

Does hiding the post-processing from a memory-limited eavesdropper raise the noise tolerance of QKD? Timelocked advantage distillation and the 20% barrier

Hikaru Wakaura

hikaruwakaura@gmail.com

QuantScape Inc

Research Article

Keywords: Quantum Key Distribution , Quantum Bit Error Rate , Advantage distillation (AD).

Posted Date: October 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11453979/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Does hiding the post-processing from a memory-limited eavesdropper raise the noise tolerance of QKD? Timelocked advantage distillation and the 20% barrier

Hikaru Wakaura^{1,*}

¹*QIRI (Quantum Integrated Research Institute Inc.),
1-16-3 Akasaka, Minato-ku, Tokyo 107-0061, Japan*

(Dated: October 19, 2026)

Advantage distillation (AD) raises the finite-key threshold QBER of decoy-state BB84 from 9.5% to 17.3% [1]; asymptotically two-way post-processing tolerates 20%, and Bae and Acín showed that an adversary who stores her probes until the AD parities are announced attains this limit, while one forced to measure each probe beforehand is held off until the 25% entanglement limit. We ask what a quantum computational timelock on the post-processing announcements can buy for this protocol, treating the parities, the block partition and the bases separately. Hiding the parities buys nothing: we prove, with an explicit measurement basis, that for every block size the best measurement made before the parities are known attains the adaptive Helstrom value, so the key length in the min-entropy accounting is the collective-attack value, 17.66% at 10^{14} pulses against the 17.30% certified by the entropic-uncertainty bound of Ref. [1]. Hiding the partition is different: the protocol of Ref. [1] partitions the sifted key at random after the quantum phase, so an adversary whose memory is shorter than the time between the probes of a block is reduced to role-blind per-probe measurements, for which the Helstrom benchmark is optimal among identical measurements at odd block sizes and gives 20.7–20.8% (20.6% with the post-selection factor charged), with 25% as asymptote. Whether a memory holding many probes without the partition can do better is open; at up to six probes a joint measurement recovers most of the adaptive value, so the gain rests on locality rather than on hiding alone. Hiding the basis does not help BB84 in guessing probability and does not add to AD for six-state at small blocks. We state the open question precisely and record two by-products for implementers of Ref. [1].

I. INTRODUCTION

Quantum key distribution promises keys whose secrecy does not rest on computational assumptions, but the promise is only redeemed when the channel is quiet enough. Every QKD protocol has a maximum QBER above which no key can be certified, and for the most deployed protocol, decoy-state BB84 with one-way post-processing, that maximum is close to 9–11% depending on the source model [1, 2]. Daylight free-space links, fibres shared with bright classical traffic, and long satellite passes can approach this level, so raising the threshold without changing the optics is one of the few remaining software levers on QKD range.

Advantage distillation is the classical lever. Introduced by Maurer for classical key agreement [3] and brought into QKD by Gottesman and Lo [4], it post-selects blocks of raw bits whose announced parities agree, so that the surviving bits are far better correlated between Alice and Bob than between Alice and Eve. Chau showed that BB84 tolerates 20.0% and six-state QKD 27.6% asymptotically with iterated two-way post-processing [5]; Acín and co-workers proved these values tight for the repetition-code protocol under collective attacks, with an explicit attack in which Eve stores her probes until the parities are public and then discriminates two product states [6, 7]. A finite-key treatment for sending-or-not-sending twin-field QKD was given by Jiang *et al.* [10],

sharper finite-blocklength entropy bounds for two-way distillation have recently been obtained from hypothesis testing [9], and for decoy-state BB84 the first complete finite-key analysis is that of Treplin, Kleinpaß and Orsucci [1], who proved that a b -bit repetition-code AD step raises the finite-key threshold from 9.5% to 17.3% at $N = 10^{14}$ pulses. Their formula bounds Eve’s smooth min-entropy about the distilled bits through an entropic uncertainty relation on the post-selected state, and it allows the adversary to hold a purification of everything until the last message of the protocol has been announced.

A different line of work asks what changes if the adversary cannot wait. Bae and Acín already observed that a delay between distribution and reconciliation, longer than the decoherence time of Eve’s memory, forces her to measure before the parities are announced, and that for the parity-announcing form of AD this restores the entanglement limit of 25% for BB84 as the security condition, whereas for the form in which Alice announces positions of equal bits it restores nothing [7]. In the cryptographic literature the same assumption is the bounded- and noisy-storage model [11–13], under which two-party primitives become possible. Watanabe, Matsumoto and Uyematsu encrypted the parities of a two-way pre-processing step with a pre-shared one-time pad and obtained a higher net key rate in the Shor–Preskill framework [14], at the price of a pre-shared key whose provenance the analysis leaves open, as Novak later noted [17]. Pastushenko and Kronberg encrypted the error-correction data and argued that, if Eve’s quantum memory decoheres before she learns the codeword set, her

* h.wakaura@qiri.co.jp

information is bounded by the additive accessible information rather than by the Holevo quantity [15]; they did not claim composable security. Novak then proved, in the non-asymptotic framework of Tomamichel and Leverrier [16], that one-time-pad protection of one-way error-correction data drawn from the same session leaves the key rate exactly unchanged at fixed security level [17], and Fung, Ma, Chau and Cai had shown that delaying privacy amplification leaves security and key rate unchanged [18]. Neither encryption nor delay by itself buys anything; whatever gain exists must come from an assumption about the adversary's memory. Vyas and Alléaume made that assumption explicit as the quantum computational hybrid model, in which a computationally secure cipher cannot be broken before a time t_{comp} and any quantum memory decoheres within $t_{\text{coh}} < t_{\text{comp}}$; their MUB-QCT protocol timelocks the basis information of a high-dimensional encoding and achieves everlasting security at rates beyond QKD [19, 20].

This paper asks what the same assumption can buy for the most deployed protocol when the hardware is left untouched and only post-processing announcements are timelocked, and it sharpens the question raised by the remark of Ref. [7]: what exactly must be assumed about the adversary's memory, and about which announcements are hidden, for the parity-announcing form of AD to beat 20%? Under the timelock the adversary must measure before she learns the hidden announcements, so her side information is classical and non-adaptive, and the problem becomes one of state discrimination with post-measurement information [25–27]. Our contributions are: (i) a threat model that lists the announcements of Protocol 1 of Ref. [1] with their timing, states which of them a timelock can hide, and identifies the adversary class that each choice enforces; (ii) Theorem 1: for every block size the best measurement made before the parities are announced equals the adaptive Helstrom measurement, with an explicit basis, so hiding the parities does not change the adversary's guessing probability, and the per-probe reading of the remark of Ref. [7] requires more than a hidden parity; (iii) a finite-key key length that reuses the decoy-state statistics of Ref. [1] and replaces only the adversary term, maximised over Bell-diagonal collective attacks, validated by reproducing all four published thresholds of that work, and evaluated in two accountings whose roles differ: the min-entropy accounting is the collective-attack value and does not depend on the timelock (it certifies 17.66% at $N = 10^{14}$ against 17.30% for the entropic-uncertainty bound), while the Shannon accounting is the only place where a timelock enters, worth at most 0.7 percentage points for block-local measurements; (iv) for a timelock on the block partition, the reduction of an adversary with short memory to role-blind per-probe measurements, for which the Helstrom-plus-tally benchmark is optimal among identical measurements at odd block sizes and gives 20.7–20.8% at $N = 10^{14}$ with the 25% limit of Gisin and Wolf [21] as asymptote, together with semidefinite programmes show-

ing that an adversary who can hold several probes recovers most of the adaptive value even without the partition, so that the gain rests on locality; and (v) numerical evidence that hiding the basis does not add to AD at small block sizes. We state the remaining open question in Sec. VII and record two by-products for implementations of Ref. [1].

II. MODEL AND ASSUMPTIONS

A1 (quantum computational timelock). A cipher E_k shared by Alice and Bob cannot be broken by the adversary before time t_{comp} after its use; any quantum memory available to the adversary decoheres within $t_{\text{coh}} < t_{\text{comp}}$ [19]. The technology assessment in Ref. [19] places t_{coh} of optical quantum memories below one second, whereas t_{comp} for a post-quantum cipher is measured in years; the assumption is only required to hold during the protocol run, which is what makes the resulting key everlasting.

A2 (announcements and their timing). Protocol 1 of Ref. [1] runs for a time T_{run} (10^5 s for $N = 10^{14}$ pulses at 1 GHz) and then announces, in this order: the detection record, the bases and intensities, a random reordering of the sifted strings, the test set, a random partition of the remaining string into blocks of length b , the AD parity strings c_A, c_B , and the reconciliation syndrome. A message m sent at time t_m may be hidden for a delay Δ_m , realised as a plain delay before publication or as encryption under E_k (then $\Delta_m = t_{\text{comp}} - t_m$); an adversary who intercepted a probe at t_{rx} must measure it before $t_{\text{rx}} + t_{\text{coh}}$, so the message is hidden from that measurement iff $\Delta_m > t_{\text{coh}} - (t_m - t_{\text{rx}})$. Since every announcement is made after the run, an adversary with $t_{\text{coh}} \ll T_{\text{run}}$ measures every probe before any of them, whether or not it is additionally hidden; the timelock matters for an adversary with $t_{\text{coh}} \gtrsim T_{\text{run}}$, for whom it decides which announcements she sees before measuring. We therefore distinguish three adversary classes by what is known at measurement time.

A3 (physical layer). As in Ref. [1]: phase-randomised coherent states with three intensities $\mu_1 > \mu_2 + \mu_3$, $\mu_3 \approx 0$, active basis choice, threshold detectors with random assignment of double clicks, and an authenticated classical channel.

A4 (adversary classes). Each pulse interacts with an independent ancilla through a collective attack, which after the Pauli twirl of Appendix B is Bell-diagonal and, in the symmetric case, is the attack of Refs. [21, 22]. Before her memory expires the adversary may apply any measurement to the ancillas she holds; afterwards her side information is classical. *Class J (joint, block-local).* She knows which b ancillas form a block and holds them together, but not the parities: this is the class an adversary with $t_{\text{coh}} \gtrsim T_{\text{run}}$ faces when only the parities and the syndrome are hidden, and it is the class that Theorem 1 settles. *Class W (window).* She does not know the partition when she measures, and her mem-

ory holds only the probes that arrived within a window of t_{coh} . For $t_{\text{coh}} \ll T_{\text{run}}$ the probes of one block fall into distinct windows except for a fraction of order $\binom{b}{2} f t_{\text{coh}} / N$ of the blocks (4×10^{-4} for $b = 9$, $f = 1$ GHz, $t_{\text{coh}} = 1$ s, $N = 10^{14}$), and for a block whose probes are in distinct windows the measurement record factorises into a per-probe instrument that may depend on the outcomes of other probes but not on which of them are its block mates. Up to the feed-forward on unrelated probes, which we neglect, this is the class of role-blind per-probe measurements: identical or randomised single-probe POVMs followed by classical decoding once the partition and the parities are announced. *Class H (Helstrom benchmark)*. Each ancilla is measured by the Helstrom measurement of its two possible states and the outcomes are decoded optimally; this is the model of Ref. [21] and the “individual attack” of Ref. [7]. Class H is a sub-class of W; in Sec. III we show numerically that for odd b it is optimal among identical per-probe measurements, and that for even b it is not.

A5 (statistics). Finite-size fluctuations of the decoy estimates, the block post-selection and the error-rate extrapolation are handled exactly as in Ref. [1]; we also verify the headline numbers with the virtual-i.i.d. reduction of Ref. [10].

The role of the cipher follows from A2. A plain delay $\tau_{\text{ann}} > t_{\text{coh}}$ is sufficient and needs no cryptography, at the price of a key-delivery latency of τ_{ann} . Encryption under a post-quantum cipher removes that latency while keeping the adversary’s delay at t_{comp} . The cipher compresses latency; it does not create the gain, which is the reading of Ref. [15] that survives the equivalence theorem of Ref. [17].

III. THE ADVERSARY’S INFORMATION ON A DISTILLED BIT

Under A1 and A2 the adversary’s optimal strategy is an immediate measurement followed by classical decoding once the delayed messages become available (Ref. [19], Proposition 1). In the symmetric attack each pulse interacts as

$$U|\uparrow\rangle|0\rangle = \sqrt{F}|\uparrow\rangle\psi_{\uparrow} + \sqrt{D}|\downarrow\rangle\varphi_{\uparrow}, \quad (1)$$

$$U|\downarrow\rangle|0\rangle = \sqrt{F}|\downarrow\rangle\psi_{\downarrow} + \sqrt{D}|\uparrow\rangle\varphi_{\downarrow}, \quad (2)$$

(Eqs. (1)–(2) of Ref. [21]), with disturbance D equal to the single-photon QBER and, by the symmetry of the attack, $F = 1 - D = (1 + c_{\varphi})/(2 - c_{\psi} + c_{\varphi})$ (Eq. (3) of Ref. [21]), where $c_{\psi} = \langle\psi_{\uparrow}|\psi_{\downarrow}\rangle$ and $c_{\varphi} = \langle\varphi_{\uparrow}|\varphi_{\downarrow}\rangle$ are real. The ψ and φ states span orthogonal subspaces, so the adversary can always learn whether a bit was disturbed ($Z_1 = X \oplus Y$ in the notation of Ref. [21], p. 2) without affecting anything else. After a b -bit block is accepted and both parity strings are announced, she knows whether the block was accepted with all bits correct (probability $p_u = (1 - D)^b/p_a$, $p_a = (1 - D)^b + D^b$)

or all bits wrong, and in the first case her b ancillas are in the product state $\psi_{a_1} \otimes \cdots \otimes \psi_{a_b}$, where the parities fix all a_i relative to the distilled bit a_1 .

Per-probe measurement (class H). This is the per-probe adversary of Ref. [7] and the model of Ref. [21]: the Helstrom measurement of each ancilla followed by optimal decoding. The Helstrom measurement distinguishes two equiprobable pure states of overlap c with success probability $\frac{1}{2} + \frac{1}{2}\sqrt{1 - c^2}$ [28] (Eq. (3.12) of Ref. [32]), so each ancilla yields a guess of its bit that is correct with probability $\delta_0 = \frac{1}{2} + \frac{1}{2}\sqrt{1 - c_{\psi}^2}$ on undisturbed bits and $\delta_1 = \frac{1}{2} + \frac{1}{2}\sqrt{1 - c_{\varphi}^2}$ on disturbed ones; with the parities the guesses become b votes on a_1 . The tally k of votes for 0 is a sufficient statistic, the majority vote is its maximum-likelihood decision with error $\gamma_b(\delta)$, and the adversary’s information on the distilled bit is

$$I_E^{\text{ind}}(D, b) = \max_{c_{\varphi}} \left[p_u J_b(\delta_0) + (1 - p_u) J_b(\delta_1) \right], \quad (3)$$

$$J_b(\delta) = 1 - \sum_{k=0}^b \Pr(k) h(\Pr(a_1 = 0 | k)),$$

with $\Pr(a_1 = 0 | k) = \delta^k (1 - \delta)^{b-k} / [\delta^k (1 - \delta)^{b-k} + (1 - \delta)^k \delta^{b-k}]$. We stress that $J_b(\delta) \geq 1 - h(\gamma_b(\delta))$ (Fano’s inequality; Appendix C.2): the expression built from the majority-vote error, which is what the asymptotic analysis of Ref. [21] needs, is only a lower bound on the adversary’s information and must not be used in a key-length formula. At the error rates relevant for key generation the maximum is at $c_{\varphi} = 1$, where Eq. (3) of Ref. [21] gives $c_{\psi} = (1 - 3D)/(1 - D)$ (Appendix C.1); close to $D = 1/4$ the maximiser moves below $c_{\varphi} = 1$ (for $b = 2$ at $D \geq 0.206$, for $b = 9$ at $D \geq 0.239$), and the numerical results below maximise over c_{φ} throughout. This is the incoherent-attack model of Ref. [21]; its Theorem 1 (Eq. (11) there) states that key agreement by AD is possible iff $D/(1 - D) < 2\sqrt{\delta_0(1 - \delta_0)} = |c_{\psi}|$, which for the 4-state protocol is $D < 1/4$ (Eq. (8) there), and it is the individual-attack security condition for the parity-announcing protocol in Ref. [7].

Joint measurement (class J). If the adversary may measure the b ancillas together, but still before the parities are announced, her task is to discriminate the 2^b product states $\psi_{a_1} \otimes \cdots \otimes \psi_{a_b}$ with the post-measurement information c , which partitions them into 2^{b-1} pairs $\{\psi_a, \psi_{\bar{a}}\}$ [25, 27]. Her success probability is the value of the semidefinite programme

$$P_b^{\text{joint}} = \max_{\{M_f\}} 2^{-b} \sum_a \text{Tr}[\rho_a E_a], \quad (4)$$

$$E_a = \sum_{f: f(c(a))=a_1} M_f, \quad M_f \geq 0, \quad \sum_f M_f = \mathbb{I},$$

over POVMs indexed by the functions $f: \{0, 1\}^{b-1} \rightarrow \{0, 1\}$ that assign a guess to every possible parity string;

the programme is the standard semidefinite formulation of minimum-error discrimination [29] with the outcome alphabet enlarged to the functions f [25]. It is upper-bounded by the adaptive value $P_b^{\text{ad}} = \frac{1}{2} + \frac{1}{2}\sqrt{1 - c_\psi^{2b}}$, the Helstrom probability for the two product states of a known pair, since the overlap of two product states is the product of the single-copy overlaps, c_ψ^b (this is Eq. (17) of Ref. [6] and Eq. (30) of Ref. [7]); this is the attack of Refs. [6, 7], which an adversary with memory implements after the announcement by an adaptive sequence of single-copy measurements [8]. The bound is attained for every b :

Theorem 1. Let ψ_0, ψ_1 be pure states with real overlap c , $|c| < 1$, and write $\Psi_a = \psi_{a_1} \otimes \cdots \otimes \psi_{a_b}$. In a basis with $\psi_1 = Z\psi_0$ let $\Psi = \psi_0^{\otimes b}$, $G = Z^{\otimes b}$, $u \propto (1 + G)\Psi$, $v \propto (1 - G)\Psi$ normalised, and $g_\pm = (u \pm v)/\sqrt{2}$. For $y' \in \{0, 1\}^{b-1}$ put $\tilde{y} = (1 \oplus |y'|, y')$, let L be the prefix-parity map $(Ly')_j = \bigoplus_{i \leq j} y'_i$, and $P_{y'} = X^{\tilde{y}} Z^{(0, Ly')}$. Then $\{P_{y'} g_+, P_{y'} g_-\}_{y'}$ is an orthonormal basis, and the projective measurement in it, read after the parities with the rule “ $a_1 = 0$ iff the outcome sign is $(-1)^{y' \cdot (Ly' \oplus c)}$ ”, guesses a_1 with probability $P_b^{\text{ad}} = \frac{1}{2} + \frac{1}{2}\sqrt{1 - c^{2b}}$ for every parity class c .

The proof (Appendix A) is short: the basis is a Pauli orbit of the class-0 Helstrom pair, orthonormality follows from $\langle \psi_0 | XZ | \psi_0 \rangle = 0$ for real states together with a property of the prefix-parity map, and the Helstrom projector of every class is a sum of basis projectors because $X^{\tilde{y}}$ anticommutes with G and commutes with the class translations $Z^{(0, c)}$ up to a sign. In the picture in which the probes carry bit flips, the outcome $(\tilde{y}, s)$ is a guess of the odd-weight parity $\tilde{y} \cdot a$, which distinguishes a from $\bar{a}$ for every class; the post-measurement information only tells the adversary how to translate that parity into a_1 . Independently, solving the semidefinite programme for $b = 2, 3, 4$ and $D \in [0.05, 0.24]$ gives $P_b^{\text{joint}} = P_b^{\text{ad}}$ to within 10^{-7} (Table II), and the construction of Theorem 1 reproduces P_b^{ad} to 10^{-14} for $b \leq 9$ in our implementation. Hiding the parities therefore costs the adversary nothing in guessing probability; the compatibility of the class-wise optimal measurements that the theorem establishes is the property studied in general by Carmeli, Heinosaari and Toigo [26]. For the adversary’s Shannon information we use the fact that for two equiprobable pure states the measurement that maximises the mutual information coincides with the Helstrom measurement, a result of Levitin [31] reproduced in Chapter 3 of Ref. [32] (Eq. (3.263) there gives the accessible information as a function of the overlap), so that $I_E^{\text{joint}}(D, b) = \max_{c_\varphi} [p_u(1 - h(1 - P_b^{\text{ad}}(c_\psi))) + (1 - p_u)(1 - h(1 - P_b^{\text{ad}}(c_\varphi)))]$ bounds the accessible information of any within-block strategy, adaptive or not; the guessing probability P_b^{ad} bounds the min-entropy version without further assumptions.

Is the Helstrom benchmark the best per-probe measurement? Class W allows any identical single-probe POVM. The probe states span a qubit, so an ex-

tremal POVM has at most four rank-one outcomes, and the block information for b identical measurements is a function of the single-probe outcome statistics (Appendix C.2). Maximising it numerically over real four-outcome POVMs (`perprobe.opt.py`) reproduces the Helstrom-plus-tally values to 10^{-6} in both J_b and the guessing probability for $b = 3, 5, 7, 9$ at $D \in [0.10, 0.23]$, whereas for even b a tilted measurement that avoids the tie of the majority vote does better in guessing probability (0.8454 against 0.8143 at $b = 2$, $D = 0.10$; 0.9246 against 0.9093 at $b = 4$) and marginally in information. Role-dependent measurements, which do better still at even b , are not available in class W because the roles are fixed by the partition. We therefore use the Helstrom benchmark for odd b only.

What hiding the partition is worth to an adversary with memory. If the adversary can hold n probes but does not know which of them will share a block, the partition joins the parities as post-measurement information. For $b = 2$ we solved the corresponding semidefinite programme, with the block containing probe 1 chosen uniformly among the $n - 1$ candidates (`hidden.partition.sdp.py`). The optimal success probability is 0.8938 at $n = 3$, 0.8913 at $n = 4$, 0.8897 at $n = 5$, 0.8884 at $n = 6$ at $D = 0.10$, against the adaptive value 0.8981 and the per-probe benchmark 0.8143 (Table III). An adversary holding six probes without the partition thus loses less than one percentage point of guessing probability, and the loss grows slowly with n . Hiding the partition therefore protects only through locality: the probes of a block must lie in different memory windows, which is what the random partition of Ref. [1] combined with $t_{\text{coh}} \ll T_{\text{run}}$ provides.

The consequence for the asymptotic threshold is immediate (Appendix C.3). Bob’s error on an accepted block is $D^b/p_a \approx (D/(1 - D))^b$. The adversary’s error is $\approx |c_\psi|^b$ for individual measurement (the Chernoff exponent of the majority vote, $2\sqrt{\delta_0(1 - \delta_0)} = |c_\psi|$) and $\approx c_\psi^{2b}/4$ for joint measurement (expanding P_b^{ad}). Key agreement by the repetition code requires the adversary’s error to decay more slowly than Bob’s, the argument of the proof of Theorem 1 of Ref. [21], i.e. $|c_\psi| > D/(1 - D)$ in the first case and $c_\psi^2 > D/(1 - D)$ in the second. With $c_\psi = (1 - 3D)/(1 - D)$ at $c_\varphi = 1$ these conditions read $D < 1/4$ and $10D^2 - 7D + 1 > 0$, i.e. $D < 1/5$. The second condition is the security condition of Ref. [7] (their Eq. (68) with $d = 2$), which that work derived for an adversary who measures after the announcement; Theorem 1 shows that the same condition governs an adversary who measures before it. The same algebra for the six-state protocol, whose symmetric attack has $c_\varphi = 0$ (Ref. [21], p. 2) and hence $c_\psi = (1 - 2D)/(1 - D)$ by Eq. (3) of that work, gives $5D^2 - 5D + 1 > 0$, i.e. $D < (5 - \sqrt{5})/10 = 27.6\%$ [5, 6]. The 25% limit of Ref. [21] is thus the limit for an adversary whose memory is shorter than the spacing between the pulses of a block, not for one whose memory is shorter than the announcement delay.

General Bell-diagonal attacks. The symmetric attack has equal error rates in the two bases. After a random Pauli twirl applied and announced by Alice and Bob, any collective attack on the single-photon rounds is Bell-diagonal [34]: Alice, Bob and the adversary share $\sum_i \sqrt{\lambda_i} |\beta_i\rangle |e_i\rangle$ with the four Bell states β_i , orthonormal e_i , and weights with $\phi_Z = \lambda_3 + \lambda_4$ and $\phi_X = \lambda_2 + \lambda_4$. Conditioned on Alice's and Bob's Z outcomes the adversary holds $\psi_a \propto \sqrt{\lambda_1} e_1 + (-1)^a \sqrt{\lambda_2} e_2$ when the bit was undisturbed and $\varphi_a \propto \sqrt{\lambda_3} e_3 + (-1)^a \sqrt{\lambda_4} e_4$ when it was flipped, so the structure of the symmetric attack is preserved with $c_\psi = (\lambda_1 - \lambda_2)/(1 - \phi_Z)$ and $c_\varphi = (\lambda_3 - \lambda_4)/\phi_Z$, and Theorem 1 applies. Eliminating $\lambda_4 = \phi_Z(1 - c_\varphi)/2$,

$$c_\psi = \frac{1 - 2\phi_X - \phi_Z c_\varphi}{1 - \phi_Z}, \quad (5)$$

$$0 \leq \lambda_4 \leq \min(\phi_X, \phi_Z), \quad \lambda_1 = 1 - \phi_Z - \phi_X + \lambda_4 \geq 0,$$

which reduces to $c_\psi = (1 - 2D - Dc_\varphi)/(1 - D)$ at $\phi_X = \phi_Z = D$. The derivation is given in Appendix C.4. The adversary's information $I_E(\phi_Z, \phi_X, b)$ is the maximum over c_φ of the expressions above with this c_ψ ; it increases with ϕ_X , because a larger phase error lets her choose a smaller $|c_\psi|$, but it is not monotone in ϕ_Z at every ϕ_X , so the key length below takes the supremum over the true Z error below its estimate.

Hiding the basis as well. A protocol that keys from both bases and delays the sifting announcement beyond t_{coh} confronts the adversary with the basis as additional post-measurement information. Without advantage distillation ($b = 1$) this is the memoryless-adversary setting solved by Bocquet, Alléaume and Leverrier [24], who optimised Eve's interaction and POVM and found thresholds of 15.4% for BB84 (against 14.6% with memory), 17.6% for SARG04 and 20.4% for six-state (against 15.6%); in guessing probability the gain is confined to protocols with more than two bases for the probe states considered here (our semidefinite programme below), while in Shannon information Ref. [24] still finds a gain of 0.8 percentage points for BB84. We extended the semidefinite programme to hidden bases and parities, with Eve's probe states conditioned on Bob's outcome and the accepted-block mixture of all-correct and all-flipped blocks. For BB84 the worst-case attack ($c_\varphi = 1$) gives $P^{\text{joint}} = P^{\text{ad}}$ at $b = 1$ and $b = 2$ for every D ; the gap that exists at $c_\varphi = 0$ is irrelevant because the adversary chooses c_φ . For six-state at $b = 1$ we reproduce the qualitative gain of Ref. [24] (e.g. $P^{\text{joint}} = 0.751$ against $P^{\text{ad}} = 0.835$ at $D = 0.168$). But as soon as blocks of two bits are formed the gain collapses: for blocks whose two pulses share an unknown basis, $P^{\text{joint}} = P^{\text{ad}}$ to 10^{-6} (0.888644 against 0.888642 at $D = 0.167$, verified with two solvers). For the realistic blocks with independently chosen hidden bases, where the full programme has 2^{18} elements, we bracket the adversary's success: the relaxation in which only the bases are hidden and the parities are known gives $P^{\text{joint}} \leq 0.8878$ at $D = 0.167$ and




FIG. 1. Adversary's uncertainty per distilled bit, $1 - I_E$, versus QBER for block sizes $b = 1, 3, 5, 9, 15$: individual measurement (solid, Eq. (3) with tally decoding), joint block measurement (dash-dotted, adaptive Helstrom value), and the entropic-uncertainty term of Ref. [1] (dashed, same colour). Vertical lines: 20% and 25%.

≤ 0.9556 at $D = 0.259$, and an explicit 16-outcome projective measurement with both hidden, found by first-order optimisation and therefore possibly short of the true optimum, achieves $P \geq 0.8826$ and 0.9395 , against $P^{\text{ad}} = 0.8886$ and 0.9583 . Hiding both therefore costs the adversary between 0.0008 and 0.006 in guessing probability at $D = 0.167$ and between 0.003 and 0.019 at $D = 0.259$, already at $b = 2$, against 0.084 at $b = 1$ and $D = 0.168$; in terms of her error probability the ratio to the adaptive value falls from 1.5 at $b = 1$ to at most 1.05 at $b = 2$ and at most 1.04 at $b = 3$ at $D = 0.167$; at $D = 0.259$, close to the six-state limit, our achievable bounds are looser and only give 1.46 and 1.43 at $b = 2, 3$, while the relaxation shows that the true ratio at $b = 2$ is at least 1.06. The construction of Theorem 1 does not extend to hidden bases, because the class translations are then the basis-change unitaries on the adversary's probe rather than Pauli operators, which is why a residual gain survives at finite b . Hiding the basis and hiding the parities therefore do not compound: the one-way memoryless gain of Ref. [24] (20.4% for six-state) is dominated by advantage distillation without any assumption (27.6%), and adding the timelock to advantage distillation is worth at most 0.03 bits of min-entropy per distilled bit at $b = 2$. The threshold is set by the error exponents of Bob and of the adversary; whether the error ratio stays bounded for all b is open, the trend from $b = 1$ to $b = 3$ being evidence rather than proof, so we make no claim about the six-state threshold with hidden bases.

Fig. 1 compares $1 - I_E$ in the two models with the entropic-uncertainty term $1 - h(\Phi_{\bar{X}})$ of Ref. [1] at the same D and b .

IV. FINITE-KEY SECURE KEY LENGTH

The protocol is Protocol 1 of Ref. [1] with two modifications: the parity strings c_A, c_B and the reconciliation syndrome are sent under E_k (or after τ_{ann}), and privacy amplification uses the bound below. Let $S_{\bar{K}}^-$ be the lower bound on the number of accepted blocks consisting only of single-photon events (Eq. (18) of Ref. [1]), ϕ_Z^+ and ϕ_X^+ the upper bounds on the single-photon error rates (Eqs. (17a)–(17b) there), $N_{\bar{K}}$ and $\Phi_{\bar{K}}$ the expected number and error rate of the distilled bits (Eq. (20) there), and $\lambda_{\text{IR}} = f N_{\bar{K}} h(\Phi_{\bar{K}})$ the reconciliation leakage with efficiency f (Sec. 2 of Ref. [1]). Then

$$\begin{aligned} \ell_{\text{QCT}} \leq & S_{\bar{K}}^- [1 - \bar{I}_E] \\ & - \Delta_{\text{AEP}}(S_{\bar{K}}^-, \epsilon) - \lambda_{\text{IR}} - \log_2 \frac{2}{\epsilon_{\text{cor}}} \\ & - 4 \log_2 \frac{2^{7/4}}{\epsilon_{\text{sec}} - \epsilon_{\text{sec}}^2/8}, \end{aligned} \quad (6)$$

with $\bar{I}_E = \sup\{I_E(\phi_Z', \phi_X^+, b) : 0 < \phi_Z' \leq \phi_Z^+\}$ for the individual or the joint model, where $\Delta_{\text{AEP}} = 4 \log_2(2 + \sqrt{2}) \sqrt{S \log_2(2/\epsilon^2)}$ is the finite-size correction of the asymptotic equipartition property (Theorem 9 of Ref. [23], with $\eta \leq 2 + \sqrt{2}$ for a binary variable and classical side information, valid for $S \geq \frac{8}{5} \log_2(2/\epsilon^2)$), evaluated with smoothing parameter $\epsilon = \epsilon_{\text{sec}}/4$, and the last two terms are the correctness and leftover-hash terms of Eq. (93) of Ref. [1] (Lemma 4 there, with the statistical failure budget $\bar{\epsilon} = \epsilon_{\text{sec}}^2/368$ chosen in Sec. 6 of that work). The only change with respect to Eq. (93) of Ref. [1] is the replacement of the entropic-uncertainty term $S_{\bar{K}}^- [1 - h(\Phi_{\bar{K}}^+)]$ by $S_{\bar{K}}^- [1 - \bar{I}_E] - \Delta_{\text{AEP}}$; Appendix C.5 traces every term to its source. We keep the parameter-estimation and hashing budget of Ref. [1] unchanged and charge the smoothing to an additional $\epsilon_{\text{sec}}/4$, so the total secrecy parameter of the stated numbers is $5\epsilon_{\text{sec}}/4$. A more conservative variant replaces $S_{\bar{K}}^- [1 - \bar{I}_E] - \Delta_{\text{AEP}}$ by $-S_{\bar{K}}^- \log_2 \bar{p}_{\text{guess}}$ with $\bar{p}_{\text{guess}}$ the same supremum of the adversary's per-block guessing probability ($1 - \gamma_b$ or P_b^{ad} , mixed over the two block types), which needs no equipartition step because the min-entropy of a classical bit given classical side information is $-\log_2 p_{\text{guess}}$ (Theorem 1 of Ref. [30]) and is additive over the independent blocks of a collective attack. Both suprema are taken separately over c_φ and over the true Z error, so that each accounting is evaluated at its own worst case; ϕ_X^+ can be used directly because both quantities increase with ϕ_X .

Three remarks delimit what this formula is. First, the two accountings play different roles. In the min-entropy accounting with P_b^{ad} the bound is the conditional min-entropy of the distilled bits given the adversary's quantum ancillas and all announcements, because the min-entropy of a product of classical-quantum states is additive (it is a semidefinite programme whose feasible solutions tensorise [30]) and P_b^{ad} is her guessing probability

with full quantum memory; this column therefore holds against any measurement, joint across blocks or not, and does not use the timelock at all. It is a collective-attack bound in its own right, and by Theorem 1 the timelock cannot improve it. In the Shannon accounting the term $1 - I_E^{\text{joint}}$ exceeds the Holevo term $1 - \chi$ of a quantum-memory adversary, and this excess is the only place where the timelock enters; it is valid for adversaries whose measurement is a product over blocks, since the asymptotic equipartition property is applied to independent classical outcomes, and we do not claim it against joint measurements across blocks, for which accessible information need not be additive. Second, the adversary term is evaluated over all Bell-diagonal collective attacks, which after the twirl of Appendix B are all collective attacks on the untagged rounds; the extension to coherent attacks is discussed in Appendix B and is not included in the numbers, except in Table III where we charge the post-selection factor to every ϵ it affects. Third, the total secrecy parameter of every number in this paper is $5\epsilon_{\text{sec}}/4$ with $\epsilon_{\text{sec}} = 10^{-9}$, that of Ref. [1] being ϵ_{sec} ; the difference is worth less than one bit of key. Second, if A2 fails, the adversary may hold her ancillas until the parities are public and the entropic-uncertainty term of Ref. [1] must be restored; the formula then coincides with Eq. (93) of that work because every other term is shared. The key length vanishes for $\phi_Z^+ \geq 1/4$ by construction. Within the model, once the adversary's side information is classical the leftover hash lemma applies in its classical form, the trace distance of the final key from uniform given her classical record is bounded by $5\epsilon_{\text{sec}}/4$ conditional on A1 holding during the run, and that trace distance bounds the accessible-information criterion of Ref. [19] through Pinsker's inequality; the key is everlasting in the sense of Ref. [19] because A1 is needed only then.

V. NUMERICAL RESULTS

We implemented Eqs. (6)–(20), (93) and (94) of Ref. [1] and the formula above in a deterministic Python module, optimising $\{\mu_1, \mu_2, p_{\mu_1}, p_{\mu_2}, p_Z, q_T\}$ by the Nelder–Mead method and locating thresholds by bisection in the misalignment angle for a lossless, noiseless channel ($\eta = 1$, $p_{\text{noise}} = 0$, QBER from misalignment only), with $\epsilon_{\text{sec}} = 10^{-9}$, $\epsilon_{\text{cor}} = 10^{-15}$, $f = 1.2$ and $\mu_3 = 0$, the settings of Ref. [1]. The implementation reproduces the four thresholds reported there: 8.67% and 11.37% at $N = 10^8$ (published 8.7% and 11.3%, $b = 3$) and 9.53% and 17.30% at $N = 10^{14}$ (published 9.5% and 17.3%, $b = 9$).

Table III reports two further sets of numbers. The first charges the post-selection factor of Appendix B, $(N+1)^{15}$, to every term it affects: the leftover-hash term grows by $60 \log_2(N+1) \approx 2.8 \times 10^3$ bits, the equipartition term Δ_{AEP} grows by a factor $\sqrt{1460/64} \approx 4.8$ (about 7×10^7 bits at the optimum), and the parameter-estimation intervals, which we do not shrink, would add a further loss; the min-entropy columns barely move,

TABLE I. Threshold QBER (sifted-key QBER ϕ_K), maximised over Bell-diagonal collective attacks, with the reference values of Ref. [1]. The per-probe columns are the Helstrom benchmark of class H, reported for odd b only (for $b = 2$ it is not the best per-probe measurement, Sec. III); the collective min-entropy column does not depend on the timelock; the joint Shannon column is the timelock-dependent accounting for block-local measurements. Secrecy parameter $5\epsilon_{\text{sec}}/4$, no post-selection factor. Replacing the block-count statistic of Ref. [1] by the virtual-i.i.d. Chernoff bound of Ref. [10] changes the $N = 10^{14}$ thresholds by less than 10^{-5} .

N	b	BB84 [T25]	AD [T25]	per-probe H, min-ent	per-probe H, Shannon	collective (joint), min-ent	joint, Shannon
10^8	2	8.67%	10.76%	12.75%	13.89%	9.96%	13.24%
10^8	3		11.37%	13.56%	14.44%	12.01%	13.33%
10^{14}	5	9.53%	16.40%	18.16%	20.18%	15.66%	18.21%
10^{14}	7		17.07%	19.70%	20.91%	16.71%	18.35%
10^{14}	9		17.30%	20.66%	20.81%	17.30%	17.40%
10^{14}	11		16.67%	21.24%	19.20%	17.66%	15.29%
10^{14}	15		13.23%	21.83%	15.45%	17.48%	11.86%

TABLE II. Guessing probability of the distilled bit of an accepted undisturbed block: per-probe Helstrom plus majority vote (for even b ties are broken at random, so $b = 4$ equals $b = 3$; the best identical per-probe measurement does better at even b , Sec. III), the optimal non-adaptive joint measurement (semidefinite programme over 2^{b-1} POVM elements, 256 elements on 16 dimensions at $b = 4$), and the adaptive Helstrom bound. The joint and adaptive columns agree to 10^{-8} for $b = 2, 3$ and to 2×10^{-8} for $b = 4$ at all 21 points computed.

b	D	c_ψ	individual $1 - \gamma_b$	joint (SDP)	adaptive P_b^{ad}
2	0.10	0.778	0.8143	0.8981	0.8981
2	0.20	0.500	0.9330	0.9841	0.9841
3	0.10	0.778	0.9093	0.9412	0.9412
3	0.20	0.500	0.9871	0.9961	0.9961
3	0.24	0.368	0.9964	0.9994	0.9994
4	0.10	0.778	0.9093	0.9653	0.9653
4	0.20	0.500	0.9871	0.9990	0.9990
4	0.24	0.368	0.9964	0.99992	0.99992

TABLE III. Finite-key thresholds with the post-selection factor charged to the smoothing and hashing parameters (rows marked PS), and the no-timelock Holevo accounting (last row).

model ($N = 10^{14}$)	$b = 5$	$b = 7$	$b = 9$
joint Shannon (PS)	—	17.80%	15.86%
collective min-ent (PS)	—	16.67%	17.25%
per-probe H Shannon (PS)	—	20.52%	19.39%
per-probe H min-ent (PS)	—	19.70%	20.61%
Holevo Shannon, no timelock	16.40%	16.76%	15.59%

the Shannon columns drop by 0.6–1.5 percentage points. The second replaces $1 - I_E^{\text{joint}}$ by the Holevo term $1 - \chi$ of a quantum-memory adversary in the same machinery, which is the timelock-free counterpart of the joint Shannon column.

TABLE IV.

n	D	per-probe H	hidden partition, n probes held	adaptive P_2^{ad}
3	0.10	0.8143	0.8938	0.8981
3	0.20	0.9330	0.9825	0.9841
4	0.10	0.8143	0.8913	0.8981
4	0.20	0.9330	0.9813	0.9841
5	0.10	0.8143	0.8897	0.8981
5	0.20	0.9330	0.9804	0.9841
6	0.10	0.8143	0.8884	0.8981
6	0.20	0.9330	0.9797	0.9841

VI. IMPLEMENTATION CONSIDERATIONS

The AD stage adds $b - 1$ exclusive-or operations and one comparison per block, fewer than one operation per raw bit, orders of magnitude below the per-bit cost of iterative LDPC decoding; hardware key-distillation engines [36–38] therefore remain decoder-limited after an AD stage is inserted, whether or not the announcements are delayed. The timelock itself is a buffer of depth τ_{ann} or an authenticated-encryption wrapper under a post-quantum key-encapsulation mechanism. Given Sec. V, neither is worth paying for in order to hide the parities; a timelock on the partition is a buffer of the same depth applied to one further message.

VII. DISCUSSION

Three announcements of the protocol can be hidden, and the results above treat them separately. Hiding the parities and the syndrome is settled by Theorem 1: for an adversary who knows which probes form a block, the parity-conditioned discrimination problems are compatible, so the delay does not change her guessing probability, and the collective-attack min-entropy bound, 17.66% at $N = 10^{14}$, is what any analysis within the model can certify. That this bound exceeds the 17.30% of Ref. [1] is a statement about the entropic-uncertainty technique, not




FIG. 2. Secure key rate ℓ/N versus QBER for $N = 10^8$ (left) and $N = 10^{14}$ (right): BB84 one-way (Ref. [1], Eq. (94)), AD (Ref. [1], Eq. (93)) and the per-probe benchmark in both entropy accountings, for $b = 1, 3, 5, 9$.




FIG. 3. Threshold QBER versus block size b : finite-key values at $N = 10^8$ and 10^{14} for Ref. [1], the per-probe benchmark and the joint model, and the asymptotic single-photon thresholds (60-digit arithmetic); horizontal lines at 20% and 25%.

about the timelock; the only timelock-dependent quantity is the excess of $1 - I_E^{\text{joint}}$ over $1 - \chi$ in the Shannon accounting, worth 0.7 percentage points for block-local measurements and open against joint measurements across blocks.

Hiding the partition is where the question of Ref. [7] lives. That work distinguishes CAD1, in which Alice announces the positions of b equal bits, from CAD2, the parity-announcing form used here and in Ref. [1]; against an adversary with memory both are secure exactly up to 20%, and against an adversary forced to measure before the announcements Ref. [7] notes that CAD1 is still

broken at 20%, because Eve knows in advance that a block carries b copies of one bit [8], whereas CAD2 is held off until the entanglement condition $D < 1/4$. The present work locates the assumption behind the second statement: not that the parities are hidden, which Theorem 1 shows to be irrelevant, but that the adversary cannot bring the probes of a block together before she measures. In Protocol 1 of Ref. [1] the partition is drawn at random after the run, so this holds for any adversary with $t_{\text{coh}} \ll T_{\text{run}}$ (class W), and the per-probe benchmark of Table I is then what she concedes, with the proviso that Helstrom-plus-tally is optimal among identical per-probe measurements only for odd b and that feed-forward on unrelated probes was neglected. The semidefinite programme of Table IV shows that the protection comes from locality rather than from hiding alone: an adversary who holds six probes without the partition still reaches within one percentage point of the adaptive value. *Open question.* For a window adversary who holds the $w = ft_{\text{coh}}$ probes of each window and learns the partition, the bases and the parities afterwards, what are the optimal guessing probability and accessible information of a distilled bit, as a function of w , b and D ? At $w = 1$ the answer is the per-probe benchmark for odd b ; at $w \geq n_K$ it is the adaptive value by Theorem 1; Table IV covers $w \leq 6$ for $b = 2$. If the per-probe value persists for $w \ll n_K/b$, a timelock on the partition raises the finite-key threshold of Ref. [1] from 17.3% to about 20.6%; if a window adversary can approach the adaptive value with w far below the block spacing, it raises nothing. Two further steps would be needed for a proof even in the first case: a bound on the fraction of blocks with two probes in one window, which the random partition makes hypergeometric, and the reduction to coherent attacks, which conflicts with block locality because permu-

tation invariance now holds only at the level of blocks (Appendix B).

Hiding the basis is the third lever. For a Z -keyed protocol it changes nothing (Section 2). For protocols that key from both bases the memoryless-adversary thresholds of Ref. [24] apply without AD; our semidefinite programmes show that for BB84 the worst-case attack is unaffected in guessing probability, that Ref. [24] nevertheless finds 0.8 percentage points in Shannon information, and that for six-state the gain of one-bit blocks shrinks by an order of magnitude at two- and three-bit blocks. Whether it vanishes for large b is open. MUB-QCT [19] is not covered by any of this, since it uses $d + 1$ bases in dimension d and no parity-announcing step; we make no claim about it beyond noting that for qubit protocols the basis and the partition are the only announcements whose hiding is not settled by Theorem 1.

Two observations about the reference formula arose in the course of the work. The bound on the logical phase-error rate of the distilled bits, Eq. (19) of Ref. [1], is proven monotone in the single-photon error rates only on the domain $\phi_X + \phi_Z < 1/2$ (Appendix B of that work), because blocks accepted with all b bits flipped carry no phase error: at $\phi_X = 0.3$ and $b = 7$ the expression falls from 0.499 to 0.25 as ϕ_Z rises from 0.3 to 0.5. An unconstrained optimisation of the protocol parameters leaves the domain by driving the estimated single-photon Z error toward one half through a vanishing test set, and then certifies key at QBERs above the asymptotic limit. Any implementation must enforce $\phi_X^+ + \phi_Z^+ < 1/2$ as an acceptance condition; we do so, and all four thresholds of Ref. [1] lie inside the domain. Separately, the Fano expression $1 - h(\gamma_b)$ built from the majority-vote error, which suffices for the asymptotic threshold of Ref. [21], underestimates a per-probe adversary's Shannon information and must be replaced by the tally information J_b of Sec. III in any finite-key use; the difference is 1.0 percentage points at $N = 10^{14}$, $b = 9$ (21.89% against 20.86% in the symmetric family).

VIII. CONCLUSION

A quantum computational timelock on the post-processing of decoy-state BB84 with advantage distillation can hide the parities, the block partition and the bases. Hiding the parities buys nothing: Theorem 1 gives, for every block size, an explicit measurement made before the parities are known that attains the adaptive Helstrom value, so the adversary's guessing probability is unchanged and the key length in the min-entropy accounting is the collective-attack value, 17.66% at 10^{14} pulses against the 17.30% certified by the entropic-uncertainty bound of Treplin *et al.*; only the Shannon accounting for block-local measurements retains a timelock-dependent excess of 0.7 percentage points. Hiding the partition is the lever that the remark of Bae and Acín requires: an adversary whose memory is shorter

than the time between the probes of a block, which the random partition of Treplin *et al.* makes a sizeable fraction of the run, is reduced to role-blind per-probe measurements, for which the Helstrom benchmark is optimal among identical measurements at odd block sizes and gives 20.7–20.8% at 10^{14} pulses, 20.6% with the post-selection factor charged, with the 25% limit of Gisin and Wolf as asymptote. Whether an adversary who holds many probes without the partition can do better is open; at six probes she recovers most of the adaptive value, so the gain rests on locality. Hiding the basis does not help BB84 in guessing probability and does not add to advantage distillation for six-state at small blocks. The key-length formula of Treplin *et al.* must be used within its stated domain, and the per-probe adversary's information must be computed from the vote tally rather than from the majority-vote error.

Appendix A: Proof of Theorem 1

Choose the qubit basis $|0\rangle \propto \psi_0 + \psi_1$, $|1\rangle \propto \psi_0 - \psi_1$, so that $\psi_0 = \cos\theta|0\rangle + \sin\theta|1\rangle$, $\psi_1 = Z\psi_0$ and $c = \cos 2\theta$; we assume $|c| < 1$. Then $\Psi_a = Z^a\Psi$ with $\Psi = \psi_0^{\otimes b}$, and $\Psi_{\pm} = \frac{1}{2}(1 \pm G)\Psi$ have norms $\|\Psi_{\pm}\|^2 = \frac{1}{2}(1 \pm c^b)$. The class c consists of $Q_c\Psi$ and $GQ_c\Psi$ with $Q_c = Z^{(0,c)}$, and since Q_c commutes with G its Helstrom vectors are $g_{\pm}^{(c)} = Q_c g_{\pm}$. For any projector E with $Eg_{+}^{(c)} = g_{+}^{(c)}$ and $Eg_{-}^{(c)} = 0$ the success probability on class c is $\frac{1}{2}[\langle Q_c\Psi|E|Q_c\Psi\rangle + \langle GQ_c\Psi|1-E|GQ_c\Psi\rangle] = \frac{1}{2}(\|\Psi_{+}\| + \|\Psi_{-}\|)^2 = P_b^{\text{ad}}$, because $Q_c\Psi = \alpha g_{+}^{(c)} + \beta g_{-}^{(c)}$ and $GQ_c\Psi = \alpha g_{-}^{(c)} + \beta g_{+}^{(c)}$ with $\alpha = (\|\Psi_{+}\| + \|\Psi_{-}\|)/\sqrt{2}$. It therefore suffices to show that the basis of Theorem 1 is orthonormal (Lemma 1) and that for each class the vectors selected by the rule span $g_{+}^{(c)}$ and are orthogonal to $g_{-}^{(c)}$ (Lemma 2).

For a product state and real ψ_0 one has $\langle\Psi|X^w Z^s|\Psi\rangle = \prod_i \langle\psi_0|X^{w_i} Z^{s_i}|\psi_0\rangle$ with $\langle\psi_0|Z|\psi_0\rangle = c$, $\langle\psi_0|X|\psi_0\rangle = \sin 2\theta \neq 0$ and $\langle\psi_0|XZ|\psi_0\rangle = 0$ (XZ is antisymmetric); hence $\langle\Psi|X^w Z^s|\Psi\rangle = 0$ if and only if w and s share a position.

Lemma 1. For $y' \neq z'$ and all signs, $\langle P_{y'} g_s | P_{z'} g_{s'} \rangle = 0$. Indeed $P_{y'}^\dagger P_{z'} = \pm X^w Z^{(0,r)}$ with $w = \tilde{y} \oplus \tilde{z} = (|d| \bmod 2, d)$, $d = y' \oplus z' \neq 0$ and $r = Ld$. As $|w|$ is even, $X^w Z^{(0,r)}$ preserves the eigenspaces of G , so the overlap equals $\frac{1}{2}[\langle u|X^w Z^{(0,r)}|u\rangle + ss'\langle v|X^w Z^{(0,r)}|v\rangle]$, and $\langle\Psi_{\pm}|X^w Z^{(0,r)}|\Psi_{\pm}\rangle = \frac{1}{2}[\langle\Psi|X^w Z^{(0,r)}|\Psi\rangle \pm \langle\Psi|X^w Z^{(1,r)}|\Psi\rangle]$. Both expectation values vanish, for every s, s' , provided (i) d and Ld share a position and (ii) either $|d|$ is odd or $d \not\subseteq Ld$. For the prefix-parity map, $(Ld)_{i_0} = 1$ at the first position i_0 of d , which gives (i); and if $|d|$ is even, $(Ld)_{i_1} = |d| \bmod 2 = 0$ at the last position i_1 of d , which gives (ii). Within one y' , $\langle g_{+}|g_{-}\rangle = 0$. The 2^b vectors are therefore orthonormal.

Lemma 2. $\langle P_{y'} g_s | g_{+}^{(c)} \rangle = 0$ unless $s = \epsilon(y', c) :=$

$(-1)^{y' \cdot (Ly' \oplus c)}$, and $\langle P_{y'} g_s | g_-^{(c)} \rangle = 0$ unless $s = -\epsilon(y', c)$. Using $X^{\tilde{y}} Z^{(0,c)} = (-1)^{y' \cdot c} Z^{(0,c)} X^{\tilde{y}}$, $\langle P_{y'} g_s | Q_c g_+ \rangle = (-1)^{y' \cdot c} \langle g_s | Z^{(0,k)} X^{\tilde{y}} | g_+ \rangle$ with $k = Ly' \oplus c$. Since $|\tilde{y}\rangle$ is odd, $Z^{(0,k)} X^{\tilde{y}}$ exchanges the eigenspaces of G and only the cross terms survive: the overlap is $(-1)^{y' \cdot c} \frac{1}{2} A [1 + s(-1)^{y' \cdot k}]$ with $A = \langle u | Z^{(0,k)} X^{\tilde{y}} | v \rangle$, where we used $\langle v | Z^{(0,k)} X^{\tilde{y}} | u \rangle = \langle u | X^{\tilde{y}} Z^{(0,k)} | v \rangle = (-1)^{y' \cdot k} A$ for real vectors. This vanishes unless $s = (-1)^{y' \cdot k}$. For $g_-^{(c)} = G g_+^{(c)}$, the relations $G X^{\tilde{y}} = -X^{\tilde{y}} G$, $[G, Z^{(0,t)}] = 0$ and $G g_s = g_{-s}$ give $\langle P_{y'} g_s | g_-^{(c)} \rangle = -\langle P_{y'} g_{-s} | g_+^{(c)} \rangle$, which vanishes unless $s = -\epsilon(y', c)$.

With $E_c = \sum_{y'} |P_{y'} g_{\epsilon(y', c)}\rangle \langle P_{y'} g_{\epsilon(y', c)}|$, Lemma 2 and the completeness of the basis give $E_c g_+^{(c)} = g_+^{(c)}$ and $E_c g_-^{(c)} = 0$, which proves the theorem. The same construction applies to the all-flipped blocks with φ in place of ψ , and to the six-state probes once the bases are announced.

Appendix B: Twirling, post-selection and block locality

The key length of Sec. IV is evaluated over Bell-diagonal collective attacks on the untagged rounds. (1) Tagging. As in Ref. [1], rounds in which Alice emits more than one photon or Bob's detector does not fire are tagged and the adversary is credited with full knowledge of them; the decoy-state bounds $S_{\tilde{K}}^-$ and ϕ_Z^+ are statements about Bob's click record and hold against any attack. (2) Symmetrisation. Alice and Bob may apply a random Pauli twirl to the untagged rounds and announce it at once; it commutes with their measurements, costs no key and reduces any collective attack to a Bell-diagonal one with the observed error rates ϕ_X, ϕ_Z [34, 35], whose purification has the structure of Sec. III with $c_\psi = (1 - 2\phi_X - \phi_Z c_\varphi)/(1 - \phi_Z)$ (Appendix C.4); an adversary holding the purification is at least as strong as one who must measure early, so the reduction is valid in the security direction. (3) Coherent attacks. For a permutation-invariant protocol the post-selection technique [33] lifts security against every i.i.d. attack with parameter ϵ to security against arbitrary attacks with parameter $\epsilon(n+1)^{d^2-1}$, $d = 4$ for a pair of qubits. Two points prevent us from applying it here as more than a comment. First, the factor must be charged to every parameter of the i.i.d. analysis that it affects: the leftover-hash term grows by $4 \log_2(n+1)^{15}$ bits and, in the Shannon accounting, the smoothing parameter of the equipartition term shrinks by the same factor, which is why Table III reports losses of up to 1.5 percentage points rather than the few hundred bits that the hashing term alone would suggest. Second, the block structure breaks round-wise permutation invariance for the classes J and W, since which rounds share a block matters; the technique then applies at the level of blocks, with local dimension 4^b and exponent $16^b - 1$, which is useless for $b \geq 2$.

For the collective min-entropy column, which does not depend on the timelock, the entropic-uncertainty proof of Ref. [1] already covers coherent attacks and the question is only whether the sharper collective bound extends; for the timelock-dependent columns, a proof against coherent attacks would need a tool compatible with block locality, such as entropy accumulation, and we leave it open.

Appendix C: Derivations of the expressions in the main text

This appendix collects the algebra behind the expressions of the main text and attributes every step that is not ours to its source, so that the reader can follow the chain from the cited equations to the numbers in Table I.

C.1. Per-ancilla quantities. Eqs. (1)–(2) of Ref. [21] define the symmetric attack, and Eq. (3) there expresses the fidelity through the overlaps, $F = (1 + c_\varphi)/(2 - c_\psi + c_\varphi)$. Since $F = 1 - D$ by definition of the disturbance, solving for c_ψ gives

$$c_\psi = 2 + c_\varphi - \frac{1 + c_\varphi}{1 - D},$$

which is $c_\psi = (1 - 3D)/(1 - D)$ at $c_\varphi = 1$ (the BB84 worst case at the error rates of interest, Sec. III) and $c_\psi = (1 - 2D)/(1 - D)$ at $c_\varphi = 0$, the value that Ref. [21] (p. 2) attributes to the six-state attack. The Helstrom measurement for two equiprobable pure states with overlap c succeeds with probability $\frac{1}{2} + \frac{1}{2}\sqrt{1 - c^2}$ [28], rederived as Eq. (3.12) of Ref. [32]; applied to $\{\psi_\uparrow, \psi_\downarrow\}$ and $\{\varphi_\uparrow, \varphi_\downarrow\}$ it gives δ_0 and δ_1 . The identity $4\delta_0(1 - \delta_0) = 1 - (1 - c_\psi^2) = c_\varphi^2$ turns the condition $D/(1 - D) < 2\sqrt{\delta_0(1 - \delta_0)}$ of Theorem 1 (Eq. (11)) of Ref. [21] into $|c_\psi| > D/(1 - D)$; with $c_\psi = (1 - 3D)/(1 - D)$ this is $1 - 3D > D$, i.e. $D < 1/4$, which is Eq. (8) of Ref. [21]. That the adversary learns whether a bit was disturbed is the statement $Z_1 = X \oplus Y$ of Ref. [21] (p. 2), a consequence of the orthogonality of the ψ and φ subspaces.

C.2. Majority vote, tally, and general per-probe measurements. After the parities c are announced, the adversary's b per-ancilla guesses $\hat{a}_i$ become b guesses $\hat{a}_i \oplus c_{i-1}$ of the distilled bit a_1 , each correct with probability δ (δ_0 or δ_1 according to the block type, which she knows from Z_1) and independent given a_1 because the ancillas are in a product state. Let k be the number of guesses equal to 0. Under $a_1 = 0$ the likelihood of k is $\binom{b}{k} \delta^k (1 - \delta)^{b-k}$ and under $a_1 = 1$ it is $\binom{b}{k} (1 - \delta)^k \delta^{b-k}$; their ratio, $(\delta/(1 - \delta))^{2k-b}$, depends on the guesses only through k , so the tally is a sufficient statistic and the maximum-likelihood decision is the majority vote, with error

$$\gamma_b(\delta) = \sum_{k > b/2} \binom{b}{k} (1 - \delta)^k \delta^{b-k} + \frac{1}{2} \binom{b}{b/2} [\delta(1 - \delta)]^{b/2},$$

where the last term is present only for even b and accounts for ties broken at random. The adversary's information is the mutual information between a_1 and the sufficient statistic, $J_b(\delta) = 1 - \sum_k \Pr(k) h(\Pr(a_1 = 0|k))$ with the Bayes posterior $\Pr(a_1 = 0|k) = \delta^k(1 - \delta)^{b-k} / [\delta^k(1 - \delta)^{b-k} + (1 - \delta)^k \delta^{b-k}]$, as stated in Sec. III. Fano's inequality for a binary variable, $H(a_1|k) \leq h(\Pr[\text{error}])$, gives $J_b \geq 1 - h(\gamma_b)$, which is why the majority-vote expression underestimates the adversary. The Chernoff bound for the binomial tail gives $\gamma_b \leq [2\sqrt{\delta(1-\delta)}]^b = |c_\psi|^b$ with the exponent tight, which is the decay rate used in the proof of Theorem 1 of Ref. [21]. For a general identical single-probe POVM $\{E_k\}$ with outcome probabilities $p_k(a) = \langle \psi_a | E_k | \psi_a \rangle$ the same reasoning applies with the pair (k_i, c_i) as the per-probe datum: given a_1 , each pair is drawn independently from $\Pr(k, c|a_1) = \frac{1}{2} p_k(a_1 \oplus c)$, so the histogram of the $2K$ categories is a sufficient statistic and both J_b and the guessing probability are explicit multinomial sums; `perprobe_opt.py` maximises them over real rank-one four-outcome POVMs in the Naimark parametrisation $E_k = V^T |m_k\rangle \langle m_k| V$ with $\{m_k\}$ an orthonormal basis of $\mathbb{R}^4$.

C.3. Asymptotic conditions. An accepted block is either all-correct, with probability $(1 - D)^b$, or all-flipped, with probability D^b (Eq. (20) of Ref. [1] with $\phi_K = D$), so Bob's error on the distilled bit is $\Phi_K = D^b / [(1 - D)^b + D^b] \approx (D/(1 - D))^b$. For the adversary, the individual model gives the majority-vote error $\approx |c_\psi|^b$ of C.2, and the joint model gives $1 - P_b^{\text{ad}} = \frac{1}{2} - \frac{1}{2} \sqrt{1 - c_\psi^{2b}} \approx c_\psi^{2b}/4$. The repetition code yields a positive rate iff the adversary's error on the distilled bit decays more slowly than Bob's as $b \rightarrow \infty$ (the comparison of exponents in the proof of Theorem 1 of Ref. [21], which relies on the classical results cited there); this gives $|c_\psi| > D/(1 - D)$ and $c_\psi^2 > D/(1 - D)$ respectively. In the joint case with $c_\psi = (1 - 3D)/(1 - D)$ the condition is $(1 - 3D)^2 > D(1 - D)$, i.e. $10D^2 - 7D + 1 > 0$, whose roots are $D = 1/5$ and $D = 1/2$, so $D < 1/5$. For the six-state protocol, $c_\psi = (1 - 2D)/(1 - D)$ gives $(1 - 2D)^2 > D(1 - D)$, i.e. $5D^2 - 5D + 1 > 0$, with roots $(5 \pm \sqrt{5})/10$, so $D < (5 - \sqrt{5})/10$. The condition $c_\psi^2 > D/(1 - D)$ is Eq. (68) of Ref. [7] specialised to $d = 2$, where it was derived for an adversary who measures after the announcement; Ref. [6] obtains the six-state value in its Eq. (23).

C.4. Bell-diagonal attacks. After a random Pauli twirl applied by Alice and Bob to their qubits the two-qubit state is Bell-diagonal, $\rho_{AB} = \sum_{i=1}^4 \lambda_i |\beta_i\rangle \langle \beta_i|$ with $\beta_{1,2} = (|00\rangle \pm |11\rangle)/\sqrt{2}$ and $\beta_{3,4} = (|01\rangle \pm |10\rangle)/\sqrt{2}$, and the attack is characterised by the weights [34, 35]; Eq. (7) of Ref. [21] is the special case of the symmetric attack. A Z -basis error occurs for β_3, β_4 and an X -basis error (a phase flip) for β_2, β_4 , so $\phi_Z = \lambda_3 + \lambda_4$

and $\phi_X = \lambda_2 + \lambda_4$. The adversary holds the purification $\sum_i \sqrt{\lambda_i} |\beta_i\rangle |e_i\rangle$ with orthonormal e_i . Projecting Alice and Bob on $|a\rangle|y\rangle$ in the Z basis gives, up to the factor $1/\sqrt{2}$, the unnormalised adversary states $\sqrt{\lambda_1} e_1 + (-1)^a \sqrt{\lambda_2} e_2$ for $y = a$ and $\sqrt{\lambda_3} e_3 + (-1)^a \sqrt{\lambda_4} e_4$ for $y \neq a$. Normalising, $\psi_a \propto \sqrt{\lambda_1} e_1 + (-1)^a \sqrt{\lambda_2} e_2$ and $\varphi_a \propto \sqrt{\lambda_3} e_3 + (-1)^a \sqrt{\lambda_4} e_4$, with overlaps $c_\psi = (\lambda_1 - \lambda_2)/(\lambda_1 + \lambda_2) = (\lambda_1 - \lambda_2)/(1 - \phi_Z)$ and $c_\varphi = (\lambda_3 - \lambda_4)/\phi_Z$, and $\psi_a \perp \varphi_{a'}$ because the e_i are orthonormal; this is the structure of Sec. III. Writing $\lambda_4 = \phi_Z(1 - c_\varphi)/2$, $\lambda_3 = \phi_Z - \lambda_4$, $\lambda_2 = \phi_X - \lambda_4$ and $\lambda_1 = 1 - \phi_Z - \phi_X + \lambda_4$ gives $\lambda_1 - \lambda_2 = 1 - \phi_Z - 2\phi_X + 2\lambda_4 = 1 - 2\phi_X - \phi_Z c_\varphi$, hence the expression for c_ψ quoted in Sec. III; the constraints $\lambda_i \geq 0$ are those stated there. At $\phi_X = \phi_Z = D$ the expression reduces to $c_\psi = (1 - 2D - D c_\varphi)/(1 - D)$, which coincides with C.1.

C.5. The key-length formula. Eq. (93) of Ref. [1] reads $\ell_{\text{AD}} \leq S_{\bar{K}}[1 - h(\Phi_X^+)] - \lambda_{\text{IR}} - \log_2(2/\epsilon_{\text{cor}}) - 4 \log_2(2^{7/4}/(\epsilon_{\text{sec}} - \epsilon_{\text{sec}}^2/8))$ when the acceptance thresholds are set to the computed bounds (Eq. (21) there). Its first term is the entropic-uncertainty lower bound on the smooth min-entropy of the distilled bits of the single-photon blocks, $S_{\bar{K}}$ being Eq. (18) and Φ_X^+ Eq. (19) of Ref. [1]; the second is the reconciliation leakage; the third is the error-verification cost (Theorem 2 of the reference cited as [42] in Ref. [1]); the fourth is the leftover-hash-lemma term (Lemma 4 of Ref. [1]) after the statistical failure probabilities are set to $\bar{\epsilon} = \epsilon_{\text{sec}}^2/368$ (Sec. 6 of Ref. [1]). Under the timelock the adversary's side information on each single-photon block is classical, so we replace the first term by a bound on the smooth min-entropy of $S_{\bar{K}}$ i.i.d. classical bits given classical side information. The fully quantum asymptotic equipartition property, Theorem 9 (Eq. (30)) of Ref. [23], gives $H_{\min}^\epsilon(A^S|E^S) \geq S H(A|E) - \sqrt{S} \delta(\epsilon, \eta)$ with $\delta(\epsilon, \eta) = 4 \log_2 \eta \sqrt{\log_2(2/\epsilon^2)}$ for $S \geq \frac{8}{5} \log_2(2/\epsilon^2)$, where η is the convergence parameter of Definition 4 (Eq. (24)) of Ref. [23], bounded there by $\sqrt{2^{-H_{\min}(A|E)}} + \sqrt{2^{H_{\max}(A|E)}} + 1$. For a binary A with classical E one has $H_{\min}(A|E) \geq 0$ and $H_{\max}(A|E) \leq 1$, hence $\eta \leq 2 + \sqrt{2}$, which is the constant in Δ_{AEP} . With a_1 uniform, $H(A|E) = 1 - I_E$, and the worst case over the attack parameters and the true error rates below their estimates gives $\bar{I}_E$. The smoothing parameter $\epsilon = \epsilon_{\text{sec}}/4$ is charged in addition to the budget of Ref. [1], giving the total $5\epsilon_{\text{sec}}/4$ stated in Sec. IV. The min-entropy variant uses instead $H_{\min}(A|E) = -\log_2 p_{\text{guess}}$ for classical side information (Theorem 1 of Ref. [30]), additive over the independent blocks of a collective attack, so that no smoothing step and no Δ_{AEP} term are needed. The leftover hash lemma against classical side information then yields a key that is ϵ_{sec} -secret in the trace-distance sense, which for classical adversary information coincides with the accessible-information criterion of Ref. [19].

-
- [1] J. Treplin, P. Kleinpaß, D. Orsucci, Finite size analysis of decoy-state BB84 with advantage distillation, arXiv:2511.21665 (2025).
 - [2] C. C. W. Lim, M. Curty, N. Walenta, F. Xu, H. Zbinden, Concise security bounds for practical decoy-state quantum key distribution, *Phys. Rev. A* 89, 022307 (2014).
 - [3] U. M. Maurer, Secret key agreement by public discussion from common information, *IEEE Trans. Inf. Theory* 39, 733 (1993).
 - [4] D. Gottesman, H.-K. Lo, Proof of security of QKD with two-way classical communications, *IEEE Trans. Inf. Theory* 49, 457 (2003).
 - [5] H. F. Chau, Practical scheme to share a secret key through a quantum channel with a 27.6% bit error rate, *Phys. Rev. A* 66, 060302 (2002).
 - [6] A. Acín, J. Bae, E. Bagan, M. Baig, Ll. Masanes, R. Muñoz-Tapia, Secrecy properties of quantum channels, *Phys. Rev. A* 73, 012327 (2006); arXiv:quant-ph/0411092.
 - [7] J. Bae, A. Acín, Key distillation from quantum channels using two-way communication protocols, *Phys. Rev. A* 75, 012334 (2007); arXiv:quant-ph/0610048.
 - [8] A. Acín, E. Bagan, M. Baig, Ll. Masanes, R. Muñoz-Tapia, Multiple-copy two-state discrimination with individual measurements, *Phys. Rev. A* 71, 032338 (2005).
 - [9] R. Bhavsar, J. Moon, J. Bae, Entropy bounds via hypothesis testing and its applications to two-way key distillation in quantum cryptography, arXiv:2602.05870 (2026).
 - [10] C. Jiang, X.-L. Hu, Z.-W. Yu, X.-B. Wang, Composable security for practical QKD with two way classical communication, arXiv:2102.00739 (2021).
 - [11] I. B. Damgård, S. Fehr, L. Salvail, C. Schaffner, Cryptography in the bounded quantum-storage model, *Proc. 46th IEEE FOCS*, 449 (2005).
 - [12] S. Wehner, C. Schaffner, B. M. Terhal, Cryptography from noisy storage, *Phys. Rev. Lett.* 100, 220502 (2008).
 - [13] R. König, S. Wehner, J. Wullschleger, Unconditional security from noisy quantum storage, *IEEE Trans. Inf. Theory* 58, 1962 (2012).
 - [14] S. Watanabe, R. Matsumoto, T. Uyematsu, Security of QKD protocol with two-way classical communication assisted by one-time pad encryption, arXiv:quant-ph/0608030 (2006).
 - [15] V. A. Pastushenko, D. A. Kronberg, Improving the performance of quantum cryptography by using the encryption of the error correction data, *Entropy* 25, 956 (2023).
 - [16] M. Tomamichel, A. Leverrier, A largely self-contained and complete security proof for QKD, *Quantum* 1, 14 (2017).
 - [17] R. Novak, Security of QKD with one-time-pad-protected error correction and its performance benefits, *Entropy* 27, 1032 (2025).
 - [18] C.-H. F. Fung, X. Ma, H. F. Chau, Q.-Y. Cai, Quantum key distribution with delayed privacy amplification and its application to the security proof of a two-way deterministic protocol, *Phys. Rev. A* 85, 032308 (2012).
 - [19] N. Vyas, R. Alléaume, Everlasting secure key agreement with performance beyond QKD in a quantum computational hybrid security model, arXiv:2004.10173 (2020).
 - [20] F. Mazzoncini, B. Bauer, P. Brown, R. Alléaume, Hybrid quantum cryptography from communication complexity, *Quantum* 9, 1862 (2025).
 - [21] N. Gisin, S. Wolf, Quantum cryptography on noisy channels: quantum versus classical key-agreement protocols, *Phys. Rev. Lett.* 83, 4200 (1999).
 - [22] C. A. Fuchs, N. Gisin, R. B. Griffiths, C.-S. Niu, A. Peres, Optimal eavesdropping in quantum cryptography. I, *Phys. Rev. A* 56, 1163 (1997).
 - [23] M. Tomamichel, R. Colbeck, R. Renner, A fully quantum asymptotic equipartition property, *IEEE Trans. Inf. Theory* 55, 5840 (2009).
 - [24] A. Bocquet, R. Alléaume, A. Leverrier, Optimal eavesdropping on QKD without quantum memory, *J. Phys. A* 45, 025305 (2012); arXiv:1106.0329.
 - [25] M. A. Ballester, S. Wehner, A. Winter, State discrimination with post-measurement information, *IEEE Trans. Inf. Theory* 54, 4183 (2008).
 - [26] C. Carmeli, T. Heinosaari, A. Toigo, State discrimination with postmeasurement information and incompatibility of quantum measurements, *Phys. Rev. A* 98, 012126 (2018).
 - [27] D. Gopal, S. Wehner, Using post-measurement information in state discrimination, *Phys. Rev. A* 82, 022326 (2010).
 - [28] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, New York, 1976).
 - [29] Y. C. Eldar, A. Megretski, G. C. Verghese, Designing optimal quantum detectors via semidefinite programming, *IEEE Trans. Inf. Theory* 49, 1007 (2003).
 - [30] R. König, R. Renner, C. Schaffner, The operational meaning of min- and max-entropy, *IEEE Trans. Inf. Theory* 55, 4337 (2009).
 - [31] L. B. Levitin, Optimal quantum measurements for two pure and mixed states, in *Quantum Communications and Measurement*, edited by V. P. Belavkin, O. Hirota, R. L. Hudson (Plenum, New York, 1995), p. 439.
 - [32] C. A. Fuchs, Distinguishability and accessible information in quantum theory, PhD thesis, University of New Mexico (1996); arXiv:quant-ph/9601020.
 - [33] M. Christandl, R. König, R. Renner, Postselection technique for quantum channels with applications to quantum cryptography, *Phys. Rev. Lett.* 102, 020504 (2009).
 - [34] B. Kraus, N. Gisin, R. Renner, Lower and upper bounds on the secret-key rate for quantum key distribution protocols using one-way classical communication, *Phys. Rev. Lett.* 95, 080501 (2005).
 - [35] R. Renner, N. Gisin, B. Kraus, An information-theoretic security proof for QKD protocols, *Phys. Rev. A* 72, 012332 (2005).
 - [36] N. Walenta et al., A fast and versatile QKD system with hardware key distillation and wavelength multiplexing, *New J. Phys.* 16, 013047 (2014).
 - [37] J. Constantin et al., An FPGA-based 4 Mbps secret key distillation engine for QKD systems, *J. Signal Process. Syst.* 86, 1 (2017).
 - [38] X. Cheng, H.-K. Mao, H.-W. Xu, Q. Li, Large-scale FPGA-based privacy amplification exceeding 10^8 bits for QKD, arXiv:2503.09331 (2025).