geq \frac{1}{2} \left(- \left(\frac{\alpha}{\alpha-1} + \frac{\alpha'}{\alpha'-1} \right) \log \frac{1}{\Pr[\Omega'_{\text{PE}}]} + 2 \log \frac{1}{\varepsilon_{\text{PA}}} + \left(\frac{\alpha}{\alpha-1} + \frac{\alpha'}{\alpha'-1} - 2 \right) \log \frac{1}{\varepsilon_{\text{EA}}} \right) + \log \frac{1}{\Pr[\Omega'_{\text{PE}}]} \\
& = \log \frac{1}{\varepsilon_{\text{PA}}} + \frac{1}{2} \left(\frac{\alpha}{\alpha-1} + \frac{\alpha'}{\alpha'-1} - 2 \right) \left(\log \frac{1}{\varepsilon_{\text{EA}}} - \log \frac{1}{\Pr[\Omega'_{\text{PE}}]} \right) \\
& \geq \log \frac{1}{\varepsilon_{\text{PA}}}, \tag{54}
\end{aligned}$$

where the last line holds because $\Pr[\Omega'_{\text{PE}}] \geq \varepsilon_{\text{EA}}$ in case 3 (and also $\frac{\alpha}{\alpha-1} + \frac{\alpha'}{\alpha'-1} - 2 \geq 2 > 0$ since $\alpha, \alpha' \in (1, 2)$).

This finally allows us to bound Eq. (47), since we have $\varepsilon_s^2 \leq \Pr[\Omega_g \wedge \Omega_h | \Omega'_{\text{PE}}] = \text{Tr}[(\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h]$ in case 3, and hence we can apply the Leftover Hashing Lemma to see that¹⁰

$$\begin{aligned}
& \frac{1}{2} \left\| \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge \Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}})_{E'} \right\|_1 \\
& = \Pr[\Omega'_{\text{PE}}] \frac{1}{2} \left\| \mathcal{M}_{\text{PA}}((\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h)_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}((\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h)_{E'} \right\|_1 \\
& \leq \Pr[\Omega'_{\text{PE}}] \left(\frac{1}{2} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{(\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h} - \ell_{\text{key}} + 2 \right) + 2\varepsilon_s \right) \\
& = \frac{1}{2} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{(\rho_{|\Omega'_{\text{PE}}}) \wedge \Omega_g \wedge \Omega_h} - \ell_{\text{key}} + 2 \right) - \log(1/\Pr[\Omega'_{\text{PE}}]) + 2\varepsilon_s \Pr[\Omega'_{\text{PE}}] \\
& \leq \varepsilon_{\text{PA}} + 2\varepsilon_s. \tag{55}
\end{aligned}$$

Since the three possible cases are exhaustive, we conclude that the secrecy condition is satisfied by choosing

$$\varepsilon_{\text{QKD}}^{\text{sec}} = \max\{\varepsilon_s^2, \varepsilon_{\text{EA}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_h = \max\{\varepsilon_{\text{EA}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_h. \tag{56}$$

Recalling that we have already shown the protocol is ε_h -correct, we finally conclude that it is $(\max\{\varepsilon_{\text{EA}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + 2\varepsilon_h)$ -sound.

⁹Keeping the $\log(1/\Pr[\Omega'_{\text{PE}}])$ term here yields a slightly tighter result as compared to [AFRV19, MDR+19], which instead used $\Pr[\Omega'_{\text{PE}}] \leq 1$ to write an inequality in place of the equality in the second-last line of Eq. (55).

¹⁰We remark that the events Ω'_{PE} and $\Omega_g \wedge \Omega_h$ were treated somewhat differently in this analysis, in that while we introduced a parameter ε_{EA} to divide the analysis of Ω'_{PE} into separate cases, there is no analogous condition for $\Omega_g \wedge \Omega_h$ (the condition for case 1 is not quite analogous; it is merely a technicality to allow us to apply Eq. (52) and the Leftover Hashing Lemma). This is fundamentally because conditioning on Ω'_{PE} via Eq. (51) has a “larger” effect on min-entropy compared to conditioning on $\Omega_g \wedge \Omega_h$ (via Eq. (52) modified for normalized conditional states). When substituting these effects into the final security parameter, one finds that the former worsens the security parameter by an amount that depends on $\log(1/\Pr[\Omega'_{\text{PE}}])$, while the latter does not, and hence we required a bound on $\Pr[\Omega'_{\text{PE}}]$ in the former.

4.3 Entropy accumulation

This section is devoted to the proof of Theorem 2. The key theoretical tool in this proof is the entropy accumulation theorem, which we shall now briefly outline in the form stated in [DF19]. To do so, we shall first introduce *EAT channels* and *tradeoff functions*.

Definition 10. A *sequence of EAT channels* is a sequence $\{\mathcal{M}_j\}_{j \in [n]}$ where each \mathcal{M}_j is a channel from a register R_{j-1} to registers $D_j S_j T_j R_j$, which satisfies the following properties:

- All D_j are classical registers with a common alphabet \mathcal{D} , and all S_j have the same finite dimension.
- For each \mathcal{M}_j , the value of D_j is determined from the registers $S_j T_j$ alone. Formally, this means \mathcal{M}_j is of the form $\mathcal{P}_j \circ \mathcal{M}'_j$, where \mathcal{M}'_j is a channel from R_{j-1} to $S_j T_j R_j$, and \mathcal{P}_j is a channel from $S_j T_j$ to $D_j S_j T_j$ of the form

$$\mathcal{P}_j(\rho_{S_j T_j}) = \sum_{s \in \mathcal{S}, t \in \mathcal{T}} (\Pi_{S_j, s} \otimes \Pi_{T_j, t}) \rho_{S_j T_j} (\Pi_{S_j, s} \otimes \Pi_{T_j, t}) \otimes |d(s, t)\rangle \langle d(s, t)|_{D_j}, \quad (57)$$

where $\{\Pi_{S_j, s}\}_{s \in \mathcal{S}}$ and $\{\Pi_{T_j, t}\}_{t \in \mathcal{T}}$ are families of orthogonal projectors on S_j and T_j respectively, and $d : \mathcal{S} \times \mathcal{T} \rightarrow \mathcal{D}$ is a deterministic function.

Definition 11. Let f_{\min} be a real-valued affine function defined on probability distributions over the alphabet \mathcal{D} . It is called a *min-tradeoff function* for a sequence of EAT channels $\{\mathcal{M}_j\}_{j \in [n]}$ if for any distribution q on \mathcal{D} , we have

$$f_{\min}(q) \leq \inf_{\sigma \in \Sigma_j(q)} H(S_j | T_j R)_\sigma \quad \forall j \in [n], \quad (58)$$

where $\Sigma_j(q)$ denotes the set of states of the form $(\mathcal{M}_j \otimes \mathcal{I}_R)(\omega_{R_{j-1} R})$ such that the reduced state on D_j has distribution q . Analogously, a real-valued affine function f_{\max} defined on probability distributions over \mathcal{D} is called a *max-tradeoff function* if

$$f_{\max}(q) \geq \sup_{\sigma \in \Sigma_j(q)} H(S_j | T_j R)_\sigma \quad \forall j \in [n]. \quad (59)$$

The infimum and supremum of an empty set are defined as $+\infty$ and $-\infty$ respectively.

With these definitions, we can now state the theorem:

Proposition 3. (*Entropy accumulation theorem [DF19, LLR+21]*) Consider a sequence of EAT channels $\{\mathcal{M}_j\}_{j \in [n]}$ and a state of the form $\rho = (\mathcal{M}_n \circ \dots \circ \mathcal{M}_1 \otimes \mathcal{I}_E)(\rho_{R_0 E}^0)$ satisfying the Markov conditions

$$I(S_{[j-1]} : T_j | T_{[j-1]} E)_\rho = 0 \quad \forall j \in [n]. \quad (60)$$

Let f_{\min} be a min-tradeoff function for $\{\mathcal{M}_j\}_{j \in [n]}$ and consider any $h \in \mathbb{R}, \varepsilon \in (0, 1), \alpha \in (1, 2)$.

Then for any event $\Omega \subseteq \mathcal{D}^n$ such that $f_{\min}(\text{freq}_{\mathbf{d}}) \geq h$ for all $\mathbf{d} \in \Omega$, we have¹¹

$$H_{\min}^{\varepsilon}(\mathbf{S}|\mathbf{T}E)_{\rho|\Omega} > nh - n \frac{(\alpha - 1) \ln 2}{2} V^2 - n(\alpha - 1)^2 K_{\alpha} - \frac{\vartheta_{\varepsilon}}{\alpha - 1} - \frac{\alpha}{\alpha - 1} \log \frac{1}{\Pr[\Omega]}, \quad (61)$$

where ϑ_{ε} is as defined in Eq. (32), and

$$\begin{aligned} V &:= \sqrt{\text{Var}_{\mathcal{Q}}(f_{\min}) + 2} + \log(2 \dim(S_j)^2 + 1), \\ K_{\alpha} &:= \frac{2^{(\alpha-1)(2 \log \dim(S_j) + \text{Max}(f_{\min}) - \text{Min}_{\mathcal{Q}}(f_{\min}))}}{6(2 - \alpha)^3 \ln 2} \ln^3 \left(2^{2 \log \dim(S_j) + \text{Max}(f_{\min}) - \text{Min}_{\mathcal{Q}}(f_{\min})} + e^2 \right), \end{aligned} \quad (62)$$

with \mathcal{Q} being the set of all distributions on \mathcal{D} that could be produced by applying some EAT channel to some state.

Informally, the Markov conditions impose the requirement that the register T_j does not “leak any information” about the previous registers $S_{[j-1]}$ beyond what is already available from $T_{[j-1]}E$. (Without this Markov condition, one could for instance have channels such that T_j simply contains a copy of S_{j-1} , in which case there could be situations where a nontrivial min-tradeoff function holds but the conclusion of the entropy accumulation theorem would be completely false.) There is also an analogous EAT statement regarding the max-entropy [DFR20], using a max-tradeoff function, though we will be using it slightly differently and will elaborate further on it at that point.

We now describe how the EAT can be used to prove Theorem 2. First, note that to prove Theorem 2 it would be sufficient to consider only the registers $\mathbf{AB}'\mathbf{X}\mathbf{Y}\mathbf{C}'E$ of the state ρ (the conditioning event Ω'_{PE} is determined by \mathbf{C}' alone). The reduced state on these registers is the same as that at the point when Step 2 of Protocol 1' has finished looping over j , since the subsequent steps do not change these registers, and thus we can equivalently study that state in place of $\rho|\Omega'_{\text{PE}}$. From this point onwards, all smoothed min- or max-entropies refer to that state conditioned on Ω'_{PE} (and normalized), hence for brevity we will omit the subscript specifying the state.

Each iteration of Step 2 of Protocol 1' can be treated as a channel \mathcal{M}_j in a sequence of EAT channels, by considering it to be a channel performing the following operations:

1. Alice generates X_j as specified in Step 2. Conditioned on the value of X_j , Alice’s device performs some measurement on its share of the stored quantum state R_{j-1} (which includes any memory retained from previous rounds), then performs sifting and noisy preprocessing on the outcome, storing the final result in register A_j .
2. Bob’s device behaves analogously, producing the registers Y_j and B'_j (we will not need to consider B_j).
3. The value of C'_j is computed from $A_j B'_j X_j Y_j$.

¹¹This expression differs slightly from those in [DF19, LLR+21] because we have not performed the simplifications based on $\vartheta_{\varepsilon} \leq \log(2/\varepsilon^2)$ and $\alpha < 2$ (though the former is only a very small improvement for typical values of ε). We also remark that in [LLR+21], a modification was made to the EAT to improve its dependence on the Var term, which we have not included here as it involves a further optimization that would significantly increase the complexity of our keyrate computations. (In an earlier version of this work, we stated that this modification did not make a difference here as our choice of affine f_{\min} is essentially equal to the tight bound on $H(S_j|T_jR)$ in our scenario; however, this was in error as this equality does not hold on e.g. any parts of the domain where f_{\min} is negative.) Still, we would expect that such an optimization would not yield better keyrates than the collective-attacks scenario (shown in Fig. 1) at least, so the potential improvement is limited to some extent.

We highlight that in the above description of \mathcal{M}_j , the only “unknowns” are the measurements it performs on the input state on R_{j-1} — all other operations are taken to be performed in trusted fashion. (This is reasonable because these measurements and the stored states are the only untrusted aspects in the true protocol.) If we had simply considered completely arbitrary channels \mathcal{M}_j producing the respective registers, it would not be possible to make a nontrivial security statement about the output.

Identifying C'_j with D_j , $A_j B'_j$ with S_j , and $X_j Y_j$ with T_j in Definition 10, we see that these channels \mathcal{M}_j indeed form a valid sequence of EAT channels: C'_j is determined from $A_j B'_j X_j Y_j$ in the manner specified by Eq. (57). Additionally, the state they produce always fulfills the Markov conditions, because the values of $X_j Y_j$ in each round are generated independently of all preceding registers.

Intuitively, it seems that we could now use the EAT to bound $H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{X}\mathbf{Y}E)$. However, there is a technical issue: to apply the EAT, the event Ω'_{PE} must be defined entirely in terms of the (classical) registers that appear in the smoothed min-entropy term that we are bounding, which is not a condition satisfied by the registers $\mathbf{A}\mathbf{X}\mathbf{Y}$ alone. This is where the register \mathbf{B}' comes into play, following the same approach as [AFRV19]: by a chain rule for the min- and max-entropies ([VDT+13] or [Tom16] Eq. (6.57)), we have for any $\varepsilon'_s + 2\varepsilon''_s < \varepsilon_s$:

$$\begin{aligned} H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{X}\mathbf{Y}E) &\geq H_{\min}^{\varepsilon'_s}(\mathbf{A}\mathbf{B}'|\mathbf{X}\mathbf{Y}E) - H_{\max}^{\varepsilon''_s}(\mathbf{B}'|\mathbf{A}\mathbf{X}\mathbf{Y}E) - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s} \\ &\geq H_{\min}^{\varepsilon'_s}(\mathbf{A}\mathbf{B}'|\mathbf{X}\mathbf{Y}E) - H_{\max}^{\varepsilon''_s}(\mathbf{B}'|\mathbf{X}\mathbf{Y}E) - 3\vartheta_{\varepsilon_s - \varepsilon'_s - 2\varepsilon''_s}. \end{aligned} \quad (63)$$

The $H_{\max}^{\varepsilon''_s}(\mathbf{B}'|\mathbf{X}\mathbf{Y}E)$ term admits a fairly simple bound, as follows: consider a sequence of EAT channels $\widetilde{\mathcal{M}}_j$ that are identical to \mathcal{M}_j except that they do not produce the registers $A_j C'_j$. As before, these maps obey the required Markov conditions. In addition, recall that for every round the register B'_j is deterministically set to 0 whenever $y \in \{0, 1\}$ (which happens with probability $1 - \gamma$), hence we always have

$$H(B'_j|X_j Y_j R)_{(\widetilde{\mathcal{M}}_j \otimes \mathcal{I}_R)(\omega_{R_{j-1}R})} = \sum_y \Pr[Y_j = y] H(B'_j|X_j R; Y_j = y)_{(\widetilde{\mathcal{M}}_j \otimes \mathcal{I}_R)(\omega_{R_{j-1}R})} \leq \gamma. \quad (64)$$

This means we can apply the max-entropy version of the EAT¹² with a *constant* max-tradeoff function of value γ . Letting $V' = 2 \log(1 + 2 \dim(B')) = 2 \log 5$, this yields the following bound for any $\alpha' \in (1, 1 + 2/V')$:

$$H_{\max}^{\varepsilon''_s}(\mathbf{B}'|\mathbf{X}\mathbf{Y}E) < n\gamma + n \left(\frac{\alpha' - 1}{4} \right) V'^2 + \frac{\vartheta_{\varepsilon''_s}}{\alpha' - 1} + \frac{\alpha'}{\alpha' - 1} \log \frac{1}{\Pr[\Omega'_{\text{PE}}]}. \quad (65)$$

The bulk of our task is to bound the $H_{\min}^{\varepsilon'_s}(\mathbf{A}\mathbf{B}'|\mathbf{X}\mathbf{Y}E)$ term. To do so, we will need an appropriate min-tradeoff function, which we shall now construct.

¹²Here we shall use the results from [DFR20], because for a constant tradeoff function, this turns out to yield a slightly better bound as compared to the version of the EAT [DF19] stated here. Strictly speaking, the reasoning used here is not a direct application of the EAT, because once again, the event Ω'_{PE} is not defined on the registers $\mathbf{B}'\mathbf{X}\mathbf{Y}$ alone (attempting to address this by including \mathbf{A} in the conditioning registers could result in the Markov conditions not being fulfilled). Fortunately, the bound (64) holds for our maps $\widetilde{\mathcal{M}}_j$ even without a constraint on the output distribution. Hence the reasoning is as follows, in terms of the equations and lemmas in [DFR20]: first apply Eq. (32) *without* the event-conditioning term (this is valid since (64) holds without constraints), then condition on Ω'_{PE} using Lemma B.6 (noting that $\rho_{\mathbf{B}'\mathbf{X}\mathbf{Y}E} = \Pr[\Omega'_{\text{PE}}] (\rho_{|\Omega'_{\text{PE}}})_{\mathbf{B}'\mathbf{X}\mathbf{Y}E} + \Pr[\Omega'^c_{\text{PE}}] (\rho_{|\Omega'^c_{\text{PE}}})_{\mathbf{B}'\mathbf{X}\mathbf{Y}E}$), and finally apply Lemma B.10 to obtain Eq. (65). (Alternatively, one could use \mathbf{C}' instead of \mathbf{B}' . This was done in [MDR+19] to slightly improve the bound in the block analysis, but it does not make a difference in our analysis. However, using \mathbf{C}' would seem to make it harder to sharpen the slightly crude bound used to obtain Eq. (67).)

4.3.1 Min-tradeoff function

Consider an arbitrary state of the form $(\mathcal{M}_j \otimes \mathcal{I}_R)(\omega_{R_{j-1}R})$. In this section, all entropies will be computed with respect to this state, and hence for brevity we will omit the subscript specifying the state.

We first note that

$$\begin{aligned} H(A_j B'_j | X_j Y_j R) &= \frac{1-\gamma}{4} \sum_{z \in \{0,1\}} H(A_j | R; X_j = Y_j = z) \\ &\quad + \frac{\gamma}{4} \sum_{y \in \{2,3\}} \sum_{x \in \{0,1\}} H(A_j B'_j | R; X_j = x, Y_j = y), \end{aligned} \quad (66)$$

where we have used the fact that $H(A_j B'_j | R; X_j = x, Y_j = y) = 0$ when $(x, y) = (0, 1)$ or $(1, 0)$, and $H(A_j B'_j | R; X_j = x, Y_j = y) = H(A_j | R; X_j = x, Y_j = y)$ when $(x, y) = (0, 0)$ or $(1, 1)$.

Let w denote the probability that the state wins the CHSH game, conditioned on the game being played. Then by applying the simple but somewhat crude bound¹³ $H(A_j B'_j | R; X_j = x, Y_j = y) \geq H(A_j | R; X_j = x, Y_j = y)$ to the terms in the second sum in Eq. (66), we get the bound¹⁴

$$H(A_j B'_j | X_j Y_j R) \geq \frac{1-\gamma}{2} r_p(w) + \frac{\gamma}{2} \sum_{y \in \{2,3\}} r_0(w) = g(w). \quad (67)$$

We can now use the function g to construct a min-tradeoff function f_{\min} , with the domain of f_{\min} being distributions on C'_j (recall that this register is set to \perp if $Y_j \in \{0, 1\}$, and otherwise is set to 0 or 1 if the CHSH game is lost or won respectively). First observe that the channel is an *infrequent-sampling channel* in the sense described in [DF19, LLR+21]. By the argument in Appendix A.7 of [LLR+21], a valid min-tradeoff function f_{\min} for the channel is given by the (unique) affine function specified by the following values (in a minor abuse of notation, here we interpret g as a function of a distribution instead of a winning probability):

$$f_{\min}(\delta_c) = \begin{cases} \frac{1}{\gamma} g(\delta_c) + \left(1 - \frac{1}{\gamma}\right) \beta & \text{if } c \neq \perp \\ \beta & \text{if } c = \perp \end{cases}, \quad (68)$$

where $\beta \in [g(\delta_0), g(\delta_1)]$ is a constant that can be chosen to optimize the keyrate. (Intuitively, this function is constructed simply by noting that the maps \mathcal{M}_j can only produce distributions that lie in the slice of the probability simplex specified by the constraint $\Pr[C'_j = \perp] = 1 - \gamma$, and hence the min-tradeoff function is free to take any value for distributions outside of this slice, recalling that we take the infimum of an empty set to be $+\infty$. For distributions within this slice, we know that g is an affine lower bound on the entropy as a function of the winning probability, and hence we can just set f_{\min} equal to g (up to a domain rescaling) on this slice. Any f_{\min} constructed this way is precisely of the form described in Eq. (68), with β being a constant determining its value on all distributions outside of the $\Pr[C'_j = \perp] = 1 - \gamma$ slice.)

¹³This marks a point where the analysis could be slightly sharpened, in that if we had a tight bound on the “two-party entropies” $H(A_j B'_j | R; X_j = x, Y_j = y)$ rather than just the “one-party entropies” $H(A_j | R; X_j = x, Y_j = y)$, we could improve the second term in (67). However, note that it would only improve the keyrate by $O(\gamma)$, because of the γ prefactor on that term.

¹⁴Recall that noisy preprocessing is not applied to the rounds with $Y_j \in \{2, 3\}$.

As shown in [LLR+21], the min-tradeoff function constructed this way satisfies

$$\begin{aligned} \text{Max}(f_{\min}) &= \max \left\{ \frac{1}{\gamma} \text{Max}(g) + \left(1 - \frac{1}{\gamma}\right) \beta, \beta \right\}, & \text{Min}_{\mathcal{Q}_f}(f_{\min}) &= \text{Min}_{\mathcal{Q}_g}(g), \\ \text{Var}_{\mathcal{Q}_f}(f_{\min}) &\leq \sup_{q \in \mathcal{Q}_g} \sum_{c \in \{0,1\}} \frac{q(c)}{\gamma} (\beta - g(\delta_c))^2, \end{aligned} \quad (69)$$

where \mathcal{Q}_f denotes the set of distributions on C'_j such that $\Pr[C'_j = \perp] = 1 - \gamma$ and $\Pr[C'_j = 1] \in [\gamma(2 - \sqrt{2})/4, \gamma(2 + \sqrt{2})/4]$, while \mathcal{Q}_g denotes the set of all distributions on the alphabet $\{0, 1\}$ such that $\Pr[1] \in [(2 - \sqrt{2})/4, (2 + \sqrt{2})/4]$.

For the specific g and range of β that we consider, these expressions simplify to Eq. (33), where we have solved the optimization $\sup_{q \in \mathcal{Q}_g}$ in the bound on $\text{Var}_{\mathcal{Q}_f}(f_{\min})$ by observing that it is an affine function of the distribution q , and the set \mathcal{Q}_g we use here is essentially a line segment (in a 1-dimensional probability simplex).

4.3.2 Final min-entropy bound

The event Ω'_{PE} is defined by the conditions $\text{freq}_{\mathbf{c}'}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma$ and $\text{freq}_{\mathbf{c}'}(0) \leq (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma$. Hence for all $\mathbf{c}' \in \Omega'_{\text{PE}}$, we have (since f_{\min} is affine)

$$\begin{aligned} f_{\min}(\text{freq}_{\mathbf{c}'}) &= \text{freq}_{\mathbf{c}'}(0) \left(\frac{1}{\gamma} g(\delta_0) + \left(1 - \frac{1}{\gamma}\right) \beta \right) + \text{freq}_{\mathbf{c}'}(1) \left(\frac{1}{\gamma} g(\delta_1) + \left(1 - \frac{1}{\gamma}\right) \beta \right) + \text{freq}_{\mathbf{c}'}(\perp) \beta \\ &= \text{freq}_{\mathbf{c}'}(0) \left(\frac{1}{\gamma} g(\delta_0) - \frac{1}{\gamma} \beta \right) + \text{freq}_{\mathbf{c}'}(1) \left(\frac{1}{\gamma} g(\delta_1) - \frac{1}{\gamma} \beta \right) + \beta \\ &\geq (1 - w_{\text{exp}} + \delta_{\text{tol}}) (g(\delta_0) - \beta) + (w_{\text{exp}} - \delta_{\text{tol}}) (g(\delta_1) - \beta) + \beta \\ &= (1 - w_{\text{exp}} + \delta_{\text{tol}}) g(\delta_0) + (w_{\text{exp}} - \delta_{\text{tol}}) g(\delta_1) \\ &= g(w_{\text{exp}} - \delta_{\text{tol}}), \end{aligned} \quad (70)$$

where the inequality holds because $\beta \in [g(\delta_0), g(\delta_1)]$, and in the last line we use the fact that g is affine and revert to interpreting it as a function of winning probability. (We remark that if we fix $\beta = g(1)$, then in fact this inequality can be derived using only the $\text{freq}_{\mathbf{c}'}(0)$ condition, following [DF19]. Hence in principle one could sacrifice the option of optimizing β in exchange for reducing the number of checks to perform in the protocol, which improves the completeness parameters.)

Therefore, we can choose $h = g(w_{\text{exp}} - \delta_{\text{tol}})$ in the EAT statement (Prop. 3) to conclude that the state conditioned on Ω'_{PE} satisfies

$$\begin{aligned} H_{\min}^{\varepsilon'_s}(\mathbf{AB}' | \mathbf{XY}E) &> ng(w_{\text{exp}} - \delta_{\text{tol}}) - n \frac{(\alpha - 1) \ln 2}{2} V^2 - n(\alpha - 1)^2 K_\alpha \\ &\quad - \frac{\vartheta_{\varepsilon'_s}}{\alpha - 1} - \frac{\alpha}{\alpha - 1} \log \frac{1}{\Pr[\Omega'_{\text{PE}}]}, \end{aligned} \quad (71)$$

where $\vartheta_\varepsilon, V, K_\alpha$ are as defined in Eq. (32). Putting this together with Eqs. (63) and (65), we finally obtain the bound in Theorem 2.

4.4 Scope of applicability

In the above analysis, we have focused on a protocol that only uses the CHSH game. However, it would be possible to modify the analysis to account for arbitrary Bell inequalities, as was done in [BRC20]. Essentially, Alice and Bob would simply need to choose a different distribution of their input settings, corresponding to a different game being played. Furthermore, it is not strictly necessary for Alice to choose uniformly random inputs in the generation rounds — as noted in [SGP+21], she could instead choose some biased distribution of inputs. It would even be possible to consider applying different amounts of noisy preprocessing for the different inputs in generation rounds. All of these modifications would essentially correspond to finding an appropriate min-tradeoff function, which we describe how to do in the subsequent section. However, we show in Sec. 5.6 that for the depolarizing-noise model at least, there is little to be gained by considering these modifications.

There is a technical issue that in order to implement the above modifications, Alice and Bob may need to know which rounds are test rounds and which rounds are generation rounds, if they need to choose different input distributions in the two cases. (In Protocol 1, this was not an issue because the CHSH game is played with uniformly random inputs, and we also used uniformly random inputs for the generation rounds, so there is no difference in the input distributions between test and generation rounds.) However, this could in principle be addressed by using a short pre-shared key in order to choose the test rounds — we describe this in more detail in Sec. 6.1.

5 Single-round bound

We now explain how to derive the function r_p as described in Eq. (27). To reduce clutter, we will use slightly different notation in this section, which should not be confused with the earlier notation. In particular, we will omit the $\bar{}$ accents from the registers $\bar{A}\bar{B}\bar{E}$, and we will be using hermitian operators A_x, B_y that should not be confused with the registers A_j, B_j in the earlier notation.

As described previously, consider an arbitrary state ρ_{ABE} and possible measurements on the A and B subsystems, indexed by x and y respectively. For the purpose of bounding the entropy, all measurements can be assumed projective by considering a suitably chosen simultaneous Stinespring dilation; see e.g. [TSG+21]. Let $P_{a|x}$ for Alice (resp. $Q_{b|y}$ for Bob) denote the projector corresponding to outcome a (resp. b) from measurement x (resp. y). We first consider a somewhat more general setting than that required for the above security proof; namely, we aim to find a lower bound on the conditional entropies, knowing only that the state and measurements produce the values ν_j for some observables $\Gamma_j(P_{a|x}, Q_{b|y}) := \sum_{abxy} c_{abxy}^{(j)} P_{a|x} \otimes Q_{b|y}$ defined by fixed coefficients $c_{abxy}^{(j)} \in \mathbb{R}$. (For instance, these could be the values of some Bell expressions, or even simply the entire set of output probabilities.) Furthermore, we allow for the possibility that the inputs are not uniformly random. Making this precise, we aim to find the function

$$\check{r}_p(\vec{\nu}) := \inf_{\rho_{ABE}, P_{a|x}, Q_{b|y}} \sum_{x \in \{0,1\}} \tau_x H(\hat{A}_x | E), \quad (72)$$

s.t. $\text{Tr}[\Gamma_j(P_{a|x}, Q_{b|y})\rho_{AB}] = \nu_j \quad \forall j$

where $\tau_x \geq 0$ are coefficients that depend on the input distributions [SGP+21]. (For the security proof of Protocol 1, we would only need to consider a single $\Gamma_j(P_{a|x}, Q_{b|y})$, which describes the

probability of winning the CHSH game. Also, we simply have $\tau_0 = \tau_1 = 1/2$, since we want a bound of the form (27).) Without loss of generality, we can restrict the optimization to pure ρ_{ABE} .

By considering the effect of Eve using mixtures of strategies, it is easily seen that \check{r}_p must be convex. It can thus be shown that for any \vec{v} in the interior of the set of values achievable by quantum theory, this optimization is in fact equal to its Lagrange dual (see e.g. [TSG+21] for a detailed explanation):

$$\check{r}_p(\vec{v}) = \sup_{\vec{\lambda}} \inf_{\rho_{ABE}, P_{a|x}, Q_{b|y}} \left(\sum_{x \in \{0,1\}} \tau_x H(\hat{A}_x|E) \right) - \vec{\lambda} \cdot \left(\text{Tr} \left[\vec{\Gamma}(P_{a|x}, Q_{b|y}) \rho_{AB} \right] - \vec{v} \right). \quad (73)$$

Since the optimization over $\vec{\lambda}$ is a supremum, it follows that for any value of $\vec{\lambda}$, we have a lower bound on $\check{r}_p(\vec{v})$ of the form

$$\check{r}_p(\vec{v}) \geq \vec{\lambda} \cdot \vec{v} + c_{\vec{\lambda}}, \quad (74)$$

where

$$c_{\vec{\lambda}} := \inf_{\rho_{ABE}, P_{a|x}, Q_{b|y}} \left(\sum_{x \in \{0,1\}} \tau_x H(\hat{A}_x|E) \right) - \vec{\lambda} \cdot \text{Tr} \left[\vec{\Gamma}(P_{a|x}, Q_{b|y}) \rho_{AB} \right]. \quad (75)$$

Importantly, such a lower bound is automatically affine with respect to \vec{v} , and hence we can take it as a possible choice of r_p for the security proof. We thus see that in order to find an affine r_p for use in the security proof, it suffices to choose some $\vec{\lambda}$ and compute (or lower-bound) the value of $c_{\vec{\lambda}}$ as defined by Eq. (75). In addition, this approach yields essentially tight bounds for the asymptotic rates¹⁵, in the sense that taking the supremum of the bounds given by all possible $\vec{\lambda}$ returns the value of $\check{r}_p(\vec{v})$, by Eq. (73). We also note that when there are multiple constraints, for any given \vec{v} the corresponding optimal $\vec{\lambda}$ yields a single Bell expression that certifies the same bound on the entropy as all the original constraints. When applying this bound in the EAT, the task of choosing $\vec{\lambda}$ here is equivalent to the task in [AFRV19, LLR+21] of choosing a tangent point to the rate curves.

The $H(\hat{A}_x|E)$ term in the optimization can be rewritten based on the approach in [WLC18]. Specifically, we observe that the state produced on $\hat{A}_x E$ by performing Alice's measurement x on ρ_{ABE} and then applying noisy preprocessing could also be obtained by the following process: append an ancilla T in the state

$$|\phi_p\rangle_T := \sqrt{1-p}|0\rangle_T + \sqrt{p}|1\rangle_T, \quad (76)$$

apply a pinching channel $\mathcal{P}(\sigma_T) := \sum_t |t\rangle\langle t|_T \sigma_T |t\rangle\langle t|_T$ to T , then perform a measurement on AT described by the projectors

$$\tilde{P}_{a|x} := P_{a|x} \otimes |0\rangle\langle 0|_T + P_{a \oplus 1|x} \otimes |1\rangle\langle 1|_T, \quad (77)$$

¹⁵There is a technicality here that was overlooked in earlier versions of this work. Specifically, even though \check{r}_p is basically affine over a large range of values (see Sec. 5.5), the optimal choice of r_p for the asymptotic rates still may not yield the optimal finite-size keyrates at each n when applying the EAT — this is because choosing a bound r_p that is suboptimal with respect to the asymptotic rates can yield a min-tradeoff function with smaller variance and range, reducing the finite-size corrections. However, given the computationally intensive nature of the algorithm we describe here, we did not attempt to optimize the choice of r_p as a function of n , instead sticking to the optimal choice for the asymptotic limit. Still, as mentioned previously in Sec. 4.3, we would expect that any potential improvements by optimizing r_p would at best bring the keyrates somewhat closer to the collective-attacks values shown in Fig. 1.

and store the outcome in \hat{A}_x without further processing. However, the pinching channel on T can in fact be omitted, because the subnormalized conditional states produced on E would still be the same without it (in the following, we leave some tensor factors of \mathbb{I} implicit for brevity, and consider an arbitrary state σ_{ABET} before applying the pinching channel):

$$\begin{aligned}
& \text{Tr}_{ABT} \left[\tilde{P}_{a|x} \mathcal{P}_T(\sigma_{ABET}) \tilde{P}_{a|x} \right] \\
&= \text{Tr}_{ABT} \left[\tilde{P}_{a|x} \left(\sum_t |t\rangle\langle t|_T \sigma_{ABET} |t\rangle\langle t|_T \right) \tilde{P}_{a|x} \right] \\
&= \sum_t \text{Tr}_{ABT} \left[(P_{a\oplus t|x} \otimes |t\rangle\langle t|_T) \sigma_{ABET} (P_{a\oplus t|x} \otimes |t\rangle\langle t|_T) \right] \\
&= \sum_{t,t'} \text{Tr}_{ABT} \left[(P_{a\oplus t|x} \otimes |t\rangle\langle t|_T) \sigma_{ABET} (P_{a\oplus t'|x} \otimes |t'\rangle\langle t'|_T) \right] \\
&= \text{Tr}_{ABT} \left[\tilde{P}_{a|x} \sigma_{ABET} \tilde{P}_{a|x} \right]. \tag{78}
\end{aligned}$$

Hence we no longer consider the pinching channel on T , i.e. we simply study the situation where we immediately perform the projective measurement (77) on the state $\hat{\rho}_{ABET} := \rho_{ABE} \otimes |\phi_p\rangle\langle\phi_p|_T$ and store the outcome in register \hat{A}_x . As mentioned previously, we can take ρ_{ABE} to be pure without loss of generality, in which case $\hat{\rho}_{ABET}$ is pure as well¹⁶, and must hence obey the following relation (derived in e.g. [Col12]):

$$H(\hat{A}_x|E) = D(\hat{\rho}_{ABT} \| \mathcal{Z}_x(\hat{\rho}_{ABT})) = D(\mathcal{G}(\rho_{AB}) \| \mathcal{Z}_x(\mathcal{G}(\rho_{AB}))), \tag{79}$$

where

$$\mathcal{G}(\sigma_{AB}) := \sigma_{AB} \otimes |\phi_p\rangle\langle\phi_p|_T, \tag{80}$$

$$\mathcal{Z}_x(\sigma_{ABT}) := \sum_a (\tilde{P}_{a|x} \otimes \mathbb{I}_B) \sigma_{ABT} (\tilde{P}_{a|x} \otimes \mathbb{I}_B). \tag{81}$$

Importantly, this expression for $H(\hat{A}_x|E)$ is entirely in terms of the reduced state ρ_{AB} , and is convex with respect to ρ_{AB} . (We give an alternative expression in Appendix A, which may be of use in other situations.)

We remark that the above analysis was fairly general, in that in fact it applies to DI scenarios with arbitrary numbers of inputs and outputs, as long as Alice’s key-generating measurements still only have 2 outcomes. In particular, the approaches described in [TSG+21, BFF21] are able to yield bounds on these optimizations when there is no noisy preprocessing, and it may be useful to study whether the above map for describing noisy preprocessing allows one to apply those approaches to this scenario. However, we leave this question for future work.

We now specialize to 2-input 2-output scenarios. In such cases, we have the following “qubit reduction” [PAB+09, HST+20]: if one finds a convex function that lower-bounds the right-hand-side of Eq. (72) with its optimization restricted to states ρ_{ABE} of dimension $2 \times 2 \times 4$ and Pauli measurements, then this function is also a lower bound on \check{r}_p . We note (by following the same arguments as before, but with the optimization domain restricted) that picking some $\vec{\lambda}$ and solving

¹⁶It was important to remove the pinching channel because if $\hat{\rho}_{ABET}$ were a mixed state, then Eq. (79) would be replaced by $D(\hat{\rho}_{ABT} \| \mathcal{Z}_x(\hat{\rho}_{ABT})) = H(\hat{A}_x|EE')$, where E' purifies $\hat{\rho}_{ABET}$. While $H(\hat{A}_x|EE')$ is a valid lower bound on $H(\hat{A}_x|E)$, it also turns out to be a trivial one, because the value of T is “copied” into E' , removing the advantage of noisy preprocessing (in which Eve is not supposed to know whether the bit-flip has occurred).

the optimization (75) over such states and measurements yields such a lower bound via Eq. (74), which is trivially convex since it is affine. Hence it suffices to consider the optimization (75) restricted to such states and measurements. Furthermore, the resulting bounds are still tight in the same sense as before, in that taking the supremum over choices of $\vec{\lambda}$ yields the convex envelope of this restricted version of the optimization (72) (except possibly for \vec{v} not in the interior of the quantum-achievable values).

Since we have reduced the analysis to Pauli measurements, it is convenient to interpret the measurements as producing values in ± 1 instead of \mathbb{Z}_2 , and define the corresponding hermitian observables $A_x := \sum_a a P_{a|x}$ and $B_y := \sum_b b Q_{b|y}$. For 2-input 2-output scenarios, the full probability distribution $\Pr[ab|xy]$ (subject to the no-signalling constraints) is completely parametrized by the 4 correlators $\text{Tr}[(A_x \otimes B_y)\rho_{AB}]$ and the 4 marginals $\text{Tr}[(A_x \otimes \mathbb{I})\rho_{AB}]$, $\text{Tr}[(\mathbb{I} \otimes B_y)\rho_{AB}]$. To simplify the analysis, one can focus on the case where the marginals are zero¹⁷, in which case only the correlators $\text{Tr}[(A_x \otimes B_y)\rho_{AB}]$ remain to be considered. In terms of the optimization (72), this means that it suffices to consider situations where there are 4 constraints, corresponding to the observables

$$\Gamma_{xy}(P_{a|x}, Q_{b|y}) = A_x \otimes B_y, \quad \text{for } x, y \in \{0, 1\}. \quad (82)$$

Considering other forms of constraints in 2-input 2-output scenarios is essentially equivalent to making specific choices of Lagrange multipliers $\vec{\lambda}$ for this 4-constraint formulation. For instance, in Protocol 1 we only impose a constraint based on the CHSH value¹⁸, which is equivalent to restricting to Lagrange-multiplier combinations of the form $(\lambda_{00}, \lambda_{01}, \lambda_{10}, \lambda_{11}) = (\lambda, \lambda, \lambda, -\lambda)$ for some $\lambda \in \mathbb{R}$.

We still have the freedom to choose the basis in which to express the optimization. Following [SGP+21], we can use the measurement axes of Alice's two Pauli measurements to define the X - Z plane on her system, taking $A_0 = Z$ and $A_1 = \cos(\theta_A)Z + \sin(\theta_A)X$ for some $\theta_A \in [0, \pi]$ (values of θ_A in $[\pi, 2\pi]$ can be brought into this range by rotating our axis choice by π around the Z -axis). Analogously, we can choose a basis for Bob such that $B_0 = Z$ and $B_1 = \cos(\theta_B)Z + \sin(\theta_B)X$ with $\theta_B \in [0, \pi]$. (This is a different basis choice from the one in [PAB+09, HST+20] that allows a reduction to *Bell-diagonal* ρ_{AB} . That choice involves more parameters for the measurements, but fewer parameters for the state.)

¹⁷While more detailed explanations can be found in e.g. [PAB+09, HST+20], the outline is as follows: consider a virtual *symmetrization step* in which Alice and Bob jointly flip their outputs using a uniformly random public bit, forcing their marginals to be zero while leaving the correlators unchanged. Some calculation shows [SR08, PAB+09, HST+20] that the entropy of their original outputs conditioned on Eve's side-information is equal to the entropy of their symmetrized outputs conditioned on Eve's side-information and the publicly communicated bit (and this still holds true with noisy preprocessing). Absorbing the publicly communicated bit into Eve's side-information, this implies that if the original optimization (72) had constraints on the marginals $\text{Tr}[(A_x \otimes \mathbb{I})\rho_{AB}]$, $\text{Tr}[(\mathbb{I} \otimes B_y)\rho_{AB}]$ in addition to the correlators $\text{Tr}[(A_x \otimes B_y)\rho_{AB}]$, we could bound it by instead considering the optimization where the marginal constraints are replaced by the condition that they are zero (with the correlator constraints unchanged). (Note that the virtual symmetrization does not need to be physically performed; it merely serves as an intermediate construction in this analysis.)

¹⁸Here, by CHSH value we mean the quantity $\nu = \text{Tr}[(A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1)\rho_{AB}]$. This is related to the probability w of winning the CHSH game by the simple equation $\nu = 8w - 4$, so an affine function of ν is easily converted to an affine function of w .

With this in mind, we rewrite the optimization (75) as

$$F_{\text{obj}}(\theta_A, \theta_B, \rho_{AB}) := \left(\min_{\theta_A} \min_{\theta_B} \min_{\rho_{AB}} F_{\text{obj}}(\theta_A, \theta_B, \rho_{AB}), \text{ where} \right. \\ \left. \sum_{x \in \{0,1\}} \tau_x D(\mathcal{G}(\rho_{AB}) \| \mathcal{Z}_x(\mathcal{G}(\rho_{AB}))) \right) - \vec{\lambda} \cdot \text{Tr} \left[\vec{\Gamma}(\theta_A, \theta_B) \rho_{AB} \right]. \quad (83)$$

(The minima are attained because the objective function is continuous and the domain is compact.) Furthermore, the objective function is invariant under the substitutions $\rho_{AB} \rightarrow (Y \otimes Y) \rho_{AB} (Y \otimes Y)$ and $\rho_{AB} \rightarrow \rho_{AB}^*$, which implies [PAB+09] we can restrict the optimization to ρ_{AB} that are “almost” Bell-diagonal — specifically, with respect to the Bell basis $\{|\Phi^+\rangle, |\Psi^-\rangle, |\Phi^-\rangle, |\Psi^+\rangle\}$ where $|\Phi^\pm\rangle = (|00\rangle \pm |11\rangle)/\sqrt{2}$ and $|\Psi^\pm\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}$, we can take

$$\rho_{AB} = \begin{pmatrix} L_{\Phi^+} & \ell_1 & 0 & 0 \\ \ell_1 & L_{\Psi^-} & 0 & 0 \\ 0 & 0 & L_{\Phi^-} & \ell_2 \\ 0 & 0 & \ell_2 & L_{\Psi^+} \end{pmatrix}, \quad (84)$$

for some $L_{\Phi^+}, L_{\Psi^-}, L_{\Phi^-}, L_{\Psi^+}, \ell_1, \ell_2 \in \mathbb{R}$.

We now describe how each minimization can be tackled when treating the parameters in the other minimizations as constants, then summarize how all these algorithms can be put together in a consistent manner, and argue that this approach indeed yields arbitrarily tight bounds.

5.1 Minimization over Alice’s measurement

We tackle this minimization simply by applying a (uniform) continuity bound for θ_A . Specifically, for $\delta \in [0, \pi]$ we describe a monotone increasing function $\varepsilon_{\text{con}}(\delta)$ that bounds the change in the objective function when θ_A is replaced by $\theta_A + \delta$ (treating θ_B and ρ_{AB} as constants). Then for any set of intervals of the form $\{[\theta_j - \delta_j, \theta_j + \delta_j]\}_j$ that covers the interval $[0, \pi]$, we would have

$$\min_{\theta_A} F_{\text{obj}}(\theta_A, \theta_B, \rho_{AB}) \geq \min_j F_{\text{obj}}(\theta_j, \theta_B, \rho_{AB}) - \varepsilon_{\text{con}}(\delta_j). \quad (85)$$

We apply this in practice by starting with a fairly “coarse” choice of intervals, then iteratively applying the process of deleting the interval that currently achieves the minimization over j , and replacing it with smaller intervals that cover the deleted interval.

To derive such a continuity bound, we first analyze the entropic term in the objective function, following [SBV+21] (with a minor modification to slightly improve the bound). Take any pure initial state $|\rho\rangle_{ABE}$, and let $\sigma_{\hat{A}_1 BE}$ be the state obtained by performing a Pauli measurement along angle θ_A in the X - Z plane on the A register of $|\rho\rangle_{ABE}$, applying noisy preprocessing, then storing the result in the classical register \hat{A}_1 and tracing out A . Let $\sigma'_{\hat{A}_1 BE}$ be the analogous state with θ_A replaced by $\theta_A + \delta$ for some $\delta \in [0, \pi]$. Our goal would be to bound $\left| H(\hat{A}_1|E)_\sigma - H(\hat{A}_1|E)_{\sigma'} \right|$.

We have $H(\hat{A}_1|E)_\sigma = \sum_{\hat{a}_1} \Pr[\hat{a}_1] H(\sigma_{E|\hat{A}_1=\hat{a}_1}) + H(\hat{A}_1)_\sigma - H(E)_\sigma$, and analogously for σ' . Since the operations on A do not affect E , we have $H(E)_\sigma = H(E)_{\sigma'}$. Also, for states of the form (84), we have $\rho_A = \mathbb{I}/2$ and hence $H(\hat{A}_1)_\sigma = H(\hat{A}_1)_{\sigma'} = 1$. This gives us

$$\left| H(\hat{A}_1|E)_\sigma - H(\hat{A}_1|E)_{\sigma'} \right| \leq \max_{\hat{a}_1} \left| \left(H(\sigma_{E|\hat{A}_1=\hat{a}_1}) - H(\sigma'_{E|\hat{A}_1=\hat{a}_1}) \right) \right|, \quad (86)$$

so it suffices to bound the difference in entropies of the conditional states on E .

Now observe that exactly the same state $\sigma_{\hat{A}_1 BE}$ would have been produced if the initial state had been $(e^{i\theta_A Y_{A/2}} \otimes \mathbb{I}_{BE}) |\rho\rangle_{ABE}$ and the initial Pauli measurement were replaced by a Z measurement. Furthermore, the fact that $\rho_A = \mathbb{I}/2$ implies we can write $|\rho\rangle_{ABE} = \sum_a |a\rangle_A |a\rangle_{BE} / \sqrt{2}$, where $\{|a\rangle_A\}_a$ is the Z -eigenbasis of A and $\{|a\rangle_{BE}\}_a$ are two orthonormal states on BE . This implies

$$\left(e^{i\theta_A Y_{A/2}} \otimes \mathbb{I}_{BE} \right) |\rho\rangle_{ABE} = \left(\mathbb{I}_A \otimes \left(e^{i\theta_A Y_{BE/2}} \right)^T \right) |\rho\rangle_{ABE} = \frac{1}{\sqrt{2}} \sum_a |a\rangle_A \otimes R_{\theta_A} |a\rangle_{BE}, \quad (87)$$

where $Y_{BE} := i|-1\rangle\langle+1|_{BE} - i|+1\rangle\langle-1|_{BE}$ and $R_{\theta_A} := (e^{i\theta_A Y_{BE/2}})^T$. Performing the analogous analysis for σ' , we conclude that

$$\begin{aligned} \sigma_{BE|\hat{A}_1=\hat{a}_1} &= (1-p)R_{\theta_A} |\hat{a}_1\rangle\langle\hat{a}_1|_{BE} R_{\theta_A}^\dagger + pR_{\theta_A} |-\hat{a}_1\rangle\langle-\hat{a}_1|_{BE} R_{\theta_A}^\dagger, \\ \sigma'_{BE|\hat{A}_1=\hat{a}_1} &= (1-p)R_{\theta_{A+\delta}} |\hat{a}_1\rangle\langle\hat{a}_1|_{BE} R_{\theta_{A+\delta}}^\dagger + pR_{\theta_{A+\delta}} |-\hat{a}_1\rangle\langle-\hat{a}_1|_{BE} R_{\theta_{A+\delta}}^\dagger. \end{aligned} \quad (88)$$

Therefore, we have

$$\begin{aligned} F(\sigma_{E|\hat{A}_1=\hat{a}_1}, \sigma'_{E|\hat{A}_1=\hat{a}_1}) &\geq F(\sigma_{BE|\hat{A}_1=\hat{a}_1}, \sigma'_{BE|\hat{A}_1=\hat{a}_1}) \\ &\geq (1-p) \left| \langle\hat{a}_1| R_{\theta_A}^\dagger R_{\theta_{A+\delta}} |\hat{a}_1\rangle \right| + p \left| \langle-\hat{a}_1| R_{\theta_A}^\dagger R_{\theta_{A+\delta}} |-\hat{a}_1\rangle \right| \\ &\geq \left| \cos \frac{\delta}{2} \right|, \end{aligned} \quad (89)$$

where the second inequality holds by concavity of fidelity, and the third inequality is given by explicit calculation (see [SBV+21]).

This lets us apply a fidelity-based continuity bound [SBV+21]:

$$\max_{\hat{a}_1} \left| \left(H(\sigma_{E|\hat{A}_1=\hat{a}_1}) - H(\sigma'_{E|\hat{A}_1=\hat{a}_1}) \right) \right| \leq 4.023 \cos^{-1} F(\sigma_{E|\hat{A}_1=-1}, \sigma'_{E|\hat{A}_1=-1}) \leq 2.012\delta, \quad (90)$$

where in the second inequality we used the condition $\delta \in [0, \pi]$. (Numerical heuristics suggest that the true bound in Eq. (90) may simply be δ , so there is some potential for improvement here, though the effect would be fairly small. In Appendix B we present an approach based on trace distance instead of fidelity, but it appears to scale poorly at small δ .)

As for the $\vec{\Gamma}$ term, we note that

$$\begin{aligned} &\left| \text{Tr} \left[\left(\vec{\lambda} \cdot \vec{\Gamma}(\theta_A + \delta, \theta_B) - \vec{\lambda} \cdot \vec{\Gamma}(\theta_A, \theta_B) \right) \rho_{AB} \right] \right| \\ &\leq \left\| \vec{\lambda} \cdot \vec{\Gamma}(\theta_A + \delta, \theta_B) - \vec{\lambda} \cdot \vec{\Gamma}(\theta_A, \theta_B) \right\|_\infty \\ &= \left\| ((\cos(\theta_A + \delta)Z + \sin(\theta_A + \delta)X) - (\cos(\theta_A)Z + \sin(\theta_A)X)) \otimes (\lambda_{10}B_0 + \lambda_{11}B_1) \right\|_\infty \\ &\leq (|\lambda_{10}| + |\lambda_{11}|) \left\| (\cos(\theta_A + \delta)Z + \sin(\theta_A + \delta)X) - (\cos(\theta_A)Z + \sin(\theta_A)X) \right\|_\infty \\ &= (|\lambda_{10}| + |\lambda_{11}|) \sqrt{2 - 2\cos(\delta)}, \end{aligned} \quad (91)$$

where the last line follows from an explicit eigenvalue calculation.

Overall, this means that for $\delta \in [0, \pi]$ we can choose

$$\varepsilon_{\text{con}}(\delta) = 2.012\tau_1\delta + (|\lambda_{10}| + |\lambda_{11}|) \sqrt{2 - 2\cos(\delta)}, \quad (92)$$

accounting for the τ_1 factor on the $H(\hat{A}_1|E)$ term. This bound is monotone increasing for $\delta \in [0, \pi]$, as required (so that it also bounds the change in entropy when the measurement angle is changed from θ_A to any value in the interval $[\theta_A, \theta_A + \delta]$).

5.2 Minimization over Bob's measurement

The entropic term in the objective function has no dependence on Bob's measurement, so we only need to consider the $\vec{\Gamma}$ term. In principle, this could be approached using the same argument as above, where we would arrive at the continuity bound

$$\left| \text{Tr} \left[\left(\vec{\lambda} \cdot \vec{\Gamma}(\theta_A, \theta_B + \delta) - \vec{\lambda} \cdot \vec{\Gamma}(\theta_A, \theta_B) \right) \rho_{AB} \right] \right| \leq (|\lambda_{01}| + |\lambda_{11}|) \sqrt{2 - 2 \cos(\delta)}. \quad (93)$$

However, some heuristic experiments indicate that the following approach (used in [SGP+21]) is more efficient: we can let $r_Z := \cos(\theta_B)$ and $r_X := \sin(\theta_B)$ and write

$$\sum_{xy} \lambda_{xy} A_x \otimes B_y = \left(\sum_x \lambda_{x0} A_x \right) \otimes Z + \left(\sum_x \lambda_{x1} A_x \right) \otimes (r_Z Z + r_X X), \quad (94)$$

in which case the minimization over $\theta_B \in [0, \pi]$ is equivalent to minimizing over (r_Z, r_X) that lie on the set $S := \{(r_Z, r_X) \mid r_Z^2 + r_X^2 = 1 \text{ and } r_X \geq 0\}$ (i.e. a semicircular arc). Crucially, the objective function is affine with respect to the vector (r_Z, r_X) . Hence if V is any finite set of points such that S is contained in their convex hull $\text{Conv}(V)$, we immediately have

$$\begin{aligned} \min_{(r_Z, r_X) \in S} F_{\text{obj}}(\theta_A, (r_Z, r_X), \rho_{AB}) &\geq \min_{(r_Z, r_X) \in \text{Conv}(V)} F_{\text{obj}}(\theta_A, (r_Z, r_X), \rho_{AB}) \\ &= \min_{(r_Z, r_X) \in V} F_{\text{obj}}(\theta_A, (r_Z, r_X), \rho_{AB}), \end{aligned} \quad (95)$$

because the minimum of an affine function over the convex hull of a finite set V is always attained at an extremal point (which will be a point in V). To apply this result, we start with a simple choice of the set V (for instance, in our code we use $V = \{(1, 0), (1, 1), (-1, 1), (-1, 0)\}$) and find the point in V that yields the minimum value. We then delete this point and replace it with two other points such that S is still contained in the convex hull, and iterate this process until a sufficiently tight bound is obtained (for instance, by checking that there is a feasible point of the optimization that is sufficiently close to the lower bound we have obtained).

5.3 Minimization over states

The minimization over states can be tackled by expressing it as a convex optimization and applying the *Frank-Wolfe algorithm* [FW56], as was observed in [WLC18]. For completeness, we now describe the method, with minor modifications and clarifications for our specific scenario. We emphasize that while this procedure is numerical, the bounds it returns are *secure*, in the sense that it will never over-estimate the true value of the minimization problem.

We observe that (for fixed measurements, parametrized as described above) the minimization over states is the minimization of the convex function

$$f_{\text{obj}}(\rho) := \left(\sum_{x \in \{0,1\}} \tau_x D(\mathcal{G}(\rho) \parallel \mathcal{Z}_x(\mathcal{G}(\rho))) \right) - \vec{\lambda} \cdot \text{Tr} \left[\vec{\Gamma}(\theta_A, (r_Z, r_X)) \rho \right]. \quad (96)$$

This function is differentiable for all ρ such that $\mathcal{G}(\rho) > 0$, with [WLC18]¹⁹

$$(\nabla f_{\text{obj}}(\rho))^T = \left(\sum_{x \in \{0,1\}} \tau_x \mathcal{G}^\dagger(\log \mathcal{G}(\rho) - \log \mathcal{Z}_x(\mathcal{G}(\rho))) \right) - \vec{\lambda} \cdot \vec{\Gamma}(\theta_A, (r_Z, r_X)), \quad (97)$$

where \mathcal{G}^\dagger is the adjoint channel of \mathcal{G} . In practice, we will not need to explicitly compute \mathcal{G}^\dagger , because all subsequent arguments rely only on “inner products” $\text{Tr}[(\nabla f_{\text{obj}}(\rho))^T \sigma]$, which can be rewritten as

$$\text{Tr}[(\nabla f_{\text{obj}}(\rho))^T \sigma] = \text{Tr} \left[\left(\sum_{x \in \{0,1\}} \tau_x (\log \mathcal{G}(\rho) - \log \mathcal{Z}_x(\mathcal{G}(\rho))) \mathcal{G}(\sigma) \right) - \vec{\lambda} \cdot \vec{\Gamma}(\theta_A, (r_Z, r_X)) \sigma \right]. \quad (98)$$

We also note that the optimization domain is a set defined by PSD constraints (namely, ρ of the form (84) with $\rho \geq 0$ and $\text{Tr}[\rho] = 1$), which means that optimizing any affine function over this set is an SDP, which can be efficiently solved and bounded [BV04] via its dual value. Together with the explicit expression (97) for the derivative of the objective function, this makes the optimization problem here a prime candidate for the Frank-Wolfe algorithm [FW56], which yields *secure* lower bounds on the true minimum value of the optimization.

The Frank-Wolfe algorithm is based on a simple geometric insight: for any point in the domain of a convex function f_{obj} , the tangent hyperplane (or a supporting hyperplane, if f_{obj} is not differentiable at that point) to the graph of f_{obj} at that point yields an affine lower bound on f_{obj} . Hence if we can minimize affine functions over the optimization domain, then any such tangent hyperplane lets us obtain a lower bound on the minimum of f_{obj} on the domain. In addition, it intuitively seems that taking the tangent at points closer to the true optimal solution should yield tighter lower bounds (though there are some technical caveats). Thus in principle, one could simply perform some heuristic computations to get an estimate of where the true minimum lies, then take the tangent at that estimate and obtain the corresponding lower bound. For this optimization, however, we found that the results were fairly sensitive to deviations from the true optimum. To get the bounds to converge, we found it more efficient to use the standard Frank-Wolfe algorithm, which (under mild assumptions) can be proven to converge in order $O(1/k)$ after k iterations:

¹⁹The notation here follows that used in [WLC18], in which for a function f defined on matrices parametrized by their matrix entries ($\rho = \sum_{jk} \rho_{jk} |j\rangle\langle k|$), its derivative is $\nabla f(\rho) := \sum_{jk} (\partial f / \partial \rho_{jk}) |j\rangle\langle k|$. Denoting the tangent to the graph of f at ρ as t_ρ , this definition of ∇f satisfies $t_\rho(\rho + \Delta) = f(\rho) + \text{Tr}[(\nabla f(\rho))^T \Delta]$.

Frank-Wolfe algorithm (As presented in [WLC18])

Let the domain of optimization be \mathcal{D} and the acceptable gap between the feasible values and the lower bounds be $\varepsilon_{\text{tol}} > 0$.

1. Set $k = 0$ and heuristically find $\rho_0 := \arg \min_{\rho \in \mathcal{D}} f_{\text{obj}}(\rho)$.
2. Solve the SDP

$$\min_{\Delta} \text{Tr} [(\nabla f_{\text{obj}}(\rho_k))^T \Delta] \quad \text{s.t.} \quad \rho_k + \Delta \in \mathcal{D}, \quad (99)$$

which returns a feasible point Δ^* as well as a dual value $-\varepsilon \leq 0$ that lower-bounds the true minimum of the SDP.

3. If $\varepsilon \leq \varepsilon_{\text{tol}}$ then stop and return $f_{\text{obj}}(\rho_k) - \varepsilon$ as a lower bound on the optimization.
 4. Otherwise, heuristically find $\mu^* := \arg \min_{\mu \in (0,1)} f_{\text{obj}}(\rho_k + \mu \Delta^*)$.
 5. Set $\rho_{k+1} = \rho_k + \mu^* \Delta^*$, $k \leftarrow k + 1$ and return to Step 2.
-

Geometrically, the SDP in Eq. (99) corresponds to considering the tangent to f_{obj} at ρ_k and computing the maximum amount by which it can decrease (as compared to its value at ρ_k) over the domain \mathcal{D} . As described earlier, this bounds the maximum amount by which f_{obj} can decrease from $f_{\text{obj}}(\rho_k)$ over \mathcal{D} .

In our application of the Frank-Wolfe algorithm, there is the technical issue that if $\mathcal{G}(\rho_k)$ is singular (or has negative eigenvalues, from numerical imprecision), then the derivative at ρ_k (Eq. (97)) is ill-behaved. To cope with this, we used the heuristic solution of simply replacing ρ_k with $(1 - \delta)\rho_k + \delta \mathbf{U}$ in Step 2, where $\delta := \max(\text{eigenvalues}(-\mathcal{G}(\rho_k)) \cup \{10^{-14}\})$. Note that this does not affect the *security* of the result, since it merely corresponds to taking a tangent at a slightly different point, which still yields a valid lower bound. On the other hand, it could possibly affect the theoretical convergence rates, but in practice this did not appear to pose a significant problem in our setting. (In [WLC18], this issue is addressed by analyzing a ‘‘perturbed’’ version of the optimization and applying a continuity bound, but we found that for the level of accuracy we desired in this work, the admissible perturbation values are too small to cope with the negative eigenvalues that occur.)

5.4 Overall algorithm

Putting the above results together, we see that for any set of intervals $[\theta_j - \delta_j, \theta_j + \delta_j]$ such that $[0, \pi] \subseteq \bigcup_j [\theta_j - \delta_j, \theta_j + \delta_j]$, and set V such that $S \subseteq \text{Conv}(V)$, we have

$$\begin{aligned} \min_{\rho_{AB}} \min_{\theta_B} \min_{\theta_A} F_{\text{obj}}(\theta_A, \theta_B, \rho_{AB}) &\geq \min_{\rho_{AB}} \min_{\theta_B} \min_j F_{\text{obj}}(\theta_j, \theta_B, \rho_{AB}) - \varepsilon_{\text{con}}(\delta_j) \\ &= \min_j \min_{\rho_{AB}} \min_{(r_Z, r_X) \in S} F_{\text{obj}}(\theta_j, (r_Z, r_X), \rho_{AB}) - \varepsilon_{\text{con}}(\delta_j) \\ &\geq \min_j \min_{\rho_{AB}} \min_{(r_Z, r_X) \in V} F_{\text{obj}}(\theta_j, (r_Z, r_X), \rho_{AB}) - \varepsilon_{\text{con}}(\delta_j) \\ &= \min_j \min_{(r_Z, r_X) \in V} \min_{\rho_{AB}} F_{\text{obj}}(\theta_j, (r_Z, r_X), \rho_{AB}) - \varepsilon_{\text{con}}(\delta_j). \end{aligned} \quad (100)$$

We can refine the intervals $[\theta_j - \delta_j, \theta_j + \delta_j]$ and set V by the iterative processes described above, with the innermost minimization over ρ_{AB} being handled by the Frank-Wolfe algorithm. It is clearly possible to swap the order of \min_j and $\min_{(r_Z, r_X) \in V}$ in the last line; however, some heuristic plots of the objective function suggest that performing the optimizations in the order shown here is slightly faster.

To see that this expression can converge to a tight bound, observe that the first inequality in (100) becomes arbitrarily tight as we choose smaller values of δ_j . For the second inequality, note that the described algorithm chooses V in such a way that the minimum over V approaches the minimum over S , hence this inequality also becomes arbitrarily tight.

In general, the above approach faces the difficulty that it needs to optimize over the choice of Lagrange multipliers $\vec{\lambda}$. Given that our approach for solving (100) for a specific choice of Lagrange multipliers is already highly computationally intensive (requiring about 5000 core-hours to achieve the level of accuracy in the bounds (103) below for each value of p), it would be impractical to also optimize over the Lagrange multipliers while doing so. It is more feasible to first optimize the Lagrange multipliers while using a simple heuristic algorithm to estimate the minimizations, then certify the final result using our approach for solving (100). (This is essentially the same perspective as presented in [SBV+21].)

We remark that this approach can also yield arbitrarily tight bounds for DIRE, where the goal would typically be [AFRV19, BRC20] to find lower bounds on (weighted sums of) “two-party entropies” $H(A_x B_y | E)$. This is because by the same arguments as above, we have $H(A_x B_y | E) = D(\rho_{AB} \| \mathcal{Z}_{xy}(\rho_{AB}))$ [Col12, TSG+21], where

$$\mathcal{Z}_{xy}(\sigma_{AB}) := \sum_{ab} (P_{a|x} \otimes Q_{b|y}) \sigma_{AB} (P_{a|x} \otimes Q_{b|y}). \quad (101)$$

(Here we omit the parts corresponding to noisy preprocessing, since it is not applied in randomness expansion.) This expression can be bounded in the same way as we have just described above, though the objective function would no longer be affine with respect to (r_Z, r_X) , and hence the optimization over Bob’s measurements would also have to be approached using a continuity bound. Our approach should yield a substantial improvement over previous results, which are restricted to the CHSH inequality and only bound the entropy of one party’s outputs [LLR+21], or which consider the full distribution and bound the entropy of both outputs but use inequalities that are not tight [BRC20, TSG+21, BFF21]. In addition, the fact that it allows for the random-key-measurement approach could yield further improvements, though there are some technicalities that we address at the end of Sec. 6.2.

It would be convenient for future analysis if it were possible to develop closed-form expressions for \check{r}_p (as was done in [HST+20, WAP21, SBV+21]), rather than the computationally intensive numerical approaches shown above. However, this appears to be rather challenging, as noted in [WAP21]. In particular, we found numerical evidence against the conjecture proposed in [SGP+21] that the minimum in (72) (when restricted to qubit strategies) can always be attained by states such that ρ_{AB} is of rank 2. (A similar observation was reported in [WAP21].) More precisely, in the process of heuristically solving the optimizations in order to estimate suitable choices of λ , we discovered that if we imposed the additional restriction that ρ_{AB} has rank 2, there was a small but numerically significant difference as compared to the results without this restriction. (The states that heuristically approach the minimum in the latter indeed tend to have two very small eigenvalues, but it appears that these eigenvalues cannot be reduced exactly to zero.) This suggests that the aforementioned conjecture is not true after all, which poses a challenge for closed-form analysis because the eigenvalues involved in computing the entropy are genuinely roots of a fourth-degree polynomial (in [HST+20, SBV+21], a key element of the analysis was to argue that it suffices to consider rank-2 ρ_{AB} , simplifying the expression for the eigenvalues).

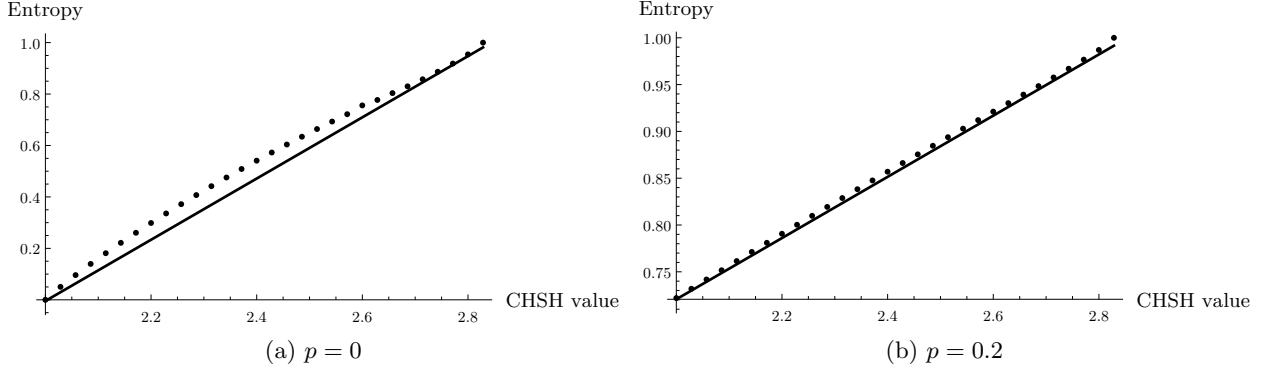


Figure 3: The solid lines are the certified lower bounds we obtained (Eq. (103)), while the points indicate the results of heuristically solving the optimization (72) over qubit states and measurements (with just the CHSH value as the constraint). As previously discussed, the tight bound in each case would be given by the convex envelope of the curve traced out by the points, assuming that the heuristics have found the true minimum. However, we can see that in each case that curve appears to be nonconvex over the interval $[2, 2.75]$ (approximately), and its convex envelope would be affine over that interval — specifically, it would be given by the linear interpolation between the feasible points at the ends of that interval. The certified bound is almost flush with this linear interpolation, indicating that it is basically tight over this interval.

5.5 Resulting bounds

For the protocols in this work, the relevant bound to compute is for the case where the only constraint imposed is the CHSH value. More precisely, we consider the described optimization with a single constraint corresponding to the operator

$$\Gamma(\theta_A, \theta_B) = A_0 \otimes B_0 + A_0 \otimes B_1 + A_1 \otimes B_0 - A_1 \otimes B_1, \quad (102)$$

with the constraint value ν being the CHSH value. (As previously mentioned, this can be implemented in the formulation where there are 4 constraint operators $\Gamma_{x,y}(\theta_A, \theta_B)$ by simply restricting to Lagrange-multiplier choices of the form $(\lambda_{00}, \lambda_{01}, \lambda_{10}, \lambda_{11}) = (\lambda, \lambda, \lambda, -\lambda)$.) Each choice of the associated Lagrange multiplier λ yields an affine lower bound of the form $\lambda\nu + c_\lambda$, as noted in Eq. (74). Importantly, some heuristic computations (also observed in [SGP+21] for the $p = 0$ case) suggest that the true bound $\check{r}_p(\nu)$ in this situation is in fact affine over a wide range of CHSH values — we show this in Fig. 3, which displays the results of heuristic minimizations compared to our certified bound in some cases. In particular, this range on which the bound is affine covers all currently experimentally reasonable values. This has the implication that there is a single λ that yields an affine lower bound which is equal to $\check{r}_p(\nu)$ (i.e. it is tight) over this entire range; specifically, it is simply the value of λ corresponding to the gradient of $\check{r}_p(\nu)$ in this range. This greatly simplifies our task since we only need to solve the optimization for this specific value of λ .

We focused on several values of noisy preprocessing, ranging from $p = 0$ to $p = 0.45$. In each case, we first solved the minimization (75) heuristically for some selection of values of λ , in order to estimate the choice of λ that yields a tight bound over the range of CHSH values in which $\check{r}_p(\nu)$ is affine. Then using our algorithm to get certified bounds on the corresponding c_λ in (75), we

²⁰The asymptotic keyrates of Protocol 1 would be given by simply halving these keyrates, which does not change their horizontal intercepts.




Figure 4: Lower bounds on asymptotic keyrates under depolarizing noise q (defined in Sec. 3.3), on a vertical log scale. The black, blue, and red solid curves show the (net) asymptotic keyrate of Protocol 2 (defined in Sec. 6.2)²⁰, for noisy-preprocessing values of $p = 0, 0.2$ and 0.3 respectively, based on the bounds we computed (Eq. (103)). For comparison, the dashed curves show the corresponding asymptotic keyrates of the protocol in [HST+20], which does not use the random-key-measurement method (the $p = 0$ case is equivalent to the [PAB+09] protocol). The solid curves intersect the horizontal axis at $q = 8.39\%$, 9.26% and 9.33% , in order of increasing p . The first value is a minor improvement over [SGP+21] (despite being effectively the same protocol), likely because our algorithm provably converges to a tight keyrate bound (for fixed p). The last value exceeds all previous bounds [HST+20, SGP+21, WAP21, SBV+21] for depolarizing-noise tolerance, by an amount comparable to their respective improvements over the original [PAB+09] protocol. It is also close to the upper bound of 9.57% that we derive in Sec. 5.6.

arrived at the final bounds

$$\begin{aligned}
 \check{r}_0(\nu) &\geq 1.190(\nu - 2) - 0.00454, & \check{r}_{0.2}(\nu) &\geq 0.327(\nu - 2) + 0.72063, \\
 \check{r}_{0.3}(\nu) &\geq 0.139(\nu - 2) + 0.88051, & \check{r}_{0.4}(\nu) &\geq 0.0341(\nu - 2) + 0.97055, \\
 \check{r}_{0.45}(\nu) &\geq 0.00855(\nu - 2) + 0.992487.
 \end{aligned} \tag{103}$$

(To express the bounds as functions of winning probability w instead, as we required for our various keyrate computations, simply replace ν in each expression with $8w - 4$.) With these bounds, we computed the asymptotic keyrates for these noisy-preprocessing values, and we show the results up to $p = 0.3$ in Fig. 4. From the $p = 0.3$ case, we obtain the depolarizing-noise threshold of 9.33% mentioned in the introduction. However, we were unable to obtain better thresholds using the higher values of p , for reasons we shall soon discuss (in Sec. 5.6 below).

For photonic experiments, we consider the [Ebe93] model described in Sec. 3.3 earlier, and the resulting keyrates for noisy-preprocessing values up to $p = 0.3$ are shown in Fig. 5 (as in the depolarizing-noise case, higher values of p gave somewhat worse results, for similar reasons that will be discussed in Sec. 5.6). As mentioned previously, in this scenario, for each value of η we have to optimize the honest states and measurements to maximize the keyrates. Since this optimization was performed heuristically and appears to be slightly numerically unstable, we are not completely certain whether these are indeed the best keyrates that could be achieved. However, for the values that we managed to obtain at least, it appears that unfortunately our protocol with $\tau_0 = \tau_1 = 1/2$ has worse detection-efficiency thresholds (for nonzero asymptotic keyrate) as compared to the results in [HST+20] based on noisy preprocessing alone, although at higher values of η our protocol




Figure 5: Lower bounds on asymptotic keyrates for limited detection efficiency η (defined in Sec. 3.3), on a vertical log scale. The black, blue, and red solid curves show the (net) asymptotic keyrate of Protocol 2 (defined in Sec. 6.2), for noisy-preprocessing values of $p = 0, 0.2$ and 0.3 respectively, based on the bounds we computed (Eq. (103)). At each value of η , the states and measurements were optimized to maximize the keyrate. For comparison, the dashed curves show the corresponding optimized asymptotic keyrates of the protocol in [HST+20], which does not use the random-key-measurement method (the $p = 0$ case is equivalent to the [PAB+09] protocol). In contrast to the depolarizing-noise case (Fig. 5), the threshold η value required for nonzero keyrates is significantly worse for our protocol than the [HST+20] protocol (although our protocol does yield slightly better keyrates at higher values of η).

can yield slightly better keyrates.²¹

The comparatively poor detection-efficiency thresholds for our protocol in this model seems to be because of the error-correction value h_{hon} — we find that for instance, the states and measurements that optimize the CHSH value yield a rather high value for at least one of the two terms $H(A_j|B_j; X_j = Y_j = 0)_{\text{hon}}$ and $H(A_j|B_j; X_j = Y_j = 1)_{\text{hon}}$ (see (17)) describing the error-correction “contributions” from the two generation inputs. Therefore, there is a significant tradeoff between maximizing the CHSH value (and hence the entropy against Eve) versus minimizing the error-correction value h_{hon} , which results in worse detection-efficiency thresholds as compared to the [HST+20] protocol which only used one input for key generation. Potential points for further keyrate improvements in the photonic model would hence be to optimize the ratio of the generation input choices or to introduce different amounts of noisy preprocessing for the two inputs, but such optimization is rather involved and we leave it for further work.

In principle, one could consider more detailed models of photonic implementations, for instance those developed in [TWF+18, HST+20] (such an analysis was performed in [SGP+21] for the random-key-measurements technique). However, since informally speaking these models account for a range of other imperfections not considered in the simple model here, they typically result

²¹Regarding the random-key-measurements technique alone, the results in [SGP+21] indicate that it does not significantly affect the detection-efficiency threshold as compared to the basic [PAB+09] protocol with optimized states and measurements, although it does somewhat improve the keyrates. We note that this comparison basically corresponds to the black curves in Fig. 5 here, although the photonic model used in [SGP+21] is a more detailed one [TWF+18] compared to the model here. Focusing on those curves in Fig. 5, we see that our findings here are roughly similar as well (in our case the random-key-measurements technique yields a slightly worse threshold, but this may just be an artifact of having considered a coarser set of data points to evaluate).

in lower keyrates than the simple model. Given that the results shown here are not entirely promising, there may be limited prospects in pursuing that direction further unless we can find some improvements within the simpler model here.

5.6 Optimality of results

For each of the bounds in Eq. (103), there is a feasible point of the optimization for c_λ which yields a value within 0.005 (or less, for higher values of p) of the certified results shown in Eq. (103), so these bounds on entropy are very close to optimal in terms of absolute error. In terms of the depolarizing-noise thresholds that they yield, taking the convex envelope of some of the feasible points shown in Fig. 3 yields the result that the thresholds for $p = 0.2$ and $p = 0.3$ cannot be improved by more than about 0.1 percentage points, so those thresholds are very close to optimal as well. However, larger values of p face the issue that the asymptotic keyrate becomes extremely low, which makes the horizontal intercept (i.e. the depolarizing-noise threshold) very sensitive to changes in \check{r}_p — even a small absolute error in this bound results in a significant change in the threshold value. Therefore, the thresholds we obtained from the certified bounds with $p = 0.4$ and $p = 0.45$ in Eq. (103) were only 9.32% and 9.10% respectively, worse than the results for $p = 0.3$. Heuristic computations suggest that the true thresholds for those cases might be approximately 9.46% and 9.50% respectively, but using our algorithm to certify these values would require it to converge to tolerances that currently appear impractical. Hence a different approach may be needed to find the true thresholds for these values of p . From a practical perspective though, such improvements may be of limited use, because the very low asymptotic rates mean that the finite-size keyrate would likely be zero until extremely large sample sizes.

In any case, we note that for depolarizing noise at least, the threshold value cannot be improved much further by *any* protocol choices within the framework we have presented in this section, e.g. by using the full distribution as constraints (which also encompasses the use of modified CHSH inequalities [WAP21, SBV+21]), or adjusting the values of τ_x . This is essentially because our bounds are very close to the linear interpolation between the points $(2, h_2(p))$ and $(2\sqrt{2}, 1)$ (as can be seen from Fig. 3). Intuitively speaking, the bound on the entropy against Eve in the depolarizing-noise scenario cannot exceed this linear interpolation (because Eve can always perform classical mixtures of strategies in order to attain every point on this linear interpolation), which means that our bounds are close to the highest bounds that are even possible in principle.

Making this reasoning quantitative, we can obtain an explicit upper bound on the depolarizing-noise threshold. (We highlight that because we will do so by constructing an extremely generic attack, this upper bound holds for *all* protocols of this form, regardless of choices of parameters such as the input distributions and noisy preprocessing. To some extent, it is surprising that there even exist any protocols that achieve thresholds close to the upper bounds implied by such a generic attack.) We first recall that (as defined in Sec. 3.3) the measurement statistics in the depolarizing-noise model are given by real projective measurements on the Werner state $\varrho_q := (1 - 2q) |\Phi^+\rangle\langle\Phi^+| + 2q \mathbb{I}/4$. It is known [AGT06, Kri79] that these measurement statistics can be reproduced by a local-hidden-variable (LHV) model as soon as the state does not violate the CHSH inequality, i.e. for $q \geq q_2 := (2 - \sqrt{2})/4 \approx 14.6\%$. (In fact, this is also true in a more general context where the honest parties have *any* number of real projective measurements on the Werner state, via the value of the second Grothendieck constant $K_G(2) = \sqrt{2}$ [AGT06, Kri79]. For general (i.e. not necessarily real) projective measurements the analogous value is known to satisfy $q_3 \lesssim 15.9\%$, following from the best known bound on the third Grothendieck constant $K_G(3)$ [HJV+17].)

In the device-independent setting, the existence of an LHV model yields an attack for Eve that gives her full knowledge of the measurement outcomes. This means that for any depolarizing-noise value q , we can construct the following attack for Eve: first, she generates a classical ancilla bit which is equal to 0 with probability $\mu := (q_2 - q)/q_2$. If the bit is equal to 0, Alice and Bob's devices simply perform the honest measurements on the maximally entangled state Φ^+ . Otherwise, the devices implement an LHV model that yields the same statistics as the honest measurements performed on the Werner state ϱ_{q_2} , but which gives Eve full knowledge of the outcomes. This attack indeed reproduces the statistics corresponding to depolarizing noise q , since $\mu |\Phi^+\rangle\langle\Phi^+| + (1 - \mu)\varrho_{q_2} = \varrho_q$ and the second strategy produces the same statistics as the honest measurements performed on ϱ_{q_2} .

In the two cases, the entropies of Alice's outputs after noisy preprocessing are $H(\hat{A}_x|\tilde{E})_{\text{triv}} = 1$ (i.e. Eve's side-information is trivial since the Alice-Bob state is pure) and $H(\hat{A}_x|\tilde{E})_{\text{LHV}} = h_2(p)$ (i.e. the uncertainty arises purely from the noisy preprocessing) respectively, where \tilde{E} denotes Eve's side-information excluding the ancilla bit. Incorporating the ancilla bit into Eve's side-information E , this attack achieves

$$H(\hat{A}_x|E) = \mu H(\hat{A}_x|\tilde{E})_{\text{triv}} + (1 - \mu)H(\hat{A}_x|\tilde{E})_{\text{LHV}} = \frac{q_2 - q}{q_2} + \frac{q}{q_2}h_2(p). \quad (104)$$

This expression hence serves as an upper bound on the best possible lower bound we could derive on the conditional entropy against Eve. Also, the conditional entropy against Bob's output (for Bob's optimal key-generation measurements on Werner states) is simply

$$H(\hat{A}_x|B_x) = h_2(p + (1 - 2p)q). \quad (105)$$

Recalling that the asymptotic keyrate is given by the Devetak-Winter bound [DW05] (up to the sifting-related weights τ_x [SGP+21]), this yields a simple upper bound on the critical value of q that still allows a positive asymptotic keyrate (for protocols of this form, i.e. applying noisy preprocessing, random key measurements, and considering the full output distribution, but only using one-way error correction):

$$\begin{aligned} q_{\text{att}}(p) &:= \max \left\{ q \left| \sum_x \tau_x (H(\hat{A}_x|E) - H(\hat{A}_x|B_x)) \geq 0 \right. \right\} \\ &= \max \left\{ q \left| \frac{q_2 - q}{q_2} + \frac{q}{q_2}h_2(p) - h_2(p + (1 - 2p)q) \geq 0 \right. \right\}, \end{aligned} \quad (106)$$

where in the first line the entropies refer to those of the attack we have described. We observe numerically that this upper bound is increasing with respect to p , so we have $q_{\text{att}}(p) \leq q_{\text{att}}(p \rightarrow 1/2)$. In the limit $p \rightarrow 1/2$ we can write $p = 1/2 - \delta$ and expand the expression in small δ to find

$$q_{\text{att}}\left(p \rightarrow \frac{1}{2}\right) = \frac{1 + 4q_2 - \sqrt{8q_2 + 1}}{8q_2} = 1 - \frac{\sqrt{7 + 4\sqrt{2}} - 1}{2\sqrt{2}} \approx 9.57\%. \quad (107)$$

Hence for protocols of the form described in this work, it is not possible for the depolarizing-noise threshold to exceed this value. For the parameter choice $p = 0.3$, we have $q_{\text{att}}(0.3) \approx 9.51\%$, which is close to the threshold of 9.33% for which we could certify a positive keyrate.²²

²²These two thresholds cannot match exactly, because the feasible points shown in Fig. 3 also imply that the tight bound $\check{r}_p(\nu)$ is not *exactly* equal to the linear interpolation between $(2, h_2(p))$ and $(2\sqrt{2}, 1)$, which essentially corresponds to the attack we describe here.

Finally, it is worth noting that while the main focus of this work is depolarizing noise applied to the statistics corresponding to the ideal CHSH measurements on Φ^+ , the above analysis in fact generalizes to a substantially larger family of scenarios. Specifically, the same analysis applies for depolarizing noise applied to the statistics from any number of real projective measurements on Φ^+ , since the results of [AGT06, Kri79] yield the required LHV models.²³ (If Bob performs suboptimal generation measurements, then Eq. (105) still holds as a lower bound, so the argument carries through.) In addition, replacing q_2 with q_3 straightforwardly provides a threshold of $q \lesssim 10.01\%$ for all possible projective measurements on that state, via the respective known bound in [HQV+17]. (Note that in all such cases where there are more than 2 possible measurements, there is no “qubit reduction” to make it easier to derive corresponding lower bounds, and hence the lower bounds have also not been very thoroughly explored.) An interesting further consideration is the case of general POVM measurements. Here, a rather loose bound $q_{\text{POVM}} \lesssim 27.3\%$ is also known [HQV+17] for the threshold that allows an LHV description. However, pure POVMs on qubits can involve up to four outcomes, opening up the possibility of more general noisy preprocessing — all doubly stochastic maps on probability vectors with four outcomes. We thus leave it for future work.

The above argument relies on the fact that depolarizing-noise statistics can be obtained by simple mixtures of “extremal” strategies. This is not necessarily the case for other noise models, such as limited detection efficiency in photonic experiments. Therefore, the same argument cannot be directly applied to obtain upper bounds on the thresholds for such forms of noise.

6 Possible modifications

6.1 Coordinating input choices by public communication

The random-key-measurement protocol has the drawback that the keyrate is effectively halved, since the generation rounds have “mismatched” inputs approximately half the time. It would be helpful to find a way to work around this issue. One possible approach could be to observe that in [AFRV19], it was assumed that the following operations can be performed in each round: the devices receive some shares of a quantum state, then²⁴ Alice and Bob publicly communicate to come to an agreement on *both* of their input choices, and finally they supply these inputs to their devices. (This was necessary in [AFRV19] because Alice and Bob’s actions in that DIQKD protocol require both of them to know whether it is a test or generation round. In fact, our analysis can be viewed as the first EAT-based security proof for a “genuinely sifting-based” DIQKD protocol, in the sense that Alice and Bob do not coordinate which rounds are test rounds, and simply choose their inputs independently.) If we assume that this is also possible in our scenario, then Alice and Bob could coordinate their inputs in the generation rounds instead of choosing them independently, thereby avoiding the sifting factor.

Unfortunately, it does not seem clear if such a proposal is entirely plausible in near-term experimental implementations. This is because it relies on the devices being able to store the quantum state for long enough for Alice and Bob to agree on their choice of inputs, which is potentially

²³A more general consideration would be the scenario of depolarizing noise applied to the statistics from an arbitrary two-qubit state, but it is not immediately obvious whether the thresholds in [AGT06, Kri79, HQV+17] apply to such states as well, so we leave this for future work.

²⁴The assumption being made here is that since the quantum states are now entirely in the possession of Alice and Bob, the test/generation decision can no longer affect the distributed state — if the distributed state could depend on whether it is a test or generation round, the protocol would be trivially insecure.

challenging for current Bell-test implementations. As an alternative, we propose the following potential modification to the DIQKD protocol in [AFRV19] — instead of agreeing on the test rounds via public communication, Alice and Bob could use a small amount of pre-shared key to choose which rounds are test rounds, in the same way as in DIRE (for details on the amount of pre-shared key required, see the DIRE protocols in [AFRV19, BRC20] or the discussion in Sec. 6.2 below). This approach would essentially be a “key expansion” protocol that requires a small amount of pre-shared key to initialize. We remark that this is not a dramatic change in perspective, because a common method to authenticate channels (namely, message authentication codes) relies on having a small amount of pre-shared key, so the assumed existence of authenticated channels in the DIQKD protocol is likely to require some pre-shared key in any case.

However, this basic notion cannot immediately be generalized to Protocol 1 here, since requiring Alice and Bob to choose uniformly distributed “matching” inputs in the generation rounds would require a large amount of pre-shared key (roughly $(1 - \gamma)n$ bits). Fortunately, in the following section, we propose a variation which overcomes this difficulty by “recovering” the entropy in the pre-shared key, thereby still achieving net key expansion.

6.2 Protocol using pre-shared key

Here we describe a variant protocol that avoids the sifting factor *without* requiring the brief quantum storage described above, through the use of a fairly long pre-shared key. The limitation of this variant is that the net increase in secret key is just a (constant) fraction of the amount of pre-shared key; however, the *rate* of net key generation does not have the sifting factor of $1/2$. Informally, the idea is to simply use the pre-shared key as Alice’s input string \mathbf{X} , which allows Bob to choose his generation inputs to match Alice’s. Just as importantly, this also allows them to (almost) entirely omit the public announcement of their inputs — hence \mathbf{X} remains private, and with some care it can be incorporated into the final key without losing the entropy it “contains”.

We now describe this idea in detail as Protocol 2 below, followed by its security proof. The protocol supposes that Alice and Bob hold a pre-shared (uniform) key of n bits, which we shall simply denote as \mathbf{X} , since it will be exactly the string that Alice uses as her inputs. The appropriate value of ℓ_{key} to choose will be described later in Theorem 3.

Protocol 2

This protocol proceeds the same way as Protocol 1, except for the following changes:

- In each round, Alice’s input X_j is determined from the pre-shared key \mathbf{X} , instead of being generated randomly in that round. Bob’s input Y_j is generated as follows: with probability γ he chooses a uniformly random $Y_j \in \{2, 3\}$, otherwise Bob chooses $Y_j = X_j$. In addition, he generates another register Y'_j which equals Y_j when $Y_j \in \{2, 3\}$ and equals \perp otherwise.
 - Alice and Bob do not publicly announce the strings \mathbf{XY} . Instead, Bob only announces the string \mathbf{Y}' .²⁵ Additionally, the sifting step is unnecessary, since there will be no rounds such that $Y_j \in \{0, 1\}$ and $X_j \neq Y_j$.
 - Privacy amplification is performed on the strings \mathbf{AX} and $\tilde{\mathbf{A}}\mathbf{X}$ instead of \mathbf{A} and $\tilde{\mathbf{A}}$.
-

To prove the security of this protocol, we can simply follow almost exactly the same security proof as for Protocol 1, with some changes we shall now describe. Firstly, the value of h_{hon} (to be used when computing EC_{max}) is replaced by

$$\begin{aligned} \tilde{h}_{\text{hon}} = H(A_j|B_jX_jY_j)_{\text{hon}} &= \frac{1-\gamma}{2} \sum_{z \in \{0,1\}} H(A_j|B_j; X_j = Y_j = z)_{\text{hon}} \\ &\quad + \frac{\gamma}{4} \sum_{x \in \{0,1\}, y \in \{2,3\}} H(A_j|B_j; X_j = x, Y_j = y)_{\text{hon}}, \end{aligned} \quad (108)$$

since the probabilities of $X_j = Y_j = z$ for $z \in \{0,1\}$ are now $(1-\gamma)/2$. (Note that no error-correction information needs to be sent from Alice to Bob regarding \mathbf{X} , since both of them have a copy of that string.)

Also, since the strings used in the privacy-amplification step are now \mathbf{AX} and $\tilde{\mathbf{A}}\mathbf{X}$, this means that we need an equivalent of Eq. (71), with $H_{\min}^{\varepsilon_s}(\mathbf{AB}'\mathbf{X}|\mathbf{Y}'E)$ in place of $H_{\min}^{\varepsilon_s}(\mathbf{AB}'|\mathbf{X}\mathbf{Y}E)$. To obtain this, we note that we can simply construct a virtual protocol in the analogous way to Protocol 1', then consider the same EAT channels \mathcal{M}_j as before, but instead we shall identify C'_j with D_j , $A_jB'_jX_j$ with S_j , and Y'_j with T_j in Definition 10. The Markov conditions are again fulfilled, since Y'_j is generated by trusted randomness in each round and independent of all previous data. To find an appropriate min-tradeoff function for these channels, we note that the output $(\mathcal{M}_j \otimes \mathcal{I}_R)(\omega_{R_{j-1}R})$ of channel \mathcal{M}_j always satisfies

$$H(X_j|Y'_jR) = H(X_j) = 1, \quad (109)$$

because X_j is produced by trusted randomness independent of Y'_jR . Therefore, we can use the chain rule to write

$$\begin{aligned} H(A_jB'_jX_j|Y'_jR) &= H(X_j|Y'_jR) + H(A_jB'_j|X_jY'_jR) \\ &\geq 1 + (1-\gamma)r_p(w) + \gamma r_0(w) \\ &:= 1 + \tilde{g}(w), \end{aligned} \quad (110)$$

where the function \tilde{g} (in contrast to g) does not have the factor of 1/2 introduced by sifting, since Alice does not “erase” the outputs of any rounds. We can thus construct a new min-tradeoff function \tilde{f}_{\min} in the same way as in Sec. 4.3.1, but using $1 + \tilde{g}(w)$ in place of $g(w)$.²⁶ The rest of the proof then proceeds as before, leading to the following security statement:

Theorem 3. *Protocol 2 has the same security guarantees as those described in Theorem 1, except with the following changes (with \tilde{g} being defined in (110)):*

- β is chosen to be in $[1 + \tilde{g}(0), 1 + \tilde{g}(1)]$ instead.
- h_{hon} is replaced by \tilde{h}_{hon} as specified in Eq. (108).

²⁵It might be possible to consider a slight variant which omits this step. However, knowing \mathbf{Y}' allows Alice to compute \mathbf{Y} , which may be relevant for error correction since it allows Alice to distinguish the test and generation rounds. In any case, it seems unclear whether the entropy of \mathbf{Y}' can be usefully extracted even if it is kept secret, since Alice does not have access to it in that case.

²⁶We remark that if we think of this replacement as happening in two steps, first replacing g by \tilde{g} and then adding a “constant offset” of 1, then the latter has no effect on $\text{Var}_{\mathcal{Q}_f}(f_{\min})$ or the difference $\text{Max}(f_{\min}) - \text{Min}_{\mathcal{Q}_f}(f_{\min})$, and hence does not change the finite-size correction to the keyrate except indirectly via changing the system dimensions and the range of β . However, the first step of replacing g by \tilde{g} does slightly increase the finite-size correction (since \tilde{g} has a somewhat larger range).

- In Eq. (31) for ℓ_{key} , $g(w_{\text{exp}} - \delta_{\text{tol}})$ is replaced by $1 + \tilde{g}(w_{\text{exp}} - \delta_{\text{tol}})$, and the values of V and K_α are replaced by

$$\begin{aligned}\tilde{V} &:= \sqrt{\text{Var}_{\mathcal{Q}_f}(\tilde{f}_{\min})} + 2 + \log 129, \\ \tilde{K}_\alpha &:= \frac{2^{(\alpha-1)(2\log 8 + \text{Max}(\tilde{f}_{\min}) - \text{Min}_{\mathcal{Q}_f}(\tilde{f}_{\min}))}}{6(2-\alpha)^3 \ln 2} \ln^3 \left(2^{2\log 8 + \text{Max}(\tilde{f}_{\min}) - \text{Min}_{\mathcal{Q}_f}(\tilde{f}_{\min})} + e^2 \right),\end{aligned}\tag{111}$$

where \tilde{f}_{\min} is a function that satisfies

$$\begin{aligned}\text{Max}(\tilde{f}_{\min}) &= 1 + \frac{1}{\gamma} \tilde{g}(1) + \left(1 - \frac{1}{\gamma}\right) \beta, & \text{Min}_{\mathcal{Q}_f}(\tilde{f}_{\min}) &= 1 + \tilde{g}\left(\frac{2 - \sqrt{2}}{4}\right), \\ \text{Var}_{\mathcal{Q}_f}(\tilde{f}_{\min}) &\leq \frac{2 - \sqrt{2}}{4\gamma} \min\{\Delta_0^2, \Delta_1^2\} + \frac{2 + \sqrt{2}}{4\gamma} \max\{\Delta_0^2, \Delta_1^2\}, & \text{where } \Delta_w &:= \beta - 1 - \tilde{g}(w).\end{aligned}\tag{112}$$

Overall, recalling that Protocol 1 required n bits of pre-shared key, we see that the *net* gain of secret key bits in Protocol 2 is larger than that of Protocol 1 by a factor of approximately (ignoring the changes to the finite-size corrections)

$$\frac{n \left(\tilde{g}(w_{\text{exp}} - \delta_{\text{tol}}) - \tilde{h}_{\text{hon}} \right)}{n \left(g(w_{\text{exp}} - \delta_{\text{tol}}) - h_{\text{hon}} \right)} \approx 2,\tag{113}$$

since it avoids the sifting factor. Informally, by keeping \mathbf{X} secret and incorporating it in the privacy amplification step, we have “recovered” the entropy that was present in the original pre-shared key.²⁷

In practice, including the string \mathbf{X} in privacy amplification essentially doubles the input size for the hash function in that step, which raises its computational difficulty substantially (though not insurmountably). One might wonder whether it would be possible to bypass this aspect — for instance, by simply performing privacy amplification on \mathbf{A} and $\tilde{\mathbf{A}}$ as before, then appending \mathbf{X} to the output. At first glance, this approach might appear plausible, since \mathbf{X} is not announced in Protocol 2. Unfortunately, it seems unclear how to certify that the publicly communicated error-correction string \mathbf{L} is independent of \mathbf{X} (in fact, it seems unlikely that this is true). Hence the idea of simply appending \mathbf{X} may not be secure. By instead incorporating it in privacy amplification in the specified manner, Protocol 2 ensures that the entropy of \mathbf{X} is securely “extracted” into the final key.

As previously mentioned, the net increase in secret key given by one instance of Protocol 2 is limited to a fraction of the amount of pre-shared key. However, it is possible in principle to recursively run Protocol 2 in order to achieve unbounded key expansion — one can use the key generated by one instance of Protocol 2 to run it again with a longer pre-shared key and larger n (since the security definition is composable, the soundness parameter will only increase additively in this process [PR14]). We stress that in doing so, one must always incorporate the seed into the privacy-amplification step exactly as specified in Protocol 2 — in particular, this means that the

²⁷Note that the underlying idea here critically relies on the fact that the bound r_p is for the entropy of the output strings *conditioned on* X_j . This allowed us to use the chain rule to obtain Eq. (63), which eventually led to the result that the seed entropy contained in \mathbf{X} simply “adds on” to the entropy of the output strings in the original Protocol 1. If r_p had been a bound on, for instance, $H(A_j|Y'_j R)$ instead of $H(A_j|X_j Y'_j R)$, this argument would not have worked.

entire key changes with every iteration, instead of simply having some new bits appended. Some care is necessary regarding device memory across instances of this recursive procedure — while it does not seem to be directly vulnerable to the memory attack of [BCK13]²⁸, it is still important to ensure that the states measured in each instance of the protocol are independent of the key generated in the preceding instance, since this key is used to choose the device inputs (which must be independent of the state in order for our security arguments to hold). Again, this relies on the notion that the registers measured by the devices do not contain information about the key generated in the preceding instance.

There is another potential variant of this idea where a pre-shared key is instead used to generate *both* input strings \mathbf{X} and \mathbf{Y} , and the input-choice announcement is omitted entirely, with privacy amplification being performed on \mathbf{AXY} and $\tilde{\mathbf{AXY}}$. This can be done by using n bits to set the value of X_j in all rounds as before, then using $\kappa h_2(\gamma)n$ bits to choose Bob’s test rounds approximately according to the desired IID distribution, and finally using $\kappa'\gamma n$ bits to set the value of Y_j in the test rounds (while the generation rounds simply have Y_j set to X_j), where $\kappa, \kappa' > 1$ are constants that can be chosen such that the approximations to the desired distributions are sufficiently accurate (see the randomness-expansion protocols in [AFRV19, BRC20] for a more complete description of this process based on the *interval algorithm*).²⁹ This would hence require $(1 + \kappa h_2(\gamma) + \kappa'\gamma)n$ bits of seed randomness. A similar argument as above could then be performed by noting that (for $X_j Y_j$ generated according to the ideal distribution) we have $H(X_j Y_j) = 1 + h_2(\gamma) + \gamma$, so most of the seed entropy can be “recovered”, up to the losses from the κ, κ' factors. However, tracking the effects of using the interval algorithm to approximate the ideal distribution is cumbersome (albeit possible; see [BRC20]), and it is unclear if this variant offers any immediate advantage over Protocol 2 for DIQKD — though it may be useful for protocols that use other Bell inequalities or non-uniform input distributions, as mentioned in Sec. 4.4.

On the other hand, it appears that this variant may have potential for the purposes of DIRE instead. (As mentioned in the introduction, this idea has also been independently proposed in [BRC21].) The main reason why the random-key-measurement approach in [SGP+21] could not be easily generalized to DIRE is that in order for Alice to select a uniformly random input in every round, she requires a (local) source of n random bits, which is a free resource in DIQKD but not in DIRE — if a proposed DIRE protocol consumes more random bits than it produces, then it has failed to achieve randomness *expansion*³⁰. However, the protocol proposed in this section has the property that it “recovers” the entropy contained in the seed, which means that one can afford to use much larger seeds while still obtaining a net increase in secret key. Explicitly, the application of this idea to DIRE would hence be as follows: one begins with $2n$ uniformly random bits, which are then used as the input strings³¹ \mathbf{XY} to the devices over n rounds to obtain outputs

²⁸This is because the only public communication in Protocol 2 that can leak any information is the error-correction string (all other public communication is based on trusted randomness). In our security proof, we have bounded the min-entropy leakage at this step simply via the length of this string, without any assumptions about its structure, and hence we can still obtain a secure bound on the min-entropy of the input for privacy amplification in the final protocol instance. Note that this claim is strictly restricted to device reuse following the recursive process specified here — once any key bits have been used for any other purpose, the attack again becomes a potential concern if the devices are reused.

²⁹We break up the use of the seed into separate processes because it allows for better efficiency as compared to directly approximating the desired distribution of \mathbf{XY} — with the approach we describe, the “inefficiency” prefactors κ, κ' of the interval algorithm only appear on the $h_2(\gamma)n, \gamma n$ terms instead of the full entropy of \mathbf{XY} .

³⁰There is, however, the related but distinct task of device-independent randomness *generation* [ZSB+20] (where the goal is to produce private randomness from an unbounded supply of public randomness that is independent from the devices), in which this would not be an issue.

³¹For DIRE based on the CHSH inequality, Bob only requires two possible measurements instead of the four

AB. Modelling this process using EAT channels in a manner similar to above (see e.g. [LLR+21] for details), for each round we would have

$$\begin{aligned} H(A_j B_j X_j Y_j | R) &= H(X_j Y_j | R) + H(A_j B_j | X_j Y_j R) \\ &= 2 + H(A_j B_j | X_j Y_j R), \end{aligned} \quad (114)$$

which (given a bound on $H(A_j B_j | X_j Y_j R)$) allows one to bound the smoothed min-entropy of **ABXY** conditioned on E . By performing privacy amplification on **ABXY**, one “recovers” all the entropy in the seed, due to the $H(X_j Y_j | R)$ term in the above equation. Overall, this proposed protocol allows one to use the improved entropy rate provided by the random-key-measurement approach [SGP+21], in the context of DIRE instead of DIQKD.

6.3 Collective attacks

As a reference to compare our results against, we could consider whether a longer secure key could be obtained under the *collective-attacks* assumption, which is the assumption that the device behaviour is IID in each round of the protocol (though Eve can still store quantum information for arbitrary periods).³² To this end, we derive a security proof under this assumption. We defer the proof to Appendix C, and just state the final key length formula here. As compared to Theorem 1, the parameters involved in the formula are somewhat different: some of the previous parameters are no longer involved, though there are some new ones, which we qualitatively describe as follows (they are closely related to the notion of “ ε -secure filtering” described in [Ren05]).

ε_{IID} : Informally, a bound on the probability that the virtual parameter estimation step (in the virtual protocol) accepts when given devices that produce insufficient min-entropy

δ_{IID} : “Relaxation” parameter that slightly enlarges the set of states for which we bound the entropy, in order to derive a nontrivial value of ε_{IID}

The formal theorem statement is:

Theorem 4. *Take any $\varepsilon_{\text{EC}}^{\text{com}}, \varepsilon_{\text{PE}}^{\text{com}}, \varepsilon_{\text{PA}}, \varepsilon_{\text{h}}, \varepsilon_{\text{s}} \in (0, 1]$, $\gamma \in (0, 1)$, $p \in [0, 1/2]$, and $\delta_{\text{IID}} \in [0, w_{\text{exp}} - \delta_{\text{tol}})$. Define*

$$\varepsilon_{\text{IID}} := B_{n, 1 - (w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}})\gamma}(\lfloor (1 - (w_{\text{exp}} - \delta_{\text{tol}})\gamma)n \rfloor). \quad (115)$$

Under the collective-attacks assumption, Protocol 1 is $(\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}})$ -complete and $(\max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_{\text{s}}\} + 2\varepsilon_{\text{h}})$ -sound when performed with any choice of EC_{max} and δ_{tol} such that Eq. (14) and Eq. (30) hold, and ℓ_{key} satisfying

$$\ell_{\text{key}} \leq ng(w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}) - \sqrt{n} (2 \log 5) \sqrt{\log \frac{2}{\varepsilon_{\text{s}}^2} - \text{EC}_{\text{max}}} - \left\lceil \log \left(\frac{1}{\varepsilon_{\text{h}}} \right) \right\rceil - 2 \log \frac{1}{\varepsilon_{\text{PA}}} + 2. \quad (116)$$

required for the DIQKD protocol here.

³²To be precise, we mean that the part of the state held by Alice and Bob’s devices is IID across the rounds, and in each round the devices have the same set of possible measurements. Since all purifications are isometrically equivalent, without loss of generality we can suppose that Eve also holds an IID purification of the Alice-Bob state.

However, the above theorem is simply a statement for Protocol 1 under the assumption of collective attacks, and that protocol does not fully exploit some implications of that assumption. For instance, in Theorem 4 there is implicitly an $O(\gamma)$ subtractive penalty to the keyrates (which was also present in Theorem 1³³) caused by having to include the test-round data in the EC_{\max} term. Yet under the collective-attacks assumption, the test rounds are completely independent of the generation rounds, which implies that the effect of γ should instead be to reduce the keyrate by a *multiplicative* factor of $(1 - \gamma)$. Importantly, in the latter case it is possible to choose arbitrarily large test probabilities γ without necessarily making the keyrates negative, which can dramatically improve the statistical bounds for parameter estimation. To formalize this idea, we consider Protocol 3 below, which attempts to minimize the finite-size correction as much as possible using the most optimistic assumptions that have been discussed thus far.

Protocol 3

This protocol proceeds the same way as Protocol 1, except for the following changes:

- Instead of independently choosing whether each round is a test or generation round, Alice chooses a uniformly random subset of size m as test rounds before the protocol begins, and we define γ as the value m/n . Alice also prepares the strings \mathbf{XY} in advance, by choosing $X_j = Y_j \in \{0, 1\}$ uniformly at random in the generation rounds, and choosing $X_j \in \{0, 1\}, Y_j \in \{2, 3\}$ uniformly at random in the test rounds.
 - In each round, Alice and Bob briefly store their received quantum states instead of immediately measuring them. Alice then publicly announces $X_j Y_j$, which Alice and Bob then use as the inputs to their devices.³⁴
 - In the error-correction step, Alice does not send error-correction data (and a corresponding hash) for the full string \mathbf{A} , but rather only the subset of it consisting of the generation rounds, denoted as \mathbf{A}_g . Bob's guess for this string will be denoted as $\tilde{\mathbf{A}}_g$. The values of \mathbf{A} in the test rounds, denoted as \mathbf{A}_t , are sent directly to Bob without compression or encryption, and Bob uses this string for parameter estimation.
 - Bob's accept condition is instead to check that $\text{hash}(\mathbf{A}_g) = \text{hash}(\tilde{\mathbf{A}}_g)$ and $\text{freq}_{\mathbf{C}_t}(1) \geq w_{\text{exp}} - \delta_{\text{tol}}$ hold, where \mathbf{C}_t denotes the substring of \mathbf{C} corresponding to the test rounds (in particular, this means the frequencies are computed with respect to a string of length γn , not n).
 - Privacy amplification is performed only on the strings \mathbf{A}_g and $\tilde{\mathbf{A}}_g$.
-

³³In fact, Theorem 1 has another $O(\gamma)$ subtractive penalty from the use of the bound (63), but this was due to the technical limitations of the EAT and the fact that the bounds (64) and (67) are slightly suboptimal.

³⁴Here, in our attempt to minimize the finite-size effects, we are following the [AFRV19] assumption mentioned previously: Alice and Bob can briefly store their received quantum states, in a manner such that the public communication cannot affect the stored states.

For this protocol, the value of EC_{\max} is to be computed based only on the number of generation rounds, since error correction is performed on the string \mathbf{A}_g rather than \mathbf{A} . Focusing on the best possible theoretical bounds from Sec. 3.1, this means we take EC_{\max} to be given by Eq. (15) with

$$H_{\max}^{\varepsilon_s}(\mathbf{A}|\mathbf{BXY})_{\text{hon}} \leq (1 - \gamma)nh_{\text{hon}} + \sqrt{(1 - \gamma)n} (2 \log 5) \sqrt{\log \frac{2}{\varepsilon_s^2}}, \quad (117)$$

where

$$h_{\text{hon}} = \sum_{z \in \{0,1\}} \frac{1}{2} H(A_j|B_j; X_j = Y_j = z)_{\text{hon}}, \quad (118)$$

since the test rounds are excluded. With this value of EC_{\max} in mind, we can state the security guarantees of this protocol, with the proof given in Appendix D (we highlight that the dependence on several security parameters here is different as compared to the previous theorems — qualitatively, this is because for instance we no longer consider a virtual protocol when introducing parameters such as ε_{IID} , and also the various binomial distributions are different since the test-round subset now has a fixed size):

Theorem 5. *Take any $\varepsilon_{\text{EC}}^{\text{com}}, \varepsilon_{\text{PE}}^{\text{com}}, \varepsilon_{\text{PA}}, \varepsilon_{\text{h}}, \varepsilon_s \in (0, 1]$, $\gamma \in (0, 1)$, $p \in [0, 1/2]$, and $\delta_{\text{IID}} \in [0, w_{\text{exp}} - \delta_{\text{tol}})$. Define*

$$\varepsilon_{\text{IID}} := B_{\gamma n, 1 - w_{\text{exp}} + \delta_{\text{tol}} + \delta_{\text{IID}}}(\lfloor (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma n \rfloor). \quad (119)$$

Under the collective-attacks assumption, Protocol 3 is $(\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}})$ -complete and $(\max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_{\text{h}})$ -sound when performed with EC_{\max} defined in terms of $\varepsilon_{\text{EC}}^{\text{com}}$ as described above, and $\delta_{\text{tol}}, \ell_{\text{key}}$ satisfying

$$\varepsilon_{\text{PE}}^{\text{com}} \geq B_{\gamma n, w_{\text{exp}}}(\lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor), \quad (120)$$

$$\begin{aligned} \ell_{\text{key}} \leq & (1 - \gamma)nr_p(w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}) - \sqrt{(1 - \gamma)n} (2 \log 5) \sqrt{\log \frac{2}{\varepsilon_s^2}} \\ & - \text{EC}_{\max} - \left\lceil \log \left(\frac{1}{\varepsilon_{\text{h}}} \right) \right\rceil - 2 \log \frac{1}{\varepsilon_{\text{PA}}} + 2. \end{aligned} \quad (121)$$

In Fig. 6, we plot the results of Theorem 5, focusing on the [RBG+17] experiment. This protocol has improved finite-size performance as compared to the original Protocol 1 (under the collective-attacks assumption) due to at least two factors. Firstly, we can potentially use larger γ values, as previously mentioned (some of the points shown in the figure correspond to values ranging up to $\gamma \approx 0.3$). Secondly, we find that for fixed values of $\gamma, n, w_{\text{exp}}, \delta_{\text{tol}}, \delta_{\text{IID}}$, the binomial-distribution bounds for this protocol (Eqs. (119) and (120)) are typically several orders of magnitude better than their counterparts for Protocol 1 (Eqs. (130) and (42)). Intuitively, this arises because in Protocol 1, the number of test rounds is itself a random variable, hence increasing the variance in e.g. the number of rounds where $C_j = 1$. Practically speaking, this means Protocol 1 requires noticeably larger values of δ_{tol} and δ_{IID} in order to achieve given completeness and soundness parameters, hence reducing the keyrate by a nontrivial amount.

However, we see that even with the optimistic assumptions that yield Theorem 5, the keyrate for the estimated experimental parameters we consider only becomes positive at fairly large n . This indicates that substantial further work is necessary in order to achieve a demonstration of positive finite-size keyrates.




Figure 6: Finite-size keyrates as a function of number of rounds in Protocol 3 (see Theorem 5), using honest devices following the estimated parameters in [MDR+19] for the loophole-free Bell test in [HBD+15], for $p = 0$ and $p = 0.03$ (the latter is a rough estimate of the choice of p which yields the highest asymptotic keyrate for these experimental parameters). Note that the latter graph is computed using a heuristic estimate of r_p rather than a certified bound. The error-correction protocol is taken to satisfy Eqs. (15) and (117). The colours correspond to soundness parameters of $\epsilon^{\text{sou}} = 10^{-3}$, 10^{-6} , and 10^{-9} for black, blue, and red respectively, and the completeness parameter is $\epsilon^{\text{com}} = 10^{-2}$ in all cases. The horizontal line denotes the asymptotic keyrate. All other parameters in Theorem 5 were numerically optimized. The required number of rounds to achieve positive keyrate is substantially lower than Protocol 1 (see Fig. 1).

7 Conclusion and further work

In this work, we have performed a finite-size analysis for a protocol that combines several of the most promising approaches towards improving keyrates for DIQKD. Furthermore, we develop an algorithm that computes arbitrarily tight lower bounds on the asymptotic keyrates for protocols of this form (i.e. allowing for noisy preprocessing and random key measurements, but restricted to one-way error correction), which applies to all 2-input 2-output scenarios. This allows us to prove a new threshold of 9.33% for noise tolerance in the depolarizing-noise model, and we show (see Sec. 5.6) that for one-way protocols, any further improvements on this threshold can only be fairly small. Finally, we propose a modified protocol, based on a pre-shared key, that overcomes the disadvantage of the sifting factor in random-key-measurement protocols.

We remark that while the finite-size analysis shown here is for a protocol based on the CHSH inequality, our algorithm for the asymptotic keyrates applies more generally to 2-input 2-output scenarios. If some exploration with this algorithm suggests that an improvement can be obtained by using an inequality other than CHSH, then it would not be difficult to modify the finite-size analysis for such an inequality, as was done in e.g. [BRC20] for DIRE (essentially, it would just correspond to having a different bound r_p).

However, our results show that for the NV-centre experiment in [HBD+15] and the cold-atom experiment in [RBG+17], an impractically large sample size (for those implementations) would still be needed in order to achieve a positive finite-size keyrate, even if one makes the optimistic assumption of collective attacks. A significant question that remains to be addressed is that of photonic experiments [SMSC+15, GVW+15], which achieve lower CHSH values but much larger sample sizes (for instance, the photonic DIRE demonstration in [LLR+21] implemented one run with $n = 1.3824 \times 10^{11}$ over 19.2 hours, and one run with $n = 3.168 \times 10^{12}$ over 220 hours). Un-

fortunately, for photonic models our heuristic results (Sec. 5.5) suggest that incorporating random key measurements is less helpful in improving the keyrate, because the experimental parameters that achieve maximal CHSH value also tend to result in higher error probability (with respect to Bob’s outputs) for one of Alice’s measurements than the other. Despite this challenge, we note that there is much freedom in parameter optimization for photonic experiments [MSS20], and given the algorithm we developed here for bounding the asymptotic keyrates, it is now possible to analyze variants such as choosing different ratios for the two key-generating measurements, or different amounts of noisy preprocessing for each basis. We aim to continue studying this in future work.

We also note that for photonic implementations, a protocol modification termed *random postselection* was recently proposed in [XZZ+22], where it was shown to provide significant improvements in the detection-efficiency thresholds required for positive asymptotic keyrates, under the assumption of collective attacks. However, that modified protocol involves public announcements in each round that violate the Markov condition in the EAT, and hence the EAT cannot be directly applied to prove security of that protocol against general attacks. It remains to be seen whether this can be proven using some other approach, although one challenge to overcome would be the issue that for other protocols with a slightly similar public-announcement structure, it has been shown [TTB+16] that the achievable keyrates against general attacks are strictly lower than for collective attacks. Any approach for proving the security of the [XZZ+22] protocol against general attacks (with the same asymptotic keyrates as for collective attacks) would have to rely on some specific difference between that protocol and the ones analyzed in [TTB+16].

As another extension of our work, our algorithm for bounding the asymptotic DIQKD keyrates also applies to DIRE. Since the results it returns are arbitrarily tight and easily incorporated into the EAT, this could be combined with the finite-size analysis of [BRC20, LLR+21] to improve the results of those works — the proof methods used there yield slightly suboptimal keyrates, in that they either bound the min-entropy rather than the von Neumann entropy [BRC20], or they only bound the entropy of one party’s outputs and are restricted to the CHSH inequality [LLR+21]. Our algorithm would yield tight bounds on the two-party von Neumann entropy for 2-input 2-output Bell inequalities; furthermore, it allows the possibility of using the random-key-measurement approach (by applying the pre-shared key proposal) to improve the keyrates.

Acknowledgements

We are very grateful to Peter Brown for detailed information on the finite-size analysis in [BRC20, LLR+21], as well as Omar Fawzi and Frédéric Dupuis for discussions on the entropy accumulation theorem. We also thank Koon Tong Goh and Ignatius W. Primaatmaja for helpful discussion and feedback. Finally, we thank the reviewers and editor for many helpful suggestions in improving the manuscript.

E. Y.-Z. T. and R. R. are supported by the Swiss National Science Foundation (SNSF) grant number 20QT21.187724 via the National Center for Competence in Research for Quantum Science and Technology (QSIT), the Air Force Office of Scientific Research (AFOSR) via grant FA9550-19-1-0202, and the QuantERA project eDICT. P. S. is supported by the Swiss National Science Foundation (SNSF). C. C.-W. L is supported by the National Research Foundation (NRF) Singapore, under its NRF Fellowship programme (NRFF11-2019-0001) and Quantum Engineering Programme 1.0 (QEP-P2).

Computational platform and code

The min-tradeoff function computations were performed using the MATLAB package YALMIP [Löf04] with the solver MOSEK [MOS19], while optimization of the finite-size keyrates was performed in Mathematica and MATLAB. Some of the calculations reported here were performed using the Euler cluster at ETH Zürich. The code used to compute the certified lower bounds is available at the following URL:

<https://github.com/ernesttyz/qbitent>

References

- [AFRV19] R. Arnon-Friedman, R. Renner, and T. Vidick. [Simple and Tight Device-Independent Security Proofs](#). *SIAM Journal on Computing* 48.1 (2019), pp. 181–225.
- [AGT06] A. Acín, N. Gisin, and B. Toner. [Grothendieck’s constant and local models for noisy entangled quantum states](#). *Physical Review A* 73 (2006), p. 062105.
- [BCK13] J. Barrett, R. Colbeck, and A. Kent. [Memory Attacks on Device-Independent Quantum Cryptography](#). *Physical Review Letters* 110 (2013), p. 010503.
- [BFF21] P. Brown, H. Fawzi, and O. Fawzi. [Computing conditional entropies for quantum correlations](#). *Nature Communications* 12.1 (2021).
- [BHK05] J. Barrett, L. Hardy, and A. Kent. [No Signaling and Quantum Key Distribution](#). *Physical Review Letters* 95 (2005), p. 010503.
- [BRC20] P. J. Brown, S. Ragy, and R. Colbeck. [A Framework for Quantum-Secure Device-Independent Randomness Expansion](#). *IEEE Transactions on Information Theory* 66.5 (2020), pp. 2964–2987.
- [BRC21] R. Bhavsar, S. Ragy, and R. Colbeck. [Improved device-independent randomness expansion rates from tight bounds on the two sided randomness using CHSH tests](#). *arXiv:2103.07504v2 [quant-ph]* (2021).
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [CMA+13] B. G. Christensen, K. T. McCusker, J. B. Altepeter, B. Calkins, T. Gerrits, A. E. Lita, A. Miller, L. K. Shalm, Y. Zhang, S. W. Nam, N. Brunner, C. C. W. Lim, N. Gisin, and P. G. Kwiat. [Detection-Loophole-Free Test of Quantum Nonlocality, and Applications](#). *Physical Review Letters* 111 (2013), p. 130406.
- [Col06] R. Colbeck. [Quantum And Relativistic Protocols For Secure Multi-Party Computation](#). *arXiv:0911.3814v2 [quant-ph]* (2006).
- [Col12] P. J. Coles. [Unification of different views of decoherence and discord](#). *Physical Review A* 85.4 (2012), p. 042103.
- [DF19] F. Dupuis and O. Fawzi. [Entropy accumulation with improved second-order term](#). *IEEE Transactions on Information Theory* (2019), pp. 1–1. ISSN: 0018-9448.
- [DFR20] F. Dupuis, O. Fawzi, and R. Renner. [Entropy Accumulation](#). *Communications in Mathematical Physics* 379.3 (2020), pp. 867–913.

- [DW05] I. Devetak and A. Winter. [Distillation of secret key and entanglement from quantum states](#). *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 461.2053 (2005), pp. 207–235.
- [Ebe93] P. H. Eberhard. [Background level and counter efficiencies required for a loophole-free Einstein-Podolsky-Rosen experiment](#). *Physical Review A* 47 (1993), R747–R750.
- [FW56] M. Frank and P. Wolfe. [An algorithm for quadratic programming](#). *Naval Research Logistics Quarterly* 3.1-2 (1956), pp. 95–110.
- [GMR+13] M. Giustina, A. Mech, S. Ramelow, B. Wittmann, J. Kofler, J. Beyer, A. Lita, B. Calkins, T. Gerrits, S. W. Nam, R. Ursin, and A. Zeilinger. [Bell violation using entangled photons without the fair-sampling assumption](#). *Nature* 497.7448 (2013), pp. 227–230.
- [GVW+15] M. Giustina, M. A. M. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-A. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, and A. Zeilinger. [Significant-Loophole-Free Test of Bell’s Theorem with Entangled Photons](#). *Physical Review Letters* 115 (2015), p. 250401.
- [HBD+15] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiou, and R. Hanson. [Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres](#). *Nature* 526 (2015), pp. 682–686.
- [HQV+17] F. Hirsch, M. T. Quintino, T. Vértesi, M. Navascués, and N. Brunner. [Better local hidden variable models for two-qubit Werner states and an upper bound on the Grothendieck constant \$K_G\(3\)\$](#) . *Quantum* 1 (2017), p. 3. ISSN: 2521-327X.
- [HST+20] M. Ho, P. Sekatski, E. Y.-Z. Tan, R. Renner, J.-D. Bancal, and N. Sangouard. [Noisy Preprocessing Facilitates a Photonic Realization of Device-Independent Quantum Key Distribution](#). *Physical Review Letters* 124.23 (2020).
- [JMS20] R. Jain, C. A. Miller, and Y. Shi. [Parallel Device-Independent Quantum Key Distribution](#). *IEEE Transactions on Information Theory* 66.9 (2020), pp. 5567–5584. ISSN: 1557-9654.
- [Kri79] J. Krivine. [Constantes de Grothendieck et fonctions de type positif sur les sphères](#). *Advances in Mathematics* 31.1 (1979), pp. 16–30.
- [LLR+21] W.-Z. Liu, M.-H. Li, S. Ragy, S.-R. Zhao, B. Bai, Y. Liu, P. J. Brown, J. Zhang, R. Colbeck, J. Fan, Q. Zhang, and J.-W. Pan. [Device-independent randomness expansion against quantum side information](#). *Nature Physics* 17.4 (2021), pp. 448–451.
- [Löf04] J. Löfberg. [YALMIP : A toolbox for modeling and optimization in MATLAB](#). *Proceedings of the CACSD Conference*. Taipei, Taiwan, 2004.
- [LZL+18] Y. Liu, Q. Zhao, M.-H. Li, J.-Y. Guan, Y. Zhang, B. Bai, W. Zhang, W.-Z. Liu, C. Wu, X. Yuan, H. Li, W. J. Munro, Z. Wang, L. You, J. Zhang, X. Ma, J. Fan, Q. Zhang, and J.-W. Pan. [Device-independent quantum random-number generation](#). *Nature* 562.7728 (2018), pp. 548–551.

- [MDR+19] G. Murta, S. B. van Dam, J. Ribeiro, R. Hanson, and S. Wehner. Towards a realization of device-independent quantum key distribution. *Quantum Science and Technology* 4.3 (2019), p. 035011.
- [ML12] X. Ma and N. Lütkenhaus. Improved Data Post-Processing in Quantum Key Distribution and Application to Loss Thresholds in Device Independent QKD. *Quantum Information and Computation* 12.3–4 (2012), pp. 203–214. ISSN: 1533–7146.
- [MOS19] MOSEK ApS. The MOSEK optimization toolbox for MATLAB manual. Version 8.1. 2019.
- [MSS20] A. A. Melnikov, P. Sekatski, and N. Sangouard. Setting Up Experimental Bell Tests with Reinforcement Learning. *Physical Review Letters* 125 (2020), p. 160401.
- [NSPS14] O. Nieto-Silleras, S. Pironio, and J. Silman. Using complete measurement statistics for optimal device-independent randomness evaluation. *New Journal of Physics* 16.1 (2014), p. 013035.
- [PAB+09] S. Pironio, A. Acín, N. Brunner, N. Gisin, S. Massar, and V. Scarani. Device-independent quantum key distribution secure against collective attacks. *New Journal of Physics* 11.4 (2009), p. 045021. ISSN: 1367-2630.
- [PAM+10] S. Pironio, A. Acín, S. Massar, A. B. de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by Bell’s theorem. *Nature* 464.7291 (2010), pp. 1021–1024.
- [PR14] C. Portmann and R. Renner. Cryptographic security of quantum key distribution. *arXiv:1409.3525v1 [quant-ph]* (2014).
- [RBG+17] W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortegel, M. Rau, and H. Weinfurter. Event-Ready Bell Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes. *Physical Review Letters* 119.1 (2017).
- [Ren05] R. Renner. Security of Quantum Key Distribution. ETH Zürich, 2005.
- [RR12] J. M. Renes and R. 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 (2012), pp. 1985–1991. ISSN: 0018-9448.
- [RW05] R. Renner and S. Wolf. Simple and Tight Bounds for Information Reconciliation and Privacy Amplification. *Advances in Cryptology — ASIACRYPT 2005*. Springer, 2005, pp. 199–216.
- [SBPC+09] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev. The security of practical quantum key distribution. *Reviews of Modern Physics* 81 (2009), pp. 1301–1350.
- [SBV+21] P. Sekatski, J.-D. Bancal, X. Valcarce, E. Y.-Z. Tan, R. Renner, and N. Sangouard. Device-independent quantum key distribution from generalized CHSH inequalities. *Quantum* 5 (2021), p. 444.
- [Sca13] V. Scarani. The device-independent outlook on quantum physics (lecture notes on the power of Bell’s theorem). *arXiv:1303.3081v4 [quant-ph]* (2013).
- [SGP+21] R. Schwonnek, K. T. Goh, I. W. Primaatmaja, E. Y.-Z. Tan, R. Wolf, V. Scarani, and C. C.-W. Lim. Device-independent quantum key distribution with random key basis. *Nature Communications* 12.1 (2021).

- [SLT+18] L. Shen, J. Lee, L. P. Thinh, J.-D. Bancal, A. Cerè, A. Lamas-Linares, A. Lita, T. Gerrits, S. W. Nam, V. Scarani, and C. Kurtsiefer. [Randomness Extraction from Bell Violation with Continuous Parametric Down-Conversion](#). *Physical Review Letters* 121 (2018), p. 150402.
- [SMSC+15] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici, A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam. [Strong Loophole-Free Test of Local Realism](#). *Physical Review Letters* 115 (2015), p. 250402.
- [SR08] V. Scarani and R. Renner. [Security Bounds for Quantum Cryptography with Finite Resources](#). *Theory of Quantum Computation, Communication, and Cryptography*. Ed. by Y. Kawano and M. Mosca. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 83–95. ISBN: 978-3-540-89304-2.
- [TCR09] M. Tomamichel, R. Colbeck, and R. Renner. [A Fully Quantum Asymptotic Equipartition Property](#). *IEEE Transactions on Information Theory* 55.12 (2009), pp. 5840–5847. ISSN: 0018-9448.
- [TL17] M. Tomamichel and A. Leverrier. [A largely self-contained and complete security proof for quantum key distribution](#). *Quantum* 1 (2017), p. 14. ISSN: 2521-327X.
- [TMMP+17] M. Tomamichel, J. Martinez-Mateo, C. Pacher, and D. Elkouss. [Fundamental finite key limits for one-way information reconciliation in quantum key distribution](#). *Quantum Information Processing* 16.11 (2017).
- [Tom16] M. Tomamichel. *Quantum Information Processing with Finite Resources*. Springer International Publishing, 2016.
- [TSG+21] E. Y.-Z. Tan, R. Schwonnek, K. T. Goh, I. W. Primaatmaja, and C. C.-W. Lim. [Computing secure key rates for quantum cryptography with untrusted devices](#). *npj Quantum Information* 7.158 (2021).
- [TTB+16] L. P. Thinh, G. de la Torre, J.-D. Bancal, S. Pironio, and V. Scarani. [Randomness in post-selected events](#). *New Journal of Physics* 18.3 (2016), p. 035007.
- [TWF+18] Y. Tsujimoto, K. Wakui, M. Fujiwara, K. Hayasaka, S. Miki, H. Terai, M. Sasaki, and M. Takeoka. [Optimal conditions for the Bell test using spontaneous parametric down-conversion sources](#). *Physical Review A* 98 (2018), p. 063842.
- [VDT+13] A. Vitanov, F. Dupuis, M. Tomamichel, and R. Renner. [Chain Rules for Smooth Min- and Max-Entropies](#). *IEEE Transactions on Information Theory* 59.5 (2013), pp. 2603–2612.
- [VV14] U. Vazirani and T. Vidick. [Fully Device-Independent Quantum Key Distribution](#). *Physical Review Letters* 113 (2014), p. 140501.
- [WAP21] E. Woodhead, A. Acín, and S. Pironio. [Device-independent quantum key distribution with asymmetric CHSH inequalities](#). *Quantum* 5 (2021), p. 443.
- [WLC18] A. Winick, N. Lütkenhaus, and P. J. Coles. [Reliable numerical key rates for quantum key distribution](#). *Quantum* 2 (2018), p. 77.
- [WTH+11] S. Winkler, M. Tomamichel, S. Hengl, and R. Renner. [Impossibility of Growing Quantum Bit Commitments](#). *Physical Review Letters* 107 (2011), p. 090502.

- [XZZ+22] F. Xu, Y.-Z. Zhang, Q. Zhang, and J.-W. Pan. [Device-Independent Quantum Key Distribution with Random Postselection](#). *Physical Review Letters* 128 (2022), p. 110506.
- [ZKB18] Y. Zhang, E. Knill, and P. Bierhorst. [Certifying quantum randomness by probability estimation](#). *Physical Review A* 98 (2018), p. 040304.
- [ZS13] A. M. Zubkov and A. A. Serov. [A Complete Proof of Universal Inequalities for the Distribution Function of the Binomial Law](#). *Theory of Probability & Its Applications* 57.3 (2013), pp. 539–544.
- [ZSB+20] Y. Zhang, L. K. Shalm, J. C. Bienfang, M. J. Stevens, M. D. Mazurek, S. W. Nam, C. Abellán, W. Amaya, M. W. Mitchell, H. Fu, C. A. Miller, A. Mink, and E. Knill. [Experimental Low-Latency Device-Independent Quantum Randomness](#). *Physical Review Letters* 124 (2020), p. 010505.

A Alternative map to describe noisy preprocessing

For Pauli measurements in the X - Z plane, applying the Pauli Y operator to the state before the measurement causes the measurement outcome to flip. Therefore, it can be seen that if one first subjects the state ρ_{ABE} to the isometry $V_A \otimes \mathbb{I}_{BE}$, where

$$V_A |\psi\rangle_A = \sqrt{1-p} |\psi\rangle_A |0\rangle_T + \sqrt{p} Y |\psi\rangle_A |1\rangle_T, \quad (122)$$

then measures A_x and directly stores the outcome in \hat{A}_x , this produces exactly the same reduced state on the registers $\hat{A}_x E$ as the noisy-preprocessing procedure. Then since $|\hat{\rho}\rangle_{ABET} := (V_A \otimes \mathbb{I}_{BE}) |\rho\rangle_{ABE}$ is a pure state, by the same reasoning as before we have

$$H(\hat{A}_x|E) = D(\hat{\rho}_{ABT} \| \tilde{\mathcal{Z}}_x(\hat{\rho}_{ABT})) = D(\tilde{\mathcal{G}}(\rho_{AB}) \| \tilde{\mathcal{Z}}_x(\tilde{\mathcal{G}}(\rho_{AB}))), \quad (123)$$

where

$$\tilde{\mathcal{G}}(\sigma_{AB}) := (V_A \otimes \mathbb{I}_B) \sigma_{AB} (V_A \otimes \mathbb{I}_B)^\dagger, \quad (124)$$

$$\tilde{\mathcal{Z}}_x(\sigma_{ABT}) := \sum_a (P_{a|x} \otimes \mathbb{I}_{BT}) \sigma_{ABT} (P_{a|x} \otimes \mathbb{I}_{BT}). \quad (125)$$

This description of noisy preprocessing appears to more specifically exploit the reduction to Pauli measurements in the 2-input 2-output scenario, and hence may be less general.

B Alternative continuity bound

In principle we could use the following approach: we first note that an intermediate step in the derivation of (123) is the relation [Col12]

$$H(\hat{A}_x|E) = H(\tilde{\mathcal{Z}}_x(\hat{\rho}_{ABT})) - H(\hat{\rho}_{ABT}), \quad (126)$$

and hence bounding how much $H(\hat{A}_1|E)$ changes in terms of the measurement angle is equivalent to bounding how much $H(\tilde{\mathcal{Z}}_1(\hat{\rho}_{ABT}))$ changes (since the $H(\hat{\rho}_{ABT})$ term is independent of the measurement). We introduce the qubit channel

$$\tilde{\mathcal{Z}}^\theta(\sigma) := \sum_a P_a(\theta) \sigma P_a(\theta), \quad (127)$$

where the projectors $P_a(\theta)$ describe a Pauli measurement in the X - Z plane along angle θ , in which case $\tilde{\mathcal{Z}}_1$ (for measurement angle θ_A) is basically the channel $\tilde{\mathcal{Z}}^{\theta_A}$ applied only to the A system. By the properties of the diamond norm, we have

$$\left\| \tilde{\mathcal{Z}}_A^{\theta_A}(\hat{\rho}_{ABT}) - \tilde{\mathcal{Z}}_A^{\theta_A+\delta}(\hat{\rho}_{ABT}) \right\|_1 \leq \left\| \tilde{\mathcal{Z}}^{\theta_A} - \tilde{\mathcal{Z}}^{\theta_A+\delta} \right\|_\diamond \leq \dim(A) \left\| J(\tilde{\mathcal{Z}}^{\theta_A}) - J(\tilde{\mathcal{Z}}^{\theta_A+\delta}) \right\|_1, \quad (128)$$

where $J(\mathcal{E})$ denotes the (normalized) Choi state of the channel \mathcal{E} . In this case we have $\dim(A) = 2$, and the trace distance between the Choi states can be computed in closed form (exploiting the fact that two of the Bell states are simultaneous eigenstates of the two Choi states). We could then apply the Fannes-Audenaert continuity bound to bound how much $H(\tilde{\mathcal{Z}}_1(\hat{\rho}_{ABT}))$ (and hence $H(\hat{A}_1|E)$) changes.³⁵ Unfortunately, the Fannes-Audenaert bound scales poorly at small values of trace distance, and hence this approach yields a much worse continuity bound than (92) (in fact, the resulting bound has infinite derivative with respect to δ at $\delta = 0$).

³⁵It was also possible to instead bound the change (in terms of trace distance) of the state on registers $\hat{A}_x E$ directly, using a slightly different channel, but it turns out to yield a slightly worse bound.

C Proof of Theorem 4

We observe that the completeness proof from Sec. 4.1 directly applies, since we have not changed the honest behaviour. As for the soundness proof, we follow the argument in Sec. 4.2 up until Eq. (47), upon which we again require a bound on the first term in that equation. To do so, we note that since the device behaviour is identical in every round, the true CHSH winning probability for the devices can be denoted as a constant value w_{true} (though this value is not directly known to Alice and Bob). We can thus define the following exhaustive possibilities for the device behaviour (again, ρ is the state at the end of Protocol 1’):

Case 1: For the state ρ , $\Pr[\Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega'_{\text{PE}}] \leq \varepsilon_s^2$.

Case 2: $w_{\text{true}} < w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}$.

Case 3: Neither of the above are true.

In case 1, the first term of Eq. (47) is immediately bounded by

$$\Pr[\Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega'_{\text{PE}}] \leq \varepsilon_s^2. \quad (129)$$

In case 2, we observe that in each round, the probability of the CHSH game being played and the devices winning is $w_{\text{true}}\gamma$. Therefore, we have

$$\begin{aligned} \Pr[\text{freq}_{\mathbf{c}'}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma] &= \Pr[\text{freq}_{\mathbf{c}'}(-1) < 1 - (w_{\text{exp}} - \delta_{\text{tol}})\gamma] \\ &\leq \Pr[\text{freq}_{\mathbf{c}'}(-1)n \leq \lfloor (1 - (w_{\text{exp}} - \delta_{\text{tol}})\gamma)n \rfloor] \\ &= B_{n, 1-w_{\text{true}}\gamma}(\lfloor (1 - (w_{\text{exp}} - \delta_{\text{tol}})\gamma)n \rfloor) \\ &\leq \varepsilon_{\text{IID}}, \end{aligned} \quad (130)$$

where the last inequality follows from the fact that we have $1 - w_{\text{true}}\gamma > 1 - (w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}})\gamma$, and $B_{n,p}(k) \leq B_{n,p'}(k)$ for $p \geq p'$. (Again, one could obtain a simpler expression using the multiplicative Chernoff bound,

$$\Pr[\text{freq}_{\mathbf{c}'}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma] \leq e^{-\frac{n\gamma\delta_{\text{IID}}^2}{3(w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}})}}, \quad (131)$$

but this gives significantly poorer results.) Overall, this allows us to bound the first term of Eq. (47) in this case by

$$\Pr[\Omega_{\text{g}} \wedge \Omega_{\text{h}} \wedge \Omega'_{\text{PE}}] \leq \Pr[\Omega'_{\text{PE}}] \leq \Pr[\text{freq}_{\mathbf{c}'}(1) \geq (w_{\text{exp}} - \delta_{\text{tol}})\gamma] \leq \varepsilon_{\text{IID}}. \quad (132)$$

As for case 3, we can directly apply the AEP (Corollary 4.10 of [DFR20]) to obtain the following bound in place of Theorem 2, with the slight difference that we have not conditioned on Ω'_{PE} yet:

$$H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{X}\mathbf{Y}E)_{\rho} > ng(w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}) - \sqrt{n}(2 \log 5) \sqrt{\log \frac{2}{\varepsilon_s^2}}. \quad (133)$$

The remainder of the analysis proceeds very similarly to Sec. 4.2. Specifically, since we have $\Pr[\Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}] \geq \varepsilon_s^2$ in this case, we can conclude

$$\begin{aligned} H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{\rho_{\wedge\Omega_g\wedge\Omega_h\wedge\Omega'_{\text{PE}}}} &\geq H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{\rho} \\ &\geq H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYE})_{\rho} - \text{EC}_{\max} - \left\lceil \log \left(\frac{1}{\varepsilon_h} \right) \right\rceil. \end{aligned} \quad (134)$$

Putting this together with Eq. (133), we see that as long as we choose ℓ_{key} satisfying Eq. (116), we will have

$$\frac{1}{2} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{\rho_{\wedge\Omega_g\wedge\Omega_h\wedge\Omega'_{\text{PE}}}} - \ell_{\text{key}} + 2 \right) \geq \log \frac{1}{\varepsilon_{\text{PA}}}. \quad (135)$$

Again noting that $\Pr[\Omega_g \wedge \Omega_h \wedge \Omega'_{\text{PE}}] \geq \varepsilon_s^2$ in this case, the Leftover Hashing Lemma then implies the first term of Eq. (47) is bounded by

$$\begin{aligned} &\frac{1}{2} \left\| \mathcal{M}_{\text{PA}}(\rho_{\wedge\Omega_g\wedge\Omega_h\wedge\Omega'_{\text{PE}}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho_{\wedge\Omega_g\wedge\Omega_h\wedge\Omega'_{\text{PE}}})_{E'} \right\|_1 \\ &\leq 2^{-\frac{1}{2}} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}|\mathbf{XYLE})_{\rho_{\wedge\Omega_g\wedge\Omega_h\wedge\Omega'_{\text{PE}}}}^{-\ell_{\text{key}}+2} \right) + 2\varepsilon_s \\ &\leq \varepsilon_{\text{PA}} + 2\varepsilon_s. \end{aligned} \quad (136)$$

Therefore, we finally conclude that under the collective-attacks assumption, the secrecy condition is satisfied by choosing

$$\varepsilon_{\text{QKD}}^{\text{sec}} = \max\{\varepsilon_s^2, \varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_h = \max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_h. \quad (137)$$

As before, the protocol is ε_h -correct, so we conclude that it is $(\max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + 2\varepsilon_h)$ -sound when performed with ℓ_{key} satisfying Eq. (116).

D Proof of Theorem 5

We first show completeness by following a similar approach as in Sec. 4.1, except that here we do not need a conversion to a virtual parameter estimation procedure, because in this protocol Bob has access to the exact value of \mathbf{A}_t . Explicitly, observe that the abort condition for this protocol is the event $(\text{hash}(\mathbf{A}) \neq \text{hash}(\tilde{\mathbf{A}})) \vee (\text{freq}_{\mathbf{c}_t}(1) < w_{\text{exp}} - \delta_{\text{tol}})$. Again, the probability of this event is upper bounded by the probability that $(\mathbf{A} \neq \tilde{\mathbf{A}}) \vee (\text{freq}_{\mathbf{c}_t}(1) < w_{\text{exp}} - \delta_{\text{tol}})$. For the honest behaviour, the probability that $\mathbf{A} \neq \tilde{\mathbf{A}}$ is again at most $\varepsilon_{\text{EC}}^{\text{com}}$ (by the guarantees of the error-correction procedure), while the probability that $\text{freq}_{\mathbf{c}_t}(1) < w_{\text{exp}} - \delta_{\text{tol}}$ is upper bounded by

$$\begin{aligned} \Pr[\text{freq}_{\mathbf{c}_t}(1) < w_{\text{exp}} - \delta_{\text{tol}}]_{\text{hon}} &\leq \Pr[\text{freq}_{\mathbf{c}_t}(1)\gamma n \leq \lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor]_{\text{hon}} \\ &= B_{\gamma n, w_{\text{exp}}}(\lfloor (w_{\text{exp}} - \delta_{\text{tol}})\gamma n \rfloor). \end{aligned} \quad (138)$$

Hence by similar reasoning as in Sec. 4.1 (although without needing the virtual parameter estimation, so one could also instead just use the union bound here), we see that the probability of the honest protocol aborting is at most $\varepsilon_{\text{EC}}^{\text{com}} + \varepsilon_{\text{PE}}^{\text{com}}$ with $\varepsilon_{\text{PE}}^{\text{com}}$ defined as in (120).

As for soundness, we follow the argument in Sec. 4 up until Eq. (46), though it is also no longer necessary to introduce the virtual parameter estimation. We also need to slightly modify the event definitions (and it is no longer necessary to consider the event Ω'_{PE}):

$$\begin{aligned}
\Omega_g: \mathbf{A}_g &= \tilde{\mathbf{A}}_g \\
\Omega_h: \text{hash}(\mathbf{A}_g) &= \text{hash}(\tilde{\mathbf{A}}_g) \\
\Omega_{\text{PE}}: \text{freq}_{\mathbf{c}_t}(1) &\geq w_{\text{exp}} - \delta_{\text{tol}}
\end{aligned}$$

Again, the accept condition for this protocol can be stated as the event $\Omega_h \wedge \Omega_{\text{PE}}$. For this protocol, we can prove a bound of the form (46) directly, without splitting it into the terms in Eq. (47). To do so, denote the state after the parameter-estimation step in Protocol 3 as ρ , and consider the following exhaustive possibilities:

Case 1: For the state ρ , $\Pr[\Omega_h \wedge \Omega_{\text{PE}}] \leq \varepsilon_s^2$.

Case 2: $w_{\text{true}} < w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}$.

Case 3: Neither of the above are true.

In case 1, the left-hand-side of Eq. (46) is immediately bounded by ε_s^2 . In case 2, we observe that in each of the γn test rounds (note that now we focus only on the test rounds, instead of the full output strings as we did in the previous proofs), the probability of winning the CHSH game is w_{true} . Therefore, we have

$$\begin{aligned}
\Pr[\text{freq}_{\mathbf{c}_t}(1) \geq w_{\text{exp}} - \delta_{\text{tol}}] &= \Pr[\text{freq}_{\mathbf{c}_t}(0) < 1 - w_{\text{exp}} + \delta_{\text{tol}}] \\
&= \Pr[\text{freq}_{\mathbf{c}_t}(0)\gamma n < (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma n] \\
&\leq \Pr[\text{freq}_{\mathbf{c}_t}(0)\gamma n \leq \lfloor (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma n \rfloor] \\
&= B_{\gamma n, 1-w_{\text{true}}}(\lfloor (1 - w_{\text{exp}} + \delta_{\text{tol}})\gamma n \rfloor) \\
&\leq \varepsilon_{\text{IID}}, \tag{139}
\end{aligned}$$

which allows us to bound the left-hand-side of Eq. (46) in this case by

$$\Pr[\Omega_h \wedge \Omega_{\text{PE}}] \leq \Pr[\Omega_{\text{PE}}] = \Pr[\text{freq}_{\mathbf{c}_t}(1) \geq w_{\text{exp}} - \delta_{\text{tol}}] \leq \varepsilon_{\text{IID}}. \tag{140}$$

Finally, for case 3, we note that since Eve's quantum side-information E is genuinely in (IID) tensor-product form under the collective-attacks assumption, we can denote the subsets corresponding to test and generation rounds as \mathbf{E}_t and \mathbf{E}_g respectively. The AEP in [DFR20] then gives

$$H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{X}_g \mathbf{Y}_g \mathbf{E}_g)_\rho > (1 - \gamma)nr_p(w_{\text{exp}} - \delta_{\text{tol}} - \delta_{\text{IID}}) - \sqrt{(1 - \gamma)n} (2 \log 5) \sqrt{\log \frac{2}{\varepsilon_s^2}}, \tag{141}$$

and we proceed similarly: since we have $\Pr[\Omega_h \wedge \Omega_{\text{PE}}] \geq \varepsilon_s^2$ in case 3, we can conclude

$$\begin{aligned}
H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{A}_t \mathbf{X} \mathbf{Y} \mathbf{L} \mathbf{E})_{\rho \wedge \Omega_h \wedge \Omega_{\text{PE}}} &\geq H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{A}_t \mathbf{X} \mathbf{Y} \mathbf{L} \mathbf{E})_\rho \\
&\geq H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{A}_t \mathbf{X} \mathbf{Y} \mathbf{E})_\rho - \text{EC}_{\text{max}} - \left\lceil \log \left(\frac{1}{\varepsilon_h} \right) \right\rceil \\
&= H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{X}_g \mathbf{Y}_g \mathbf{E}_g)_\rho - \text{EC}_{\text{max}} - \left\lceil \log \left(\frac{1}{\varepsilon_h} \right) \right\rceil, \tag{142}
\end{aligned}$$

where the last line holds because the test rounds are independent of the generation rounds. Hence as long as we choose ℓ_{key} satisfying Eq. (121), we will have

$$\frac{1}{2} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{A}_t \mathbf{XYLE})_{\rho \wedge \Omega_h \wedge \Omega_{\text{PE}}} - \ell_{\text{key}} + 2 \right) \geq \log \frac{1}{\varepsilon_{\text{PA}}}. \quad (143)$$

Again noting that $\Pr[\Omega_h \wedge \Omega_{\text{PE}}] \geq \varepsilon_s^2$ in this case, the Leftover Hashing Lemma then implies

$$\begin{aligned} & \frac{1}{2} \|\mathcal{M}_{\text{PA}}(\rho \wedge \Omega_h \wedge \Omega_{\text{PE}})_{K_A E'} - \mathbb{U}_{K_A} \otimes \mathcal{M}_{\text{PA}}(\rho \wedge \Omega_h \wedge \Omega_{\text{PE}})_{E'}\|_1 \\ & \leq 2^{-\frac{1}{2}} \left(H_{\min}^{\varepsilon_s}(\mathbf{A}_g | \mathbf{A}_t \mathbf{XYLE})_{\rho \wedge \Omega_h \wedge \Omega_{\text{PE}}} - \ell_{\text{key}} + 2 \right) + 2\varepsilon_s \\ & \leq \varepsilon_{\text{PA}} + 2\varepsilon_s. \end{aligned} \quad (144)$$

Therefore, we finally conclude that under the collective-attacks assumption, the secrecy condition is satisfied by choosing

$$\varepsilon_{\text{QKD}}^{\text{sec}} = \max\{\varepsilon_s^2, \varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} = \max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\}. \quad (145)$$

As before, the protocol is ε_h -correct, so we conclude that it is $(\max\{\varepsilon_{\text{IID}}, \varepsilon_{\text{PA}} + 2\varepsilon_s\} + \varepsilon_h)$ -sound when performed with ℓ_{key} satisfying Eq. (121).