Generation of quantum randomness by probability estimation with classical side information

Emanuel Knill,^{1,2,*} Yanbao Zhang^{3,4,†} and Peter Bierhorst^{1,5}

¹National Institute of Standards and Technology, Boulder, Colorado 80305, USA

²Center for Theory of Quantum Matter, University of Colorado, Boulder Colorado 80309, USA

³NTT Basic Research Laboratories, NTT Corporation, 3-1 Morinosato-Wakamiya, Atsugi, Kanagawa 243-0198, Japan

⁴NTT Research Center for Theoretical Quantum Physics, NTT Corporation, 3-1 Morinosato-Wakamiya, Atsugi, Kanagawa 243-0198, Japan

⁵Mathematics Department, University of New Orleans, New Orleans, Louisiana 70148, USA

(Received 6 January 2020; accepted 17 August 2020; published 22 September 2020)

We develop the framework of probability estimation for certifying randomness with respect to classical side information from a sequence of Bell-test or other randomness-generating trials. The framework is based on directly estimating the probability of measurement outcomes conditional on settings choices and classical side information with adaptive test supermartingales. Accordingly, the number of trials needs not to be predetermined, and one can stop performing trials early, as soon as the desired amount of randomness is extractable. It can be used with arbitrary, partially known and time-varying probabilities for the random settings choices. It can also adapt to other time-varying experimental parameters. Furthermore, it is suitable for application to experiments with low Bell violation per trial, such as current optical loophole-free Bell tests. Compared with our previous work [Phys. Rev. A **98** 040304(R) (2018)], here we formulate the framework for the general situation where the randomness can be extracted from a sequence of private data determined in an arbitrary way by the measurement outcomes of the trials. Trial-wise probability estimators can be adapted using all accessible, private information in addition to the results of previous trials. We prove that probability estimation achieves the asymptotically optimal rate for certified randomness generation and makes possible the exponential expansion of settings entropy. We implement probability estimation numerically and apply it to a representative settings-conditional distribution of the measurement outcomes from an atomic loophole-free Bell test [W. Rosenfeld *et al.*, Phys. Rev. Lett. **119**, 010402 (2017)] to illustrate trade-offs between the amount of randomness, error, settings entropy, adversarial settings bias, and number of trials. We then show that probability estimation yields more randomness from the optical loophole-free Bell-test data analyzed in [P. Bierhorst *et al.*, arXiv:1702.05178v1] and tolerates adversarial settings biases.

DOI: [10.1103/PhysRevResearch.2.033465](https://doi.org/10.1103/PhysRevResearch.2.033465)

I. INTRODUCTION

Device-independent quantum randomness generation exploits the fact that there are quantum correlations with measurement outcomes that are necessarily nondeterministic with respect to prior side information. This nondeterminism is ensured by nonsignaling constraints that can be enforced by causal separation of the relevant events and certified by tests based on freely chosen measurement settings. The certified randomness can then be extracted by means of classical algorithms. Device independence means that the physical devices producing and measuring the quantum correlations can be obtained from untrusted manufacturers, without affecting the desired randomness properties. Device-independent quantum randomness generation was introduced in Colbeck's thesis [1] and has since been studied by many researchers. See Refs. [2,3] for reviews.

The types of randomness generation protocols are diverse and depend on the context and assumptions. In general, a randomness generation protocol produces a string of σ bits certified to be uniformly random within a total variation distance ϵ . It requires devices whose behaviors are known to satisfy restrictions based on physical principles such as speed-of-light limits on communication. Device independence also requires some random input bits, usually used for measurement settings choices, where different assumptions can be made on their quality and source. The certification is conditional on the device-behavior restrictions as well as assumptions on outside entities with respect to which the string of generated bits is to be random. For example, it matters whether these entities are classical or quantum and what type of access they had to the protocol devices in the past. This work considers entities holding classical side information. That outside entities have only classical information can be justified if they are not quantum-capable, if we built the devices and the entities never had access to them, or if we verified the absence of long-term quantum memory in the devices. See Sec. III C. We assume that the source of input bits is random relative to outside entities at the time they last interacted with the protocol devices.

For device-independent quantum randomness generation, we use experiments modelled on Bell tests [4]. These tests are designed to show that there are quantum correlations

*Corresponding author: knill@nist.gov

†Corresponding author: yanbao.zhang.xf@hco.ntt.co.jp

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

that cannot be mixtures of locally determined probability distributions referred to as “local realistic” (LR) distributions. See Ref. [5] for a review. A Bell test consists of a sequence of trials (sometimes called “rounds”) where two or more stations make measurements on a shared physical state with randomly chosen measurement settings. For full device independence, each trial needs to satisfy that the different stations’ time intervals between applying the settings to the devices and determining the measurement outcomes are spacelike separated, preventing any communication between them. Furthermore, trials must be committed to in advance, so that it is not possible to postselect them on a success criterion. Because these conditions require large separation and/or fast devices as well as high-efficiency measurements, only recently has it become possible to perform successful Bell tests satisfying these criteria. Such Bell tests may be referred to as “loophole-free” and the list of successful loophole-free Bell tests includes ones based on heralded atom entanglement [6,7] and ones utilizing entangled photon-pairs with high-efficiency detectors [8,9].

Experimental certified device-independent quantum randomness was first demonstrated in Ref. [10] (see also Refs. [11,12]) with pairs of ions located in separate traps. This demonstration claimed the presence of 42 bits of entropy with an error of 0.01 in the string of measurement outcomes, with respect to classical side information and restricting correlations to quantum achievable ones. Extraction would have reduced the number of bits produced and increased the error by an amount that was not determined. Recently, we and our collaborators demonstrated end-to-end device-independent quantum randomness extraction [13], producing 256 bits within 0.001 of uniform from one of the data sets from the loophole-free Bell test reported in Ref. [9]. These bits are certified with respect to classical side information and nonsignaling assumptions, which in principle allows for super-quantum correlations. Extracting randomness from today’s optical loophole-free Bell tests required the theoretical advances in Ref. [13] to deal with the fact that each trial demonstrates very little violation of Bell inequalities. Previous works were not sufficient for certifying entropy without increasing the number of trials by orders of magnitude. A specific comparison is in Ref. [13].

A benefit of the theory developed in Ref. [13] is that it allows for an adaptive protocol that can track changes in the trial statistics during the protocol. This is helpful in current experiments, where we find that measurable drifts in parameters can wipe out a randomness certificate if not accounted for. The fact that the protocol can adapt is inherited from its use of the “probability-based ratio” protocol for obtaining p -value bounds against local realism (LR) in Bell tests [14,15]. Here we develop a different class of randomness generation protocols based on “probability estimation.” Probability estimation involves obtaining high-confidence-level upper bounds on the actual probability of the measurement outcomes given the known constraints on the distributions. We show that randomness generation can be reduced to probability estimation. Since probability estimation is a statistical estimation problem, we then take advantage of the theory of test supermartingales [16] to bypass the framework of Bell inequalities and directly determine probability estimators ex-

pressed as products of *probability estimation factors* (PEFs). PEFs are functions of a trial’s settings and outcomes and provide a way of multiplicatively accumulating probability estimates trial-by-trial. While relationships between device-independent quantum randomness and Bell violations exist, they are inequivalent quantities: a stronger violation of a fixed Bell inequality does not necessarily certify a larger amount of device-independent quantum randomness [17]. We develop tools for obtaining PEFs. In particular, we show that when the distributions of settings and outcomes are constrained to a convex polytope, PEFs can be effectively optimized with convex optimization over the polytope given one parameter, the “power” (see definition 22). The optimization can explicitly take into account the number of trials and the error goal. In the limit where the power parameter goes to zero, an asymptotic rate is obtained that can be interpreted as the optimal rate for producing entropy for random bits. This generalizes and improves on min-entropy estimators for Bell configurations described in works such as Refs. [18–20], which are optimal for single-trial min-entropy estimation.

In a large class of situations including the standard Bell-test configurations, PEFs directly lead to exponential expansion of input randomness, as expected from previous works [21–23], which prove exponential expansion with more trials or worse settings entropy than PEFs but secure against quantum side information. We prove that asymptotically, the settings entropy can be logarithmic in the output entropy. This is the best result so far for randomness expansion without using a cross-feeding protocol, which can accomplish infinite expansion [24,25]. To accomplish exponential expansion, we use highly biased settings distributions. We point out that it is not necessary to have independent and identically distributed (i.i.d.) settings choices. In particular, if the settings are obtained by choosing, ahead of time, a random “test” trial among a block of 2^k trials, we can achieve good expansion while eliminating the need for decompressing uniformly random input bits into a stream of highly biased and independent ones.

As a demonstration of the power of probability estimation, we show how it would perform on a representative example for distributions achieved in loophole-free Bell experiments with atoms based on heralded entanglement [7]. We then apply it to the main data set from Ref. [9] analyzed in Ref. [13], showing that we could improve the amount of device-independent quantum randomness extracted substantially, while ensuring that the certificates are valid even if the input randomness is biased, a problem that was noticed and accounted for in the report on the loophole-free Bell test for this experiment [9]. Finally, we reanalyze the data from the first demonstration of certified experimental randomness in a Bell test free of the detection loophole but subject to the locality loophole, which was based on ions [10]. We demonstrate significantly more randomness than reported in this reference. These examples demonstrate that probability estimation is a practical way for implementing entropy production in randomness generation.

Our framework is in the spirit of the entropy accumulation framework of Ref. [26], but takes advantage of the simplifications possible for randomness generation with respect to classical side information. In particular, the outside entities have no interaction with the protocol devices after the protocol

starts, and the framework can be cast purely in terms of random variables without invoking quantum states. This avoids the complications of a full representation of the protocol in terms of quantum processes. With these simplifications, our framework applies to any situation with known constraints on device behavior, not only device-independent situations. The only requirement is that the constraints can be formulated as constraints on past-conditional probability distributions and are sufficiently strong to allow for randomness certification. In this work, we choose to introduce our framework with a focus on device-independent applications.

In the entropy accumulation framework, the relevant estimators, called min-tradeoff functions, must be chosen before the protocol, and the final certificate is based on the sum of statistics derived from these functions. Finding suitable min-tradeoff functions is in general difficult. In the probability estimation framework, probability estimators can be adapted and accumulate multiplicatively. For relevant situations, PEFs are readily obtained and the tradeoff between randomness and error can be effectively optimized.

The analog of min-tradeoff functions in the probability estimation framework are entropy estimators. We show that logarithms of PEFs are proportional to entropy estimators, and essentially all entropy estimators are related to PEFs in this way. In this sense, there is no difficulty in finding entropy estimators. However, PEFs are more informative for applications, so except for illuminating asymptotic behavior, there is little to be gained by seeking entropy estimators directly.

A feature of entropy accumulation is optimality of asymptotic rates for min-tradeoff functions. Probability estimation also achieves optimal asymptotic rates. In both cases, the tradeoff between error and amount of randomness makes these asymptotic rates less relevant, which we demonstrate for probability estimation on the Bell-test examples mentioned above.

The remainder of the paper is structured as follows. We summarize the main results in Sec. II. We lay out our notation and define the basic concepts required for the probability estimation framework in Sec. III. This section includes introductions to less familiar material on classical smooth min-entropies, test martingales and the construction of test martingales from Bell inequalities. In Sec. IV, we define exact and soft probability estimation and show how randomness generation can be reduced to probability estimation. The measurement outcomes can be fed into appropriate randomness extractors, where the number of near-uniform random bits is naturally related to the probability estimate. We give three protocols that compose probability estimators with randomness extractors. The first is based on general relationships between probability estimation and smooth min-entropy and reprises techniques from Refs. [12,13,27]. The second relies on banked randomness to avoid the possibility of protocol failure. The third requires linear strong extractors to enable a direct analysis of the composition. Although we do not demonstrate an end-to-end randomness-generation protocol including extraction here, our goal is to provide all the information needed for implementing such a protocol in future work, with all relevant constants given explicitly. Sec. V shows how to perform probability estimation for a sequence of trials by means of implicit test supermartingales defined adaptively. The main tool involves PEFs to successively

accumulate probability estimates. The main results involve theorems showing how PEFs can be “chained” to form probability estimates. PEFs are readily constructed for distribution constraints relevant in Bell-test configurations. We proceed to an exploration of basic PEF properties in Sec. VI, where we find that there is a close relationship between the rates for PEFs and those for a class of functions called “entropy estimators”, which are the analog of min-tradeoff functions in our framework. We establish that for error that is $e^{-o(n)}$, the achievable asymptotic rates are optimal. Next, in Sec. VII, we consider a family of PEFs constructed from Bell functions whose expectations bound maximum conditional probabilities in one trial. In Sec. VIII, we show that this family can be used for exponential expansion by means of highly biased settings choices. Given a constant error bound, the settings distribution can be interpreted as a random choice of a constant (on average) number of test trials with uniformly random settings, where the remaining trials have fixed settings. We note that this is a theoretical proof-of-principle of exponential expansion. In practice, we prefer to numerically optimize the PEFs with respect to the desired error and calibrated experimental distribution. The final section Sec. IX explores the three examples mentioned above.

We remark that the results for certifying randomness with exact probability estimation were already presented in Ref. [28]. Here, we formulate the results for the general case with soft probability estimation. Recently, we also generalized exact probability estimation for certifying randomness with respect to quantum side information, see Refs. [29,30].

II. SUMMARY OF MAIN RESULTS

This work aims to establish the foundations of probability estimation and contains a large number of mathematical results based on mathematical concepts introduced here. In this section, we summarize the main results without precise definitions, in more familiar terms and with less generality than the probability estimation framework established later.

The context for our work consists of experiments where a sequence of trials is performed. In each trial, settings are chosen according to a random variable (RV) Z and outcomes are obtained according to an RV C . The sequences of outcomes and settings obtained in the experiment are denoted by $\mathbf{C}$ and $\mathbf{Z}$, where for n trials, $\mathbf{C} = (C_i)_{i=1}^n$ and $\mathbf{Z} = (Z_i)_{i=1}^n$, with C_i and Z_i the i th trial’s outcomes and settings. The main example of such an experiment is the standard Bell test, where there are two physically separated stations. In a trial, the stations randomly select measurement settings X and Y (respectively) and the stations’ devices produce outcomes A and B (respectively). In this case, $Z = XY$ and $C = AB$. The physical separation of the stations and the physics of the measurement devices constrain the distributions of $ABXY$ to be nonsignaling once loopholes are accounted for. Classical devices are further constrained by LR, which can be violated by quantum devices. This violation is associated with randomness that can be exploited for randomness generation [1], see the introduction above. Here we consider randomness generation where the generated bits are random relative to any external entity holding classical side information E .

The traditional approach to randomness generation is to first derive a bound on the smooth min-entropy of the out-

comes conditional on settings and E from the statistics of the observed value $\mathbf{cz}$ of $\mathbf{CZ}$ given constraints on the joint distribution of $\mathbf{CZ}$ and E . The smooth conditional min-entropy $H_{\min, \mu}^{\epsilon}(\mathbf{C}|ZE)$ for the joint distribution μ is given by the negative natural logarithm of the maximum probability of $\mathbf{C}$ given $\mathbf{Z}$ and E , averaged over $\mathbf{Z}$ and E , up to an error bound ϵ , which is the smoothness parameter. It can be formally defined as the maximum $\lambda \geq 0$ such that there exist μ' and ν where μ' is within total variation distance ϵ of μ and $\mu'(\mathbf{c}ze) \leq \exp(-\lambda)\nu(\mathbf{z}e)$ for all $\mathbf{c}ze$. (A more involved but equivalent definition based on maximum probabilities is given in Sec. III D, definition 3.) The smooth conditional min-entropy characterizes the number of near-uniform random bits that can be extracted from the outcomes with a randomness extractor applied after obtaining the min-entropy bound. The distance from uniform of the random bits obtained is parametrized by an error bound that, in the simplest case, is the sum of the error bound of the smooth min-entropy and an error parameter of the extractor used to obtain the random bits. Here we reduce the problem of obtaining a smooth conditional min-entropy bound to that of estimating the conditional probability $\mathbb{P}(\mathbf{c}|\mathbf{z}e)$ for the observed values $\mathbf{cz}$, independent of the value e of E .

Constraints on the joint distribution of $\mathbf{CZ}$ and E are determined by a statistical model $\mathcal{H}$ consisting of the allowed joint probability distributions, which may enforce nonsignaling conditional on E and other constraints such as that the conditional distributions are quantum achievable by causally separated devices sharing an initial state. Given $\mathcal{H}$, a level- ϵ (conditional) probability estimator for $\mathcal{H}$ and $\mathbf{C}|Z$ is a function $U : \mathbf{cz} \mapsto U(\mathbf{cz}) \in [0, 1]$ such that for all $\mu \in \mathcal{H}$ and all values e of E , the probability that $\mathbf{CZ}$ takes a value $\mathbf{cz}$ for which $U(\mathbf{cz}) \geq \mathbb{P}_{\mu}(\mathbf{C} = \mathbf{c}|\mathbf{Z} = \mathbf{z}, E = e)$ is at least $1 - \epsilon$. The probability estimate $U(\mathbf{cz})$ differs from a smooth min-entropy estimate in that the quantity being estimated depends on the data $\mathbf{cz}$, while the smooth conditional min-entropy is a characteristic of the overall distribution μ . Our first result is that one can obtain a smooth min-entropy estimate from a probability estimate.

Lemma (Lemma 18). Consider a level- ϵ probability estimator U for $\mathcal{H}$ and $\mathbf{C}|Z$, $\mu \in \mathcal{H}$ and $\kappa = \mathbb{P}_{\mu}(U \leq p)$. Define the probability distribution μ' by $\mu'(\mathbf{c}ze) = \mu(\mathbf{c}ze|U \leq p)$ for all $\mathbf{c}ze$. Then $H_{\min, \mu'}^{\epsilon/\kappa}(\mathbf{C}|ZE) \geq -\ln(p/\kappa)$.

We establish this lemma for the larger class of soft probability estimators, which provide extensions that may be useful in some applications. In particular, softness enhances adaptability and enables use of trial information not determined by $\mathbf{CZ}$. For simplicity, we do not consider softening in this section.

Here is a sketch of one way to generate randomness from $\mathcal{C}$ using the lemma above. First determine a level- ϵ^2 probability estimator U , then run an experiment to obtain an instance $\mathbf{cz}$ of $\mathbf{CZ}$. If $U(\mathbf{cz}) > p$, the protocol failed. If not, apply a classical-proof extractor $\mathcal{E}$ to $\mathbf{c}$ with input min-entropy $-\log_2(p/\epsilon)$ to produce random bits. The number of random bits produced can be close to that determined by the input min-entropy. The definition and properties of extractors are summarized in Sec. III D. The parameters chosen ensure that if the probability of success is at least ϵ , then conditional on success, the random bits produced are uniform within TV distance $\epsilon + \epsilon_x$, where ϵ_x is the extractor error. In Sec. IV C, we

provide details for three protocols for randomness generation from probability estimators that take advantage of features of probability estimation to improve on the sketch just given.

We previously developed a powerful method for constructing “probability-based ratios” such that the resulting p -value bounds for testing LR are asymptotically tight [14,15]. The method can be seen as an application of the theory of test supermartingales [16]. The definitions and basic properties of test supermartingales are given in Sec. III F. Here we show that this theory can be applied to the problem of constructing probability estimators. The basic tool is to construct *probability estimation factors* (PEFs) that are computed for each trial of an experiment. Let $\mathcal{C}$ be the trial model. A PEF with power $\beta > 0$ for $\mathcal{C}$ and $\mathbf{C}|Z$ is a non-negative function $F : \mathbf{cz} \mapsto F(\mathbf{cz})$ such that for all $\mu \in \mathcal{C}$ we have $\mathbb{E}_{\mu}(F(\mathbf{CZ})\mu(\mathbf{C}|Z)^{\beta}) \leq 1$. The fundamental theorem of PEFs is that they can be “chained” over trials to construct probability estimators. When chaining trials, the experimental model $\mathcal{H}$ is constructed from individual trial models $\mathcal{C}$ by requiring that each trial’s probability distribution conditional on the past is in $\mathcal{C}$. The trial models and PEFs may depend on past settings and outcomes. Conditioning on settings requires an additional conditional independence property as specified in Sec. V A. A simplified version of the fundamental theorem for the case where all trial models and PEFs are the same is the following:

Theorem (Theorem 23). Let F be a PEF with power β for $\mathcal{C}$ and $\mathbf{C}|Z$ and define $T(\mathbf{CZ}) = \prod_{i=1}^n F(\mathbf{C}_i Z_i)$. Then $1/(\epsilon T(\mathbf{CZ}))^{1/\beta}$ is a level- ϵ probability estimator for $\mathcal{H}$ and $\mathbf{C}|Z$.

The proof is enabled by martingale theory and requires constructing a test supermartingale.

From the fundamental theorem for PEFs, we can see that up to adjustments for probability of success and the error bound, the min-entropy per trial witnessed by a PEF F with power β at distribution μ is expected to be the “log-prob rate” $\mathbb{E}_{\mu}(\ln(F(\mathbf{cz})))/\beta$. In a sequence of trials with devices behaving according to specifications, we can expect the trial distributions to be approximately i.i.d. with known distribution μ . In this case, the RVs $\ln(F(\mathbf{C}_i Z_i))/\beta$ are also approximately i.i.d., and their sum is typically close to $n\mathbb{E}_{\mu}(\ln(F(\mathbf{cz})))/\beta$. Since the sum determines the conditional min-entropy witnessed by the probability estimator obtained from F , a goal of PEF construction is to maximize the log-prob rate. We show that PEF construction and optimization reduces to the problem of maximizing a concave function over a convex domain. If the trial model is a convex polytope, the convex domain is defined by finitely many extreme points and PEF optimization has an effective implementation (theorem 28). For the standard Bell-test configuration, convex polytopes that include the model and that have a manageable number of extreme points exist. We implement and apply PEF optimization in Sec. IX.

Given a trial model and a distribution μ in the model, the maximum number of near-uniform random bits that can be produced per trial in an asymptotically long sequence of trials is given by the minimum conditional entropy of the outcomes given the settings and E . The minimum is over all distributions ν of $\mathbf{CZE}$ such that $\nu(\mathbf{CZ}|e) \in \mathcal{C}$ for all e and $\nu(\mathbf{CZ}) = \mu(\mathbf{CZ})$. This is a consequence of the asymptotic equipartition property [31]. We prove that PEFs can achieve the maximum rate in the asymptotic limit.

Theorem (Theorem 43). For any trial model $\mathcal{C}$ and distribution $\mu \in \mathcal{C}$ with minimum conditional entropy g , the supremum of the log-prob rates of PEFs for $\mathcal{C}$ is g .

To prove optimality we define entropy estimators for model $\mathcal{C}$ as real-valued functions $K(CZ)$ such that $\mathbb{E}_\mu(K(CZ))$ is a lower bound on the conditional entropy for all $\mu \in \mathcal{C}$. We then show that there are entropy estimators for which the entropy estimate $\mathbb{E}_\mu(K(CZ))$ approaches the minimum conditional entropy, and that for every entropy estimator, there are PEFs whose log-prob rates approach the entropy estimate.

It is desirable to minimize the settings entropy used for randomness generation. For this, we consider *maximum probability estimators* for a model $\mathcal{C}$, which are defined as functions $G(CZ)$ such that $\mathbb{E}_\mu(G(CZ)) \geq \max_{c,z} \mu(c|z)$ for all $\mu \in \mathcal{C}$. Nontrivial maximum probability estimators exist. For example, every Bell inequality for the standard two-settings, two-outcomes Bell-test configuration for two parties has associated maximum probability estimators $G(CZ)$, and a distribution that violates the Bell inequality satisfies $\mathbb{E}_\mu(G(CZ)) < 1$. For every maximum probability estimator $G(CZ)$, we construct a family of PEFs for which the log-prob rates at $\mu \in \mathcal{C}$ approach $-\ln(\mathbb{E}_\mu(G(CZ)))$. We analyze this family of PEFs to determine how log-prob rates depend on number of trials and power, and find that with this family it is possible to achieve exponential expansion of settings entropy. For this, we consider models $\mathcal{C}_{C|Z}$ of distributions of C conditional on Z . An unconditional model $\mathcal{C}$ is obtained by specifying a probability distribution for the settings.

Theorem (Theorem 52). Let $G(CZ)$ be a maximum probability estimator for $\mathcal{C}$ where $\mathcal{C}$ is determined by the conditional model $\mathcal{C}_{C|Z}$ with the uniform settings distribution. Suppose that for a given $\mu \in \mathcal{C}$, $\mathbb{E}_\mu(G(CZ)) < 1$. Then for a constant error bound, there exists a family of PEFs and settings probability distributions determined by the number of trials n such that the settings entropy is $O(\ln(n))$ and the smooth conditional min-entropy of the outcomes is $\Omega(n)$.

The constants in the construction for exponential expansion are not of excessive size, but we consider the construction a proof of principle, not a practical proposal. Given that for relevant configurations, we can optimize PEFs directly, for finite experiments and the best expansion, it is preferable to use directly optimized PEFs. We determine expansion opportunities for the trial distribution observed in an atomic Bell test in Sec. IX.

III. BASIC CONCEPTS

A. Notation

Much of this work concerns stochastic sequences of random variables (RVs). RVs are functions on an underlying probability space. The range of an RV is called its value space. Here, all RVs have finite value spaces. We truncate sequences of RVs so that we only consider finitely many RVs at a time. With this we may assume that the underlying probability space is finite too. We use upper-case letters such as $A, B, \dots, X, Y, \dots$ to denote RVs. The value space of an RV such as X is denoted by $\text{Rng}(X)$. The cardinality of the value space of X is $|\text{Rng}(X)|$. Values of RVs are denoted by the corresponding lower-case letters. Thus x is a value of X , often thought of as the particular value realized in an experiment. In

the same spirit, we use Ω to denote the *universal* RV defined as the identity function on the set of the underlying probability space. Values of Ω are denoted by ω . When using symbols for values of RVs, they are implicitly assumed to be members of the range of the corresponding RV. In many cases, the value space is a set of letters or a set of strings of a given length. We use juxtaposition to denote concatenation of letters and strings. For a string s , $|s|$ denotes its length. Unless stated otherwise, a string-valued RV S produces fixed length strings. $|S|$ denotes the length of the strings, S_i is the i th letter of the string, and $S_{\leq i}$ is the length i prefix of the string. By default, strings are binary, which implies, for example, $|S| = \log_2(|\text{Rng}(S)|)$. Sequence RVs (stochastic sequences) are denoted by capital bold-face letters, with the corresponding lower-case bold-face letters for their values. For example, we write $\mathbf{A} = (A_i)_{i=1}^N$ and $\mathbf{A}_{\leq m} = (A_i)_{i=1}^m$. Our conventions for indices are that we generically use N to denote a large upper bound on sequence lengths, n to denote the available length and i, j, k, l, m as running indices. By convention, $\mathbf{A}_{\leq 0}$ is the empty sequence of RVs. Its value is constant, independent of Ω . When multiple stochastic sequences are in play, we refer to the collection of i th RVs in the sequences as the data from the i th *trial*. We typically imagine the trials as happening in time and being performed by an experimenter. We refer to the data from the trials preceding the upcoming one as the *past*. The past can also include initial conditions and any additional information that may have been obtained. These are normally implicit when referring to or conditioning on the past.

Probabilities are denoted by $\mathbb{P}(\dots)$. If there are multiple probability distributions involved, we disambiguate with a subscript such as in $\mathbb{P}_\nu(\dots)$ or simply $\nu(\dots)$, where ν is a probability distribution. We generally reserve the symbol μ for the global, implicit probability distribution, and may write $\mu(\dots)$ instead of $\mathbb{P}(\dots)$. Expectations are similarly denoted by $\mathbb{E}(\dots)$ or $\mathbb{E}_\mu(\dots)$. If ϕ is a logical expression involving RVs, then $\{\phi\}$ denotes the event where ϕ is true for the values realized by the RVs. For example, $\{f(X) > 0\}$ is the event $\{\omega : f(X(\omega)) > 0\}$ written in full set notation. The brackets $\{\dots\}$ are omitted for events inside $\mathbb{P}(\dots)$ or $\mathbb{E}(\dots)$. As is conventional, commas separating logical expressions are interpreted as conjunction. When the capital/lower-case convention can be unambiguously interpreted, we abbreviate “ $X = x$ ” by “ x ”. For example, with this convention, $\mathbb{P}(x, y) = \mathbb{P}(X = x, Y = y)$. Furthermore, we omit commas in the abbreviated notation, so $\mathbb{P}(xy) = \mathbb{P}(x, y)$. RVs or functions of RVs appearing outside an event but inside $\mathbb{P}(\dots)$ or after the conditioner in $\mathbb{E}(\dots | \dots)$ result in an expression that is itself an RV. We can define these without complications because of our assumption that the event space is finite. Here are two examples. $\mathbb{P}(f(X)|Y)$ is the RV whose value at ω is $\mathbb{P}(f(X) = f(X(\omega))|Y = Y(\omega))$. This is a function of the RVs X and Y and can be described as the RV whose value is $\mathbb{P}(f(X) = f(x)|Y = y)$ whenever the values of X and Y are x and y , respectively. Similarly $\mathbb{E}(X|Y)$ is the RV defined as a function of Y , with value $\mathbb{E}(X|Y = y)$ whenever Y has value y . Note that X plays a different role before the conditioners in $\mathbb{E}(\dots)$ than it does in $\mathbb{P}(\dots)$, as $\mathbb{E}(X|Y)$ is not a function of X , but only of Y . We comment that conditional probabilities with conditioners having probability zero are not well-defined, but in most cases can be defined arbitrarily. Typically, they occur

in a context where they are multiplied by the probability of the conditioner and thereby contribute zero regardless. An important context involves expectations, where we use the convention that when expanding an expectation over a finite set of values as a sum, zero-probability values are omitted. We do so without explicitly adding the constraints to the summation variables. We generally use conditional probabilities without explicitly checking for probability-zero conditioners, but it is necessary to monitor for well-definedness of the expressions obtained.

To denote general probability distributions, usually on the joint value spaces of RVs, we use symbols such as μ, ν, ρ , with modifiers as necessary. As mentioned, we reserve the unmodified μ for the distinguished global distribution under consideration, if there is one. Other symbols typically refer to probability distributions defined on the joint range of some subset of the available RVs. The set of probability distributions on $\text{Rng}(X)$ is denoted by $\mathcal{S}_X$. We usually just say “distribution” instead of “probability distribution.” The terms “distributions on $\text{Rng}(X)$ ” and “distributions of X ” are synonymous. The support of a distribution ν on X is denoted by $\text{Supp}(\nu) = \{x | \nu(x) > 0\}$. When multiple RVs are involved we denote marginal and conditional distributions by expressions such as $\mu[X|Y = y]$ for the distribution of X on its value space conditional on $\{Y = y\}$. The probability of x for this distribution can be denoted by $\mu[X|Y = y](x)$, which is well-defined for $\mathbb{P}(Y = y) > 0$ and therefore well-defined with probability one because of our finiteness assumptions. Note the use of square brackets to distinguish the distribution specification from the argument determining the probability at a point. If ν is a joint distribution of RVs, then we extend the conventions for arguments of $\mathbb{P}(\dots)$ to arguments of ν ,

as long as all the arguments are determined by the RVs for which ν is defined. For example, if ν is a joint distribution of X, Y , and Z , then $\nu(x|y)$ has the expected meaning, as does the RV $\nu(X|Y)$ in contexts requiring no other RVs. We denote the uniform distribution on $\text{Rng}(X)$ by Unif_X , omitting the subscript if the value space is clear from context. If R and S are independent RVs with marginal distributions $\nu = \mu[R]$ and $\nu' = \mu[S]$ on their ranges, then their joint distribution is denoted by $\mu[R, S] = \nu \otimes \nu'$.

In our work, probability distributions are constrained by a statistical *model*, which is defined as a set of distributions and denoted by letters such as $\mathcal{H}$ or $\mathcal{C}$. The models for trials to be considered here are usually convex and closed. For a model $\mathcal{C}$, we write $\text{Extr}(\mathcal{C})$ for the set of extreme points of $\mathcal{C}$ and $\text{Cvx}(\mathcal{C})$ for the convex closure of $\mathcal{C}$ defined as the smallest closed convex set containing $\mathcal{C}$.

The total variation (TV) distance between ν and ν' is defined as

$$\begin{aligned} \text{TV}(\nu, \nu') &= \sum_x (\nu(x) - \nu'(x)) \llbracket \nu(x) \geq \nu'(x) \rrbracket \\ &= \frac{1}{2} \sum_x |\nu(x) - \nu'(x)|, \end{aligned} \quad (1)$$

where $\llbracket \phi \rrbracket$ for a logical expression ϕ denotes the $\{0, 1\}$ -valued function evaluating to 1 iff ϕ is true. True to its name, the TV distance satisfies the triangle inequality. Here are three other useful properties. First, if ν and ν' are joint distributions of X and Y and the marginals satisfy $\nu[Y] = \nu'[Y]$, then the TV distance between ν and ν' is the average of the TV distances of the Y -conditional distributions:

$$\begin{aligned} \text{TV}(\nu, \nu') &= \sum_y \sum_x (\nu(x, y) - \nu'(x, y)) \llbracket \nu(x, y) \geq \nu'(x, y) \rrbracket \\ &= \sum_y \sum_x (\nu(x|y)\nu(y) - \nu'(x|y)\nu'(y)) \llbracket \nu(x|y)\nu(y) \geq \nu'(x|y)\nu'(y) \rrbracket \\ &= \sum_y \sum_x (\nu(x|y)\nu(y) - \nu'(x|y)\nu(y)) \llbracket \nu(x|y)\nu(y) \geq \nu'(x|y)\nu(y) \rrbracket \\ &= \sum_y \nu(y) \sum_x (\nu(x|y) - \nu'(x|y)) \llbracket \nu(x|y) \geq \nu'(x|y) \rrbracket \\ &= \sum_y \nu(y) \text{TV}(\nu[X|y], \nu'[X|y]). \end{aligned} \quad (2)$$

Second, if for all y , $\nu[X|y] = \nu'[X|y]$, then the TV distance between ν and ν' is given by the TV distance between the marginals on Y :

$$\begin{aligned} \text{TV}(\nu, \nu') &= \sum_y \sum_x (\nu(x, y) - \nu'(x, y)) \llbracket \nu(x, y) \geq \nu'(x, y) \rrbracket \\ &= \sum_y \sum_x (\nu(x|y)\nu(y) - \nu'(x|y)\nu'(y)) \llbracket \nu(x|y)\nu(y) \geq \nu'(x|y)\nu'(y) \rrbracket \\ &= \sum_y \sum_x (\nu(x|y)\nu(y) - \nu(x|y)\nu'(y)) \llbracket \nu(x|y)\nu(y) \geq \nu(x|y)\nu'(y) \rrbracket \\ &= \sum_y \sum_x \nu(x|y)(\nu(y) - \nu'(y)) \llbracket \nu(y) \geq \nu'(y) \rrbracket \end{aligned}$$

$$\begin{aligned}
&= \sum_y (v(y) - v'(y)) \mathbb{I}[v(y) \geq v'(y)] \\
&= \text{TV}(v[Y], v'[Y]).
\end{aligned} \tag{3}$$

Third, the TV distance satisfies the data-processing inequality. That is, for any stochastic process $\mathcal{E}$ on $\text{Rng}(X)$ and distributions v and v' of X , $\text{TV}(\mathcal{E}(v), \mathcal{E}(v')) \leq \text{TV}(v, v')$. We use this property only for functions $\mathcal{E}$, but for general forms of this result, see Ref. [32].

When constructing distributions close to a given one in TV distance, it is often convenient to work with subprobability distributions. A subprobability distribution of X is a subnormalized non-negative measure on $\text{Rng}(X)$, which in our case is simply a non-negative function $\tilde{v}$ on $\text{Rng}(X)$ with *weight* $w(\tilde{v}) = \sum_x \tilde{v}(x) \leq 1$. For expressions not involving conditionals, we use the same conventions for subprobability distributions as for probability distributions. When comparing subprobability distributions, $\tilde{v} \leq \tilde{v}'$ means that for all x , $\tilde{v}(x) \leq \tilde{v}'(x)$, and we say that $\tilde{v}'$ *dominates* $\tilde{v}$.

Lemma 1. Let $\tilde{v}$ be a subprobability distribution of X of weight $w = 1 - \epsilon$. Let v and v' be distributions of X satisfying $\tilde{v} \leq v$ and $\tilde{v} \leq v'$. Then $\text{TV}(v, v') \leq \epsilon$.

Proof. Calculate

$$\begin{aligned}
\text{TV}(v, v') &= \sum_x (v(x) - v'(x)) \mathbb{I}[v(x) \geq v'(x)] \\
&\leq \sum_x (v(x) - \tilde{v}(x)) \mathbb{I}[v(x) \geq \tilde{v}(x)] \\
&= \sum_x (v(x) - \tilde{v}(x)) \\
&= 1 - w = \epsilon.
\end{aligned} \tag{4}$$

Lemma 2. Assume that $p \geq 1/|\text{Rng}(X)|$. Let v be a distribution of X and $\tilde{v} \leq v$ a subprobability distribution of X with weight $w = 1 - \epsilon$ and $\tilde{v} \leq p$. Then there exists a distribution v' of X with $v' \geq \tilde{v}$, $v' \leq p$, and $\text{TV}(v, v') \leq \epsilon$.

Proof. Because $p \geq 1/|\text{Rng}(X)|$ and $\sum_x (p - \tilde{v}(x)) = |\text{Rng}(X)|p - w \geq 1 - w$, there exists a distribution $v' \geq \tilde{v}$ with $v' \leq p$. Since $v' \leq p$ and v are distributions dominating $\tilde{v}$ and by lemma 1, $\text{TV}(v, v') \leq \epsilon$. $\blacksquare$

B. Bell-test configurations

The standard example of a Bell-test configuration involves a pair of devices located at two stations, **A** and **B**. The stations engage in a sequence of trials. In a trial, each station chooses one of two settings and obtains a measurement outcome, either 0 or 1, from their device. A Bell test with this configuration is called a (2,2,2) Bell test, where the numbers indicate the number of stations, the number of settings available to each station and the number of outcomes at each setting. The test produces two stochastic sequences from each station. The sequence of settings choices is denoted by **X** for **A** and **Y** for **B**, and the sequence of measurement outcomes is denoted by **A** and **B**, respectively. In using this notation, we allow for arbitrary numbers of settings and measurement outcomes. For

a $(2, k, l)$ Bell-test configuration, $|\text{Rng}(X_i)| = |\text{Rng}(Y_i)| = k$ and $|\text{Rng}(A_i)| = |\text{Rng}(B_i)| = l$. The main role of separating the configuration by station is to justify the assumptions, which can be generalized to more stations if desired. The assumptions constrain the models consistent with the configuration. In this work, once the constraints are determined, the separation into stations plays little role, but we continue to identify settings and outcomes RVs. We use **Z** and **C** to denote the corresponding stochastic sequences. For the case of two stations, $\mathbf{Z} = (X_i Y_i)_{i=1}^N$ and $\mathbf{C} = (A_i B_i)_{i=1}^N$.

We also refer to RVs **D** and **R**. For **D**, we assume that D_i is determined by (that is a function of) C_i . We use **D** to study protocols where the randomness is extracted from **D** instead of **C**. For example, this includes protocols where only **A**'s measurement outcomes are used, in which case we set $\mathbf{D} = \mathbf{A}$. We assume that C_i is determined by R_i , and we use **R** to contain additional information accumulated during the experiment that can be used to adapt the protocol but is kept private. We may write $R_i = C_i R'_i$, where C_i and R'_i are the outcomes observable and the additional information obtainable in the i th trial, respectively. For example, if some calibration data is obtained after the i th trial and before the $(i + 1)$ th trial, we can use this calibration data to adapt the PEF for the $(i + 1)$ th trial, and so R'_i will include the information about this calibration data. We note that the special case where $D_i = C_i = R_i$ was already studied in our previous work [28]. In general, we need to differentiate the variables D_i , C_i and R_i as D_i is used for extracting randomness, C_i is for constraining the model for the i th trial (see Sec. V A), and R_i is for adapting the PEF construction. The main purpose of the RV **Z** is to contain data that may become public. This applies to the settings if the source of the random settings choices is public or otherwise not sufficiently trusted. In situations where there is only one trial under consideration or we do not need the structure of the stochastic sequences, we use the nonboldface variants of the RVs, that is D, C, R , and Z .

C. Assumptions for Bell tests

We are interested in limiting the information held by an external entity. This work is concerned with external entities holding classical side information, so the external entity **E**'s state is characterized by an RV E . Physical entities and systems are denoted by capital sans-serif letters, a convention we already used to refer to stations. All our results are proven for finite event spaces. While this is reasonable with respect to the RVs representing information explicitly used by the protocols, it is an unverifiable restriction on E . For countable E , one can use the observation that conditioning on large-probability finite subsets of E gives distributions that are arbitrarily close in TV distance to the unconditioned distribution.

Our theory is formulated in the untrusted device model [1] where **E** may have had the devices in possession before the trials of the protocol of interest start. Once the protocol starts,

E cannot receive information from the devices. Note that this does not preclude that E sends information to the devices via one-way channels, see below. This simplification is possible for local randomness-generation protocols where the participating parties need no remote classical communication. In many other applications such as quantum key distribution, the protocols involve both quantum and classical communication, in which case E can gain additional information at any time.

To ensure that E cannot receive information from the devices after the start of the protocol requires physical security to isolate the joint location of the devices and the stations. Formalizing this context requires a physical framework for which subsystems and interactions can be defined, with classical subsystems playing a special role. In the case of applications to Bell-test configurations, we motivate the constraints on the models with physical arguments, but we do not prove the constraints with respect to a formal physical framework. Let D be the pair of devices with which A and B implement the protocol. The absence of information passing from ABD to E is ensured if there is no physical interactions between the two sets of systems after the protocol starts. While Z is private, we can time shift any one-way communication from E to ABD , or noninteracting dynamics of E to before the beginning of the protocol. A communication from E to ABD can be described by adding the communication as a system C to E , which interacts with E then becomes isolated from E and later becomes part of the devices. This makes time shifting possible, after which we can assume that any physical processes relevant to the protocol act only on ABD . Of course, we insist that after time-shifting, the state of E is classical. This justifies the use of a single RV E to determine the state of E . In this work, we make few assumptions on the physics of D and formulate constraints on distributions purely in terms of the RVs produced by the protocol, conditional on E .

The formal restriction that E holds only classical side information allows for quantum E provided that there is an additional quantum system H independent of E such that the joint state at the start of the protocol of the systems ABD , H and E forms a quantum Markov chain $ABD \leftrightarrow H \leftrightarrow E$ [33]. This is a quantum generalization of stating that E and ABD are independent conditionally on H . The system H needs to have no interaction with E and ABD after the start of the protocol. For example, H can be part of a generic decohering environment. The Markov chain property implies that after an extension of E , without loss of generality, the information held by E can be assumed to be classical. Operationally, this situation can be enforced if we trust or ensure that the devices have no long-term quantum memory.

Two assumptions constraining the possible distributions are required for randomness generation with Bell tests. The settings constraints restrict the settings distributions, and the nonsignaling constraints enforce the absence of communication between the stations during a trial. There are different ways in which the settings distributions can be constrained. The strongest constraint we use is that Z_{i+1} is uniform, independent of $Z_{\leq i}$, $C_{\leq i}$ and E . In general, we can consider weaker constraints on Z_{i+1} . For example, Z_{i+1} may have any distribution that is independent of E but determined in a known way by the past. Most generally, Z_{i+1} has a distribution belonging to a specified set conditional on E and the past.

If we later allow Z to become public, or if Z is from a public source of randomness, then Z_{i+1} must be conditionally independent of $R_{\leq i}$ given $Z_{\leq i}$ and E . This constraint can be avoided if we trust that the data that normally contributes to Z remains private, in which case we can exploit that our framework is flexible in how the experimental data relates to the random variables: If normally Z contains settings choices, we can instead let Z be empty, add the settings choices into the RV C and extract randomness from all or some of C . With this assignment of RVs, conditional independence is not required. However, for nontrivial randomness generation, it is then desirable to extract strictly more random bits than were required for the settings choices. Otherwise the randomness obtained may just be due to the input randomness with no other contribution.

To state the nonsignaling constraints requires explicitly referencing each station's RVs. The nonsignaling constraints are assumed to be satisfied for trial i conditional on the past and E . With conditioning on these RVs implicit and for two stations, the nonsignaling constraints consist of the identities

$$\mathbb{P}(A_i|X_iY_i) = \mathbb{P}(A_i|X_i), \quad \mathbb{P}(B_i|X_iY_i) = \mathbb{P}(B_i|Y_i), \quad (5)$$

which assert remote settings independence of measurement outcomes. The nonsignaling constraints can be strengthened using the semidefinite-programming hierarchy for quantum distributions [34] when the devices are assumed to be quantum.

In any given implementation of randomness generation, the settings and nonsignaling constraints must be justified by physical arguments. The settings constraints require a source of random bits that are sufficiently independent of external entities at the last point in time where they had interactions with the devices and the stations. For protocols such as those discussed here, where the output's randomness is assured conditionally on the settings, the settings randomness can become known after the last interaction, so it is possible to use public sources of randomness. However, if the same public string is used by many randomness generators, a scattershot attack may be possible. To justify the nonsignaling constraints, this randomness must not leak to the stations' devices before it is used to apply the device settings. Thus we need to trust that it is possible to generate bits that are, before the time of use, unpredictable by the stations' devices. For the nonsignaling constraints, we also rely on physical isolation of the stations during a trial, preferably enforced by relativistic causality. It usually goes without saying that we trust the timing, recording and computing devices that monitor and analyze the data according to the protocol. However, given the complexity of modern electronics and the prevalence of reliability and security issues, caution is advised.

D. Min-entropy and randomness extraction

To be able to extract some amount of randomness from a bit-string RV D , it suffices to establish a min-entropy bound for the distribution of D . The min-entropy of D is $H_{\min}(D) = -\ln(p_{\max}(D))$, where $p_{\max}(D) = \max_d \mathbb{P}(d)$. By default, logarithms are base e , so entropy is expressed in *nits* (or *e*-bits), which simplifies calculus with entropic quantities. We convert to bits to determine string lengths as needed. We switch to

logarithms base 2 in Sec. IX for quantitative comparisons. Since we usually work with bounds on min-entropies rather than exact values, we say that D has min-entropy $-\ln(p)$ or D has max-prob p if $p_{\max} \leq p$ and add the adjective “exact” to refer to $-\ln(p_{\max})$ or $p_{\max}$. If D has min-entropy $\sigma \ln(2)$, then it is possible to extract close to σ near-uniform bits from D given some uniform seed bits. The actual number of near-uniform bits that can be safely extracted depends on $|D|$, how many uniform seed bits are used, the maximum acceptable TV distance from uniform of the near-uniform output bits, and the *extractor* used. An extractor $\mathcal{E}$ has as input a bit string D , a string of uniform seed bits S independent of all other relevant RVs, and the following parameters: a lower bound $\sigma_h \ln(2)$ on the min-entropy of D , the desired number σ of uniform bits, and the maximum acceptable TV distance from uniform ϵ_x of the output. We write $\mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)$ for the length σ bit string produced by the extractor. This bit string is close to uniform provided the input satisfies *extractor constraints* that depend on the specific extractor used. Formally, a *strong extractor* has the property that if the parameters $n = |D|$, $l = |S|$, σ_h , σ , ϵ_x satisfy its extractor constraints, S is an independent and uniform bit string, and D has min-entropy at least $\sigma_h \ln(2)$, then

$$\text{TV}(\mu[\mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)S], \text{Unif}) \leq \epsilon_x. \quad (6)$$

Extractors used in this work are strong by default. The statement that $(n, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints means that there exists l so that the full list of parameters $(n, l, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints. We always include $\sigma_h \leq n$, $1 \leq \sigma \leq \sigma_h$ and $0 < \epsilon_x \leq 1$ among the extractor constraints. Extractor constraints are partially determined by information theoretic bounds, but constraints for practical extractors are typically stronger. See Ref. [35].

In previous work [13], we have used an implementation of Trevisan’s strong extractor based on the framework of Maurer, Portmann, and Scholz [36] that we called the TMPS extractor $\mathcal{E}_{\text{TMPS}}$. The extractor constraints for $\mathcal{E}_{\text{TMPS}}$ are not optimized. Simplified constraints are $0 < \epsilon_x \leq 1$, $2 \leq \sigma \leq \sigma_h \leq n$, $\sigma \in \mathbb{N}$, and

$$\sigma + 4 \log_2(\sigma) \leq \sigma_h - 4 \log_2(1/\epsilon_x) - 6, \\ l \geq 36 \log_2(\sigma) (\log_2(4n\sigma^2/\epsilon_x^2))^2. \quad (7)$$

See Ref. [13] for the smaller expression for l used by the implementation. The TMPS extractor is designed to be secure against quantum side information. In general, strong extractors can be made classical-proof, that is, secure against classical side information, with a modification of the extractor constraints. Some strong extractors such as Trevisan’s can also be made quantum-proof, that is, secure against quantum side information, with a further modification of the constraints. See Ref. [36] for explanations and references to the relevant results, specifically lemma A.3, theorem B.4, lemma B.5, and lemma B.8. All our protocols assume strong extractors and are secure against classical side information. However, they do not require that the strong extractor be classical- or quantum-proof; in one case, the relevant constraint modification is implicit in our proof, and in the other cases, we take advantage of special properties of our techniques. The TMPS extractor constraints given above include the modifications

required for being quantum-proof; we did not roll back these modifications.

The TMPS extractor also satisfies that conditional on the seed bits, the output is a linear hash of D . For bits, a linear hash of D is a parity of D , computed as $\sum_i h_i D_i \bmod(2)$, where the $h_i \in \{0, 1\}$ form the parity vector. We call extractors with this property *linear extractors*.

Given that the extracted bits are allowed a nonzero distance from uniform, it is not necessary to satisfy a strict min-entropy bound on the extractor input D . It suffices for the distribution of D to have some acceptably small TV distance ϵ_h from a distribution with min-entropy $\sigma_h \ln(2)$. We say that D has ϵ_h -smooth min-entropy $\sigma_h \ln(2)$ if its distribution is within TV distance ϵ_h of one with min-entropy $\sigma_h \ln(2)$. For the situation considered so far in this section, the error bound (or *smoothness*) ϵ_h and the extractor error ϵ_x can be added when applying the extractor to D . In this work, we generally work directly with the maximum probability. We say that D has ϵ_h -smooth max-prob p if it has ϵ_h -smooth min-entropy $-\ln(p)$.

We generally want to generate bits that are near-uniform conditional on E and often other variables such as Z . For our analyses, E is not particularly an issue because our results hold uniformly for all values of E , that is, conditionally on $\{E = e\}$ for each e . However this is not the case for Z .

Definition 3. The distribution μ of DZE has ϵ -smooth ZE -conditional max-prob p if the following two conditions hold. (1) For each ze there exists a subprobability distribution $\tilde{\mu}_{ze}$ of D such that $\tilde{\mu}_{ze} \leq \mu[D|ze]$ and $\tilde{\mu}_{ze} \leq p$; and (2) the weights of these subprobability distributions satisfy $\sum_{ze} w(\tilde{\mu}_{ze}) \mu(ze) \geq 1 - \epsilon$. The minimum p for which μ has ϵ -smooth ZE -conditional max-prob p is denoted by $P_{\max, \mu}^{\epsilon, u}(D|ZE)$. (The superscript u alludes to the uniformity of the bound with respect to ZE .)

The distribution μ of DZE has ϵ -smooth average ZE -conditional max-prob p if there exists a distribution ν of DZE with $\text{TV}(\nu, \mu) \leq \epsilon$ and $\sum_{ze} \max_d (\nu(d|ze)) \nu(ze) \leq p$. The minimum p for which μ has ϵ -smooth average ZE -conditional max-prob p is denoted by $P_{\max, \mu}^{\epsilon}(D|ZE)$. The quantity $H_{\min, \mu}^{\epsilon}(D|ZE) = -\ln(P_{\max, \mu}^{\epsilon}(D|ZE))$ is the ϵ -smooth ZE -conditional min-entropy.

We refer to the smoothness parameters as *error bounds*. Observe that the definitions are monotonic in the error bound. For example, if $P_{\max, \mu}^{\epsilon} \leq p$ and $\epsilon' \geq \epsilon$, then $P_{\max, \mu}^{\epsilon'} \leq p$. The quantity $\sum_{ze} \max_d (\nu(d|ze)) \nu(ze)$ in the definition of $P_{\max, \mu}^{\epsilon}$ can be recognized as the (average) maximum guessing probability of D given Z and E (with respect to ν), whose negative logarithm is the guessing entropy defined, for example, in Ref. [27]. The relationships established below reprise results from the references. We use them to prove soundness of the first two protocols for composing probability estimation with randomness extractors (theorems 19 and 20) but bypass them for soundness of the third (theorem 21).

We focus on probabilities rather than entropies because in this work, we achieve good performance for finite data by direct estimates of probabilities of actual events, not entropies of distributions. While entropy estimates may be considered the ultimate goal for existing applications, they are not fundamental in our approach. The focus on probabilities helps us to avoid introducing logarithms unnecessarily.

A summary of the relationships between conditional min-entropies and randomness extraction is given in Ref. [37] for the quantum case and can be specialized to the classical case. When so specialized, the definition of smooth conditional min-entropy in, for example, Ref. [37] differs from the one above in that Ref. [37] uses one of the fidelity-related distances. One such distance reduces to the Hellinger distance h for probability distributions for which $h^2 \leq \text{TV} \leq \sqrt{2}h$.

Uniform or average bounds on Z -conditional max-probs with respect to $E = e$ can be lifted to the ZE -conditional max-probs, as formalized by the next lemma.

Lemma 4. If for all e , $P_{\max, \mu[DZ|e]}^{u, \epsilon}(D|Z) \leq p$, then $P_{\max, \mu[DZE]}^{u, \epsilon}(D|ZE) \leq p$. Suppose that for all e , $P_{\max, \mu[DZ|e]}^{\epsilon}(D|Z) \leq p_e$, and let $\bar{\epsilon} = \sum_e \epsilon_e \mu(e)$ and $\bar{p} = \sum_e p_e \mu(e)$. Then $P_{\max, \mu[DZE]}^{\bar{\epsilon}}(D|ZE) \leq \bar{p}$.

Proof. For the first claim, for each e , let $\tilde{\mu}_{ze}$ be subprobability distributions witnessing that $P_{\max, \mu[DZ|e]}^{u, \epsilon}(D|Z) \leq p$ according to the definition. Then $\tilde{\mu}_{ze} \leq \mu[D|ze]$, $\tilde{\mu}_{ze} \leq p$ and $\sum_z w(\tilde{\mu}_{ze})\mu(z|e) \geq 1 - \epsilon$, from which we get

$$\begin{aligned} \sum_{ze} w(\tilde{\mu}_{ze})\mu(ze) &= \sum_e \mu(e) \sum_z w(\tilde{\mu}_{ze})\mu(z|e) \\ &\geq \sum_e \mu(e)(1 - \epsilon) = (1 - \epsilon). \end{aligned} \quad (8)$$

For the second claim, for each e , let v_e witness $P_{\max, \mu[DZ|e]}^{\epsilon}(D|Z) \leq p_e$. Then $\text{TV}(v_e, \mu[DZ|e]) \leq \epsilon_e$ and $\sum_z \max_d(v_e(d|z))v_e(z) \leq p_e$. Define v by $v(dze) = v_e(dz)\mu(e)$. Then $v[E] = \mu[E]$, so we can apply Eq. (2) for

$$\text{TV}(v, \mu) = \sum_e \text{TV}(v_e, \mu[DZ|e])\mu(e) \leq \sum_e \epsilon_e \mu(e) = \bar{\epsilon}. \quad (9)$$

Furthermore,

$$\begin{aligned} \sum_{ze} \max_d(v(d|ze))v(ze) &= \sum_e \mu(e) \sum_z \max_d(v_e(d|z))v(z|e) \\ &\leq \sum_e \mu(e)p_e = \bar{p}, \end{aligned} \quad (10)$$

as required for the conclusion. ■

The main relationships between conditional and average conditional max-probs are determined by the next lemma.

Lemma 5. Let $p \geq 1/|\text{Rng}(D)|$, $\epsilon > 0$ and μ a distribution of DZE . If $P_{\max, \mu}^{u, \epsilon}(D|ZE) \leq p$, then $P_{\max, \mu}^{\epsilon}(D|ZE) \leq p$. If $P_{\max, \mu}^{\epsilon}(D|ZE) \leq p\delta$, then there exists a distribution v' of DZE with $\text{TV}(v', \mu) \leq \epsilon + \delta$ such that $P_{\max, v'}^{u, 0}(D|ZE) \leq p$. If $P_{\max, \mu}^{\epsilon}(D|ZE) \leq p\delta$, then $P_{\max, \mu}^{u, 2\epsilon + \delta}(D|ZE) \leq p$.

Proof. Suppose that $P_{\max, \mu}^{u, \epsilon}(D|ZE) \leq p$. Let $\tilde{\mu}_{ze}$ be subprobability distributions witnessing this inequality as required by the definition. Since $p \geq 1/|\text{Rng}(D)|$ and by lemma 2, there exist distributions v_{ze} with $v_{ze} \geq \tilde{\mu}_{ze}$, $v_{ze} \leq p$ and $\text{TV}(v_{ze}, \mu[D|ze]) \leq 1 - w(\tilde{\mu}_{ze})$. Define v by $v(dze) =$

$v_{ze}(d)\mu(ze)$. Then $v[ZE] = \mu[ZE]$, and according to Eq. (2)

$$\begin{aligned} \text{TV}(v, \mu) &= \sum_{ze} \text{TV}(v_{ze}, \mu[D|ze])\mu(ze) \\ &\leq \sum_{ze} (1 - w(\tilde{\mu}_{ze}))\mu(ze) \leq 1 - (1 - \epsilon) = \epsilon. \end{aligned} \quad (11)$$

Since $\sum_{ze} \max_d(v_{ze}(d))v(ze) \leq \sum_{ze} p\mu(ze) = p$, the first part of the lemma follows.

For the second part of the lemma, suppose that $P_{\max, \mu}^{\epsilon}(D|ZE) \leq p\delta$ and let v be a distribution witnessing this inequality. By definition, $\text{TV}(v, \mu) \leq \epsilon$ and $\sum_{ze} \max_d(v(d|ze))v(ze) \leq p\delta$. Let $m(ze) = \max_d(v(d|ze))$. By Markov's inequality, $v(m(ze) > p) \leq \delta$. Let $v'(dze) = v(dze)$ if $m(ze) \leq p$ and $v'(dze) = v(ze)/|\text{Rng}(D)|$ otherwise. Then $\text{TV}(v', v) \leq \delta$ and for all ze , $\max_d(v'(d|ze)) \leq p$. Note that $v'[ZE] = v[ZE]$. By the triangle inequality, $\text{TV}(v', \mu) \leq \epsilon + \delta$. By letting $\tilde{v}'_{ze} = v'[D|ze]$ in the definition for ZE -conditional max-prob, we see that $P_{\max, v'}^{u, 0}(D|ZE) \leq p$.

For the last part of the lemma, define $v''(dze) = v'(d|ze)\mu(ze)$. Then

$$\begin{aligned} \text{TV}(v', v'') &= \text{TV}(v'[ZE], v''[ZE]) \\ &= \text{TV}(v[ZE], \mu[ZE]) \leq \epsilon, \end{aligned} \quad (12)$$

where we used the identities $v'[D|ze] = v''[D|ze]$ and Eq. (3), and applied the data-processing inequality. Let $\tilde{\mu}_{ze}(d) = \min(v''(d|ze), \mu(d|ze)) = \min(v'(d|ze), \mu(d|ze))$. Then the $\tilde{\mu}_{ze}$ are subprobability distributions, $\tilde{\mu}_{ze} \leq p$, $\tilde{\mu}_{ze} \leq \mu[D|ze]$ and

$$\begin{aligned} \sum_{ze} w(\tilde{\mu}_{ze})\mu(ze) &= \sum_{zed} \min(v''(d|ze), \mu(d|ze))\mu(ze) \\ &= \sum_{zed} \min(v''(dze), \mu(dze)) \\ &= \sum_{zed} \mu(dze) - \sum_{zed} (\mu(dze) - v''(dze)) \\ &\quad \times \mathbb{I}[\mu(dze) > v''(dze)] \\ &= 1 - \text{TV}(v'', \mu) \\ &\geq 1 - \text{TV}(v'', v') - \text{TV}(v', v) - \text{TV}(v, \mu) \\ &\geq 1 - 2\epsilon - \delta, \end{aligned} \quad (13)$$

so $\tilde{\mu}_{ze}$ is as required by the definition. ■

We remark for emphasis that as noted in the proof of the lemma, the construction of v' satisfies that $v'[ZE] = v[ZE]$, where v is the witness of the assumption in the second part.

The smoothness parameter for smooth conditional max-prob composes well with strong extractors. See also Ref. [27], proposition 1.

Lemma 6. Suppose that the distribution μ of DZE satisfies $P_{\max, \mu}^{u, \epsilon_h}(D|ZE) \leq 2^{-\sigma_h}$, and S is a uniform and independent seed string with respect to μ . If $(n = |D|, l = |S|, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints, then $\text{TV}(\mu[\mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)SZE], \text{Unif} \otimes \mu[ZE]) \leq \epsilon_x + \epsilon_h$.

Proof. Let $\tilde{\mu}_{ze}$ be subprobability distributions witnessing $P_{\max, \mu}^{u, \epsilon_h}(D|ZE) \leq 2^{-\sigma_h}$ as required by the definition, and define v as we did at the beginning of the proof of the first part

of lemma 5, with $p = 2^{-\sigma_h}$ and $\epsilon = \epsilon_h$. Then $v[D|ZE] \leq p$, $v[ZE] = \mu[ZE]$, and by Eq. (11), $\text{TV}(v, \mu) \leq \epsilon_h$. According to the extractor specification, for all ze , we have $\text{TV}(v[\mathcal{E}(D, S)|ze], \text{Unif}) \leq \epsilon_x$. Since $v[ZE] = \mu[ZE]$, we can apply Eq. (2) to get

$$\begin{aligned} & \text{TV}(v[\mathcal{E}(D, S)SZE], \text{Unif} \otimes \mu[ZE]) \\ &= \sum_{ze} \text{TV}(v[\mathcal{E}(D, S)|ze], \text{Unif}) \mu(ze) \\ &\leq \sum_{ze} \epsilon_x \mu(ze) = \epsilon_x. \end{aligned} \quad (14)$$

By the data-processing inequality, $\text{TV}(v[\mathcal{E}(D, S)SZE], \mu[\mathcal{E}(D, S)SZE]) \leq \text{TV}(v, \mu) \leq \epsilon_h$. The result now follows by the triangle inequality. ■

By means of the third part of lemma 5, we can also compose smooth average conditional max-prob with a strong extractor.

Lemma 7. Suppose that the distribution μ of DZE satisfies $P_{\max, \mu}^{\epsilon_h}(D|ZE) \leq 2^{-\sigma_h} \delta$, and S is a uniform and independent seed string with respect to μ . If $(n = |D|, l = |S|, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints, then $\text{TV}(\mu[\mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)SZE], \text{Unif} \otimes \mu[ZE]) \leq \epsilon_x + 2\epsilon_h + \delta$.

Proof. From the last part of lemma 5, we have $P_{\max, \mu}^{2\epsilon_h + \delta}(D|ZE) \leq 2^{-\sigma_h}$. The result now follows from lemma 6. ■

E. Randomness generation protocols

A randomness generation protocol $\mathcal{P}$ is parameterized by σ , the requested number of uniform bits and ϵ , the protocol error bound, defined as the TV distance from uniform of these bits. Minimally, the output consists of a string of σ bits. In addition, the output can contain an additional string containing other random bits that may have been used internally for seeding the extractor or choosing settings, in case these bits can be used elsewhere. Further, the protocol may output a flag indicating success or failure. If the probability of success can be less than 1, the protocol may require as input a minimum acceptable probability of success κ . We write $\mathcal{P} = (\mathcal{P}_X, \mathcal{P}_S, \mathcal{P}_P)$, where $\mathcal{P}_X$ is the principal σ -bit output, $\mathcal{P}_S$ is the possibly empty but potentially reusable string of seed and other random bits, and $\mathcal{P}_P$ is the success flag, which is 0 (“fail”) or 1 (“pass”). Here, we have suppressed the arguments σ and ϵ (or σ, ϵ and κ) of $\mathcal{P}, \mathcal{P}_X, \mathcal{P}_S$, and $\mathcal{P}_P$, since they are clear from context. The outputs of $\mathcal{P}$ are treated as RVs, jointly distributed with E and any other RVs relevant to the situation. In this context, we constrain the joint distribution of the RVs according to a model $\mathcal{H}$. When $\mathcal{P}_S$ is nonempty, we assume that the marginal distribution $\mu[\mathcal{P}_S]$ is known to the user.

The property that a protocol output satisfies the request is called *soundness*. We formulate soundness so that the TV distance is conditional on success and ϵ absorbs the probability of success criterion in a sense to be explained.

Definition 8. The randomness generation protocol $\mathcal{P}$ is (σ, ϵ) sound with respect to E' and model $\mathcal{H}$ if for all $\mu \in \mathcal{H}$,

$|\mathcal{P}_X| = \sigma$ and there is a distribution $v_{E'}$ of E' such that

$$\begin{aligned} & \text{TV}(\mu[\mathcal{P}_X \mathcal{P}_S E' | \mathcal{P}_P = 1], \text{Unif}_{\mathcal{P}_X} \otimes \mu[\mathcal{P}_S] \\ & \otimes v_{E'}) \mathbb{P}(\mathcal{P}_P = 1) \leq \epsilon. \end{aligned} \quad (15)$$

We normally use $E' = ZE$, but $E' = E$ is an option if the settings choices are private. The definition ensures that $\mathcal{P}$ is nearly indistinguishable, namely within TV distance ϵ , from an *ideal* protocol with the same probability of success, namely, such a protocol for which $\mathcal{P}_X$ is perfectly uniform and independent of other variables conditional on success.

An alternative definition of soundness is to make the minimum acceptable probability of success κ explicit and require that either $\mathbb{P}(\mathcal{P}_P = 1) < \kappa$ or the conditional TV distance in Eq. (15) is bounded by ϵ . We refer to a protocol satisfying this property as $(\sigma, \epsilon, \kappa)$ sound. The two definitions are closely related:

Lemma 9. Fix E' and $\mathcal{H}$. Let $\kappa > 0$ and $\epsilon > 0$. Suppose that $\mathcal{P}$ is a $(\sigma, \epsilon\kappa)$ sound randomness generation protocol, then $\mathcal{P}$ is $(\sigma, \epsilon, \kappa)$ sound. If $\mathcal{P}$ is $(\sigma, \epsilon, \kappa)$ sound, then $\mathcal{P}$ is $(\sigma, \max(\epsilon, \kappa))$ sound.

Proof. Suppose that $\mathcal{P}$ is $(\sigma, \epsilon\kappa)$ sound. Consider any $\mu \in \mathcal{H}$ and let $v_{E'}$ be the distribution of E' in Eq. (15). If $\mathbb{P}(\mathcal{P}_P = 1) \geq \kappa$, then the desired TV distance is at most $\epsilon\kappa/\kappa = \epsilon$. Since $\mu \in \mathcal{H}$ is arbitrary, $\mathcal{P}$ is $(\sigma, \epsilon, \kappa)$ sound.

For the other direction, again consider any $\mu \in \mathcal{H}$. If $\mathbb{P}(\mathcal{P}_P = 1) < \kappa$, then the left-hand side of Eq. (15) is at most κ , because the TV distance is bounded by 1. If $\mathbb{P}(\mathcal{P}_P = 1) \geq \kappa$, let $v_{E'}$ be such that the TV distance in Eq. (15) is bounded by ϵ . Then the full left-hand side of this equation is bounded by $\epsilon \mathbb{P}(\mathcal{P}_P = 1) \leq \epsilon$. The conclusion follows. ■

We prefer to use (σ, ϵ) soundness, since this is often more informative than $(\sigma, \epsilon, \kappa)$ soundness, allowing for flexibility in how the protocol is used without looking inside the soundness proof. However, in both cases, soundness proofs may contain additional information about the relationship between the probability of success and the error bound that is useful when implementing the protocol in a larger context. For example, the soundness established in theorem 19 establishes a relationship between relevant error bounds that does not match either definition, but readily implies forms of either one.

In studies of specific randomness generation protocols, an important consideration is *completeness*, which requires that the actual probability of success of the protocol is nontrivially large, preferably exponentially close to 1 with respect to relevant resource parameters. Completeness is readily satisfied by our protocols for typical Bell-test configurations.

The actual probability of success should be distinguished from any relevant minimum acceptable probability of success in view of the discussion above. Experimental implementations so far demonstrate that success probabilities are acceptable, but not arbitrarily close to 1. As discussed in the introduction, theory shows that there are randomness generation protocols using quantum systems that can achieve high success probabilities. Here we emphasize protocols for which the success probability is 1 by allowing for injection of banked randomness when insufficient randomness is available from the allotted resources. Assuming completeness, we can also take advantage of the ability to stop the protocol only

when enough randomness is generated, but since this requires care when extracting near-uniform bits from potentially long strings, we do not develop this approach further here.

F. Test factors and test supermartingales

Definition 10. A test supermartingale [16] with respect to a stochastic sequence $\mathbf{R}$ and model $\mathcal{H}$ is a stochastic sequence $\mathbf{T} = (T_i)_{i=0}^N$ with the properties that $T_0 = 1$, for all i $T_i \geq 0$, T_i is determined by $\mathbf{R}_{\leq i}$ and the governing distribution, and for all distributions in $\mathcal{H}$, $\mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i}) \leq T_i$. The ratios $F_i = T_i/T_{i-1}$ with $F_i = 1$ if $T_{i-1} = 0$ are called the *test factors* of $\mathbf{T}$.

Here, $\mathbf{R}$ captures the relevant information that accumulates in a sequence of trials. It does not need to be accessible to the experimenter. The σ -algebras induced by $\mathbf{R}_{\leq i}$ define the nested sequence of σ -algebras used in more general formulations. Between trials i and $i+1$, the sequence $\mathbf{R}_{\leq i}$ is called the past. In the definition, we allow for T_i to depend on the governing distribution μ . With this, for a given μ , T_i is a function of $\mathbf{R}_{\leq i}$. Below, when stating that RVs are determined, we implicitly include the possibility of dependence on μ without mention. The μ dependence can arise through expressions such as $\mathbb{E}_\mu(G|\mathbf{R}_{\leq i})$ for some G , which is determined by $\mathbf{R}_{\leq i}$ given μ . One way to formalize this is to consider μ -parameterized families of RVs. We do not make this explicit and simply allow for our RVs to be implicitly parameterized by μ . We note that the governing distribution in a given experiment or situation is fixed but usually unknown with most of its features inaccessible. As a result, many RVs used in mathematical arguments cannot be observed even in principle. Nevertheless, they play important roles in establishing relationships between observed and inferred quantities.

Defining $F_i = 1$ when $T_{i-1} = 0$ makes sense because given $\{T_{i-1} = 0\}$, we have $\{T_i = 0\}$ with probability 1. The sequence $\mathbf{F}$ satisfies the conditions that for all i , (1) $F_i \geq 0$, (2) F_i is determined by $\mathbf{R}_{\leq i}$, and (3) for all distributions in $\mathcal{H}$, $\mathbb{E}(F_{i+1}|\mathbf{R}_{\leq i}) \leq 1$. We can define test supermartingales in terms of such sequences: Let $\mathbf{F}$ be a stochastic sequence satisfying the three conditions. Then the stochastic sequence with members $T_0 = 1$ and $T_i = \prod_{1 \leq j \leq i} F_j$ for $i \geq 1$ is a test supermartingale. It suffices to check that $\mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i}) \leq T_i$. This follows from

$$\mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i}) = \mathbb{E}(F_{i+1}T_i|\mathbf{R}_{\leq i}) = \mathbb{E}(F_{i+1}|\mathbf{R}_{\leq i})T_i \leq T_i, \quad (16)$$

where we pulled out the determined quantity T_i from the conditional expectation. In this work, we construct test supermartingales from sequences $\mathbf{F}$ with the above properties. We refer to any such sequence as a sequence of test factors, without necessarily making the associated test supermartingale explicit. We extend the terminology by calling an RV F a test factor with respect to $\mathcal{H}$ if $F \geq 0$ and $\mathbb{E}(F) \leq 1$ for all distributions in $\mathcal{H}$.

For an overview of test supermartingales and their properties, see Ref. [16]. The notion of test supermartingales and proofs of their basic properties are due to Ville [38] in the same work that introduced the notion of martingales. The name “test supermartingale” appears to have been introduced in Ref. [16]. Test supermartingales play an important theoretical role in proving many results in martingale theory, including that of proving tail bounds for large classes of

martingales. They have been studied and applied to Bell tests [14,15,39].

The definition implies that for a test supermartingale $\mathbf{T}$, for all n , $\mathbb{E}(T_n) \leq 1$. This follows inductively from $\mathbb{E}(T_{i+1}) = \mathbb{E}(\mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i})) \leq \mathbb{E}(T_i)$ and $T_0 = 1$. An application of Markov’s inequality shows that for all $\epsilon > 0$,

$$\mathbb{P}(T_n \geq 1/\epsilon) \leq \epsilon. \quad (17)$$

Thus a large final value $t = T_n$ of the test supermartingale is evidence against $\mathcal{H}$ in a hypothesis test with $\mathcal{H}$ as the (composite) null hypothesis. Specifically, the RV $1/T$ is a p -value bound against $\mathcal{H}$, where in general, the RV U is a p -value bound against $\mathcal{H}$ if for all distributions in $\mathcal{H}$, $\mathbb{P}(U \leq \epsilon) \leq \epsilon$. This result can be strengthened as follows. Define $T^* = \max_{1 \leq i \leq n} T_i$.

Theorem 11. For a test supermartingale $\mathbf{T}$ with respect to $\mathbf{R}$ and $\mathcal{H}$, the RV $1/T^*$ is a p -value bound against $\mathcal{H}$.

This theorem is due to Ville [38] Ch. V.2.2 (pp. 97 and 98).

Proof. The theorem follows from Doob’s maximal inequalities. The particular inequality we need is normally stated for a non-negative submartingale $\mathbf{T}'$ in the form

$$\mathbb{P}(\max_{1 \leq i \leq n} T'_i \geq \lambda) \leq \mathbb{E}(T'_n)/\lambda, \quad (18)$$

where $\lambda = 1/\epsilon$ for our purposes. Statements and proofs of this inequality are readily found online. A textbook treatment is in Ref. [40] Ch. V, Cor. 22. To apply the above maximal inequality to the test supermartingale $\mathbf{T}$, let $\tilde{T}_{i+1} = \mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i})$. Note that $\tilde{T}_{i+1}$ is determined by $\mathbf{R}_{\leq i}$ and by the definition of supermartingales, $\tilde{T}_{i+1} \leq T_i$. Define $T'_0 = 1$ and $T'_i = \prod_{1 \leq j \leq i} T_j/\tilde{T}_j$. By pulling out determined quantities from the conditional expectation, we get

$$\begin{aligned} \mathbb{E}(T'_{i+1}|\mathbf{R}_{\leq i}) &= T'_i \mathbb{E}(T_{i+1}/\tilde{T}_{i+1}|\mathbf{R}_{\leq i}) \\ &= T'_i \mathbb{E}(T_{i+1}|\mathbf{R}_{\leq i})/\tilde{T}_{i+1} = T'_i. \end{aligned} \quad (19)$$

Hence $\mathbf{T}'$ is a test martingale. In particular, it is a non-negative submartingale with $\mathbb{E}(T'_n) = 1$. We claim that $T'_i \geq T_i$. This holds for $i = 0$. For a proof by induction, compute $T'_{i+1} = T'_i T_{i+1}/\tilde{T}_{i+1} \geq T_i T_{i+1}/\tilde{T}_{i+1} \geq T_i T_{i+1}/T_i = T_{i+1}$. This gives $\max_{1 \leq i \leq n} T'_i \geq T^*$, so the event $\{T^* \geq \lambda\}$ implies the event in the maximal inequality, and monotonicity of probabilities implies the theorem. ■

One can produce a test supermartingale adaptively by determining the test factors F_{i+1} to be used at the next trial. If the i th trial’s data is R_i , including any incidental information obtained, F_{i+1} is expressed as a function of $\mathbf{R}_{\leq i}$ and data from the $(i+1)$ th trial (a *past-parameterized* function of R_{i+1}), and constructed to satisfy $F_{i+1} \geq 0$ and $\mathbb{E}(F_{i+1}|\mathbf{R}_{\leq i}) \leq 1$ for any distribution in the model $\mathcal{H}$. Note that in between trials, we can effectively stop the experiment by assigning all future $F_{i+1} = 1$ conditional on the past. This is equivalent to constructing the stopped process relative to a stopping rule. This argument also shows that the stopped process is still a test supermartingale. Here we consider only bounded length test supermartingales (with very large bound N if necessary), so are not concerned with questions of convergence as $N \rightarrow \infty$. But we note that since $\mathbb{E}(|F_i|) = \mathbb{E}(F_i) \leq 1$ for all i , Doob’s martingale convergence theorem applies, and $\lim_{i \rightarrow \infty} F_i$ exists almost surely and is integrable. Furthermore, this limit also

has expectation at most 1. See, for example, Ref. [40] Ch. V, theorem 28.

More generally, we use test supermartingales for estimating lower bounds on products of positive stochastic sequences $\mathbf{G}$. Such lower bounds are associated with unbounded-above confidence intervals. We need the following definition:

Definition 12. Let U, V, X be RVs and $1 \geq \epsilon \geq 0$. $I = [U, V]$ is a (conservative) *confidence interval for X at level ϵ with respect to $\mathcal{H}$* if for all distributions in $\mathcal{H}$ we have $\mathbb{P}(U \leq X \leq V) \geq 1 - \epsilon$. The quantity $\mathbb{P}(U \leq X \leq V)$ is called the *coverage probability*.

As noted above, the RVs U, V , and X may be μ -dependent. For textbook examples of confidence intervals, X is a parameter determined by μ , and U and V are calculated from an estimation error. We need the full generality of the definition above. The quantity ϵ in the definition is a significance level, which corresponds to a confidence level of $(1 - \epsilon)$.

Lemma 13. Let $\mathbf{F}$ and $\mathbf{G}$ be two stochastic sequences with $F_i \in [0, \infty)$, $G_i \in (0, \infty]$, and F_i and G_i determined by $\mathbf{R}_{\leq i}$. Define $T_0 = 1$, $T_i = \prod_{1 \leq j \leq i} F_j$ and $U_0 = 1$, $U_i = \prod_{1 \leq j \leq i} G_j$, and suppose that for all $\mu \in \mathcal{H}$, $\mathbb{E}(F_{i+1}/G_{i+1} | \mathbf{R}_{\leq i}) \leq 1$. Then $[T_n \epsilon, \infty)$ is a confidence interval for U_n at level ϵ with respect to $\mathcal{H}$. If U_i is monotone nondecreasing, then $[T^* \epsilon, \infty)$ is a confidence interval for U_n at level ϵ with respect to $\mathcal{H}$.

Proof. The assumptions imply that the F_{i+1}/G_{i+1} form a sequence of test factors with respect to $\mathcal{H}$ and generate the test supermartingale $\mathbf{T}/\mathbf{U}$, where division in this expression is term-by-term. Therefore

$$\mathbb{P}(T_n \epsilon \geq U_n) = \mathbb{P}(T_n/U_n \geq 1/\epsilon) \leq \epsilon, \quad (20)$$

So $[T_n \epsilon, \infty)$ is a confidence interval for U_n at level ϵ .

If U_i is monotone nondecreasing, observe that the event $\{\max_{1 \leq i \leq n} T_i \epsilon \geq U_n\}$ is the same as $\{(\max_{1 \leq i \leq n} T_i)/U_n \geq 1/\epsilon\}$. This event in turn implies $\{\max_{1 \leq i \leq n} T_i/U_i \geq 1/\epsilon\}$, because U_n being nondecreasing implies that for all i , $T_i/U_n \leq T_i/U_i$. Applying theorem 11, we obtain

$$\begin{aligned} \mathbb{P}(T^* \epsilon \geq U_n) &= \mathbb{P}(T^*/U_n \geq 1/\epsilon) \\ &\leq \mathbb{P}(\max_{1 \leq i \leq n} T_i/U_i \geq 1/\epsilon) \leq \epsilon. \end{aligned} \quad (21)$$

It follows that $[T^* \epsilon, \infty)$ is a confidence interval for U_n at level ϵ . ■

G. From Bell functions to test factors

Consider a specific trial subject to settings and nonsignaling constraints, where the settings distribution is fixed and known. To simplify notation, we omit trial indices and conditioning on the past. Let $\nu = \mu[Z]$ be the settings distribution. A *Bell function* B maps settings and outcomes to real values and satisfies $\mathbb{E}(B(CZ)) \leq 0$ for all *local realistic* (LR) distributions with the given settings distribution. The set of LR distributions given that the settings distribution is ν consists of mixtures of nonsignaling distributions with outcomes determined by the settings. These are distributions such that $C = AB$ is a function of $Z = XY$ for which A does not depend on Y and B does not depend on X . See Sec. IX, Eq. (158) for more detail. The expression $\mathbb{E}(B(CZ)) \leq 0$ is closely related to traditional Bell inequalities, which may be expressed in the *conditional form* $\sum_z \mathbb{E}(B'(Cz)|z) \leq 0$.

Provided $\nu(z) = 0$ implies $B'(cz) = 0$ for all c , any Bell inequality in this form can be converted to a Bell function by defining $B(cz) = B'(cz)/\nu(z)$. Conversely, a Bell function B for settings distribution ν determines a Bell inequality by defining $B'(cz) = B(cz)\nu(z)$. Bell inequalities apply to LR distributions independent of what the settings probabilities are, and these settings probabilities are then considered to be the free choice of the experimenter. We do not use this perspective here.

Let $-l$ be a lower bound for the Bell function B with $l > 0$. Then $F = (B + l)/l$ is a test factor with respect to LR distributions. Such test factors can provide an effective method for rejecting local realism (LR) at high significance levels by use of Eq. (17). As an example, consider the (2,2,2) Bell-test configuration, with the uniform settings distribution, $\nu = 1/4$. The ranges of the settings X, Y and the outcomes A, B are $\{0, 1\}$. The following is a Bell function:

$$\begin{aligned} B(xyab) &= (\llbracket xy = 11 \rrbracket - \llbracket xy = 01 \rrbracket - \llbracket xy = 10 \rrbracket \\ &\quad - \llbracket xy = 00 \rrbracket) |a - b|. \end{aligned} \quad (22)$$

The inequality $\mathbb{E}(B(CZ)) \leq 0$ is equivalent to the well-known Clauser-Horne-Shimony-Holt (CHSH) inequality [41]. We give the function in this form for simplicity and because one way to verify that B is a Bell function is to note that $d(a, b) = |a - b|$ satisfies the triangle inequality, so this Bell function belongs to the class of distance-based Bell functions. See Refs. [42,43]. Since for all arguments, exactly one of the expressions inside $\llbracket \dots \rrbracket$ is true, the minimum value of B is -1 . Thus $B + 1$ is a test factor. More generally, $1 + \lambda B$ is a test factor for any $0 \leq \lambda \leq 1$ [15]. By optimizing these test factors, asymptotically optimal statistical strength [expected $-\ln(p\text{-value})$ per trial] for rejecting LR can be achieved [14,15].

While probability estimation uses probability estimation factors (PEFs) to assign numerical values to experimental outcomes, it does not require Bell functions or even Bell-test configurations. The connection of Bell functions to test factors serves as an instructive example and can yield PEFs useful for probability estimation as witnessed implicitly by Ref. [13]. See also Sec. VII A. In general, optimal PEFs cannot be constructed from Bell functions. Therefore, in this work, we prefer to directly optimize PEFs without referencing Bell functions, see Secs. V and VII.

IV. PROBABILITY ESTIMATION AND RANDOMNESS EXTRACTION

A. Probability estimation

Consider RVs R, Z , and E whose joint distribution is in the model $\mathcal{H}$. In an experiment, we observe Z, C , and possibly other parts of R , but not E . For this section, it is not necessary to structure the RVs as stochastic sequences, so we use D, C, Z , and R in place of $\mathbf{D}, \mathbf{C}, \mathbf{Z}$, and $\mathbf{R}$, but follow the conventions introduced in Sec. III B. In particular, we assume that D is determined by C , and C is determined by R . There is no loss of generality in this assumption. It simply allows us to omit C (or D) as arguments of functions and in conditioners when R is already present.

We focus on probability estimates with coverage probabilities that do not depend on E , formalized as follows.

Definition 14. Let $1 \geq \epsilon \geq 0$. The function $F : \text{Rng}(RZ) \rightarrow [0, \infty)$ is a *level- ϵ E -uniform probability estimator* for $\mathcal{H}$ [ϵ -UPE or with specifics, ϵ -UPE($D:R|Z; E, \mathcal{H}$)] if for all e and distributions μ in $\mathcal{H}$, we have $\mu(F(RZ) \geq \mu(D|Ze)|e) \geq 1 - \epsilon$. We omit specifics such as $\mathcal{H}$ if they are clear from context. The function $\mathcal{U} : \text{Rng}(RZ) \times [0, 1] \rightarrow [0, \infty)$ is a UPE if it is monotone nondecreasing in the second argument (the confidence level) and for each ϵ , $\mathcal{U}(RZ, 1 - \epsilon)$ is an ϵ -UPE.

The level of a probability estimator relates to the smoothness parameter for smooth min-entropy via the relationships established below in Sec. IV B. We also use the term *error bound* to refer to the level of a probability estimator. The first condition on $\mathcal{U}$ ensures that the confidence upper bounds that it provides are nondecreasing with confidence level, so that the corresponding confidence intervals are consistently nested. The second is the required minimum coverage probability for confidence regions at a given confidence level. Our inclusion of the random variable E here and in the next definition is in a sense redundant: Uniformity of the estimator means that we could instead have considered the model $\mathcal{H}'$ of distributions on Z and R consisting of distributions $\mu[RZ|e]$ over all $\mu \in \mathcal{H}$ and all e . We refer to E explicitly because of the role played by external entities in this work.

Every ϵ -UPE provides a lower bound on a smooth min-entropy. For the special case $D = C = R$, this is theorem 11 of Ref. [28]. Define a family of subprobability distributions $\tilde{\mu}_{ze}$ of C by $\tilde{\mu}_{ze}(c) = \mu(c|ze)[F(cz) \geq \mu(c|ze)]$. The key observation in the proof of theorem 11 in Ref. [28] is that the subprobability distributions $\tilde{\mu}_{ze}$ have the following properties. (1) The weights of $\tilde{\mu}_{ze}$ satisfy $\sum_z w(\tilde{\mu}_{ze})\mu(z|e) \geq 1 - \epsilon$ for each e . (2) By definition $\tilde{\mu}_{ze}(c) \leq \mu(c|ze)$ and $\tilde{\mu}_{ze}(c) \leq F(cz)$ for each cze . The conclusion of theorem 11 in Ref. [28] is derived from these properties alone. This theorem directly generalizes to the case where D differs from R by replacing RVs and rewriting property 2. Specifically, for every ϵ -UPE $F(RZ)$, the family of subprobability distributions $\tilde{\mu}_{ze}$ of DR defined by $\tilde{\mu}_{ze}(dr) = \mu(dr|ze)[F(rz) \geq \mu(d|ze)]$ has the properties:

(1) The weights of $\tilde{\mu}_{ze}$ satisfy $\sum_z w(\tilde{\mu}_{ze})\mu(z|e) \geq 1 - \epsilon$ for each e ; (2) by definition $\tilde{\mu}_{ze}(dr) \leq \mu(dr|ze)$ and $\tilde{\mu}_{ze}(dr) \leq F(rz)\mu(r|dze)$ for each $drze$. To be able to certify randomness in D conditional on Z while drawing on other information in R without conditioning on it, we weaken the definition of ϵ -UPEs by directly imposing (1) and a “softened” version of (2) as in the next definition. Below we show that this suffices for establishing lower bounds on smooth min-entropy.

Definition 15. Let $1 \geq \epsilon \geq 0$. The function $F : \text{Rng}(RZ) \rightarrow [0, \infty)$ is a *level- ϵ soft E -uniform probability estimator* for $\mathcal{H}$ [ϵ -soft UPE, or with specifics, ϵ -soft UPE($D:R|Z; E, \mathcal{H}$)] if for each e and each distribution μ in $\mathcal{H}$ there exists a family of subprobability distributions $\tilde{\mu}_z \leq \mu[DR|ze]$ of DR such that the following two conditions hold: 1) The weights of $\tilde{\mu}_z$ satisfy $\sum_z w(\tilde{\mu}_z)\mu(z|e) \geq 1 - \epsilon$; 2) there exists a non-negative function $q(R|DZ)$ of DRZ such that $\sum_r q(r|dz) \leq 1$ for all dz , and

$$\tilde{\mu}_z(dr) \leq F(rz)q(r|dz), \quad (23)$$

for all drz . The function $\mathcal{U} : \text{Rng}(RZ) \times [0, 1] \rightarrow [0, \infty)$ is a soft UPE if it is monotone nondecreasing in the second argument, and for each ϵ , $\mathcal{U}(RZ, 1 - \epsilon)$ is an ϵ -soft UPE.

Note that in the definition of soft probability estimators, we have chosen to leave the dependence of $\tilde{\mu}_z$ and q on e and on μ implicit. In this definition, we explicitly consider the joint distribution of DR , even though we assume that D is determined by R . Considering the joint distribution helps simplify the notation for probabilities. However, we take this assumption into account by not explicitly including D as an argument to F . The purpose of Eq. (23) is to establish an upper bound on the subprobability $\tilde{\mu}_z(d)$ conditionally on the event $\{F(RZ) \leq p\}$, see the proofs of lemmas 17 and 18 for details. For any given e , one can interpret $F(rz)q(r|dz)$ as an implicit and distribution-dependent estimate of the probability of dr given z , where $q(r|dz)$ depends on the distribution to be estimated and accounts for the probability of r conditional on dz .

The discussions in the paragraph above definition 15 suggest that the notion of a soft UPE is weaker than that of a UPE, which is formalized by the next lemma.

Lemma 16. If F is an ϵ -UPE, then it is an ϵ -soft UPE.

Proof. For any given e , define $\tilde{\mu}_z(dr) = \mu(dr|ze)[F(rz) \geq \mu(d|ze)]$. Then by the definition of ϵ -UPEs, $\sum_z w(\tilde{\mu}_z)\mu(z|e) \geq 1 - \epsilon$. Let $q(r|dz) = \mu(r|dze)$. It suffices to check Eq. (23) for elements of $\{drz : F(rz) \geq \mu(d|ze)\}$:

$$\begin{aligned} \tilde{\mu}_z(dr) &= \mu(dr|ze) \\ &= \mu(r|dze)\mu(d|ze) \\ &\leq q(r|dz)F(rz). \end{aligned} \quad (24)$$

B. From probability estimation to min-entropy

The next lemma shows that ϵ -soft UPEs with constant upper bounds certify smooth min-entropy.

Lemma 17. Suppose that F is an ϵ -soft UPE($D:R|Z; E, \mathcal{H}$) such that $F \leq p$ for some constant p . Then for each e and $\mu \in \mathcal{H}$, $P_{\max, \mu[DZ|e]}^{u, \epsilon}(D|Z) \leq p$. Consequently, $P_{\max, \mu[DZE]}^{u, \epsilon}(D|ZE) \leq p$.

Proof. Fix e and let $\tilde{\mu}_z(DR)$ and $q(R|DZ)$ be as in the definition of soft UPEs. According to our conventions, $\tilde{\mu}_z[D](d) = \tilde{\mu}_z(d) = \sum_r \tilde{\mu}_z(dr)$. Therefore $\sum_z w(\tilde{\mu}_z[D])\mu(z|e) \geq 1 - \epsilon$. Since $\sum_r q(r|dze) \leq 1$, for all dz , we have

$$\tilde{\mu}_z(d) \leq \sum_r F(rz)q(r|dz) \leq \sum_r pq(r|dz) \leq p. \quad (25)$$

Further, $\tilde{\mu}_z(d) \leq \mu(d|ze)$ for all dz , as can be seen by summing the defining inequality for $\tilde{\mu}_z$ over r . It follows that the $\tilde{\mu}_z[D]$ witness that $P_{\max, \mu[DZ|e]}^{u, \epsilon}(D|Z) \leq p$. For the last statement, apply lemma 4. ■

A weaker relationship holds for general soft UPEs.

Lemma 18. Suppose that F is an ϵ -soft UPE($D:R|Z; E, \mathcal{H}$), $\mu \in \mathcal{H}$, $p \geq 1/|\text{Rng}(D)|$, and let $\kappa = \mathbb{P}(F \leq p)$. Then $P_{\max, \mu[DZE|F \leq p]}^{\epsilon/\kappa}(D|ZE) \leq p/\kappa$.

Proof. Let $\kappa_e = \mathbb{P}(F \leq p|e)$. Below we show that for all values e of E ,

$$P_{\max, \mu[DZ|e, F \leq p]}^{\epsilon/\kappa_e}(D|Z) \leq p/\kappa_e. \quad (26)$$

Once this is shown, we can use

$$\begin{aligned} \sum_e \frac{1}{\kappa_e} \mu(e|F \leq p) &= \sum_e \frac{1}{\mathbb{P}(F \leq p|e)} \mu(e|F \leq p) \\ &= \sum_e \frac{\mu(e)}{\mu(F \leq p)} = 1/\kappa, \end{aligned} \quad (27)$$

and lemma 4 to complete the proof. In order to prove the bound in Eq. (26), we apply the two conditions in the definition of ϵ -soft UPEs. Specifically, from the second condition in definition 15 we can prove that there exists a distribution $\nu[DZ|e]$ according to which the average max-prob satisfies $P_{\max, \nu}^0(D|Z) \leq p/\kappa_e$, and from the first condition in definition 15, we can bound the distance $\text{TV}(\nu[DZ|e], \mu[DZ|e, F \leq p])$ from above by ϵ/κ_e . For the remainder of the proof, e is fixed, so we simplify the notation by universally conditioning on $\{E = e\}$ and omitting the explicit condition. Further, we omit e from suffixes. Thus $\kappa = \kappa_e$ from here on.

Let $\chi(RZ) = \mathbb{I}[F(RZ) \leq p]$. Then $\kappa = \mathbb{E}(\chi)$. Let $\kappa_z = \mu(F \leq p|z)$. We have $\sum_z \kappa_z \mu(z) = \kappa$ and $\kappa_z = \mu(z|F \leq p)\kappa/\mu(z)$.

Let $\tilde{\mu}_z$ and q witness that F is an ϵ -soft UPE($D:R|Z; E, \mathcal{H}$). Define $\tilde{\nu}(drz) = \tilde{\mu}_z(dr)\chi(rz)\mu(z)/\kappa$. The weight of $\tilde{\nu}$ satisfies

$$\begin{aligned} w(\tilde{\nu}) &= \sum_{drz} \tilde{\mu}_z(dr)\chi(rz)\mu(z)/\kappa \\ &\leq \sum_{drz} \mu(drz)\chi(rz)/\kappa \\ &= \mathbb{P}(F \leq p)/\kappa = 1, \end{aligned} \quad (28)$$

$$\begin{aligned} w(\tilde{\nu}) &= \sum_{drz} \mu(drz)\chi(rz)/\kappa \\ &\quad - \sum_{drz} (\mu(drz) - \tilde{\mu}_z(dr)\mu(z))\chi(rz)/\kappa \\ &= 1 - \sum_{drz} (\mu(drz) - \tilde{\mu}_z(dr)\mu(z))\chi(rz)/\kappa \\ &\geq 1 - \sum_{drz} (\mu(drz) - \tilde{\mu}_z(dr)\mu(z))/\kappa \\ &\geq 1 - (1 - (1 - \epsilon))/\kappa = 1 - \epsilon/\kappa. \end{aligned} \quad (29)$$

Thus $\tilde{\nu}$ is a subprobability distribution of weight at least $1 - \epsilon/\kappa$. We use it to construct the distribution ν witnessing the conclusion of the lemma. For each drz , we bound

$$\begin{aligned} \tilde{\nu}(drz)/\mu(z|F \leq p) &= \tilde{\mu}_z(dr)\chi(rz)/\kappa_z \\ &\leq F(rz)q(r|dz)\chi(rz)/\kappa_z \leq pq(r|dz)/\kappa_z, \end{aligned} \quad (30)$$

from which we get $\tilde{\nu}(dz)/\mu(z|F \leq p) \leq p/\kappa_z$ by summing over r . Define $\tilde{\nu}[D|z](d) = \tilde{\nu}(dz)/\mu(z|F \leq p)$ and let $w_z = w(\tilde{\nu}[D|z])$, where this extends the conditional probability notation to the subprobability distribution $\tilde{\nu}$ with the understanding that the conditionals are with respect to μ given $\{F \leq p\}$. Applying the first step of Eq. (30) and continuing

from there, we have

$$\begin{aligned} \tilde{\nu}[D|z](d) &= \sum_r \tilde{\nu}(drz)/\mu(z|F \leq p) \\ &= \sum_r \tilde{\mu}_z(dr)\chi(rz)/\kappa_z \\ &\leq \sum_r \mu(dr|z)\chi(rz)/\kappa_z \\ &= \sum_r \mu(dr|z)\mathbb{I}[F(rz) \leq p]/\kappa_z \\ &= \mu(d, F \leq p|z)/\mu(F \leq p|z) = \mu(d|z, F \leq p), \end{aligned} \quad (31)$$

which also implies that $w_z \leq 1$. Since $\tilde{\nu}[D|z](d) \leq p/\kappa_z$ and $p/\kappa_z \geq p \geq 1/|\text{Rng}(D)|$, we can apply lemma 2 to obtain distributions ν_z of D such that $\nu_z \geq \tilde{\nu}[D|z]$, $\nu_z \leq p/\kappa_z$, and $\text{TV}(\nu_z, \mu[D|z, F \leq p]) \leq 1 - w_z$. Now we can define $\nu(dz) = \nu_z(d)\mu(z|F \leq p)$. From Eq. (2),

$$\begin{aligned} \text{TV}(\nu, \mu[DZ|F \leq p]) &= \sum_z \text{TV}(\nu_z, \mu[D|z, F \leq p])\mu(z|F \leq p) \\ &\leq \sum_z (1 - w_z)\mu(z|F \leq p) \\ &= 1 - \sum_z w(\tilde{\nu}[D|z])\mu(z|F \leq p) \\ &= 1 - \sum_z \sum_{dr} (\tilde{\nu}(drz)/\mu(z|F \leq p))\mu(z|F \leq p) \\ &= 1 - w(\tilde{\nu}) \leq \epsilon/\kappa. \end{aligned} \quad (32)$$

For the average maximum probability of ν , we get

$$\begin{aligned} \sum_z \max_d \nu(d|z)\nu(z) &= \sum_z \max_d \nu_z(d)\mu(z|F \leq p) \\ &\leq p \sum_z \mu(z|F \leq p)/\kappa_z \\ &= p \sum_z \mu(z)/\kappa = p/\kappa, \end{aligned} \quad (33)$$

which together with the argument at the beginning of the proof establishes the lemma. ■

We remark that the witness ν constructed in the proof of the lemma satisfies $\nu[ZE] = \mu[ZE|F \leq p]$. We refer to this ν in the next section, where to emphasize that ν is constructed conditionally on $\{F \leq p\}$ we write $\nu[DZE|F \leq p]$ instead of $\nu[DZE]$.

C. Protocols

Our goal is to construct probability estimators from test supermartingales and use them in randomness generation by composition with an extractor. The required supermartingales are introduced in Sec. V. Here we give three ways in which probability estimators can be composed with an extractor for randomness generation protocols. For the first, given a probability estimator, we can estimate smooth min-entropy and compose with an extractor by chaining lemma 18 with

lemma 7. The second draws on banked randomness to avoid failure. The third requires a linear extractor. The first protocol is given in the following theorem.

Theorem 19. Let F be an ϵ_h -soft UPE($D:R|Z; E, \mathcal{H}$) with $\mathbb{P}(F \leq p\delta) = \kappa$, D a bitstring of length n , and $p\delta \geq 1/2^n$. Write $\sigma_h = -\log_2(p)$ and suppose that $\mathcal{E}$ is a strong extractor such that $(n, l, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints. Let S be a length l uniform and independent seed string. Abbreviate $\mathcal{E} = \mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)$. Then

$$\begin{aligned} & \text{TV}(\mu[\mathcal{E}SZE|F \leq p\delta], \text{Unif}_{\mathcal{E}S} \otimes \mu[ZE|F \leq p\delta]) \\ & \leq \epsilon_x + (2\epsilon_h + \delta)/\kappa. \end{aligned} \quad (34)$$

In particular, $\mathcal{P}_X = \mathcal{E}, \mathcal{P}_S = S, \mathcal{P}_P = \llbracket F \leq p\delta \rrbracket$ defines a $(\sigma, \epsilon_x + 2\epsilon_h + \delta)$ sound randomness generation protocol.

Proof. By lemma 18, $P_{\max, \mu[DZE|F \leq p\delta]}^{\epsilon_h/\kappa}(D|ZE) \leq p\delta/\kappa$. We obtain Eq. (34) by applying lemma 7 with ϵ_h there replaced by ϵ_h/κ here and δ there with δ/κ here. The soundness statement follows by multiplying both sides by κ in consideration of $\epsilon_x \kappa \leq \epsilon_x$. ■

We can use the remarks after lemmas 5 and 18 to improve the factor of 2 on ϵ_h in the statement of the theorem. For this, let v witness that $P_{\max, \mu[DZE|F \leq p]}^{\epsilon_h/\kappa}(D|ZE) \leq p\delta/\kappa$ according to lemma 18, where from the remark after this lemma, $v[ZE|F \leq p\delta] = \mu[ZE|F \leq p\delta]$. Then apply the second part of lemma 5 and the remark thereafter to obtain a distribution v' of DZE with $\text{TV}(v', v[DZE|F \leq p\delta]) \leq \delta/\kappa$ and $v'[ZE] = v[ZE|F \leq p\delta] = \mu[ZE|F \leq p\delta]$, where $P_{\max, v'}^{u, 0}(D|ZE) \leq p$. From lemma 6, we have $\text{TV}(v'[\mathcal{E}(D, S; \sigma_h, \sigma, \epsilon_x)SZE], \text{Unif}_{\mathcal{E}S} \otimes v'[ZE]) \leq \epsilon_x$. The observation now follows from $v'[ZE] = \mu[ZE|F \leq p\delta]$ and the triangle and data-processing inequalities. Namely, we get

$$\begin{aligned} & \text{TV}(\mu[\mathcal{E}SZE|F \leq p\delta], \text{Unif}_{\mathcal{E}S} \otimes \mu[ZE|F \leq p\delta]) \\ & \leq \epsilon_x + (\epsilon_h + \delta)/\kappa. \end{aligned} \quad (35)$$

The protocol performance and soundness proof of theorem 19 is closely related to the performance and proof of the protocol soundness theorem in Ref. [13], which in turn are based on results of Refs. [12,27]. The parameters in these theorems implicitly compensate for the fact that we do not require the extractor to be classical-proof, see the comments on extractor constraints in Sec. III D.

The next protocol has no chance of failure but needs access to banked randomness. The banked randomness cannot be randomness from previous instances of the protocol involving the same devices. Let $\mathcal{U}$ be a soft UPE($D:R|Z; E$) for $\mathcal{H}$. Let $\sigma \in \mathbb{N}_+$ and $\epsilon > 0$ be the requested number of bits and error bound. Assume that we have access to a source of uniform bits S that is independent of all other relevant RVs. For banked randomness, we also have access to a second source of such random bits S_b . Let $\mathcal{E}$ be a strong extractor. We define a randomness generation protocol $\mathcal{P}(\sigma, \epsilon; \mathcal{U})$ by the following steps.

0. Choose $n, \sigma_h, \epsilon_h > 0, \epsilon_x > 0$ and l so that $(n + \sigma_h, l, \sigma_h, \sigma, \epsilon_x)$ satisfies the extractor constraints and $\epsilon_h + \epsilon_x = \epsilon$.

1. Perform trials to obtain values r and z of R and Z expressed as bit strings, with $|D| = n$. Let d be the value taken by D . (Recall that by our conventions, D is determined by R .)

2. Determine $p = \mathcal{U}(rz, 1 - \epsilon_h)$ and let $k = \max(0, \lceil \sigma_h - \log_2(1/p) \rceil)$.

3. Obtain values $(s)_{\leq l}$ and $(s_b)_{\leq k}$ from $S_{\leq l}$ and $(S_b)_{\leq k}$.

4. Let $d' = d(s_b)_{\leq k} 0^{\sigma_h - k}$, where $0^{\sigma_h - k}$ is a string of zeros of length $\sigma_h - k$.

5. If the distribution of Z is known and independent of E , let $z' = z$, otherwise let z' be the empty string.

6. Output $\mathcal{P}_X = \mathcal{E}(d', (s)_{\leq l}, \sigma_h, \sigma, \epsilon_x)$, $\mathcal{P}_S = z'(s)_{\leq l}$ and $\mathcal{P}_P = 1$.

Note that we have left the parameter choices made in the first step free, subject to the given constraints. They can be made to minimize the expected amount of banked randomness needed given information on device performance and resource bounds. It is important that the choices be made before obtaining r, z and the seed bits. That is, they are conditionally independent of these RVs given the pre-protocol past, where $\mathcal{H}$ is satisfied conditionally on this past. The length of D must be fixed before performing trials, but this does not preclude stopping trials early if an adaptive probability estimator is used. In this case, D can be zero-filled to length n . The requirement for fixed length n is imposed by the extractor properties: We do not have an extractor that can take advantage of variable- and unbounded-length inputs.

Theorem 20. The protocol $\mathcal{P}(\sigma, \epsilon; \mathcal{U})$ is a (σ, ϵ) sound and complete randomness generation protocol with respect to $\mathcal{H}$.

Proof. The protocol has zero probability of failure, so completeness is immediate. However, note that for the protocol to be nontrivial, the expected number of banked random bits used should be small compared to σ .

Let K, D' , and Z' be the RVs corresponding to the values k, d' and z' constructed in the protocol. Given the parameter choices that are made first, K is determined by R and Z . Define $S_K = (S_b)_{\leq K} 0^{\sigma_h - K}$ and $R' = RS_K$. Taking advantage of uniformity in e , we fix e and implicitly condition everything on $\{E = e\}$. Let $\tilde{\mu}_z$ and q be the functions witnessing that $p = \mathcal{U}(rz, 1 - \epsilon_h)$ is an ϵ_h -soft UPE. Now let $\tilde{v}_z$ be the family of subprobability distributions on $D'R'$ defined by $\tilde{v}_z(d'r') = \tilde{\mu}_z(dr)\mu(s_k|drz) = \tilde{\mu}_z(dr)2^{-k}$, where we used that K is determined by RZ . Then $\tilde{v}_z(d'r') \leq p2^{-k}q(r|dz) \leq 2^{-\sigma_h}q(r|dz)$. Furthermore,

$$\begin{aligned} \sum_z w(\tilde{v}_z)\mu(z) &= \sum_z \sum_{d'r'} \tilde{v}_z(d'r')\mu(z) \\ &= \sum_z \sum_{d'r'} \tilde{\mu}_z(dr)\mu(s_k|drz)\mu(z) \\ &= \sum_z \sum_{dr} \tilde{\mu}_z(dr)\mu(z) \sum_{k, s_k} \mu(s_k|drz) \\ &= \sum_z \sum_{dr} \tilde{\mu}_z(dr)\mu(z) \\ &= \sum_z \omega(\tilde{\mu}_z)\mu(z) \\ &\geq 1 - \epsilon_h. \end{aligned} \quad (36)$$

It follows that the function $d'r'z \mapsto 2^{-\sigma_h}$ is a ϵ_h -soft UPE($D':R'|Z; E, \mathcal{H}$) with a constant upper bound, witnessed by $\tilde{v}_z$ and $q(r'|d'z) = q(r|dz)$. From lemma 17, D' has ϵ -smooth Z -conditional max-prob $2^{-\sigma_h}$ given $\{E = e\}$ for each

e. The theorem now follows from the composition lemma lemma 6 after taking note of lemma 4. ■

The third protocol is the same as the second except that if the estimated probability is too large in step 2, the protocol fails. In this case, banked random bits are not used. This gives a protocol $\mathcal{Q}(\sigma, \epsilon; \mathcal{U})$ whose steps are the following.

0. Choose $n, \sigma_h, \epsilon_h > 0, \epsilon_x > 0$ and l so that $(n, l, -\log_2(2^{-\sigma_h} + 2^{-n}), \sigma, \epsilon_x)$ satisfies the extractor constraints, $\epsilon_h + \epsilon_x = \epsilon$ and $2^{-n} \leq 2^{-\sigma_h}$.

1. Perform trials to obtain values r and z of R and Z expressed as bit strings, with $|D| = n$.

2. Determine $F(rz) = \mathcal{U}(rz, 1 - \epsilon_h)$.

3. If $F(rz) > 2^{-\sigma_h}$, return with $\mathcal{Q}_X$ and $\mathcal{Q}_S$ the empty strings and $\mathcal{Q}_P = 0$ to indicate failure.

4. Obtain value $(s)_{\leq l}$ from $S_{\leq l}$.

5. Output $\mathcal{Q}_X = \mathcal{E}(d, (s)_{\leq l}, -\log_2(2^{-\sigma_h} + 2^{-n}), \sigma, \epsilon_x)$, $\mathcal{Q}_S = (s)_{\leq l}$ and $\mathcal{Q}_P = 1$.

Theorem 21. The protocol $\mathcal{Q}(\sigma, \epsilon; \mathcal{U})$ defined above when used with a linear strong extractor is a (σ, ϵ) sound randomness generation protocol with respect to $\mathcal{H}$.

Proof. Because the distribution of Z conditional on $\{\mathcal{Q}_P = 1\}$ can differ significantly from its initial distribution in a way that is not known ahead of time, Z 's randomness cannot be reused. Consequently, in this proof, Z and E always occur together, so we let Z stand for the joint RV ZE throughout.

The proof first replaces the distribution $\mu[DZS\mathcal{Q}_P]$ by $\nu[DZS\mathcal{Q}_P]$ so that $\nu[D|Z] \leq 2^{-\sigma_h} + 2^{-n}$ and conditionally on $\{\mathcal{Q}_P = 1\}$, the distribution's change is small in a sense to be defined. Since the conclusion only depends on $\{\mathcal{Q}_P = 1\}$, changes when $\{\mathcal{Q}_P = 0\}$ can be arbitrary. We therefore define TV_{pass} so that for global distributions ν and ν' satisfying $\nu[\mathcal{Q}_P] = \nu'[\mathcal{Q}_P]$ and for all RVs U ,

$$\begin{aligned} \text{TV}_{\text{pass}}(\nu[U\mathcal{Q}_P], \nu'[U\mathcal{Q}_P]) \\ = \text{TV}(\nu[U|\mathcal{Q}_P = 1], \nu'[U|\mathcal{Q}_P = 1])\nu(\mathcal{Q}_P = 1). \end{aligned} \quad (37)$$

Note that TV_{pass} is the same as the total variation distance between $\nu[\mathcal{M}(U\mathcal{Q}_P)\mathcal{Q}_P]$ and $\nu'[\mathcal{M}(U\mathcal{Q}_P)\mathcal{Q}_P]$ for any process $\mathcal{M}$ satisfying $\mathcal{M}(U1) = U$ and $\mathcal{M}(U0) = v$ for a fixed value v . Such processes forget U when $\mathcal{Q}_P = 0$. Note that (σ, ϵ) soundness is defined in terms of TV_{pass} .

Write $\kappa = \mathbb{P}(\mathcal{Q}_P = 1)$ and $p = 2^{-\sigma_h}$. We omit the subscript $\leq l$ from $S_{\leq l}$. To construct ν we refine the proof of lemma 18. Write $F(rz) = \mathcal{U}(rz, 1 - \epsilon_h)$, and let $\tilde{\mu}_z$ and q witness that F is an ϵ_h -soft UPE($D:R|Z; E, \mathcal{H}$). The event $\{\mathcal{Q}_P = 1\}$ is the same as $\{F \leq p\}$. The distribution ν to be constructed satisfies $\nu[ZS\mathcal{Q}_P] = \mu[ZS\mathcal{Q}_P]$. We maintain that S is uniform and independent of the other RVs for both ν and μ . Thus it suffices to construct $\nu[DZ\mathcal{Q}_P]$ and define $\nu = \nu[DZ\mathcal{Q}_P] \otimes \text{Unif}_S$. We start by defining $\nu[DZ|\mathcal{Q}_P = 1]$ here to be $\nu[DZ]$ as constructed in the proof of lemma 18. Consistent with the notation in that proof, we also use the notation $\nu_z = \nu[D|z, \mathcal{Q}_P = 1]$. Note that $\nu[Z|\mathcal{Q}_P = 1](z) = \mu(z|\mathcal{Q}_P = 1)$. We then set

$$\begin{aligned} \nu(dz, \mathcal{Q}_P = b) &= \nu[DZ|\mathcal{Q}_P = 1](dz)\kappa[b = 1] \\ &\quad + \text{Unif}_D(d)\mu(z, \mathcal{Q}_P = 0)[b = 0] \\ &= \nu_z(d)\mu(z, \mathcal{Q}_P = 1)[b = 1] \\ &\quad + \text{Unif}_D(d)\mu(z, \mathcal{Q}_P = 0)[b = 0]. \end{aligned} \quad (38)$$

This ensures that $\nu[ZS\mathcal{Q}_P] = \mu[ZS\mathcal{Q}_P]$ after adding in the uniform and independent RV S . The distribution ν satisfies the following additional properties by construction:

$$\begin{aligned} \nu(d|z) &= \nu(d, \mathcal{Q}_P = 1|z) + \nu(d, \mathcal{Q}_P = 0|z) \\ &= \nu_z(d)\mu(\mathcal{Q}_P = 1|z) + \mu(\mathcal{Q}_P = 0|z)/|\text{Rng}(D)| \\ &\leq p + 2^{-n} = 2^{-\sigma_h} + 2^{-n} \end{aligned} \quad (39)$$

and

$$\begin{aligned} \text{TV}_{\text{pass}}(\nu[DZS\mathcal{Q}_P], \mu[DZS\mathcal{Q}_P]) \\ = \text{TV}(\nu[DZS|\mathcal{Q}_P = 1], \mu[DZS|\mathcal{Q}_P = 1])\kappa \\ \leq \epsilon_h, \end{aligned} \quad (40)$$

in view of Eq. (32) and the preceding paragraph.

We can now apply the extractor to get

$$\text{TV}(\nu[\mathcal{E}(D, S)SZ], \text{Unif}_{\mathcal{E}S} \otimes \nu[Z]) \leq \epsilon_x. \quad (41)$$

Here we transferred the extractor guarantee conditionally on $Z = z$ for each z , noting that the marginal distribution on Z is the same for both arguments. From here on we abbreviate $\mathcal{E} = \mathcal{E}(D, S)$.

The main task for completing the proof of the theorem is to bound $\text{TV}(\nu[\mathcal{E}SZ|\mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S} \otimes \nu[Z|\mathcal{Q}_P = 1])$, which requires transferring the distance of Eq. (41) to the conditional distributions. For this, we need to show that for all z ,

$$\begin{aligned} \mu(\mathcal{Q}_P = 1|z)\text{TV}(\nu[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\ \leq \text{TV}(\nu[\mathcal{E}S|z], \text{Unif}_{\mathcal{E}S}). \end{aligned} \quad (42)$$

Before proving this inequality we use it to prove the theorem as follows:

$$\begin{aligned} \text{TV}_{\text{pass}}(\nu[\mathcal{E}SZ\mathcal{Q}_P], \text{Unif}_{\mathcal{E}S} \otimes \mu[Z\mathcal{Q}_P]) \\ = \mu(\mathcal{Q}_P = 1)\text{TV}(\nu[\mathcal{E}SZ|\mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S} \otimes \mu[Z|\mathcal{Q}_P = 1]) \\ = \mu(\mathcal{Q}_P = 1) \sum_z \mu(z|\mathcal{Q}_P = 1)\text{TV}(\nu[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\ = \sum_z \mu(z, \mathcal{Q}_P = 1)\text{TV}(\nu[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\ = \sum_z \mu(z)\mu(\mathcal{Q}_P = 1|z)\text{TV}(\nu[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\ \leq \sum_z \mu(z)\text{TV}(\nu[\mathcal{E}S|z], \text{Unif}_{\mathcal{E}S}) \\ = \text{TV}(\nu(\mathcal{E}SZ), \text{Unif}_{\mathcal{E}S} \otimes \nu(Z)) \\ \leq \epsilon_x, \end{aligned} \quad (43)$$

where we applied Eq. (2) (twice, first with $\mu[Z|\mathcal{Q}_P = 1] = \nu[Z|\mathcal{Q}_P = 1]$, then with $\mu[Z] = \nu[Z]$), Eq. (42) (to get the third-to-last line), and Eq. (41). The theorem now follows via the triangle inequality for TV_{pass} , adding the TV_{pass} -distance from $\nu[\mathcal{E}SZ\mathcal{Q}_P]$ to $\mu[\mathcal{E}SZ\mathcal{Q}_P]$. According to Eq. (40) and the data-processing inequality, this distance is bounded by ϵ_h .

It remains to establish Eq. (42). We start by taking advantage of the linearity of the extractor. We have

$$\begin{aligned} \nu(\mathcal{E} = x, s|z) &= \nu(\mathcal{E} = x, s|z, \mathcal{Q}_P = 1)\nu(\mathcal{Q}_P = 1|z) \\ &\quad + \nu(\mathcal{E} = x, s|z, \mathcal{Q}_P = 0)\nu(\mathcal{Q}_P = 0|z) \end{aligned}$$

$$\begin{aligned}
&= v(\mathcal{E} = x, s|z, \mathcal{Q}_P = 1)\mu(\mathcal{Q}_P = 1|z) \\
&\quad + v(\mathcal{E} = x, s|z, \mathcal{Q}_P = 0)\mu(\mathcal{Q}_P = 0|z). \tag{44}
\end{aligned}$$

Define $R_s = \text{Rng}(\mathcal{E}(D, s))$ and $r_s = |R_s| \leq |\text{Rng}(\mathcal{E}(D, S))| = 2^\sigma$. Since $\mathcal{E}$ is a linear extractor, $\text{Rng}(D)$ can be treated as a vector space over a finite field. For all $x \in R_s$, $N_{sx} = \{d : x = \mathcal{E}(d, s)\}$ is a translation of the null space of $d \mapsto \mathcal{E}(d, s)$. Hence $n_s = |N_{sx}|$ is independent of x , and $r_s n_s = |\text{Rng}(D)|$. Since $v[D|z, \mathcal{Q}_P = 0]$ is uniform by design, for $x \in R_s$

$$v(\mathcal{E} = x, s|z, \mathcal{Q}_P = 0) = \frac{|N_{sx}|}{|\text{Rng}(D)|} = \frac{1}{r_s} \geq 2^{-\sigma}. \tag{45}$$

For $x \notin R_s$, $v(\mathcal{E} = x, s|z, \mathcal{Q}_P = b) = 0$. In the last two equations, we can move s into the conditioner in each expression

and use independence and uniformity of S with respect to Z and $\mathcal{Q}_P$ to get

$$\begin{aligned}
v(\mathcal{E} = x|zs) &= v(\mathcal{E} = x|zs, \mathcal{Q}_P = 1)\llbracket x \in R_s \rrbracket \mu(\mathcal{Q}_P = 1|z) \\
&\quad + \frac{1}{r_s} \llbracket x \in R_s \rrbracket \mu(\mathcal{Q}_P = 0|z). \tag{46}
\end{aligned}$$

Abbreviate $v_{zs} = v[\mathcal{E}|zs]$, and $v_{zsb} = v[\mathcal{E}|zs, \mathcal{Q}_P = b]$, which are conditional distributions of $\mathcal{E}$. Then $v_{zs0}(x) = \llbracket x \in R_s \rrbracket / r_s$. Write $\kappa_z = \mu(\mathcal{Q}_P = 1|z) = \mu(\mathcal{Q}_P = 1|zs)$ for the conditional passing probabilities. With these definitions, we can write $v_{zs} = \kappa_z v_{zs1} + (1 - \kappa_z) v_{zs0}$ and calculate

$$\begin{aligned}
\text{TV}(v_{zs}, \text{Unif}_{\mathcal{E}}) &= \sum_x (v_{zs}(x) - 2^{-\sigma}) \llbracket v_{zs}(x) > 2^{-\sigma} \rrbracket \\
&= \sum_{x \in R_s} (\kappa_z v_{zs1}(x) + (1 - \kappa_z)/r_s - 2^{-\sigma}) \llbracket \kappa_z v_{zs1}(x) + (1 - \kappa_z)/r_s > 2^{-\sigma} \rrbracket \\
&= \sum_{x \in R_s} (\kappa_z (v_{zs1}(x) - 2^{-\sigma}) + (1 - \kappa_z)(1/r_s - 2^{-\sigma})) \llbracket \kappa_z (v_{zs1}(x) - 2^{-\sigma}) + (1 - \kappa_z)(1/r_s - 2^{-\sigma}) > 0 \rrbracket \\
&\geq \sum_{x \in R_s} (\kappa_z (v_{zs1}(x) - 2^{-\sigma})) \llbracket \kappa_z (v_{zs1}(x) - 2^{-\sigma}) > 0 \rrbracket \\
&= \kappa_z \sum_{x \in R_s} (v_{zs1}(x) - 2^{-\sigma}) \llbracket (v_{zs1}(x) - 2^{-\sigma}) > 0 \rrbracket \\
&= \kappa_z \text{TV}(v_{zs1}, \text{Unif}_{\mathcal{E}}). \tag{47}
\end{aligned}$$

Here, we used the fact that $\sum_x f(x) \llbracket f(x) > 0 \rrbracket$ is monotone increasing with f , noting that $1/r_s \geq 2^{-\sigma}$.

Since $v(s|z) = v(s|z, \mathcal{Q}_P = 1) = \text{Unif}_S(s)$, we can apply Eq. (2) to obtain

$$\begin{aligned}
\text{TV}(v[\mathcal{E}S|z], \text{Unif}_{\mathcal{E}S}) &= \sum_s \text{TV}(v[\mathcal{E}|zs], \text{Unif}_{\mathcal{E}}) v(s|z) \\
&= \sum_s \text{TV}(v_{zs}, \text{Unif}_{\mathcal{E}}) / |\text{Rng}(S)| \tag{48}
\end{aligned}$$

and

$$\begin{aligned}
&\text{TV}(v[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\
&= \sum_s \text{TV}(v[\mathcal{E}|zs, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}}) v(s|z, \mathcal{Q}_P = 1) \\
&= \sum_s \text{TV}(v_{zs1}, \text{Unif}_{\mathcal{E}}) / |\text{Rng}(S)|. \tag{49}
\end{aligned}$$

Applying the previous two displayed equations and the inequality before that, we conclude that for all z ,

$$\begin{aligned}
\mu(\mathcal{Q}_P = 1|z) \text{TV}(v[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\
&= \kappa_z \text{TV}(v[\mathcal{E}S|z, \mathcal{Q}_P = 1], \text{Unif}_{\mathcal{E}S}) \\
&= \kappa_z \sum_s \text{TV}(v_{zs1}, \text{Unif}_{\mathcal{E}}) / |\text{Rng}(S)|
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_s \text{TV}(v_{zs}, \text{Unif}_{\mathcal{E}}) / |\text{Rng}(S)| \\
&= \text{TV}(v[\mathcal{E}S|z], \text{Unif}_{\mathcal{E}S}). \tag{50}
\end{aligned}$$

■

V. TEST SUPERMARTINGALES FOR UNIFORM PROBABILITY ESTIMATION

A. Standard models for sequences of trials

Each of theorems 19–21 reduces the problem of randomness generation to that of probability estimation. We now consider the situation where $\mathbf{CZ}$ is a stochastic sequence of n trials, with the distributions of the i th trial RVs $C_i Z_i$ in a model C_i conditional on the past and on E . For the remainder of the paper, the conditioning on E applies universally, and relevant statements are uniform in the values of E . We therefore no longer mention E explicitly. For the most general treatment, we define $R_i = C_i R'_i$ where R'_i is additional information obtainable in a trial and $R_0 = R'_0$ is information available initially. Test supermartingales and test factors are with respect to $\mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}$. According to our convention, D_i is a function of C_i . We note that other situations can be cast into this form by changing the definition of C_i , for example by extending C_i with other information, and modifying models accordingly.

Formally, we are considering models $\mathcal{H}(\mathcal{C})$ of distributions of $\mathbf{RZ}$ defined by a family of conditionally defined models $\mathcal{C}_{i+1}|\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}$ of $C_{i+1}Z_{i+1}$. $\mathcal{H}(\mathcal{C})$ consists of the distributions μ with the following two properties. First, for all i and $\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}$,

$$\mu[C_{i+1}Z_{i+1}|\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}] \in \mathcal{C}_{i+1}|\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}. \quad (51)$$

We emphasize that the model $\mathcal{C}_{i+1}$ specifies the set of all possible distributions of $C_{i+1}Z_{i+1}$ conditional on the past and on E , although we perform probability estimation using all information contained in R_{i+1} and we extract randomness only from D_{i+1} . Second, μ satisfies that Z_{i+1} is independent of $\mathbf{R}_{\leq i}$ conditionally on $\mathbf{Z}_{\leq i}$ (and E). Models $\mathcal{H}(\mathcal{C})$ satisfying these conditions are called *standard*. The second condition is needed in order to be able to estimate $\mathbf{Z}$ -conditional probabilities of $\mathbf{D}$ and corresponds to the Markov-chain condition in the entropy accumulation framework [26].

In many cases, the sets $\mathcal{C}_{i+1}|\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}$ do not depend on $\mathbf{r}$, but we take advantage of dependence on $\mathbf{z}_{\leq i}$. In our applications, the sets capture the settings and nonsignaling constraints. For simplicity, we write $\mathcal{C}_{i+1} = \mathcal{C}_{i+1}|\mathbf{r}_{\leq i}, \mathbf{z}_{\leq i}$, leaving the conditional parameters implicit.

Normally, models for trials are convex and closed. If $\mathbf{Z}$ is nontrivial, the second condition on standard models prevents their being convex closed for $n > 1$, but we note that our results generally extend to the convex closure of the model used.

B. UPE constructions

Let F_{i+1} be non-negative functions of $C_{i+1}Z_{i+1}$, parameterized by $\mathbf{R}_{\leq i}$ and $\mathbf{Z}_{\leq i}$. We call such functions “past-parameterized”. Let $T_0 = 1$ and $T_i = \prod_{1 \leq j \leq i} F_j$ for $i \geq 1$. For $\beta > 0$, define $\mathcal{U}_\beta$ and $\mathcal{U}_\beta^*$ by

$$\mathcal{U}_\beta(\mathbf{RZ}, 1 - \epsilon) = (T_n \epsilon)^{-1/\beta}, \quad (52)$$

$$\mathcal{U}_\beta^*(\mathbf{RZ}, 1 - \epsilon) = \left(\max_{i \leq n} (T_i \epsilon) \right)^{-1/\beta}. \quad (53)$$

We choose the F_i so that they are *probability estimation factors* according to the following definition.

Definition 22. Let $\beta > 0$, and let $\mathcal{C}$ be any model, not necessarily convex. A *probability estimation factor (PEF)* with power β for $\mathcal{C}$ and $D|Z$ is a non-negative RV $F = F(\mathbf{CZ})$ such that for all $v \in \mathcal{C}$, $\mathbb{E}_v(F v(D|Z)^\beta) \leq 1$.

As usual, if the parameters are clear from context, we may omit them. In particular, unless differently named RVs are involved, PEFs are always for $D|Z$.

In the case where $\mathbf{R} = \mathbf{C} = \mathbf{D}$, we show that $\mathcal{U}_\beta$ and $\mathcal{U}_\beta^*$ are UPEs. In the general case, $\mathcal{U}_\beta$ is a soft UPE. We start with the special case, which was also studied in our previous work [28].

Theorem 23. Fix $\beta > 0$ and assume $\mathbf{R} = \mathbf{C} = \mathbf{D}$. Let $\mathbf{F}$ be a sequence of past-parameterized PEFs, with F_i a PEF with power β for $\mathcal{C}_i$. Then $\mathcal{U}_\beta(\mathbf{RZ}, 1 - \epsilon)$ and $\mathcal{U}_\beta^*(\mathbf{RZ}, 1 - \epsilon)$ as defined in Eqs. (52) and (53) are ϵ -UPE($\mathbf{C}:\mathbf{C}|\mathbf{Z}; \mathcal{H}(\mathcal{C})$).

Note that β cannot be adapted during the trials. On the other hand, before the i th trial, we can design the PEFs F_i for the particular constraints relevant to the i th trial.

Proof. We first observe that

$$\prod_{j=0}^{i-1} \mathbb{P}(C_{j+1}|Z_{j+1}\mathbf{Z}_{\leq j}\mathbf{C}_{\leq j}) = \mathbb{P}(\mathbf{C}_{\leq i}|\mathbf{Z}_{\leq i}). \quad (54)$$

This follows by induction with the identity

$$\begin{aligned} \mathbb{P}(\mathbf{C}_{\leq j+1}|\mathbf{Z}_{\leq j+1}) &= \mathbb{P}(C_{j+1}|Z_{j+1}\mathbf{Z}_{\leq j}\mathbf{C}_{\leq j})\mathbb{P}(\mathbf{C}_{\leq j}|\mathbf{Z}_{\leq j}) \\ &= \mathbb{P}(C_{j+1}|Z_{j+1}\mathbf{Z}_{\leq j}\mathbf{C}_{\leq j})\mathbb{P}(\mathbf{C}_{\leq j}|\mathbf{Z}_{\leq j}) \end{aligned} \quad (55)$$

by conditional independence of Z_{j+1} .

We claim that $F_{i+1}\mathbb{P}(C_{i+1}|Z_{i+1}\mathbf{Z}_{\leq i}\mathbf{C}_{\leq i})^\beta$ is a test factor determined by $\mathbf{C}_{\leq i+1}, \mathbf{Z}_{\leq i+1}$. To prove this claim, for all $\mathbf{c}_{\leq i}, \mathbf{z}_{\leq i}$, $v = \mu[C_{i+1}Z_{i+1}|\mathbf{c}_{\leq i}, \mathbf{z}_{\leq i}] \in \mathcal{C}_{i+1}$. With $F_{i+1} = F_{i+1}(C_{i+1}Z_{i+1}; \mathbf{c}_{\leq i}, \mathbf{z}_{\leq i})$, we obtain the bound

$$\begin{aligned} \mathbb{E}(F_{i+1}\mathbb{P}(C_{i+1}|Z_{i+1}\mathbf{z}_{\leq i}\mathbf{c}_{\leq i})^\beta | \mathbf{c}_{\leq i}, \mathbf{z}_{\leq i}) &= \mathbb{E}_v(F_{i+1}v(C_{i+1}|Z_{i+1})^\beta) \\ &\leq 1, \end{aligned} \quad (56)$$

where we invoked the assumption that F_{i+1} is a PEF with power β for $\mathcal{C}_{i+1}$. By arbitrariness of $\mathbf{c}_{\leq i}, \mathbf{z}_{\leq i}$, and because the factors are determined by $\mathbf{C}_{\leq i+1}, \mathbf{Z}_{\leq i+1}$, the claim follows. The product of these test factors is

$$\begin{aligned} &\prod_{j=0}^{i-1} F_{j+1}\mathbb{P}(C_{j+1}|Z_{j+1}\mathbf{Z}_{\leq j}\mathbf{C}_{\leq j})^\beta \\ &= T_i \prod_{j=0}^{i-1} \mathbb{P}(C_{j+1}|Z_{j+1}\mathbf{Z}_{\leq j}\mathbf{C}_{\leq j})^\beta \\ &= T_i \mathbb{P}(\mathbf{C}_{\leq i}|\mathbf{Z}_{\leq i})^\beta, \end{aligned} \quad (57)$$

with $T_i = \prod_{j=1}^i F_j$. Thus, the sequence $(T_i \mathbb{P}(\mathbf{C}_{\leq i}|\mathbf{Z}_{\leq i})^\beta)_i$ is a test supermartingale, where the inverse of the second factor is monotone nondecreasing. We remark that as a consequence, T_n is a PEF with power β for $\mathcal{H}(\mathcal{C})$, that is, for $\mathbf{R} = \mathbf{C} = \mathbf{D}$, chaining PEFs yields PEFs for standard models.

From Eqs. (20) and (21) with $U_i = \mathbb{P}(\mathbf{C}_{\leq i}|\mathbf{Z}_{\leq i})^{-\beta}$ and manipulating the inequalities inside $\mathbb{P}(\cdot)$, we get

$$\mathbb{P}(\mathbb{P}(\mathbf{C}_{\leq n}|\mathbf{Z}_{\leq n}) \geq (T_n \epsilon)^{-1/\beta}) \leq \epsilon, \quad (58)$$

$$\mathbb{P}(\mathbb{P}(\mathbf{C}_{\leq n}|\mathbf{Z}_{\leq n}) \geq (\max_i T_i \epsilon)^{-1/\beta}) \leq \epsilon. \quad (59)$$

To conclude that $\mathcal{U}_\beta$ and $\mathcal{U}_\beta^*$ are UPEs, it now suffices to refer to their defining identities Eqs. (52) and (53). ■

That F_{i+1} can be parameterized in terms of the past as $F_{i+1} = F_{i+1}(C_{i+1}Z_{i+1}; \mathbf{C}_{\leq i}, \mathbf{Z}_{\leq i})$ allows for adapting the PEFs based on $\mathbf{CZ}$, but no other information in $\mathbf{R}$ can be used. Since T_i is a function of $\mathbf{C}_{\leq i}, \mathbf{Z}_{\leq i}$, this enables stopping when a target value of T_i is achieved.

In order to use additional information available in $\mathbf{R}$, we now treat the case where $C, R \neq D$. While not necessary at this point, we do this with respect to a softening of the PEF properties.

Definition 24. Let $\beta > 0$, and let $\mathcal{C}$ be any model, not necessarily convex. A *soft probability estimation factor (soft PEF)* with power β for $\mathcal{C}$ and $D|Z$ is a non-negative RV $F = F(\mathbf{CZ})$ such that for each $v \in \mathcal{C}$, there exists a function $q(\mathbf{C}|DZ) \geq 0$ of CDZ depending on v implicitly with $\sum_c q(c|dz) \leq 1$ for

all dz and

$$\mathbb{E}_v \left(\frac{F(CZ)}{q(C|DZ)^\beta} v(C|Z)^\beta \right) \leq 1. \quad (60)$$

If $\sum_c q(c|dz) < 1$ for some dz , we can increase q to ensure $\sum_c q(c|dz) = 1$ without increasing the left-hand side of Eq. (60). Hence, without loss of generality, we can let $\sum_c q(c|dz) = 1$ in the above definition. We always set $q(c|dz) = 0$ for values cdz with $v(cdz) = 0$. This does not cause problems with Eq. (60) given the convention that when expanding an expectation over a finite set of values, probability-zero values are omitted from the sum.

The direct calculation in the proof of the next lemma shows that PEFs are soft PEFs.

Lemma 25. If F is a PEF with power β for $\mathcal{C}$, then it is a soft PEF with power β for $\mathcal{C}$.

Proof. Let F be a PEF with power β for $\mathcal{C}$ and $v \in \mathcal{C}$. Define $q(C|DZ) = v(C|DZ)$. Then q satisfies the condition in the definition of soft PEFs. Since D is a function of C , we have $v(dcz) = v(cz)[d = D(c)]$. From this identity, we can deduce

$$\begin{aligned} & \mathbb{E}_v \left(\frac{F(CZ)}{q(C|DZ)^\beta} v(C|Z)^\beta \right) \\ &= \mathbb{E}_v \left(F(CZ) \left(\frac{v(C|Z)}{q(C|DZ)} \right)^\beta \right) \\ &= \mathbb{E}_v \left(F(CZ) \left(\frac{v(CZ)/v(Z)}{v(CDZ)/v(DZ)} \right)^\beta \right) \\ &= \mathbb{E}_v \left(F(CZ) \left(\frac{v(DZ)}{v(Z)} \right)^\beta \right) \\ &= \mathbb{E}_v (F(CZ) v(D|Z)^\beta) \\ &\leq 1. \end{aligned} \quad (61)$$

Since $v \in \mathcal{C}$ is arbitrary, this verifies that every PEF F is a soft PEF. ■

Lemma 26. Let $F(CZ)$ be a soft PEF with power β for $\mathcal{H}$. Then $\mathcal{V} = (F\epsilon)^{-1/\beta}$ is an ϵ -soft UPE($D: \mathcal{C}|\mathcal{Z}; \mathcal{H}$).

Proof. Fix $v \in \mathcal{H}$ and let $q(C|DZ)$ be the function witnessing that F is a soft PEF for this v . Define

$$\tilde{v}_z(dz) = v(dz)[\mathcal{V}(cz) \geq v(cz)/q(c|dz)]. \quad (62)$$

The Markov inequality and the definition of soft PEFs implies that

$$\begin{aligned} & \mathbb{P}_v(\mathcal{V}(CZ) < v(C|Z)/q(C|DZ)) \\ &= \mathbb{P}_v((v(C|Z)/q(C|DZ))^\beta > \mathcal{V}(CZ)^\beta) \\ &= \mathbb{P}_v(F(CZ)(v(C|Z)/q(C|DZ))^\beta > 1/\epsilon) \\ &\leq \epsilon. \end{aligned} \quad (63)$$

Hence

$$\begin{aligned} \sum_{dcz} \tilde{v}_z(dz)v(z) &= \sum_{dcz} v(dcz)[\mathcal{V}(cz) \geq v(cz)/q(c|dz)] \\ &= \mathbb{P}_v(\mathcal{V}(CZ) \geq v(C|Z)/q(C|DZ)) \\ &\geq 1 - \epsilon. \end{aligned} \quad (64)$$

The definition of $\tilde{v}_z(dz)$ and $v(dz|z) = v(c|z)[d = D(c)]$ ensure that

$$\tilde{v}_z(dz) \leq \mathcal{V}(cz)q(c|dz), \quad (65)$$

as required for soft UPEs. Since v is arbitrary, the above verifies that $\mathcal{V} = (F\epsilon)^{-1/\beta}$ is an ϵ -soft UPE. ■

Theorem 27. Fix $\beta > 0$. Let $\mathbf{F}$ be a sequence of past-parameterized soft PEFs, with F_i a soft PEF with power β for $\mathcal{C}_i$. Then $\mathcal{U}_\beta(\mathbf{RZ}, 1 - \epsilon)$ as defined in Eq. (52) is an ϵ -soft UPE($\mathbf{D}:\mathbf{R}|\mathbf{Z}; \mathcal{H}(\mathcal{C})$).

Proof. Below we first show that chaining soft PEFs yields soft PEFs. Then, by applying lemma 26, we prove the theorem.

For this result, direct chaining of the probabilities fails. Instead, we decompose alternately according to $\mathcal{C}_i$ and $\mathcal{R}_i$ given $\mathcal{C}_i$, as follows:

$$\begin{aligned} & \mathbb{P}(\mathbf{C}_{\leq i+1} \mathbf{R}_{\leq i+1} | \mathbf{Z}_{\leq i+1}) \\ &= \mathbb{P}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i+1}) \mathbb{P}(C_{i+1} R_{i+1} | \mathbf{C}_{\leq i} \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i+1}) \\ &= \mathbb{P}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | Z_{i+1} \mathbf{Z}_{\leq i}) \mathbb{P}(C_{i+1} | Z_{i+1} \mathbf{C}_{\leq i} \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}) \\ &\quad \times \mathbb{P}(R_{i+1} | C_{i+1} Z_{i+1} \mathbf{C}_{\leq i} \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}) \\ &= \mathbb{P}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i}) \mathbb{P}(C_{i+1} | Z_{i+1} \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}) \\ &\quad \times \mathbb{P}(R_{i+1} | C_{i+1} Z_{i+1} \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}), \end{aligned} \quad (66)$$

where we simplified in the last step by applying the fact that Z_{i+1} is conditionally independent of $\mathbf{R}_{\leq i}$ given $\mathbf{Z}_{\leq i}$ and taking advantage of the assumption that $\mathcal{C}_i$ is determined by $\mathcal{R}_i$. The identity can be expanded recursively, replacing the first term in the last expression obtained each time. We identify two products after the expansion, namely,

$$P_{C,i}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i}) = \prod_{j=0}^{i-1} \mathbb{P}(C_{j+1} | Z_{j+1} \mathbf{R}_{\leq j} \mathbf{Z}_{\leq j}) \quad (67)$$

and

$$P_{R,i}(\mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i}) = \prod_{j=0}^{i-1} \mathbb{P}(R_{j+1} | C_{j+1} Z_{j+1} \mathbf{R}_{\leq j} \mathbf{Z}_{\leq j}), \quad (68)$$

which satisfy

$$\mathbb{P}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i}) = \mathbb{P}(\mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i}) = P_{C,i} P_{R,i}, \quad (69)$$

whenever $\mathbb{P}(\mathbf{C}_{\leq i} \mathbf{R}_{\leq i} | \mathbf{Z}_{\leq i})$ is not zero. Again, we took advantage of our assumption that $\mathcal{C}_i$ is a function of $\mathcal{R}_i$ to omit unnecessary arguments.

Let μ be an arbitrary distribution in $\mathcal{H}(\mathcal{C})$ and note that for all $\mathbf{r}_{\leq i} \mathbf{z}_{\leq i}$, $v_{i+1} = \mu[R_{i+1} Z_{i+1} | \mathbf{r}_{\leq i} \mathbf{z}_{\leq i}]$ satisfies that $v_{i+1}[C_{i+1} Z_{i+1}] \in \mathcal{C}_{i+1}$. According to the definition of soft PEFs, there exists $q_{i+1}(C_{i+1} | D_{i+1} Z_{i+1}; \mathbf{r}_{\leq i} \mathbf{z}_{\leq i})$ such that

$$\mathbb{E}_{v_{i+1}}(F_{i+1}(v_{i+1}(C_{i+1} | Z_{i+1})/q_{i+1}(C_{i+1} | D_{i+1} Z_{i+1}))^\beta) \leq 1 \quad (70)$$

and $\sum_{c_{i+1}} q_{i+1}(c_{i+1}|d_{i+1}z_{i+1}) \leq 1$ for all $d_{i+1}z_{i+1}$, where the past parameters are implicit. We define the additional product

$$Q_i(C_{\leq i}|D_{\leq i}Z_{\leq i}; \mathbf{R}_{\leq i-1}) = \prod_{j=0}^{i-1} q_{j+1}(C_{j+1}|D_{j+1}Z_{j+1}; \mathbf{R}_{\leq j}Z_{\leq j}). \quad (71)$$

Later we find that the function $q(\mathbf{R}|\mathbf{DZ})$ needed for the final soft UPE construction according to definition 15 is the $i = n$ instance of

$$p_{\leq i}(\mathbf{R}_{\leq i}|D_{\leq i}Z_{\leq i}) = Q_i(C_{\leq i}|D_{\leq i}Z_{\leq i}; \mathbf{R}_{\leq i-1})P_{R,i}(\mathbf{R}_{\leq i}Z_{\leq i}).$$

Write

$$\begin{aligned} p_{i+1}(R_{i+1}|D_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i}) \\ = q_{i+1}(C_{i+1}|D_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i})\mathbb{P}(R_{i+1}|C_{i+1}Z_{i+1}\mathbf{R}_{\leq i}Z_{\leq i}) \end{aligned} \quad (72)$$

for the incremental factor defining $p_{\leq i+1}$, so that $p_{\leq i+1} = p_{i+1}p_{\leq i}$. We show by induction that $p_{\leq i}$ satisfies the required normalization condition. The initial value is $p_{\leq 0} = 1$ (empty products evaluate to 1), and for the induction step, for all $\mathbf{d}_{\leq i+1}z_{\leq i+1}$,

$$\begin{aligned} \sum_{\mathbf{r}_{\leq i+1}} p_{\leq i+1}(\mathbf{r}_{\leq i+1}|\mathbf{d}_{\leq i+1}z_{\leq i+1}) \\ = \sum_{\mathbf{r}_{\leq i}} p_{\leq i}(\mathbf{r}_{\leq i}|\mathbf{d}_{\leq i}z_{\leq i}) \sum_{r_{i+1}} q_{i+1}(c_{i+1}|d_{i+1}z_{i+1}; \mathbf{r}_{\leq i}z_{\leq i}) \\ \times \mathbb{P}(r_{i+1}|c_{i+1}z_{i+1}\mathbf{r}_{\leq i}z_{\leq i}). \end{aligned} \quad (73)$$

where to get the first line we used Eq. (72) and the fact $\mathbb{P}(R_{i+1}|Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i}) = \mathbb{P}(R_{i+1}|C_{i+1}Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i})$ because C_{i+1} is determined by R_{i+1} , and to get the last line we used the definition of soft PEFs, particularly Eq. (70). By arbitrariness of $\mathbf{r}_{\leq i}z_{\leq i}$, the claim follows.

Since $\prod_{j=0}^{n-1} p_{j+1} = p_{\leq n}$, if we set the function q in the definition 24 to $p_{\leq n}$, we conclude that $T_n(\mathbf{RZ}) = \prod_{j=0}^{n-1} F_{j+1}$ is a soft PEF with power β for $\mathbf{D}|\mathbf{Z}$ and $\mathcal{H}(\mathcal{C})$. Hence, chaining soft PEFs yields soft PEFs for standard models. Since $\mathcal{U}_\beta(\mathbf{RZ}, (1 - \epsilon)) = (T_n \epsilon)^{-1/\beta}$, to finish the proof of the theorem, we apply lemma 26 with C there replaced by $\mathbf{R}$ here, D there with $\mathbf{D}$ here and Z there with $\mathbf{Z}$ here. The proof of lemma 26 shows that the softness witness $q(\mathbf{R}|\mathbf{DZ})$ for $\mathcal{U}_\beta(\mathbf{RZ}, 1 - \epsilon)$ is also given by $p_{\leq n}$. ■

C. Effectiveness of PEF Constraints

The constraints on PEFs F_i are linear, but it is not obvious that this set of linear constraints has a practical realization as a linear or semidefinite program. Theorem 28 below shows

The inner sum on the right-hand side evaluates to

$$\begin{aligned} \sum_{r_{i+1}} q_{i+1}(c_{i+1}|d_{i+1}z_{i+1}; \mathbf{r}_{\leq i}z_{\leq i})\mathbb{P}(r_{i+1}|c_{i+1}z_{i+1}\mathbf{r}_{\leq i}z_{\leq i}) \\ = \sum_{c_{i+1}} \left(\sum_{r_{i+1}: C_{i+1}(r_{i+1})=c_{i+1}} \mathbb{P}(r_{i+1}|c_{i+1}z_{i+1}\mathbf{r}_{\leq i}z_{\leq i}) \right) \\ \times q_{i+1}(c_{i+1}|d_{i+1}z_{i+1}; \mathbf{r}_{\leq i}z_{\leq i}) \\ = \sum_{c_{i+1}} q_{i+1}(c_{i+1}|d_{i+1}z_{i+1}; \mathbf{r}_{\leq i}z_{\leq i}) \\ \leq 1. \end{aligned} \quad (74)$$

Substituting back in Eq. (73) and applying the induction hypothesis gives the normalization condition.

We claim that

$$\begin{aligned} G_{i+1}(R_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i}) \\ = \frac{F_{i+1}(C_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i})}{p_{i+1}(R_{i+1}|D_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i})^\beta} \mathbb{P}(R_{i+1}|Z_{i+1}\mathbf{R}_{\leq i}Z_{\leq i})^\beta \end{aligned} \quad (75)$$

is a test factor determined by $\mathbf{R}_{\leq i+1}Z_{\leq i+1}$. To prove this claim, with $F_{i+1} = F_{i+1}(C_{i+1}Z_{i+1}; \mathbf{R}_{\leq i}Z_{\leq i})$ and omitting the past parameters of p_{i+1} , for each $\mathbf{r}_{\leq i}z_{\leq i}$, we compute

$$\begin{aligned} \mathbb{E} \left(F_{i+1} \left(\frac{\mathbb{P}(R_{i+1}|Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i})}{p_{i+1}(R_{i+1}|D_{i+1}Z_{i+1})} \right)^\beta \middle| \mathbf{r}_{\leq i}z_{\leq i} \right) &= \mathbb{E} \left(F_{i+1} \left(\frac{\mathbb{P}(R_{i+1}|C_{i+1}Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i})\mathbb{P}(C_{i+1}|Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i})}{\mathbb{P}(R_{i+1}|C_{i+1}Z_{i+1}\mathbf{r}_{\leq i}z_{\leq i})q_{i+1}(C_{i+1}|D_{i+1}Z_{i+1})} \right)^\beta \middle| \mathbf{r}_{\leq i}z_{\leq i} \right) \\ &= \mathbb{E}_{v_{i+1}} \left(F_{i+1} \left(\frac{\mathbb{P}_{v_{i+1}}(C_{i+1}|Z_{i+1})}{q_{i+1}(C_{i+1}|D_{i+1}Z_{i+1})} \right)^\beta \right) \\ &\leq 1, \end{aligned} \quad (76)$$

that when $C = D$, it suffices to use the extreme points $\text{Extr}(\mathcal{C})$ of the convex set $\mathcal{C}$, which implies that if $\mathcal{C}$ is a known, convex polytope, then a finite number of constraints suffice. This includes the cases where $\mathcal{C}$ is characterized by typical settings and nonsignaling constraints for (k, l, m) Bell-test configurations and generalizes to cases where the settings distributions are not fixed but constrained to be a convex polytope. See Sec. VIII A. For the case where $C \neq D$, we show that soft PEFs for $\text{Extr}(\mathcal{C})$ are soft PEFs for $\mathcal{C}$. In particular PEFs for $\text{Extr}(\mathcal{C})$ yield soft PEFs for $\mathcal{C}$. When $\mathcal{C}$ is a convex polytope this yields an effective construction for soft PEFs.

Theorem 28. Let $\mathcal{C}$ be a convex set of distributions of CZ . Then for $\beta > 0$ the family of inequalities

$$\mathbb{E}_v(F(CZ)v(C|Z)^\beta) \begin{matrix} \geq 0 \\ \leq 1 \end{matrix} \quad \text{for } v \in \mathcal{C}, \quad (77)$$

is implied by the subset with $v \in \text{Extr}(\mathcal{C})$.

Proof. For a given $\nu \in \mathcal{C}$, we can write the constraint on F as

$$\mathbb{E}_\nu(F(CZ)\nu(C|Z)^\beta) = \sum_{cz} F(cz)\nu(cz)^{1+\beta}/\nu(z)^\beta \leq 1. \quad (78)$$

Thus each ν determines a linear inequality for F . For distributions σ on CZ and ρ on Z with $\text{Supp}(\sigma[Z]) \subseteq \text{Supp}(\rho)$ define

$$l_{\sigma|\rho}(cz) = \sigma(cz)^{1+\beta}/\rho(z)^\beta, \quad (79)$$

so that $l_{\nu|\nu}(cz)$ are the coefficients of F in the inequality above. If $\rho(z) = 0$, then $\sigma(cz) \leq \sigma(z) = 0$, so we set $l_{\nu|\rho}(cz) = 0$, which is consistent since the power in the numerator is larger than that in the denominator. Since F is non-negative, it suffices to show that for each cz , the coefficients $l_{\nu|\rho}(cz)$ are biconvex functions of ν and ρ . If $\nu = \lambda\nu_1 + (1-\lambda)\nu_2$ with $0 \leq \lambda \leq 1$ and $\nu_i \in \mathcal{C}$, then

$$l_{\nu|\nu}(cz) \leq \lambda l_{\nu_1|\nu_1}(cz) + (1-\lambda)l_{\nu_2|\nu_2}(cz). \quad (80)$$

Given that $\sum_{cz} F(cz)l_{\nu'|\nu'}(cz) \leq 1$ for $\nu' = \nu_1$ and $\nu' = \nu_2$, we find

$$\begin{aligned} \sum_{cz} F(cz)l_{\nu|\nu}(cz) &\leq \sum_{cz} F(cz)(\lambda l_{\nu_1|\nu_1}(cz) + (1-\lambda)l_{\nu_2|\nu_2}(cz)) \\ &\leq \lambda \sum_{cz} F(cz)l_{\nu_1|\nu_1}(cz) + (1-\lambda) \\ &\quad \times \sum_{cz} F(cz)l_{\nu_2|\nu_2}(cz) \\ &\leq \lambda + (1-\lambda) = 1. \end{aligned} \quad (81)$$

This observation extends to arbitrary finite convex combinations by reduction to the above case. Consequently the constraints associated with extremal distributions suffice.

Let $r(x|y) = x^{1+\beta}y^{-\beta}$, where we define $r(0|0) = 0$. The desired convexity expressed in Eq. (80) is implied by the fact that $r(x|y)$ is biconvex in its arguments for relevant values, which is lemma 29 below. ■

Lemma 29. Define $r(x|y) = x^{1+\beta}y^{-\beta}$ with $\beta > 0$. If (x, y) is expressed as a finite convex combination $(x, y) = \sum_i \lambda_i(x_i, y_i)$ where $x_i \geq 0$, $y_i \geq 0$ and $y_i = 0$ implies $x_i = 0$, then $r(x|y) \leq \sum_i \lambda_i r(x_i|y_i)$.

Proof. This is a special case of the biconvexity of sandwiched Rényi powers, see Refs. [44,45]. To prove biconvexity of $r(x|y)$ directly, it suffices to consider convex combinations of two arguments, so let $x = \lambda x_1 + (1-\lambda)x_2$ and $y = \lambda y_1 + (1-\lambda)y_2$ be convex combinations as expressed in the theorem statement. Consider

$$f(\lambda) = r(\lambda x_1 + (1-\lambda)x_2 | \lambda y_1 + (1-\lambda)y_2). \quad (82)$$

If $y_1 = y_2 = 0$, then $f(\lambda) = 0$ for all λ , which is convex. If $y_1 = 0 < y_2$, then $f(\lambda) = (1-\lambda)x_2^{1+\beta}/y_2^\beta$, which again is convex. By symmetry, so is $f(\lambda)$ for the case where $y_2 = 0 < y_1$. It therefore suffices to consider the case where $y_1 > 0$ and $y_2 > 0$. Now $f(\lambda)$ is continuous for $\lambda \in [0, 1]$ and smooth for $\lambda \in (0, 1)$, so it suffices to show that $f''(\lambda) \geq 0$ as follows:

$$\begin{aligned} f'(\lambda) &= (\lambda x_1 + (1-\lambda)x_2)^\beta (\lambda y_1 + (1-\lambda)y_2)^{-\beta-1} \\ &\quad \times ((1+\beta)(x_1 - x_2)(\lambda y_1 + (1-\lambda)y_2) \\ &\quad + (-\beta)(\lambda x_1 + (1-\lambda)x_2)(y_1 - y_2)) \end{aligned} \quad (83)$$

$$\begin{aligned} f''(\lambda) &= (\lambda x_1 + (1-\lambda)x_2)^{\beta-1} (\lambda y_1 + (1-\lambda)y_2)^{-\beta-2} \\ &\quad \times \left(\beta(1+\beta)(x_1 - x_2)^2 (\lambda y_1 + (1-\lambda)y_2)^2 \right. \\ &\quad + 2(-\beta)(1+\beta)(x_1 - x_2)(y_1 - y_2)(\lambda x_1 \\ &\quad + (1-\lambda)x_2)(\lambda y_1 + (1-\lambda)y_2) \\ &\quad \left. + (-\beta)(-1-\beta)(y_1 - y_2)^2 (\lambda x_1 + (1-\lambda)x_2)^2 \right) \\ &= (\lambda x_1 + (1-\lambda)x_2)^{\beta-1} (\lambda y_1 + (1-\lambda)y_2)^{-\beta-2} \\ &\quad \times \beta(1+\beta)((x_1 - x_2)(\lambda y_1 + (1-\lambda)y_2) \\ &\quad - (y_1 - y_2)(\lambda x_1 + (1-\lambda)x_2))^2, \end{aligned} \quad (84)$$

which is a non-negative multiple of a square. Consequently $f(\lambda)$ is convex in λ , and biconvexity of $r(x|y)$ follows. ■

Corollary 30. If $F(CZ)$ is a PEF with power β for $C|Z$ and $\text{Extr}(\mathcal{C})$, then $F(CZ)$ is a PEF with power β for $C|Z$ and the convex closure of $\mathcal{C}$.

The same observation holds for soft PEFs.

Theorem 31. If $F(CZ)$ is a soft PEF with power β for $D|Z$ and $\text{Extr}(\mathcal{C})$, then $F(CZ)$ is a soft PEF with power β for $D|Z$ and the convex closure of $\mathcal{C}$.

Proof. Let $\nu_i \in \text{Extr}(\mathcal{C})$ and $\nu = \sum_i \lambda_i \nu_i$ be a finite convex combination of the ν_i . Let $q_i(C|DZ)$ witness that F is a soft PEF at ν_i . From the joint convexity of $\nu^{1+\beta}/\sigma^\beta$ according to lemma 29, for all cdz

$$\frac{\nu(cz)^{1+\beta}}{(\sum_i \lambda_i q_i(c|dz)\nu_i(z))^\beta} \leq \sum_i \lambda_i \frac{\nu_i(cz)^{1+\beta}}{(q_i(c|dz)\nu_i(z))^\beta}. \quad (85)$$

To apply lemma 29, note that the definition of soft PEFs implies that $q_i(c|dz) > 0$ whenever $\nu_i(cz) > 0$ as otherwise the defining inequality cannot be satisfied. Define $q(c|dz) = \sum_i \lambda_i q_i(c|dz)\nu_i(z)/\nu(z)$ if $\nu(z) > 0$ and $q(c|dz) = 0$ otherwise. We have

$$\begin{aligned} \sum_c q(c|dz) &= \sum_i \lambda_i \sum_c q_i(c|dz)\nu_i(z)/\nu(z) \\ &\leq \sum_i \lambda_i \nu_i(z)/\nu(z) = \nu(z)/\nu(z) = 1. \end{aligned} \quad (86)$$

Now

$$\begin{aligned} &\frac{F(cz)}{q(c|dz)^\beta} \nu(c|z)^\beta \nu(cz) \\ &= F(cz) \frac{\nu(cz)^{1+\beta}}{q(c|dz)^\beta \nu(z)^\beta} \\ &= F(cz) \frac{\nu(cz)^{1+\beta}}{(\sum_i \lambda_i q_i(c|dz)\nu_i(z))^\beta} \\ &\leq \sum_i \lambda_i F(cz) \nu_i(cz)^{1+\beta} / (q_i(c|dz)\nu_i(z))^\beta \\ &= \sum_i \lambda_i \frac{F(cz)}{q_i(c|dz)^\beta} \nu_i(c|z)^\beta \nu_i(cz). \end{aligned} \quad (87)$$

Summing both sides over dcz gives the desired inequality

$$\mathbb{E}_v \left(\frac{F(CZ)}{q(C|DZ)^\beta} v(C|Z)^\beta \right) \leq \sum_i \lambda_i \mathbb{E}_{v_i} \left(\frac{F(CZ)}{q_i(C|DZ)^\beta} v_i(C|Z)^\beta \right) \leq 1. \quad (88)$$

Corollary 32. If $F(CZ)$ is a PEF with power β for $D|Z$ and $\text{Extr}(\mathcal{C})$, then $F(CZ)$ is a soft PEF with power β for $D|Z$ and the convex closure of $\mathcal{C}$.

Proof. By lemma 25, $F(CZ)$ is a soft PEF with power β for $D|Z$ and $\text{Extr}(\mathcal{C})$, so we can apply theorem 31. ■

D. Log-probs and PEF optimization strategy

Given the probability estimation framework, the number of random bits that can be extracted is determined by the negative logarithm of the final probability estimate after adjustments for error bounds and extractor constraints. We focus on the quantities obtained from Eq. (52), which reflect the final values of the soft PEF products.

Definition 33. Let F_i be soft PEFs with power $\beta > 0$ for C_i . The *log-prob* of $\mathbf{F}$ is $\sum_i \ln(F_i(c_i z_i))/\beta$. Given an error bound ϵ , the *net log-prob* is given by $\sum_i \ln(F_i(c_i z_i))/\beta - \ln(1/\epsilon)/\beta$.

In the randomness generation protocols of Sec. IV C with the soft UPEs of Sec. V B, the net log-prob is effectively a “raw” entropy base e available for extraction. For the protocol to succeed, the net log-prob must exceed $\sigma_h \ln(2)$, where σ_h is the extractor parameter for the input min-entropy in bits. In the protocol with banked randomness, the net log-prob contributes almost fully to the input min-entropy provided to the extractor. In Sec. IX, we further adjust the net log-prob by subtracting the entropy used for settings choices and call the resulting quantity the *net entropy*. This does not yet take into account the entropy needed for the extractor seed, nor the fact that the number of output bits is smaller than the input min-entropy according to the extractor constraints.

The log-prob and net log-prob are empirical quantities. For the average quantities with respect to a distribution, we consider expected contributions per trial.

Definition 34. If $v \in \mathcal{C}$ is the probability distribution of CZ at a trial and F is a soft PEF with power $\beta > 0$ for $\mathcal{C}$, then we define the *log-prob rate* of F at v as

$$\mathcal{O}_v(F; \beta) = \sum_{cz} v(cz) \ln(F(cz))/\beta = \mathbb{E}_v(\ln(F))/\beta. \quad (89)$$

Let v be a distribution of CZ . If $\mu(c_{i+1} z_{i+1} | \mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}) = v(C = c_{i+1}, Z = z_{i+1})$ for each i and c_{i+1}, z_{i+1} , then the trials are *empirically i.i.d.* with distribution v . That is, they are i.i.d. from the point of view of the user, but not necessarily from the point of view of $\mathbf{E}$. In this case, the expected log-prob after n trials is given by $n\mathcal{O}_v(F; \beta)$. When using the term “rate,” we implicitly assume a situation where we expect many trials with the same distribution. We omit the parameters β or v when they are clear from context. The quantity $\mathcal{O}(F)$ could be called the trial’s nominal min-entropy increase. We refer to it as “nominal” because the final min-entropy provided to the extractor has to be reduced by $\ln(1/\epsilon)/\beta$ (see definition 33). This reduction diverges as β goes to zero.

The log-prob rate $\mathcal{O}(F)$ is a concave function of F , where F is constrained to a convex set through the inequalities determined by $\mathcal{C}$. The problem of maximizing $\mathcal{O}(F)$ is therefore that of maximizing a concave function over a convex domain. For (k, l, m) Bell-test configurations with k, l, m small and typical settings and nonsignaling constraints, it is feasible to solve this numerically. However, the complexity of the full problem grows very rapidly with k, l, m .

E. Schema for probability estimation

The considerations above lead to the following schematic probability estimation protocol. We are given the maximum number of available trials n , a probability estimation goal q and the error bound ϵ_h . In general, the protocol yields a soft UPE. The protocol can take advantage of experimentally relevant operational details to improve probability estimation by adapting soft PEFs while it is running. In view of theorem 23 (if $C_i = D_i$) or theorem 27 (if not), each trial’s soft PEF can be adapted by taking advantage of the values of the R_j and Z_j from previous trials as well as any initial information that may be part of E . However, the values of the parameters n, q, ϵ_h and β in the protocol must be fully determined by initial information and cannot be adapted later.

Schema for probability estimation is as follows:

Given: n , the maximum number of trials, q , the probability estimation goal, and ϵ_h , the error bound.

1. Set $T_0 = 1$. Choose $\beta > 0$ based on prior data from device commissioning or other earlier uses of the devices.

2. For each $i = 0, \dots, n - 1$:

2.1. After the i th and before the $(i + 1)$ th trial:

2.1.1. Determine the set $\mathcal{C}_{i+1}$ constraining the distributions at the $(i + 1)$ th trial.

2.1.2. Estimate the distribution $v \in \mathcal{C}_{i+1}$ of $C_{i+1}Z_{i+1}$ from previous measurements (see the remark after the schema).

2.1.3. Based on the estimate v and other information available in $\mathbf{R}_{\leq i}$ and $\mathbf{Z}_{\leq i}$, determine a soft PEF F_{i+1} with power β for $D_{i+1}|Z_{i+1}$ and $\mathcal{C}_{i+1}$, optimized for log-prob rate at v .

2.2. Obtain the $(i + 1)$ th trial result $c_{i+1}z_{i+1}$.

2.3. After the $(i + 1)$ th trial, compute $T_{i+1} = T_i F_{i+1}(c_{i+1}z_{i+1})$.

2.4. If $U = (T_{i+1}\epsilon_h)^{-1/\beta} \leq q$, stop and return U .

3. Return $U = (T_n\epsilon_h)^{-1/\beta}$.

In this schema, we have taken advantage of the ability to stop early while preserving the required coverage properties. According to theorem 27, an implementation of the schema returns the value of an ϵ_h -soft UPE suitable for use in a randomness generation protocol with theorem 19, 20, or 21.

We remark that there are different ways of estimating the distribution $v[C_{i+1}Z_{i+1}]$. Our method works as long as the estimated distribution is determined by $\mathbf{R}_{\leq i} \mathbf{Z}_{\leq i}$ and initial information. In practice, we estimate the distribution $v \in \mathcal{C}_{i+1}$ by maximum likelihood using the past trial results $\mathbf{C}_{\leq i} \mathbf{Z}_{\leq i}$, where the likelihood is computed with the assumption that these trial results are i.i.d. [see Eq. (163) for details]. We emphasize that the use of the i.i.d. assumption is only for performing the estimate. The probability estimation protocol and the derived randomness generation protocol are sound regardless of how we obtain the estimated distribution.

VI. SOME PEF PROPERTIES AND PEF OPTIMALITY

A. Dependence on power

Lemma 35. If F is a soft PEF with power β for $\mathcal{C}$, then for all $0 < \gamma \leq 1$, F^γ is a soft PEF with power $\gamma\beta$ for $\mathcal{C}$. Furthermore, given a distribution $\rho \in \mathcal{C}$, the log-prob rate at ρ is the same for F and F^γ .

We refer to the transformation of soft PEFs in the lemma as “power reduction” by γ .

Proof. We have, for all $\nu \in \mathcal{C}$,

$$\begin{aligned} \mathbb{E}_\nu(F^\gamma(CZ)(\nu(C|Z)/q(C|DZ))^{\gamma\beta}) \\ = \mathbb{E}_\nu((F(CZ)(\nu(C|Z)/q(C|DZ))^\beta)^\gamma) \\ \leq (\mathbb{E}_\nu(F(CZ)(\nu(C|Z)/q(C|DZ))^\beta))^\gamma \leq 1, \end{aligned} \quad (90)$$

by concavity of $x \mapsto x^\gamma$ with $0 < \gamma \leq 1$. Hence F^γ is a soft PEF with power $\gamma\beta$ for $\mathcal{C}$. Accordingly,

$$\begin{aligned} \mathcal{O}_\rho(F; \beta) &= \mathbb{E}_\rho(\ln(F))/\beta = \mathbb{E}_\rho(\gamma \ln(F))/(\gamma\beta) \\ &= \mathbb{E}_\rho(\ln(F^\gamma))/(\gamma\beta) = \mathcal{O}_\rho(F^\gamma; \gamma\beta). \end{aligned} \quad (91)$$

■

The lemma implies the following corollary.

Corollary 36. The supremum of the log-prob rates at ρ for all soft PEFs with power β is nonincreasing in β .

Hence, to determine the best log-prob rate without regard to the error bound, one can analyze soft PEFs in the limit where the power goes to 0. In particular, we determine the best log-prob rate achievable by all PEFs with any power in Sec. VIB, and we show the asymptotic optimality of this log-prob rate in Sec. VID. Accordingly, the best log-prob rate achievable by all PEFs is equal to the best log-prob rate achievable by all soft PEFs. We also note the following property: If F is a PEF with power β for $\mathcal{C}$, then for all $0 < \gamma \leq 1$, F is a PEF with power β/γ for $\mathcal{C}$, see lemma 14 of Ref. [28]. This property is useful for determining the finite-data performance of PEFs, as detailed in Ref. [28].

B. PEFs and entropy estimation

Definition 37. An entropy estimator for $D|C$ and $\mathcal{C}$ is a real-valued function K of CZ with the property that for all $\nu \in \mathcal{C}$,

$$\mathbb{E}_\nu(K(CZ)) \leq \mathbb{E}_\nu(-\ln(\nu(D|Z))). \quad (92)$$

The entropy estimate of K at ν is $\mathbb{E}_\nu(K(CZ))$.

Entropy estimators are the analog of affine min-tradeoff functions in the entropy accumulation framework of Ref. [26]. The following results establish a close relationship between entropy estimators and PEFs. First we show that the set of entropy estimators contains the log- β -roots of positive PEFs with power $\beta > 0$.

Theorem 38. Let F be a function of CZ such that $F \geq 0$. If there exists $\beta > 0$ such that F^β is a PEF with power β for $\mathcal{C}$, then for all $\nu \in \mathcal{C}$, $\mathbb{E}_\nu(\ln(F(CZ))) \leq \mathbb{E}_\nu(-\ln(\nu(D|Z)))$.

Proof. Suppose that F^β is a PEF with power $\beta > 0$ for $\mathcal{C}$. For any distribution $\nu \in \mathcal{C}$,

$$\begin{aligned} 1 &\geq \mathbb{E}_\nu(F^\beta(CZ)\nu(D|Z)^\beta) \\ &= \mathbb{E}_\nu(e^{\beta \ln(F(CZ)) + \beta \ln(\nu(D|Z))}) \\ &\geq \mathbb{E}_\nu(1 + \beta \ln(F(CZ)) + \beta \ln(\nu(D|Z))) \\ &= 1 + \beta[\mathbb{E}_\nu(\ln(F(CZ))) - \mathbb{E}_\nu(-\ln(\nu(D|Z)))]. \end{aligned} \quad (93)$$

Here we used that $e^x \geq 1 + x$. It follows that $\mathbb{E}_\nu(\ln(F(CZ))) \leq \mathbb{E}_\nu(-\ln(\nu(D|Z)))$. ■

According to theorem 38, positive PEFs F^β with power β define entropy estimators $\ln(F)$ whose entropy estimate at a distribution $\rho \in \mathcal{C}$ is the log-prob rate of the PEF. Our next goal is to relate log-prob rates for PEFs with power approaching 0 to entropy estimates. For this, we need the following definition.

Definition 39. The asymptotic gain rate at ρ of $\mathcal{C}$ is the supremum of the log-prob rates at ρ achievable by PEFs with any positive power for $\mathcal{C}$.

By the properties of PEFs, the asymptotic gain rate at ρ of $\mathcal{C}$ is the same as that of $\text{Cvx}(\mathcal{C})$.

The asymptotic gain rate is defined in terms of a supremum over all PEFs, including those whose support is properly contained in $\text{Rng}(CZ)$. The next lemma shows that the same supremum is obtained for positive PEFs.

Lemma 40. Let g be the asymptotic gain rate at ρ of $\mathcal{C}$ and define

$$\begin{aligned} g_+ &= \sup\{\mathcal{O}_\rho(F^\beta; \beta) : F > 0 \\ &\text{and } F^\beta \text{ is a PEF with power } \beta > 0 \text{ for } \mathcal{C}\}. \end{aligned} \quad (94)$$

Then $g = g_+$ and the supremum of the entropy estimates at ρ of entropy estimators for $\mathcal{C}$ is at least g .

Proof. It is clear that $g \geq g_+$. Both g and g_+ are non-negative because we can always choose $F = 1$ (see theorem 46 below). It suffices to consider the nontrivial case $g > 0$. Suppose that F contributes nontrivially to the supremum defining the asymptotic gain rate, namely, F^β is a PEF with power $\beta > 0$ and log-prob rate $g_F = \mathcal{O}_\rho(F^\beta; \beta) > 0$. We can choose F so that g_F is arbitrarily close to g . Consider $G^\beta = (1 - \delta)F^\beta + \delta$ with $\delta > 0$ sufficiently small. Then G^β is a positive PEF with power β (see theorem 46) and has log-prob rate

$$\begin{aligned} g_G &= \mathcal{O}_\rho(G^\beta; \beta) = \mathbb{E}_\rho(\ln(G^\beta))/\beta \\ &= \mathbb{E}_\rho(\ln((1 - \delta)F^\beta + \delta))/\beta \\ &\geq \mathbb{E}_\rho(\ln((1 - \delta)F^\beta))/\beta \\ &= g_F + \ln(1 - \delta)/\beta, \end{aligned} \quad (95)$$

by monotonicity of the logarithm. We have $\ln(1 - \delta)/\beta = O(\delta)$. Consequently,

$$g_G \geq g_F + O(\delta). \quad (96)$$

Since δ is arbitrary, we can find positive G so that G^β is a PEF with power β and log-prob rate g_G at ρ bounded below by a quantity arbitrarily close to g_F . Since g_F is itself arbitrarily close to g , we have $g = g_+$. The last statement of the lemma follows from theorem 38 and the definitions of asymptotic gain rate and entropy estimators. ■

According to the lemma above, the supremum of the entropy estimates at ρ by entropy estimators is an upper bound for the asymptotic gain rate. The upper bound could be strict. However, the next theorem implies that in fact, the supremum of the entropy estimates gives exactly the asymptotic gain rate.

Theorem 41. Suppose K is an entropy estimator for $\mathcal{C}$. Then the asymptotic gain rate at ρ is at least the entropy estimate of K at ρ given by $\sigma = \mathbb{E}_\rho(K(CZ))$.

Proof. To prove the theorem requires constructing families of PEFs with small powers whose log-prob rates approach the entropy estimate of K .

Let $k_{\max} = \max(K)$ and $k_{\min} = \min(K)$. We may assume that $k_{\max} > 0$ as otherwise, the entropy estimate is not positive. For sufficiently small $\epsilon > 0$, we determine $\gamma > 0$ such that $G(CZ)^\gamma = (e^{-\epsilon+K(CZ)})^\gamma$ is a PEF with power γ for $\mathcal{C}$. We require that $\epsilon < 1/2$. Consider $v \in \mathcal{C}$ and define

$$f(\gamma) = \mathbb{E}_v(G(CZ)^\gamma v(D|Z)^\gamma) = \sum_{cz} (e^{-\epsilon+K(cz)} v(d|z))^\gamma v(cz). \quad (97)$$

G^γ is a PEF with power γ provided that $f(\gamma) \leq 1$ for all v . We Taylor-expand $f(\gamma)$ at $\gamma = 0$ for $\gamma \geq 0$ with a second-order remainder.

$$\begin{aligned} f(0) &= 1, \\ \frac{d}{d\gamma} f(\gamma) &= \sum_{cz} \ln(e^{-\epsilon+K(cz)} v(d|z)) (e^{-\epsilon+K(cz)} \\ &\quad \times v(d|z))^\gamma v(cz) \\ &= \sum_{cz} (K(cz) + \ln(v(d|z)) - \epsilon) \\ &\quad \times (e^{-\epsilon+K(cz)} v(d|z))^\gamma v(cz), \\ \left. \frac{d}{d\gamma} f(\gamma) \right|_{\gamma=0} &= \sum_{cz} (K(cz) + \ln(v(d|z)) - \epsilon) v(cz) \\ &= \mathbb{E}_v(K(CZ) + \ln(v(D|Z))) - \epsilon \\ &\leq -\epsilon, \end{aligned} \quad (98)$$

$$\begin{aligned} \frac{d^2}{d\gamma^2} f(\gamma) &= \sum_{cz} (K(cz) + \ln(v(d|z)) - \epsilon)^2 \\ &\quad \times (e^{-\epsilon+K(cz)} v(d|z))^\gamma v(cz). \end{aligned} \quad (99)$$

The second derivative is positive, so $f(\gamma)$ is convex and

$$f(\gamma) \leq 1 - \gamma\epsilon + \int_0^\gamma \frac{d^2}{d\gamma^2} f(\gamma) \Big|_{\gamma=\gamma'} (\gamma - \gamma') d\gamma'. \quad (100)$$

To bound the remainder term, we use $v(cz) \leq v(dz) = v(d|z)v(z)$, $e^{-\epsilon} \leq 1$ and $e^{K(CZ)\gamma} \leq e^{k_{\max}\gamma}$ for

$$\begin{aligned} \frac{d^2}{d\gamma^2} f(\gamma) &\leq \sum_{cz} (K(cz) + \ln(v(d|z)) - \epsilon)^2 v(d|z)^{1+\gamma} \\ &\quad \times e^{k_{\max}\gamma} v(z). \end{aligned} \quad (101)$$

As a function of $v(d|z)$ (which is bounded by 1) the summand has a maximum value determined by the maximum of $g(x) = (a + \ln(x))^2 x^{1+\gamma}$ for $0 < x \leq 1$ with $a = K(cz) - \epsilon$. The critical points of $g(x)$ are obtained from solving

$$0 = \frac{d}{dx} g(x) = x^\gamma (a + \ln(x))((1 + \gamma)(a + \ln(x)) + 2). \quad (102)$$

The solutions $x = 0$ and $\ln(x) = -a$ are minima with $g(x) = 0$. The remaining solution is obtained from $\ln(x_0) = -(a +$

$2/(1 + \gamma))$. The candidate maxima are at this solution and at the boundary $x_1 = 1$. The values of g at these points are

$$\begin{aligned} g(x_0) &= \frac{4}{(1 + \gamma)^2} e^{-(a(1+\gamma)+2)}, \\ g(x_1) &= a^2. \end{aligned} \quad (103)$$

If $x_0 > 1$, the solution at x_0 is irrelevant. The condition $x_0 \leq 1$ is equivalent to $-(a(1 + \gamma) + 2) \leq 0$, in which case $g(x_0) \leq 4$. We have $a^2 \leq \max((k_{\max} - \epsilon)^2, (k_{\min} - \epsilon)^2)$, so let $u = \max(4, k_{\max}^2, k_{\min}^2 + |k_{\min}| + 1/4)$, which is a loose upper bound on the maximum of $g(x)$ for $0 < x \leq 1$. For the bound, we used that $k_{\max} > 0$ and $0 < \epsilon < 1/2$ imply $(k_{\max} - \epsilon)^2 \leq \max(1/4, k_{\max}^2)$. Returning to bounding f , we get

$$\begin{aligned} \frac{d^2}{d\gamma^2} f(\gamma) &\leq \sum_{cz} u e^{k_{\max}\gamma} v(z) \\ &= u e^{k_{\max}\gamma} \sum_c 1 = u e^{k_{\max}\gamma} |\text{Rng}(\mathcal{C})|, \end{aligned} \quad (104)$$

from which

$$\begin{aligned} f(\gamma) &\leq 1 - \epsilon\gamma + \int_0^\gamma u e^{k_{\max}\gamma'} |\text{Rng}(\mathcal{C})| (\gamma - \gamma') d\gamma' \\ &\leq 1 - \epsilon\gamma + u |\text{Rng}(\mathcal{C})| e^{k_{\max}\gamma} \int_0^\gamma (\gamma - \gamma') d\gamma' \\ &= 1 - \epsilon\gamma + u |\text{Rng}(\mathcal{C})| e^{k_{\max}\gamma} \gamma^2 / 2. \end{aligned} \quad (105)$$

To ensure that $f(\gamma) \leq 1$, it suffices to satisfy

$$\gamma \leq \frac{2\epsilon e^{-k_{\max}\gamma}}{u |\text{Rng}(\mathcal{C})|}. \quad (106)$$

Let w be the right-hand side of this inequality. To remove the dependence of w on γ while maintaining $f(\gamma) \leq 1$, we reduce the upper bound on γ in a few steps. Since $e^{-k_{\max}\gamma} \geq 1 - k_{\max}\gamma$,

$$w \geq w' = \frac{2\epsilon(1 - k_{\max}\gamma)}{u |\text{Rng}(\mathcal{C})|}. \quad (107)$$

If $\gamma \leq w'$, then $\gamma \leq 2\epsilon/(u |\text{Rng}(\mathcal{C})|)$, so we can substitute this in the right-hand side for

$$w' \geq w'' = \frac{2\epsilon(1 - 2k_{\max}\epsilon/(u |\text{Rng}(\mathcal{C})|))}{u |\text{Rng}(\mathcal{C})|}. \quad (108)$$

For $\epsilon < u |\text{Rng}(\mathcal{C})|/(4k_{\max})$, this simplifies further to $w'' \geq w''' = \epsilon/(u |\text{Rng}(\mathcal{C})|)$ by substituting the bound on ϵ for the second occurrence of ϵ in the expression for w'' . The bound on ϵ is satisfied since $\epsilon < 1/2$ is already assumed and from $u \geq \max(4, k_{\max}^2, u/(4k_{\max})) \geq 1/2$. We now require

$$\begin{aligned} \gamma &\leq w''' = \frac{\epsilon}{u |\text{Rng}(\mathcal{C})|} \\ &= \frac{\epsilon}{\max(4, k_{\max}^2, k_{\min}^2 + |k_{\min}| + 1/4) |\text{Rng}(\mathcal{C})|}, \end{aligned} \quad (109)$$

which ensures that $f(\gamma) \leq 1$. Since the bound on $f(\gamma)$ is independent of $v \in \mathcal{C}$, G^γ is a PEF with power γ and the

log-prob rate of G^γ at ρ is

$$\begin{aligned}\mathcal{O}_\rho(G^\gamma; \gamma) &= \mathcal{O}_\rho(e^{-\epsilon^\gamma} e^{K^\gamma}; \gamma) \\ &= \mathbb{E}_\rho(K(CZ) - \epsilon) = \sigma - \epsilon,\end{aligned}\quad (110)$$

which approaches σ as ϵ goes to zero, as required to complete the proof of the theorem. ■

C. Error bound tradeoffs

So far, the results of this section have ignored the contribution of the error bound ϵ_h to the probability estimate, giving the appearance that arbitrarily small powers are optimal. For a given finite number of trials, or if ϵ_h is intended to decrease exponentially with output min-entropy, the optimal power for soft PEFs is bounded away from zero. This is because ϵ_h increases the probability estimate by a factor of $\epsilon_h^{-1/\beta}$, which diverges as β goes to zero. We can analyze the situation for the case where $\epsilon_h = e^{-\kappa\sigma n}$, where σ is the log-prob rate $\mathcal{O}_\rho(F; \beta)$ at a given distribution $\rho \in \mathcal{C}$. Per trial, we get a *net log-prob rate* of

$$\begin{aligned}\mathcal{O}_\rho(F; \kappa; \beta) &= \mathcal{O}_\rho(F; \beta) - \kappa \mathcal{O}_\rho(F; \beta)/\beta \\ &= (1 - \kappa/\beta) \mathbb{E}_\rho(\ln(F(CZ)))/\beta.\end{aligned}\quad (111)$$

We can consider the family F^γ of soft PEFs with power $\gamma\beta$ obtained by power reduction of F . The net log-prob rate at power $\gamma\beta$ is

$$\mathcal{O}_\rho(F^\gamma, \kappa; \gamma\beta) = \left(1 - \frac{\kappa}{\gamma\beta}\right) \mathcal{O}_\rho(F; \beta). \quad (112)$$

Consequently, the net log-prob rate is never improved by power reduction of a given soft PEF. However, there are usually soft PEFs with higher net log-prob rates but lower powers. There is therefore a tradeoff between the log-prob rate and the term κ/β that we expect to be optimized at some finite nonzero β . This effect is demonstrated by example in Sec. IX, see Fig. 1.

We remark that in the above definitions, we may consider $\kappa = \kappa(n)$ as a function of n .

D. Optimality of PEFs

Consider the situation where $R_i Z_i E_i$ are i.i.d. with a distribution ν such that for all e , $\nu[C_i Z_i | E_i = e]$ is in $\mathcal{C}$. Here we have structured the RV E as a sequence RV, $\mathbf{E}$. We call this the *i.i.d. scenario* with respect to ν , where we specify ν as a distribution of RZE without indices. We consider the error bound to be $e^{-o(n)}$. Our results so far establish that if the asymptotic gain rate at $\rho = \nu[CZ]$ is g , we can certify smooth conditional min-entropy and extract near-uniform random bits at an asymptotic rate arbitrarily close to g . This is formalized by the next theorem.

Theorem 42. Let g be the asymptotic gain rate at ρ of $\mathcal{C}$ and $\epsilon_n = e^{-\kappa(n)n}$ with $\kappa(n) = o(1)$. Assume the i.i.d. scenario. Then for any $\delta > 0$, there exists a PEF G^β with power β for $\mathcal{C}$

such that with asymptotic probability one, the net log-prob after n trials with PEFs $F_i = G(C_i Z_i)^\beta$ and error bound ϵ_n exceeds $(g - \delta)n$.

Proof. By the definition of asymptotic gain rate, there exists a PEF G^β with power β such that $\mathbb{E}(\ln(G)) = g - \delta/3$. We may assume $g - \delta/3 > 0$. The RV $S_n = \sum_{i=1}^n \ln(F_i)/\beta = \sum_{i=1}^n \ln(G(C_i Z_i))$ is a sum of bounded i.i.d. RVs, so according to the weak law of large numbers, the probability that $S_n \geq (g - 2\delta/3)n$ goes to 1 as $n \rightarrow \infty$. The net log-prob is given by $S_n - \ln(1/\epsilon)/\beta = S_n - \kappa(n)n/\beta$. Since $\kappa(n) = o(1)$, for n sufficiently large, $\kappa(n)/\beta \leq \delta/3$. Thus with asymptotic probability 1, the net log-prob after n trials exceeds $(g - \delta)n$. ■

We claim that the asymptotic gain rate g at ρ is equal to the minimum of the conditional entropies $H(D|ZE; \nu)$ over all distributions ν of CZE such that $\nu[CZ] = \rho$, where the conditional entropy base e is defined by

$$\begin{aligned}H(D|ZE; \nu) &= \sum_{ze} H(\nu(D|ze)) \nu(ze) \\ &= \mathbb{E}_\nu(-\ln(\nu(D|ZE))).\end{aligned}\quad (113)$$

According to the asymptotic equipartition property (AEP) [31] specialized to the case of finite classical-classical states, the infimum of the conditional entropies is the optimal randomness-generation rate in the asymptotic limit. In this sense, our method is optimal. This also implies that the supremum of the log-prob rates at ρ achievable by soft PEFs with any power is the same as g , the supremum of the log-prob rates at ρ achievable by PEFs with any power. The claim is established by the next theorem.

Theorem 43. Let g be the asymptotic gain rate at ρ of $\mathcal{C}$. For all distributions ν of CZE with $\nu[CZ] = \rho[CZ]$, we have $g \leq H(D|ZE; \nu)$. For $|\text{Rng}(E)|$ sufficiently large, there is a distribution ν of CZE with $\nu[CZ] = \rho[CZ]$ such that $g = H(D|ZE; \nu)$.

Proof. Suppose that ν is a distribution of CZE such that $\nu[CZ] = \rho$. Define $\rho_e = \nu[CZ|E = e] \in \mathcal{C}$ and $\lambda_e = \nu(e)$. Then, $\rho = \sum_e \lambda_e \rho_e$. Consider an arbitrary entropy estimator K for $\mathcal{C}$ and write $f(\sigma) = \mathbb{E}_\sigma(K(CZ))$ for its entropy estimate at σ . By definition, $f(\sigma) \leq H(D|Z; \sigma)$ for all $\sigma \in \mathcal{C}$. Since $f(\sigma)$ is linear in σ , we have

$$f(\rho) = \sum_e \lambda_e f(\rho_e) \leq \sum_e \lambda_e H(D|Z; \rho_e) = H(D|ZE; \nu). \quad (114)$$

By arbitrariness of the entropy estimator, the supremum of the entropy estimates at ρ is bounded by $H(D|ZE; \nu)$, so by lemma 40, we have $g \leq H(D|ZE; \nu)$.

For the second part of the theorem, since the asymptotic gain rate of $\mathcal{C}$ is the same as that of $\overline{\text{Conv}}(\mathcal{C})$, without loss of generality assume that $\mathcal{C}$ is convex closed. For any distribution $\sigma \in \mathcal{C}$ define

$$h_{\min}(\sigma) = \inf \left\{ \sum_e \lambda_e H(D|Z; \sigma_e) : \sigma_e \in \mathcal{C}, \lambda_e \geq 0, \sum_e \lambda_e = 1, \sum_e \lambda_e \sigma_e = \sigma \right\}. \quad (115)$$

We claim that $h_{\min}(\rho)$ is the supremum of entropy estimates at ρ for $\mathcal{C}$ and that the infimum in the definition of $h_{\min}(\rho)$ is achieved by a sum involving a bounded number of terms. The conditional entropy is concave in the joint distribution of its variables (see, for example, Ref. [46], Cor. 11.13, which is readily specialized to the classical case). It follows that if one of the σ_e contributing to the sum defining $h_{\min}(\sigma)$ is not extremal, we can replace it by a convex combination of extremal distributions to decrease the value of the sum. Thus, we only have to consider $\sigma_e \in \text{Extr}(\mathcal{C})$ for defining $h_{\min}$. It follows that $h_{\min}(\sigma)$ is the convex roof extension of the function $\sigma \mapsto H(D|Z; \sigma)$ on $\text{Extr}(\mathcal{C})$. Convex roof extensions are defined in Ref. [47]. See this reference for a proof that $h_{\min}$ is convex. In fact, the graph of $h_{\min}$ on $\mathcal{C}$ is the lower boundary of the convex closure of the set $\{(\sigma, H(D|Z; \sigma)) : \sigma \in \text{Extr}(\mathcal{C})\}$. Specializing $h_{\min}(\sigma)$ to the case $\sigma = \rho$, since the dimension is finite we can apply Carathéodory's theorem and express $h_{\min}(\rho)$ as finite convex combinations $\sum_e \lambda_e H(D|Z; \rho_e)$ with $\rho_e \in \text{Extr}(\mathcal{C})$. The number of terms required is at most $d + 2$, where d is the dimension of $\mathcal{C}$. Since $h_{\min}$ is convex and has a closed epigraph, for any $\epsilon > 0$ there exists an affine function f on $\mathcal{C}$ such that $f(\rho) \geq h_{\min}(\rho) - \epsilon$ and the graph of f is below that of $h_{\min}$ (see Ref. [48], Sect. 3.3.2 and Exercise 3.28). We can extend f to all distributions σ of CZ and express $f(\sigma)$ as an expectation $f(\sigma) = \mathbb{E}_\sigma(K(CZ))$ for some real-valued function K of CZ . Relevant existence and extension theorems can be found in textbooks on convex analysis, topological vector spaces or operator theory. For example, see Ref. [49], Ch. 1. We now have that for all $\sigma \in \mathcal{C}$,

$$\mathbb{E}_\sigma(K(CZ)) \leq h_{\min}(\sigma) \leq H(D|Z; \sigma) = \mathbb{E}_\sigma(-\ln(\sigma(D|Z))). \quad (116)$$

It follows that K is an entropy estimator and that its entropy estimate at ρ is $h_{\min}(\rho) - \epsilon$. By arbitrariness of ϵ , $h_{\min}(\rho)$ is the supremum of entropy estimates at ρ for $\mathcal{C}$, and hence $g = h_{\min}(\rho)$ by lemma 40 and theorem 41. From above, we can write $h_{\min}(\rho) = \sum_{e=1}^{d+2} \lambda_e H(D|Z; \rho_e)$, with $\lambda_e \geq 0$, $\sum_{e=1}^{d+2} \lambda_e = 1$ and $\rho_e \in \text{Extr}(\mathcal{C})$. We can set $\text{Rng}(E) = \{1, \dots, d+2\}$ and define v on CZE by $v(cze) = \rho_e(cz)\lambda_e$. Then $H(D|ZE; v) = \sum_e \lambda_e H(D|Z; \rho_e) = h_{\min}(\rho) = g$, which is what we aimed to prove. ■

VII. ADDITIONAL PEF CONSTRUCTIONS

Effectiveness of PEF constraints as discussed in Sec. V C provides a practical way to construct soft PEFs when the convex closure of the set $\mathcal{C}$ has a finite number of extreme points. We demonstrate this construction for the case of PEFs in Sec. IX. In addition, we can construct PEFs by the strategies discussed in the following two subsections.

A. PEFs from maximum probability estimators

A strategy for constructing PEFs is to determine functions F that solve the inequalities

$$\begin{aligned} F &\geq 0, \\ \mathbb{E}_v(F(CZ) \max_{dz} v(d|z)^\beta) &\leq 1 \quad \text{for } v \in \text{Extr}(\mathcal{C}). \end{aligned} \quad (117)$$

Any such F is a PEF with power β for $\mathcal{C}$. The next theorem provides families of PEFs satisfying these inequalities. Since

the expectations of such PEFs witness the maximum conditional probability $\max_{dz} v(d|z)$ in a trial, the corresponding probability estimators may be referred to as *maximum probability estimators*.

Theorem 44. Suppose that B is a function of CZ such that

$$1 - \mathbb{E}_v(B(CZ)) \geq \max_{dz} v(d|z) \quad \text{for all } v \in \mathcal{C}. \quad (118)$$

Let $\alpha < 1$, $\beta > 0$, $0 \leq \lambda \leq 1$ and define

$$F(CZ) = (1 - \alpha)^{-\beta} \left(1 + \beta \frac{B(CZ) - \alpha}{1 - \alpha} \right). \quad (119)$$

If $F' = 1 + \lambda(F - 1) \geq 0$, then F' is a PEF with power β for $\mathcal{C}$ satisfying Eq. (117).

A reasonable choice for α in the theorem is $\alpha = \bar{b} = \mathbb{E}_\rho(B(CZ))$ with ρ our best guess for the true distribution of the trial. The inequality in Eq. (118) and the expression for F suggests that we should maximize $\bar{b}$ for the best results at ρ . To optimize the log-prob rate or the net log-prob rate, we can then vary β and α . If the condition $F' \geq 0$ is not satisfied, we can either reduce λ in the definition or replace B by γB and α by $\gamma\alpha$ for an appropriate $\gamma \in (0, 1)$. The latter replacement preserves the validity of Eq. (118). Theorem 49 shows that PEFs obtained by these methods for a given B witness an asymptotic gain rate of at least $-\ln(1 - \bar{b})$, which justifies the goal of maximizing $\bar{b}$ and is what we would hope for given the interpretation of the right-hand-side of Eq. (118) as a worst-case conditional max-prob.

Functions B satisfying the conditions in the theorem with $\bar{b} > 0$ for a given non-LR distribution are readily constructed for a large class of Bell-test configurations. See the discussions after the proofs of this and the next theorem. The family of PEFs constructed accordingly contains PEFs with good log-prob rates, as witnessed by theorem 49 below, which quantifies the performance of B in terms of $\bar{b}$. The family can also be used as a tool for proving exponential randomness expansion, see theorem 52.

Proof. We use the following general inequality: for $x < 1$,

$$\begin{aligned} (1 - x)^{-\beta} &= (1 - \alpha - (x - \alpha))^{-\beta} \\ &= (1 - \alpha)^{-\beta} \left(1 - \frac{x - \alpha}{1 - \alpha} \right)^{-\beta} \\ &\geq (1 - \alpha)^{-\beta} \left(1 + \beta \frac{x - \alpha}{1 - \alpha} \right), \end{aligned} \quad (120)$$

since the right-hand-side defines the tangent line of the graph of the convex function $(1 - x)^{-\beta}$ at $x = \alpha$. Provided $F' \geq 0$, given any $v \in \mathcal{C}$, we can compute

$$\begin{aligned} \mathbb{E}_v(F'(CZ) \max_{dz} v(d|z)^\beta) &\leq \mathbb{E}_v(F'(CZ)) \max_{dz} v(d|z)^\beta \\ &= \mathbb{E}_v(1 + \lambda(F(CZ) - 1)) \max_{dz} v(d|z)^\beta \\ &= (1 - \lambda) \max_{dz} v(d|z)^\beta \\ &\quad + \lambda \mathbb{E}_v(F(CZ)) \max_{dz} v(d|z)^\beta \\ &\leq (1 - \lambda) + \lambda \mathbb{E}_v(F(CZ)) \max_{dz} v(d|z)^\beta. \end{aligned} \quad (121)$$

To continue, with Eq. (120) we bound

$$\begin{aligned}\mathbb{E}_v(F(CZ)) &= \mathbb{E}_v\left((1-\alpha)^{-\beta}\left(1+\beta\frac{B(CZ)-\alpha}{1-\alpha}\right)\right) \\ &= \left[(1-\alpha)^{-\beta}\left(1+\beta\frac{\mathbb{E}_v(B(CZ))-\alpha}{1-\alpha}\right)\right] \\ &\leq [1-\mathbb{E}_v(B(CZ))]^{-\beta}.\end{aligned}\quad (122)$$

To apply Eq. (120), note that Eq. (118) implies $\mathbb{E}_v(B(CZ)) < 1$. Substituting in Eq. (121) and applying Eq. (118) gives

$$\begin{aligned}\mathbb{E}_v(F'(CZ)v(D|Z)^\beta) \\ \leq (1-\lambda) + \lambda[1-\mathbb{E}_v(B(CZ))]^{-\beta} \max_{dz} v(d|z)^\beta \leq 1.\end{aligned}\quad (123)$$

Since v is an arbitrary distribution in $\mathcal{C}$, we conclude that if $F' \geq 0$, it is a PEF with power β and, in consideration of the first line of Eq. (121), it satisfies Eq. (117). ■

For (2,2,2) Bell tests with known settings probabilities, starting from a Bell function B_0 with positive expectation at ρ , one can determine $m > 0$ such that $B = B_0/2m$ satisfies Eq. (118). Computationally, m can be found by checking the constraints of Eq. (118) at extremal $v[C|Z]$. This construction was exploited in Ref. [13]. See the lemma in the proof of the “Entropy production theorem” in this reference, where setting $B = (T-1)/2m$ with respect to the notation there defines a function satisfying Eq. (118). This observation about nontrivial Bell functions is a consequence of the fact that positive expectations of such Bell functions witness the presence of a Popescu-Rohrlich (PR) box in the distribution, and such a box has maximum outcome probability $1/2$ for each setting z . (See Sec. IX for the definition of PR boxes.)

For $\lambda = 1$, theorem 44 can be generalized to a weighted form with minor modifications to the proof.

Theorem 45. Let γ be a positive function of DZ . Suppose that B is a function of CZ such that

$$1 - \mathbb{E}_v(B(CZ)) \geq \max_{dz} \gamma(dz)v(d|z) \text{ for all } v \in \mathcal{C}. \quad (124)$$

Let $\alpha < 1$, $\beta > 0$ and define

$$F(CZ) = \gamma(DZ)^\beta (1-\alpha)^{-\beta} \left(1 + \beta \frac{B(CZ)-\alpha}{1-\alpha}\right). \quad (125)$$

If $F \geq 0$, then F is a PEF with power β for $\mathcal{C}$.

Proof. It suffices to adjust the sequence of inequalities leading to Eq. (123) as follows

$$\begin{aligned}\mathbb{E}_v(F(CZ)v(D|Z)^\beta) \\ &= \mathbb{E}_v\left((1-\alpha)^{-\beta}\left(1+\beta\frac{B(CZ)-\alpha}{1-\alpha}\right)(\gamma(DZ)v(D|Z))^\beta\right) \\ &\leq \mathbb{E}_v\left((1-\alpha)^{-\beta}\left(1+\beta\frac{B(CZ)-\alpha}{1-\alpha}\right)\right) \max_{dz} (\gamma(dz)v(d|z))^\beta \\ &= \left[(1-\alpha)^{-\beta}\left(1+\beta\frac{\mathbb{E}_v(B(CZ))-\alpha}{1-\alpha}\right)\right] \max_{dz} (\gamma(dz)v(d|z))^\beta \\ &\leq [1-\mathbb{E}_v(B(CZ))]^{-\beta} (\max_{dz} \gamma(dz)v(d|z))^\beta \\ &\leq 1.\end{aligned}\quad (126)$$

Equation (124) defines a set of convex constraints on B . Namely,

$$1 - \mathbb{E}_v(B(CZ)) \geq \gamma(dz)v(d|z) \text{ for all } dz \text{ and } v \in \mathcal{C}. \quad (127)$$

If we use B to construct PEFs according to theorem 45, a reasonable goal is to maximize $\bar{b} = \mathbb{E}_\rho(B(CZ))$ subject to these constraints, where ρ is an estimate of the true distribution. We remark that the asymptotic gain rate witnessed by PEFs constructed according to theorem 45 from a given B is at least $-\ln(1-\bar{b}) + \mathbb{E}_\rho \ln(\gamma(DZ))$, see theorem 49.

To maximize $\bar{b}$, we can define $\tilde{B} = 1 - B$ and solve the following problem:

$$\begin{aligned}\text{Minimize: } & \mathbb{E}_\rho(\tilde{B}(CZ)) \\ \text{Subject to: } & \mathbb{E}_v(\tilde{B}(CZ)) \geq \gamma(dz)v(d|z) \text{ for all } dz \text{ and } v \in \mathcal{C}.\end{aligned}\quad (128)$$

Given an optimal B , one can use Eq. (125) to define PEFs, choosing parameters to optimize the log-prob rate at ρ . If F as constructed does not satisfy $F \geq 0$, we can replace B by $\gamma'B$ and α by $\gamma'\alpha$ for an appropriate $\gamma' \in (0, 1)$. Alternatively, we can reduce the power β .

For general $\mathcal{C}$, solving Eq. (128) may be difficult. But suppose that the measurement settings distribution $v[Z]$ is fixed, $v[Z] = \rho[Z]$ for all $v \in \mathcal{C}$, and the conditional distributions $v[C|Z]$ belong to a convex set $\mathcal{C}_{C|Z}$ determined by semidefinite constraints. Now $\mathcal{C} = \{v[C|Z]\rho(Z) : v[C|Z] \in \mathcal{C}_{C|Z}\}$, which is a special case of the sets free for Z defined in Sec. VIII A. This is the standard situation for Bell configurations, in which case Eq. (128) is related to the optimization problems described in Refs. [18–20]. These references define convex programs that determine the maximum available min-entropy for one trial with distribution ρ . In fact, the program given in Eq. (8) of Ref. [19] is related to the dual of Eq. (128) when $\mathcal{C}_{C|Z}$ is the set of quantum realizable conditional probability distributions. To make the relationship explicit and show that with the given assumptions, Eq. (128) is effectively solvable, define $\hat{B}$ by $\hat{B}(cz) = \tilde{B}(cz)\rho(z)$. With explicit sums, Eq. (128) becomes

$$\begin{aligned}\text{Minimize: } & \sum_{c'z'} \hat{B}(c'z')\rho(c'|z') \\ \text{Subject to: } & \sum_{c'z'} \hat{B}(c'z')v(c'|z') \geq \gamma(dz)v(d|z) \\ & \text{for all } dz \text{ and } v[C|Z] \in \mathcal{C}_{C|Z}.\end{aligned}\quad (129)$$

Note that $v[d|z]$ is a linear function of $v[C|Z]$. Let $\hat{B}_{\min}$ minimize the objective function of Eq. (129) with minimum value $1 - \bar{b}$. Such a solution exists since $\hat{B}(cz) = \max(\gamma)\rho(z)$ is feasible, which implies that the minimum value of the objective function is less than or equal to $\max(\gamma)$. Moreover, the constraints imply that the minimum is positive. By construction, $(\hat{B}_{\min}, -1)$ is in the cone $\mathcal{D}^*$ dual to the closed convex cone $\mathcal{D}$ generated by $\mathcal{D}_1 = \{(v[C|Z], \gamma(dz)v(d|z)) : v[C|Z] \in \mathcal{C}_{C|Z}, dz \in \text{Rng}(DZ)\}$. Further, $\vec{x} = (\rho(C|Z), (1 - \bar{b}))$ defines a supporting hyperplane of $\mathcal{D}^*$ at $(\hat{B}_{\min}, -1)$. Since $(\hat{B}_{\min}, -1) \cdot (\rho(C|Z), (1 - \bar{b})) = 0$, it suffices to check that $\vec{x}$ is in the dual of $\mathcal{D}^*$. Nonzero elements of $\mathcal{D}^*$ are positive linear combinations of elements of the form $(\hat{B}, -1)$ for some $\hat{B}$ and (σ, a) for some σ and $a \geq 0$. The first satisfies that $\hat{B}$ is a feasible solution of Eq. (129) and by definition of $\bar{b}$, $(\hat{B}, -1) \cdot \vec{x} = \sum_{c'z'} \hat{B}(c'z')\rho(c'|z') - (1 - \bar{b}) \geq 0$. For the second form, we have $(\sigma, a) \cdot \vec{x} = \sum_{c'z'} \sigma(c'z')\rho(c'|z') + a(1 - \bar{b})$. For all dz , $\sum_{c'z'} \sigma(c'z')\rho(c'|z') \geq -a\gamma(dz)\rho(d|z)$,

and since $(\hat{B}_{\min}, -1) \in \mathcal{D}^*$, $(1 - \bar{b}) \geq \gamma(dz)\rho(d|z)$. Hence $(\sigma, a) \cdot \bar{x} \geq -a\gamma(dz)\rho(d|z) + a\gamma(dz)\rho(d|z) = 0$, in consideration of $a \geq 0$. The dual of $\mathcal{D}^*$ is $\mathcal{D}$ again, so $\bar{x} \in \mathcal{D}$. Because of probability normalization, $\mathcal{D}_1$ is contained in several hyperplanes not containing the origin, so $\mathcal{D}$ is pointed. The intersection of any of these hyperplanes with $\mathcal{D}$ is $\mathcal{D}_1$. Since the set $\mathcal{D}_1$ is closed and bounded, it meets every extremal ray of $\mathcal{D}$ where it intersects these hyperplanes (theorem 1.4.5 of Ref. [49]). Consequently, $(\rho[C|Z], (1 - \bar{b}))$ is a convex combination of elements of $\mathcal{D}_1$. Since we are in finite dimensions we can apply Carathéodory's theorem and express $\rho[C|Z]$ as a finite convex combination $\sum_{dzk} \lambda_{dzk} v_{dzk}(c'|z') = \rho(c'|z')$ satisfying $\sum_{dzk} \lambda_{dzk} \gamma(dz) v_{dzk}(d|z) = (1 - \bar{b})$ with $v_{dzk} \in \mathcal{C}_{C|Z}$, $\sum_{dzk} \lambda_{dzk} = 1$ and $\lambda_{dzk} \geq 0$. If we define $\lambda_{dz} = \sum_k \lambda_{dzk}$ and $v_{dz} = \sum_k \lambda_{dzk} v_{dzk} / \lambda_{dz}$ if $\lambda_{dz} > 0$ (otherwise, v_{dz} can be any member of the set $\mathcal{C}_{C|Z}$), then k can be eliminated in the convex combination. By construction, $(1 - \bar{b})$ is the maximum value of $\sum_{dz} \lambda_{dz} \gamma(dz) v_{dz}(d|z)$ for any family $v_{dz} \in \mathcal{C}_{C|Z}$ and $\lambda_{dz} \geq 0$ satisfying that for all $c'z'$, $\sum_{dz} \lambda_{dz} v_{dz}(c'|z') = \rho(c'|z')$ and $\sum_{dz} \lambda_{dz} = 1$. We can write $\tilde{v}_{dz} = \lambda_{dz} v_{dz}$ to absorb the coefficients λ_{dz} . With this, the value of the following problem is $(1 - \bar{b})$:

$$\begin{aligned} \text{Maximize: } & \sum_{dz} \gamma(dz) \tilde{v}_{dz}(d|z) \\ \text{Subject to: } & \sum_{dz} \tilde{v}_{dz}(c'|z') = \rho(c'|z') \quad \text{for all } c'z', \\ & \tilde{v}_{dz} \in [0, \infty) \mathcal{C}_{C|Z} \quad \text{for all } dz. \end{aligned} \quad (130)$$

The set $[0, \infty) \mathcal{C}_{C|Z}$ is the cone generated by $\mathcal{C}_{C|Z}$. If $\mathcal{C}_{C|Z}$ is characterized by a semidefinite program, then so is $[0, \infty) \mathcal{C}_{C|Z}$ (by eliminating inhomogenous constraints in the semidefinite program's standard form, see Ref. [48], Eq. (4.51)). Equation (130) can then be cast as a semidefinite program also and solved effectively. By semidefinite-programming duality, this can then be used to obtain effective solutions of Eq. (128), provided the semidefinite program is formulated to satisfy strong duality. For related programs, Ref. [18] claim strong duality by exhibiting strictly feasible solutions.

We conclude with remarks on the relationship between the optimization problem in Eq. (130) and that in Eq. (8) of Ref. [19] (referenced as “P8” below). To relate Eq. (130) to P8, set $\gamma(dz) = 1$ and let $\mathcal{C}_{C|Z}$ be the set of quantum achievable conditional probability distributions, for which there is a hierarchy of semidefinite-programming relaxations [34]. Then identify both c and d here with ab there (so $c = d$), z here with xy there, and $\tilde{v}_{dz}(c'|z')$ here with $\rho(z) P_{cz}(c'|z')$ there, where $P_{cz}(c'|z') = \sum_{\alpha\beta: \alpha_z \beta_z = c} P_{\alpha\beta}(c'|z')$. The objective function of Eq. (130) now matches that of P8. However, unless there is only one setting, the equality constraints in Eq. (130) are a proper subset of those of P8. Observe that the equality constraints of P8 when expressed in terms of the variables $P_{cz}(c'|z')$ require that for each z and $c'z'$, $\sum_c P_{cz}(c'|z') = \rho(c'|z')$. For this, note that P8 includes the constraints $\sum_{\alpha\beta} P_{\alpha\beta}(c'|z') = \rho(c'|z')$ for all $c'z'$. For any z , the left-hand side can be written as $\sum_c \sum_{\alpha\beta: \alpha_z \beta_z = c} P_{\alpha\beta}(c'|z') = \sum_c P_{cz}(c'|z')$. These identities imply that $\sum_{cz} \rho(z) P_{cz}(c'|z') = \sum_z \rho(z) \rho(c'|z') = \rho(c'|z')$, but are stronger when $|\text{Rng}(Z)| > 1$. Furthermore, according to P8, the P_{cz} must be expressed as sums of the $P_{\alpha\beta} \in \mathcal{C}_{C|Z}$ as specified above, and this implies additional constraints. As a result the optimal value for P8 is in general smaller.

B. Convex combination

Optimization over all PEFs with a given power can be highly demanding, because both the size of the range of CZ and the number of extreme points of $\mathcal{C}$ in a general (k, l, m) Bell-test configuration are large. Furthermore, the size of the range of CZ determines the dimension of the search space, so if this size is large and the amount of data available for making an estimate of the true distribution is limited, there is a risk of overfitting when optimizing the log-prob rate for the estimated distribution at a given power β . See Eq. (162) of Sec. IX for the explicit formulation of the optimization problem. However, in many cases we have a small set of candidate PEFs expected to be helpful in a given situation, where the set was obtained in earlier studies of the experimental devices before running any protocols. Then, the optimization problem can be greatly simplified by means of the next theorem.

Theorem 46. Let $(F_i)_{i=1}^r$ be a collection of PEFs with power β and define $F_0 = 1$. Then every weighted average $F = \sum_{i=0}^r \lambda_i F_i$ with $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$ is a PEF with power β .

Proof. It suffices to observe that $F_0 = 1$ is a PEF with an arbitrary power, and the set of PEFs with power β is convex closed since it is defined by a family of linear constraints. ■

Optimizing PEFs with respect to the coefficients of weighted averages is efficient and less susceptible to overfitting issues. If ρ is estimated by empirical frequencies from past trials, one can evaluate the objective function directly on the past data, in which case the technique is not limited to configurations with computationally manageable $|\text{Rng}(CZ)|$. This strategy was proposed for optimizing test factors for rejecting LR in Ref. [15] and used in Refs. [39,43].

VIII. REDUCING SETTINGS ENTROPY

A. Sets of distributions that are free for Z

In our applications, the sets of distributions $\mathcal{C}$ are determined by constraints on the Z -conditional distributions of C and separate constraints on the marginal distributions of Z . We first formalize this class of sets. As defined in Sec. III, for any RV X , $\mathcal{S}_X = \{\rho : \rho \text{ is a distribution of } X\}$ and $\overline{\text{Cvx}}(\mathcal{X})$ is the convex closure of $\mathcal{X}$.

Definition 47. Let $\mathcal{S}_{C|Z} = \{(\rho_{C|z})_{z \in \text{Rng}(Z)} : \text{for each } z, \rho_{C|z} \text{ is a distribution of } C\}$. The set of distributions $\mathcal{C}$ of CZ is *free* for Z if there are closed convex sets $\mathcal{C}_{C|Z} \subseteq \mathcal{S}_{C|Z}$ and $\mathcal{C}_Z \subseteq \mathcal{S}_Z$ such that

$$\mathcal{C} = \overline{\text{Cvx}}(\{\nu \in \mathcal{S}_{CZ} : (\nu[C|Z] = z)_z \in \mathcal{C}_{C|Z} \text{ and } \nu[Z] \in \mathcal{C}_Z\}). \quad (131)$$

In the case of Bell tests, $\mathcal{C}_{C|Z}$ consists of the set of nonsignaling distributions, possibly satisfying additional quantum constraints. If Z is a settings choice RV with known distribution ν , then $\mathcal{C}_Z = \{\nu\}$. For $\rho_{C|Z} = (\rho_{C|z})_z \in \mathcal{S}_{C|Z}$, we define $\rho_{C|Z} \rtimes \nu$ by $(\rho_{C|Z} \rtimes \nu)(cz) = \rho_{C|z}(c) \nu(z)$. With this, $\mathcal{C}$ in Eq. (131) can be written as $\overline{\text{Cvx}}(\mathcal{C}_{C|Z} \rtimes \mathcal{C}_Z)$. If $\mathcal{C}_Z = \{\nu\}$, then $\mathcal{C}_{C|Z} \rtimes \mathcal{C}_Z = \mathcal{C}_{C|Z} \rtimes \nu$ is already convex closed, so we can omit the convex-closure operation.

The extreme points of $\mathcal{C}_{C|Z} \rtimes \nu$ consist of the set of $\rho_{C|Z} \rtimes \nu$ with $\rho_{C|Z}$ extremal in $\mathcal{C}_{C|Z}$. In general, we have the following.

Lemma 48. $\text{Extr}(\overline{\text{Cvx}}(\mathcal{C}_{C|Z} \rtimes \mathcal{C}_Z)) \subseteq \overline{\text{Extr}(\mathcal{C}_{C|Z})} \rtimes \overline{\text{Extr}(\mathcal{C}_Z)}$.

Here, the overline on the right-hand side denotes topological closure.

Proof. $\mathcal{C}_{C|Z}$ and $\mathcal{C}_Z$ are bounded closed sets with finite dimension, hence compact. The operation $\bowtie$ is continuous and therefore maps bounded closed sets to bounded closed sets. So $\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z$ is bounded and closed, as is $\overline{\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)}$, which is contained in $\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z$. From the last property, $\overline{\text{Cvx}(\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z)} \supseteq \overline{\text{Cvx}(\overline{\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)})}$.

By bilinearity of $\bowtie$, we have $\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z \subseteq \overline{\text{Cvx}(\overline{\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)})}$. Since convex closure is idempotent, $\overline{\text{Cvx}(\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z)} \subseteq \overline{\text{Cvx}(\overline{\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)})}$. We conclude that these two convex sets are identical. Every bounded closed set contains the extreme points of its convex closure (theorem 1.4.5 of Ref. [49]). Accordingly, $\text{Extr}(\overline{\text{Cvx}(\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z)}) \subseteq \overline{\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)}$. ■

We note that the members of $\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)$ are extremal. This can be seen as follows: Let the member $\rho_{C|Z} \bowtie v$ of $\text{Extr}(\mathcal{C}_{C|Z}) \bowtie \text{Extr}(\mathcal{C}_Z)$ be a finite convex combination of members of $\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z$, written as $\rho_{C|Z} \bowtie v = \sum_i \lambda_i \rho_{C|Z,i} \bowtie v_i$. By marginalizing to Z and by extremality of v , it follows that $v_i = v$ for all i . Given this, and extremality of $\rho_{C|Z}$, we can also conclude that $\rho_{C|Z,i} = \rho_{C|Z}$. Thus the inclusion in lemma 48 becomes an equality if we topologically close the left-hand side. In particular, if $\mathcal{C}_{C|Z}$ and $\mathcal{C}_Z$ are convex polytopes, then the right-hand side is finite and therefore $\overline{\text{Cvx}(\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z)}$ is also a convex polytope with the expected extreme points, a fact that we exploit for finding soft PEFs by our methods.

We fix the convex set $\mathcal{C}_{C|Z} \subseteq \mathcal{S}_{C|Z}$ for the remainder of this section. For brevity, instead of referring to properties “for $\mathcal{C} = \overline{\text{Cvx}(\mathcal{C}_{C|Z} \bowtie \mathcal{C}_Z)}$,” we just say “for $\mathcal{C}_Z$,” provided $\mathcal{C}_{C|Z}$ is clear from context.

B. Gain rates for biased settings

For randomness expansion, it is desirable to minimize the entropy of $\mu[Z]$, since this is the main contributor to the number of random bits required as input to the protocol. The other contributor is the seed required for the extractor, but good extractors already use relatively little seed. We expect that reducing the entropy of $\mu[Z]$ may reduce the asymptotic gain rate. However, in the case where we have uniform bounds on the maximum probability as required for theorem 45, we can show that the reduction is limited.

Theorem 49. Let $\rho_{C|Z} \in \mathcal{C}_{C|Z}$. Let γ be a positive function of DZ . Let B be a function of CZ satisfying Eq. (124) for $\mathcal{C}_Z = \{\text{Unif}_Z\}$. Assume that $\bar{b} = \mathbb{E}_{\rho_{C|Z} \bowtie \text{Unif}_Z}(B) > 0$. Let v be a distribution of Z with $p_{\min} = \min_z v(z) > 0$. Then the asymptotic gain rate at $\rho_{C|Z} \bowtie v$ for $\mathcal{C}_Z = \{v\}$ is at least $-\ln(1 - \bar{b}) + \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln(\gamma(DZ))$.

Proof. Let $q = 1/|\text{Rng}(Z)|$ and $w = -\min(0, \min(B))$. Define $B'(CZ) = B(CZ)\text{Unif}_Z(Z)/v(Z) = B(CZ)q/v(Z)$. For all $\sigma_{C|Z} \in \mathcal{S}_{C|Z}$,

$$\begin{aligned} \mathbb{E}_{\sigma_{C|Z} \bowtie v}(B') &= \sum_z \mathbb{E}_{\sigma_{C|Z}}(B'(Cz))v(z) \\ &= \sum_z \mathbb{E}_{\sigma_{C|Z}}(B(Cz)q/v(z))v(z) \\ &= \sum_z \mathbb{E}_{\sigma_{C|Z}}(B(Cz))q = \mathbb{E}_{\sigma_{C|Z} \bowtie \text{Unif}_Z}(B). \end{aligned} \quad (132)$$

By our assumptions on B , for all $\sigma_{C|Z} \in \mathcal{C}_{C|Z}$, $1 - \mathbb{E}_{\sigma_{C|Z} \bowtie \text{Unif}_Z} B(CZ) \geq \max_{dz} \gamma(dz) \mathbb{P}_{\sigma_{C|Z} \bowtie \text{Unif}_Z}(d|z)$, which is Eq. (124). Since $\mathbb{P}_{\sigma_{C|Z} \bowtie v}(d|z) = \mathbb{P}_{\sigma_{C|Z} \bowtie \text{Unif}_Z}(d|z)$, we can apply Eq. (132) to conclude that B' satisfies Eq. (124) for $\mathcal{C}_Z = \{v\}$. According to theorem 45,

$$F_\beta(CZ) = \gamma(DZ)^\beta (1 - \bar{b})^{-\beta} \left(1 + \beta \frac{B'(CZ) - \bar{b}}{1 - \bar{b}}\right) \quad (133)$$

is a PEF with power β for $\mathcal{C}_Z = \{v\}$ provided that $F_\beta \geq 0$. Note that Eq. (124) implies $\bar{b} < 1$. Since $B' \geq -wq/p_{\min}$, for sufficiently small β , the condition $F_\beta \geq 0$ is satisfied. Specifically, we require that $\beta < 1/a$, where we define $a = (wq/p_{\min} + \bar{b})/(1 - \bar{b})$. The asymptotic gain rate g from the theorem statement satisfies

$$g \geq \limsup_{\beta \rightarrow 0^+} \mathbb{E}_{\rho_{C|Z} \bowtie v}(\ln(F_\beta))/\beta. \quad (134)$$

The expression inside the limit is

$$\begin{aligned} \mathbb{E}_{\rho_{C|Z} \bowtie v}(\ln(F_\beta))/\beta &= -\ln(1 - \bar{b}) + \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln(\gamma(DZ)) \\ &\quad + \frac{1}{\beta} \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln \left(1 + \beta \frac{B'(CZ) - \bar{b}}{1 - \bar{b}}\right). \end{aligned} \quad (135)$$

In general, given $\beta a < 1$ and $x \geq -a$, we can approximate $\ln(1 + \beta x)$ from Taylor expansion around $x = 0$ with a second-order remainder to get

$$\ln(1 + \beta x) \geq \beta x - \frac{1}{2(1 - \beta a)^2} (\beta x)^2, \quad (136)$$

since $|\frac{d^2}{dy^2} \ln(1 + y)| \leq 1/(1 - \beta a)^2$ for $y \geq -\beta a$. Let $X = (B'(CZ) - \bar{b})/(1 - \bar{b})$. Then $\mathbb{E}_{\rho_{C|Z} \bowtie v} X = 0$. Let $v = \mathbb{E}_{\rho_{C|Z} \bowtie v}(X^2)$, which is the variance of X . Substituting X for x in the inequality of Eq. (136) and applying expectations to both sides gives

$$\begin{aligned} \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln(1 + \beta X) &\geq -\frac{\beta^2 v}{2(1 - \beta a)^2} \\ &= -O(\beta^2), \end{aligned} \quad (137)$$

where the minus sign on the order notation emphasizes that the lower bound is negative. The expression in Eq. (135) is therefore bounded below by $-\ln(1 - \bar{b}) + \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln(\gamma(DZ)) - O(\beta)$ which goes to $-\ln(1 - \bar{b}) + \mathbb{E}_{\rho_{C|Z} \bowtie v} \ln(\gamma(DZ))$ as $\beta \rightarrow 0^+$. In view of Eq. (134), the theorem follows. ■

C. Spot-checking settings distributions

To approach the asymptotic gain rate in theorem 49 requires small powers which negatively impact the net log-prob rate. We analyze the effect on the net log-prob rate in the case where μ is a mixture of a deterministic distribution δ_{z_0} and Unif_Z . This corresponds to using the same setting for most trials and randomly choosing *test* trials with uniform settings distribution. This is referred to as a *spot-checking* strategy [23]. Later we consider a protocol where one random trial out of a block of 2^k trials has Z uniformly distributed.

To simplify the analysis, we take advantage of the fact that for configurations such as those of Bell tests, we can hide the

choice of whether or not to apply a test trial from the devices. This corresponds to appending a test bit T to Z , where $T = 1$ indicates a test trial and $T = 0$ indicates a fixed one, with $Z = z_0$. The set $\mathcal{C}_{C|ZT}$ is obtained from $\mathcal{C}_{C|Z}$ by constraining $\mu[C|zt]$ to be independent of t for each z and $(\mu[C|zt])_z \in \mathcal{C}_{C|Z}$. For any $v_{C|Z} \in \mathcal{C}_{C|Z}$ there is a corresponding $\tilde{v}_{C|ZT} \in \mathcal{C}_{C|ZT}$ defined by $\tilde{v}_{C|ZT}(c) = v_{C|Z}(c)$ for all c, z and t . The map $v_{C|Z} \mapsto \tilde{v}_{C|ZT}$ is a bijection.

Let $q = 1/|\text{Rng}(Z)|$ and $\rho_{C|Z} \in \mathcal{C}_{C|Z}$. Let v_r be the probability distribution of ZT defined by $v_r(z1) = rq$ and $v_r(z0) = (1-r)\delta_{z,z_0}$ for some value z_0 of Z . Since we are analyzing the case where r is small, we assume $r < 1/2$. The entropy of the distribution v_r is given by $S(v_r) = H(r) + r \ln(1/q)$, where $H(r) = -r \ln(r) - (1-r) \ln(1-r)$. Let B be a function of CZ satisfying Eq. (118) for $\mathcal{C}_Z = \{\text{Unif}_Z\}$ with $\bar{b} = \mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z}(B) > 0$ and $\bar{v} = \mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z}((B - \bar{b})^2)$. From Eq. (118), it follows that $\bar{b} < 1$. Define $B_r(CZT)$ by

$$\begin{aligned} B_r(CZ0) &= 0, \\ B_r(CZ1) &= B(CZ) \frac{1}{r}. \end{aligned} \quad (138)$$

Let

$$F_{r,\beta}(CZT) = (1 - \bar{b})^{-\beta} \left(1 + \beta \frac{B_r(CZT) - \bar{b}}{1 - \bar{b}} \right). \quad (139)$$

Setting the function B_r to zero when $\{T = 0\}$ is convenient. Because $F_{r,\beta}$ is related to the tangent line of $(1 - x)^{-\beta}$ at $x = \bar{b}$ and $B_r = 0$ corresponds to $x = 0$, the expected value of $F_{r,\beta}$ is slightly below 1 for nontest trials.

Theorem 50. There exist constants $d > 0, d' \geq 0$ independent of r such that for $\beta \leq dr$, $F_{r,\beta}$ as defined in Eq. (139) is a PEF with power β for $\mathcal{C}_{C|ZT} \times v_r$, and its log-prob rate $g_{r,\beta}$ at $\tilde{\rho}_{C|ZT} \times v_r$ satisfies $g_{r,\beta} \geq -\ln(1 - \bar{b}) - d'\beta/r$.

In practical situations, we anticipate using numerical optimization to determine the PEFs, which we expect to improve on the bounds in the theorem. However, we believe that the constants d and d' obtained in the proof are reasonable. They are given by

$$\begin{aligned} d &= \frac{1 - \bar{b}}{2w + \bar{b}}, \\ d' &= 2 \frac{\bar{v} + \bar{b}^2}{(1 - \bar{b})^2}, \end{aligned} \quad (140)$$

with $w = -\min(0, \min(B))$.

Proof. The proof is a refinement of that of theorem 49. For any $\sigma_{C|Z} \in \mathcal{S}_{C|Z}$, consider the expectation of B_r with respect to $\tilde{\sigma}_{C|ZT} \times v_r$.

$$\begin{aligned} \mathbb{E}_{\tilde{\sigma}_{C|ZT} \times v_r}(B_r) &= \sum_z \mathbb{E}_{\tilde{\sigma}_{C|z0}} B_r(Cz0) v_r(z0) \\ &\quad + \sum_z \mathbb{E}_{\tilde{\sigma}_{C|z1}} B_r(Cz1) v_r(z1) \\ &= \sum_z \mathbb{E}_{\sigma_{C|z}} B_r(Cz1) qr \\ &= \sum_z \mathbb{E}_{\sigma_{C|z}} B(Cz) q \\ &= \mathbb{E}_{\sigma_{C|Z} \times \text{Unif}_Z} B(CZ). \end{aligned} \quad (141)$$

Since $\mathbb{P}_{\tilde{\sigma}_{C|ZT} \times v_r}(d|zt) = \mathbb{P}_{\sigma_{C|Z} \times \text{Unif}_Z}(d|z)$ and B satisfies Eq. (118) for $\mathcal{C}_Z = \{\text{Unif}_Z\}$, so does B_r but for $\mathcal{C}_{ZT} = \{v_r\}$. Thus theorem 44 applies, and we have that $F_{r,\beta}$ as defined in Eq. (139) is a PEF with power β for $\mathcal{C}_{C|ZT} \times v_r$, provided β is small enough. From Eq. (138) and the fact $\min(B) \geq -w$, we have that $B_r \geq -w/r$. From the proof of theorem 49, we can replace $wq/p_{\min}$ by w/r in the expression for a there to see that it suffices to satisfy $\beta < 1/a = (1 - \bar{b})/(w/r + \bar{b})$ in order to make sure $F_{r,\beta} \geq 0$. The upper bound can be estimated as

$$1/a = \frac{1 - \bar{b}}{w/r + \bar{b}} = r \frac{1 - \bar{b}}{w + \bar{b}r} > r \frac{1 - \bar{b}}{w + \bar{b}/2} \quad (142)$$

given our assumption that $r < 1/2$ and $\bar{b} > 0$. With foresight, we set $d = (1 - \bar{b})/(2w + \bar{b}) < 1/(2ra)$. With $\beta \leq dr$, this implies $F_{r,\beta} > 0$ and $\beta a < 1/2$.

Next we lower bound $\mathbb{E}_{\tilde{\rho}_{C|ZT} \times v_r}(\ln(F_{r,\beta}))/\beta$ by the same strategy that we used in the proof of theorem 49. The result is

$$g_{r,\beta} = \mathbb{E}_{\tilde{\rho}_{C|ZT} \times v_r}(\ln(F_{r,\beta}))/\beta \geq -\ln(1 - \bar{b}) - \frac{\beta v_r}{2(1 - \beta a)^2}, \quad (143)$$

where here we define $v_r = \mathbb{E}_{\tilde{\rho}_{C|ZT} \times v_r} X_r^2$ with $X_r = (B_r - \bar{b})/(1 - \bar{b})$. Note that from Eq. (141) and the definition of $\bar{b}$, $\mathbb{E}_{\tilde{\rho}_{C|ZT} \times v_r} X_r = 0$. An explicit expression for v_r is obtained as follows:

$$\begin{aligned} v_r &= \mathbb{E}_{\tilde{\rho}_{C|ZT} \times v_r} (X_r(CZT))^2 \\ &= \sum_{zt} \mathbb{E}_{\tilde{\rho}_{C|zt}} (X_r(Czt))^2 v_r(zt) \\ &= \sum_{zt} \mathbb{E}_{\rho_{C|z}} (X_r(Czt))^2 v_r(zt) \\ &= \sum_z \mathbb{E}_{\rho_{C|z}} (X_r(Cz0))^2 v_r(z0) + \sum_z \mathbb{E}_{\rho_{C|z}} (X_r(Cz1))^2 v_r(z1) \\ &= \mathbb{E}_{\rho_{C|z_0}} (X_r(Cz_00))^2 (1 - r) + \sum_z \mathbb{E}_{\rho_{C|z}} (X_r(Cz1))^2 qr \\ &= \frac{\bar{b}^2}{(1 - \bar{b})^2} (1 - r) + \mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z} (X_r(CZ1))^2 r. \end{aligned} \quad (144)$$

For the second term in the last line, compute

$$\begin{aligned} &\mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z} (X_r(CZ1))^2 \\ &= \frac{1}{(1 - \bar{b})^2} \mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z} ((B_r(CZ1) - \bar{b})^2) \\ &= \frac{1}{(1 - \bar{b})^2} \mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z} ((B(CZ)/r - \bar{b})^2) \\ &= \frac{1}{(1 - \bar{b})^2} [\mathbb{E}_{\rho_{C|Z} \times \text{Unif}_Z} ((B(CZ)/r - \bar{b}/r)^2) + (\bar{b}/r - \bar{b})^2] \\ &= \frac{1}{(1 - \bar{b})^2} (\bar{v}/r^2 + \bar{b}^2(1/r - 1)^2) \\ &= \frac{\bar{v} + \bar{b}^2(1 - r)^2}{r^2(1 - \bar{b})^2}, \end{aligned} \quad (145)$$

where for the fourth line we applied the identity $\mathbb{E}(U^2) = \mathbb{E}((U - \mathbb{E}(U))^2) + \mathbb{E}(U)^2$ relating the expectation of the

square to the variance. Combining this with Eq. (144) gives

$$\begin{aligned} v_r r &= \frac{\bar{v} + ((1-r)^2 + r(1-r))\bar{b}^2}{(1-\bar{b})^2} = \frac{\bar{v} + (1-r)\bar{b}^2}{(1-\bar{b})^2} \\ &\leq \frac{\bar{v} + \bar{b}^2}{(1-\bar{b})^2}. \end{aligned} \quad (146)$$

Write c for the right-hand side of this inequality, which is independent of r . Substituting in Eq. (143), we get

$$\begin{aligned} \mathbb{E}_{\tilde{\rho}_{C|ZT} \propto v_r}(\ln(F_{r,\beta}))/\beta &\geq -\ln(1-\bar{b}) - \frac{\beta c/r}{2(1-\beta a)^2} \\ &\geq -\ln(1-\bar{b}) - 2\beta c/r, \end{aligned} \quad (147)$$

because our earlier choice for d implies $\beta a < 1/2$. We now set $d' = 2c$ to complete the proof of the theorem. ■

From the proof above, we can extract the variance of $F_{r,\beta}$ at $\tilde{\rho}_{C|ZT} \propto v_r$, from which we can obtain a bound on the variance of $\ln(F_{r,\beta})$. Since we refer to it below when discussing the statistical performance of exponential randomness expansion, we record the result here.

Lemma 51. The variance $v_{r,\beta}$ of $\ln(F_{r,\beta})$ with $\beta \leq dr$ at $\tilde{\rho}_{C|ZT} \propto v_r$ satisfies

$$v_{r,\beta} \leq 4\ln(2)^2 \frac{\beta^2}{r} \frac{\bar{v} + \bar{b}^2}{(1-\bar{b})^2}. \quad (148)$$

Proof. The variance of $F_{r,\beta}$ at $\tilde{\rho}_{C|ZT} \propto v_r$ can be determined from v_r in the proof of theorem 50. It satisfies

$$\begin{aligned} \text{Var}(F_{r,\beta}) &= (1-\bar{b})^{-2\beta} \beta^2 v_r \\ &\leq \frac{\beta^2}{r} \frac{\bar{v} + \bar{b}^2}{(1-\bar{b})^{2(1+\beta)}}. \end{aligned} \quad (149)$$

By construction of $F_{r,\beta}$, the mean of $F_{r,\beta}$ is $(1-\bar{b})^{-\beta}$, and $F_{r,\beta} > (1-\bar{b})^{-\beta}/2$. The latter follows from $\beta a < 1/2$, noted where we defined d after Eq. (142). In general, for a positive RV U with mean $\bar{u}$ and $U > \bar{u}/2$,

$$\begin{aligned} \text{Var}(\ln(U)) &\leq \mathbb{E}((\ln(U) - \ln(\bar{u}))^2) \\ &= \mathbb{E}((\ln(U/\bar{u}))^2) \\ &= \mathbb{E}((\ln(1 + (U - \bar{u})/\bar{u}))^2) \\ &\leq 4\ln(2)^2 \mathbb{E}(((U - \bar{u})/\bar{u})^2) \\ &= 4\ln(2)^2 \text{Var}(U)/(\bar{u})^2. \end{aligned} \quad (150)$$

Here, the first step follows from $\text{Var}(Y) \leq \mathbb{E}((Y - y_0)^2)$ for any y_0 , where we substituted $\ln(U)$ and $\ln(\bar{u})$ for Y and y_0 . In the second-last step, we used $|\ln(1+y)| \leq 2\ln(2)|y|$ for $y \geq -1/2$, substituting $(U - \bar{u})/\bar{u}$ for y . Setting $U = F_{r,\beta}$ and $\bar{u} = \mathbb{E}(F_{r,\beta}) = (1-\bar{b})^{-\beta}$, we get the inequality in the lemma. ■

D. Exponential randomness expansion

For PEFs satisfying the conditions for theorem 50, it is possible to achieve exponential expansion of the entropy of Z .

Theorem 52. Let $F_{r,\beta}$ be the family of PEFs from theorem 50. Consider n trials with model $\mathcal{H}(C_{C|ZT} \propto v_r)$ and a fixed error bound ϵ_h . Consider $\rho_{C|Z} \in \mathcal{C}_{C|Z}$ and let $g_{r,\beta}$ be the log-prob rate of $F_{r,\beta}$ at $\tilde{\rho}_{C|ZT} \propto v_r$, and assume

that the trials are i.i.d. for CZT with this distribution. We can choose $r = r_n$ and $\beta = \beta_n$ as functions of n so that the expected net log-prob is $g_{\text{net}} = n g_{r_n, \beta_n} - \ln(1/\epsilon_h)/\beta_n \geq n(-\ln(1-\bar{b}))/3 = e^{\Omega(nS(v_{r_n}))}$, where $nS(v_{r_n})$ is the total input entropy of $\mathbf{ZT}$.

The performance in the theorem statement may be compared to that in Ref. [23], where the settings entropy for n trials is $\Omega(\ln^3(n))$ (see the end of Sect. 1.1 of the reference). With our terminology, this implies a net log-prob $e^{O(S^{1/3})}$ where S is the total input entropy, which is subexponential in input entropy. It is also possible to match the soundness performance of the “one-shot” protocol of Ref. [22], Cor. 1.2, by modifying the proof below with r_n and β_n proportional to $n^{-\omega}$ with $0 < \omega < 1$ instead of $\omega = 1$. However, the referenced results are valid for quantum side information.

The proof of theorem 52 given below computes explicit constants for the orders given in the theorem statement. The constants depend on the PEF construction in the previous section, and the performance can always be improved by direct PEF optimization. As a result we do not expect the constants to be needed in practice.

The i.i.d. assumption is a completeness requirement that we expect to hold approximately if the devices perform as designed. Explicit constants expressed in terms of those in theorem 50 are given with Eq. (153) below. Notably, the expected number of test trials is independent of n and proportional to $\ln(1/\epsilon_h)$, which is what one would hope for in the absence of special strategies allowing feed-forward of generated randomness from independent devices, an example of which is the cross-feeding protocol of Ref. [25]. However, since nontest trials contribute negatively to the log-prob, this means that the full gain and the error bound are witnessed by a constant number of test trials. One should therefore expect a large instance-to-instance variation in the actual log-prob obtained, raising the question of whether the success probability, or the expected number of trials before reaching a log-prob goal, are sufficiently well behaved. We consider this question after the proof.

Proof. Let d and d' be the constants in theorem 50. From this theorem, the expected net log-prob is bounded by

$$\begin{aligned} g_{\text{net}} &= n g_{r_n, \beta_n} - \frac{\ln(1/\epsilon_h)}{\beta_n} \\ &\geq n \left(-\ln(1-\bar{b}) - \frac{d'\beta_n}{r_n} \right) - \frac{\ln(1/\epsilon_h)}{\beta_n}, \end{aligned} \quad (151)$$

with $\bar{b}$ as defined for theorem 50. The input entropy per trial $S(v_r) = H(r) + r \ln(1/q)$ is bounded above by $-2r \ln(r)$, provided we take $r \leq q/e$. For this, note that $r \ln(1/q) \leq r \ln(1/(re)) = -r \ln(r) - r$ and $-(1-r) \ln(1-r) \leq r$ since the function $x \mapsto -(1-x) \ln(1-x)$ is concave for $0 \leq x < 1$, is equal to 0 at $x = 1$, and has slope 1 at $x = 0$. If we choose $r_n = c'/n$ with c' to be determined later, the expected number of test trials is c' and the total input entropy satisfies $nS(v_{r_n}) \leq 2c' \ln(n/c')$, or equivalently $n \geq c' e^{nS(v_{r_n})/(2c')}$. For sufficiently large n , $r \leq q/e$ is satisfied. If we then choose $\beta_n = cr_n = cc'/n$ with $c \leq d$ such that g_{net} grows linearly with n , we have accomplished our goal. Write $g_0 = -\ln(1-\bar{b})$. Substituting

the expressions for β_n and r_n in Eq. (151),

$$g_{\text{net}} \geq ng_0 \left(1 - \frac{d'c}{g_0} - \frac{\ln(1/\epsilon_h)}{cc'g_0} \right). \quad (152)$$

We first set $c = \min(d, g_0/(3d'))$, which ensures that $d'c/g_0 \leq 1/3$. We then set $c' = 3 \ln(1/\epsilon_h)/(cg_0)$. This gives the inequality

$$g_{\text{net}} \geq ng_0/3 \geq (c'g_0/3)e^{nS(v_{r_n})/(2c')}, \quad (153)$$

which implies the theorem. ■

According to lemma 51, the variance at $\tilde{\rho}_{C|ZT} \propto v_r$ of the quantity $n \ln(F_{r_n, \beta_n})/\beta$ whose expectation is the expected log-prob for the parameters chosen in the proof above is given by

$$\begin{aligned} v &= n \frac{v_{r_n, \beta_n}}{\beta_n^2} \leq 4 \ln(2)^2 \frac{n}{r_n} \frac{\bar{v} + \bar{b}^2}{(1 - \bar{b})^2} \\ &= 4 \ln(2)^2 \frac{n^2}{c'} \frac{\bar{v} + \bar{b}^2}{(1 - \bar{b})^2} \\ &= 2 \ln(2)^2 \frac{n^2 d'}{c'}. \end{aligned} \quad (154)$$

Here, to get the last line, we used Eq. (140). In view of Eq. (152), the expected log-prob is at least $2ng_0/3$, with up to $1/2$ of that taken away by the error-bound requirement. Here we are interested in the distribution of the net log-prob, which depends on the actual trial results. The amount that is subtracted from the log-prob to get the net log-prob for the error bound is independent of the trial results and given by $f = n \ln(1/\epsilon_h)/(cc')$. By increasing c' by a constant factor while holding c fixed, we can reduce f and still achieve exponential expansion, albeit with a smaller constant in the exponent. The first two terms on the right-hand side of Eq. (152) are independent of c' and provide a lower bound on the expected log-prob. These considerations imply that if the variance is small enough compared to $(ng_0/6)^2$ (for example), with correspondingly high probability the net log-prob exceeds $ng_0/6$. Small variance is implied by large c' , because according to Eq. (154), the variance of the log-prob is reduced if c' is increased. To be specific, let $\kappa \geq \ln(1/\epsilon_h)$. We can set

$$\begin{aligned} c' &= \frac{3\kappa}{cg_0} \\ &= \frac{3\kappa}{g_0 \min(d, g_0/(3d'))} \\ &= \frac{9d'\kappa}{g_0^2 \min(3dd'/g_0, 1)}, \end{aligned} \quad (155)$$

from which

$$v \leq 2 \ln(2)^2 \frac{n^2 g_0^2 \min(3dd'/g_0, 1)}{9\kappa} \leq 2 \ln(2)^2 \frac{n^2 g_0^2}{9\kappa}. \quad (156)$$

By Chebyshev's inequality, a conservative upper bound on the probability that $n \ln(F_{r_n, \beta_n})/\beta < 2ng_0/3 - \sqrt{v}/\lambda$ is λ^2 . Since at most $ng_0/3$ is subtracted to obtain the net log-prob, if we require that $\sqrt{v}/\lambda \leq ng_0/6$, then with probability at least $1 - \lambda^2$ we get a net log-prob of at least $ng_0/6$, which is sufficient for exponential expansion. From Eq. (156), we can achieve

this if we have

$$2 \ln(2)^2 \frac{n^2 g_0^2}{9\kappa \lambda^2} \leq \frac{n^2 g_0^2}{6^2}, \quad (157)$$

that is, if $\kappa \geq 8 \ln(2)^2/\lambda^2$. In terms of c' , this means that c' needs to be set to $\max(3 \ln(1/\epsilon_h), 24 \ln(2)^2/\lambda^2)/(cg_0)$ for a probability of at least $1 - \lambda^2$ of having a net log-prob of $ng_0/6$.

Better bounds on the probabilities above can be obtained by taking into consideration the i.i.d. assumptions. For example, one can apply Chernoff-Hoeffding bounds to improve the estimates. Alternatively, we can use the option to acquire trials until a desired net log-prob is obtained with a conservative upper bound on the maximum number of trials.

E. Blockwise spot checking

For randomness expansion without having to decompress uniform bits for generating biased distributions of Z , we suggest the following spot-checking strategy consisting of blocks of trials. Each block nominally contains 2^k trials. One of them is chosen uniformly at random as a test trial, for which we use the random variable T_s , now valued between 0 and $2^k - 1$ with a uniform distribution. The l th trial in a block uses $Z = z_0$ (distribution δ_{z_0}), unless $l = T_s + 1$ in which case the distribution of Z is Unif_Z . For analyzing these trials, observe that the distribution of Z at the l th trial conditional on not having encountered the test trial yet, that is, given that $\{l - 1 \leq T_s\}$, is a mixture of δ_{z_0} and Unif_Z , with the probability of the latter being $1/(2^k - l + 1)$. Thus we can process the trial using an adaptive strategy, designing the PEF for the l th trial according to theorem 50, where the distribution of T has probability $r_l = 1/(2^k - l + 1)$ of $T = 1$. We call this the *blockwise spot-checking* strategy. Here we must choose a single power for the PEFs independent of l . Once the test trial is encountered, the remaining trials in the block do not need to be performed as we can, equivalently, set the PEFs to be trivial from this point on. Of course, if there is verifiable randomness even when the setting is fixed and known by the external entity and devices, we can continue. The average number of trials per block is therefore $(2^k + 1)/2$. Note that $r_{2^k-l+1} = 1/l$ and the probability of reaching the $2^k - l + 1$ th trial is $l/2^k$. When we choose the common power of the PEFs, the goal is to optimize the expected log-prob per block in consideration of the way in which the settings probabilities change as well as the anticipated trial statistics. More details on randomness expansion with the blockwise spot-checking strategy are available in Ref. [50], where PEFs are constructed by direct optimization.

IX. APPLICATIONS TO BELL TESTS

A. General considerations

All our applications are for the (2,2,2) Bell-test configuration. In this configuration, $Z = XY$ and $D = C = AB$ with $X, Y, A, B \in \{0, 1\}$, so we can effectively construct the PEFs for a trial according to theorem 28. Therefore, in this section, we choose to illustrate the performance of PEFs as opposed to that of soft PEFs. The distributions are free for Z . The set of conditional distributions consists of those satisfying the

nonsignaling constraints Eq. (5). We denote this set by $\mathcal{N}_{C|Z}$. This includes distributions not achievable by two-party quantum states, so additional quantum constraints can be considered. There is a hierarchy of semidefinite programs to generate such constraints due to Navascués, Pironio, and Acín [34] (the NPA hierarchy). Because semidefinite programs define convex sets that are not convex polytopes and have a continuum of extreme points, we cannot make direct use of the NPA hierarchy. In the absence of a semidefinite hierarchy for the PEF constraints, we require the conditional distributions to be a convex polytope defined by extreme points. In general, we can use the NPA hierarchy to generate finitely many additional linear constraints to add to the nonsignaling constraints. The set of extreme points can then be obtained by available linear programming tools and used for PEF optimization. To keep the number of extreme points low, it helps to add only constraints relevant to the actual distribution expected at a trial. We do not have a systematic way to identify relevant constraints. Below, we consider additional quantum constraints based on Tsirelson's bound [51]. We note that one can consider constructing PEFs via theorem 44 as discussed after the theorem. With this method one can take advantage of the NPA hierarchy directly but is restricted to PEFs satisfying Eq. (117).

The extreme points of $\mathcal{N}_{C|Z}$ are given by the deterministic LR distributions and variations of Popescu-Rohrlich (PR) boxes [52]. The deterministic LR distributions are parameterized by two functions $f_A : x \mapsto f_A(x) \in \{0, 1\}$ and $f_B : y \mapsto f_B(y) \in \{0, 1\}$ defining each station's settings-dependent outcome, giving distributions

$$v_{f_A, f_B}(ab|xy) = \llbracket f_A(x) = a \rrbracket \llbracket f_B(y) = b \rrbracket. \quad (158)$$

This yields the 16 LR extreme points. The PR boxes have distributions v_g defined by functions $g : xy \mapsto \{0, 1\}$, where $|\{xy|g(xy) = 1\}|$ is odd, according to

$$v_g(ab|xy) = \llbracket a = b \oplus g(xy) \rrbracket / 2, \quad (159)$$

with addition of values in $\{0, 1\}$ modulo 2. The function g indicates for which settings pairs the stations' outcomes are perfectly anticorrelated. There are eight extreme points in this class.

The simplest quantum constraint to add is Tsirelson's bound [51], which can be expressed in the form $\mathbb{E}(B(CZ)) \leq (\sqrt{2} - 1)/4$ for all distributions compatible with quantum mechanics, with $B(CZ)$ given by Eq. (22) and fixing the settings distribution to be uniform. There are 8 variants of Tsirelson's bound expressed as

$$\mathbb{E}(B_g(CZ)) \leq (\sqrt{2} - 1)/4 \quad (160)$$

corresponding to the eight PR boxes labeled by g , with $B_g(CZ)$ as defined in the next paragraph. Let $\mathcal{Q}_{C|Z}$ be the set of conditional distributions obtained by adding the eight conditional forms of Eq. (160) to the constraints.

For any PR box v_g , let $p(g) = 0$ or $p(g) = 1$ according to whether $|g^{-1}(1)| = 1$ or $|g^{-1}(1)| = 3$ and define

$$B_g(xyab) = -(-1)^{g(xy)+p(g)} \llbracket a \neq b \oplus p(g) \rrbracket. \quad (161)$$

This implies that the Bell function $B(CZ)$ in Eq. (22) can be expressed as $B(CZ) = B_{g, xy \mapsto x \times y}(CZ)$, with $x \times y$ the numeric product of x and y . For any distribution $\rho \in \mathcal{N}_{C|Z} \rtimes \{\text{Unif}_Z\}$,

the expectation of $\mathbb{E}_\rho(B_g(CZ))$ bounds the contribution of the PR box v_g to ρ in the sense that if $\rho = (1 - p)\rho_{\text{LR}} + pv_g$, with ρ_{LR} being an LR distribution, then $\mathbb{E}_\rho(B_g(CZ)) \leq p/4$. From Ref. [53], one determines that the bound is tight in the sense that if $\mathbb{E}_\rho(B_g(CZ)) = p/4$, then there exists an LR distribution ρ_{LR} such that $\rho = (1 - p)\rho_{\text{LR}} + pv_g$. These LR distributions are convex combinations of the eight deterministic LR distributions for which $\mathbb{E}_{v_{f_A, f_B}}(B_g(CZ)) = 0$. We observe that these are extremal, and all nonsignaling distributions can be expressed as convex combinations of extreme points involving at most one PR box. This makes it possible to explicitly determine the extreme points of $\mathcal{Q}_{C|Z}$. They are the convex combinations $(1 - q)v_{f_A, f_B} + qv_g$ with $q = \sqrt{2} - 1$ that satisfy Tsirelson's bound with equality on the one-dimensional line connecting one of the above eight deterministic LR distributions to the corresponding PR box. Note that the above argument can be generalized to determine the extreme points of the convex polytopes obtained using other Bell inequalities, see Ref. [54].

We emphasize that probability estimation is not based on the Bell-inequality framework. However, any PEF F for a Bell-test configuration has the property that $F - 1$ is a Bell function and therefore associated with a Bell inequality. This follows because the conditional probabilities for the extremal LR distributions are either 0 or 1. Therefore, regardless of the power β , the PEF constraints from these extreme points are equivalent to the constraints for Bell functions after subtracting 1. One can construct PEFs from Bell functions by reversing this argument and rescaling the resulting factor to satisfy the constraints induced by non-LR extreme points. But the relationship between measures of quality of Bell functions (such as violation signal-to-noise or winning probability) and those of related PEFs is not clear.

B. Applications and methods

We illustrate probability estimation with three examples. In the first, we estimate the conditional probability distribution of the outcomes observed in the recent loophole-free Bell test with entangled atoms reported in Ref. [7]. In the second, we reanalyze the data used for extracting 256 random bits uniform within 0.001 in Ref. [13]. Finally, we reanalyze the data from the ion experiment of Ref. [10], which was the first to demonstrate min-entropy estimation from a Bell test. For the first example, we explore the log-prob rates and net log-prob rates over a range of choices of parameters for PEF power, error bound, unknown settings bias and test-trial probability. The log-prob rates and net log-prob rates are determined directly from the inferred probabilities of the outcomes, from which the expected net log-prob can be inferred under the assumption that the trials are i.i.d. For the other two examples, we explicitly apply the probability estimation schema and obtain the total log-prob for different scenarios. In each scenario, the parameter choices are made based on training data, following the protocol in Ref. [13]. We show the certifiable min-entropy for the reported data as a function of the error bound and compare to the min-entropies reported in Refs. [10, 13].

For determining the best PEFs, we can optimize the log-prob rate at the estimated distribution ν given the power β .

TABLE I. Inferred settings-conditional distribution $\rho_{\text{atoms}}(ab|xy)$. Shown are the settings-conditional probabilities of the outcomes estimated from all trials with the prepared state Ψ^- reported in Ref. [7]. This is the $\mathcal{Q}_{C|Z}$ -satisfying distribution based on the outcomes from 27683 heralded trials. Each column is for one combination of outcomes ab , and the rows correspond to the settings combination xy . Because the probabilities are settings-conditional, each row's probabilities add to 1, except for rounding error. We give the numbers at numerical precision for traceability. They are used for determining PEFs, not to make a statement about the actual distribution of the outcomes in the experiment. Statistical error does not affect the validity of the PEFs obtained [see the paragraph after Eq. (162)].

ab	00	10	01	11
xy				
00	0.114583230563265	0.408949785618886	0.369310344143205	0.107156639674644
10	0.399140705802719	0.124392310379432	0.111262723543957	0.365204260273892
01	0.102208313465938	0.403210760398438	0.381685261240533	0.112895664895092
11	0.127756153189431	0.377662920674945	0.382647276157245	0.111933649978380

No matter which of $\mathcal{N}_{C|Z}$ or $\mathcal{Q}_{C|Z}$ is considered, the free-for- Z sets $\mathcal{C}$ obtained accordingly have a finite number of extreme points. We can therefore take advantage of the results in Sec. VC to solve the optimization problem effectively. Given the list of extreme points $(\rho_i)_{i=1}^q$ of the set $\mathcal{C}$ and the estimated distribution ν in $\mathcal{C}$, the optimization problem can be stated as follows:

$$\begin{aligned} &\text{Maximize: } \sum_{cz} \ln(F(cz))\nu(cz) \\ &\text{Subject to: } F(cz) \geq 0, \quad \forall cz, \\ &\quad \sum_{cz} F(cz)\rho_i(c|z)^\beta \rho_i(cz) \leq 1, \quad \forall i. \end{aligned} \quad (162)$$

Since the objective function is concave with respect to F and the constraints are linear, the problem can be solved by any algorithm capable of optimizing nonlinear functions with linear constraints on the arguments. In our implementation, we use sequential quadratic programming. Due to numerical imprecision, it is possible that the returned numerical solution does not satisfy the second constraint in Eq. (162) and the corresponding PEF is not valid. In this case, we can multiply the returned numerical solution by a positive factor smaller than 1, whose value is given by the reciprocal of the maximal left-hand side of the above second constraint at the extreme points of $\mathcal{C}$. Then, the rescaled solution is a valid PEF.

All examples involve inferring an experimental probability distribution, in the first example from the published data, and in the other two from an initial set of trials referred to as the *training set*. Due to finite statistics, the data's empirical frequencies are unlikely to satisfy the constraints used for probability estimation. It is necessary to ensure that the distribution for which log-prob rates and corresponding PEFs are calculated satisfies the constraints, because otherwise, the distribution is not consistent with the physical model underlying the randomness generation process, and so unrealistically high predicted log-prob rates may be obtained. To ensure self-consistency, before solving the optimization in Eq. (162), we determine the closest constraints-satisfying conditional distributions. For this, we take as input the empirical frequencies $f(cz)$ and find the most likely distribution ν of CZ with $(\nu[C|z])_z \in \mathcal{Q}_{C|Z}$ and $\nu(z) = f(z)$ for each z . That is, we solve the following optimization problem:

$$\begin{aligned} &\text{Maximize: } \sum_{cz} f(cz) \ln(\nu'(c|z)/f(c|z)) \\ &\text{Subject to: } (\nu'[C|z])_z \in \mathcal{Q}_{C|Z}. \end{aligned} \quad (163)$$

The objective function is proportional to the logarithm of the ratio of the likelihood of the observed results with respect to the constrained distribution ν' and that with respect to the empirical frequencies f . This optimization problem involves maximizing a concave function over a convex domain and can be solved by widely available tools. Note that in all cases, we ensure that the solution satisfies Tsirelson's bounds. Once $\nu(C|Z)$ has been determined by solving the optimization problem in Eq. (163), we use $(\nu[C|z])_z \in \mathcal{C}_{C|Z}$ for the predictions in the first example, and for determining the best PEFs for the data to be analyzed in the other two examples. In each case, we use the model for the settings probability distribution relevant to the situation, not the empirical one $f(z)$.

An issue that comes up when using physical random number sources to choose the settings is that these sources are almost but not quite uniform. Thus allowances for possible

FIG. 1. $\log_2$ -prob rates (labelled with $\kappa = 0$) at ρ_{atoms} for different powers β . The rates are shown for nonsignaling $\mathcal{N}_{C|Z}$ and quantum $\mathcal{Q}_{C|Z}$ (Tsirelson's bounds only) constraints. Also shown are the net $\log_2$ -prob rates according to Eq. (111) for error bounds $\epsilon_h = 2^{-\kappa\sigma\eta}$, where σ is the $\log_2$ -prob rate and for three positive values of κ . Note that the constant κ can be interpreted as a $\log_2$ -error-bound rate with respect to the nominally available entropy for randomness generation.

FIG. 2. Expected net $\log_2$ -probs for 27683 i.i.d. trials at ρ_{atoms} as a function of error bound. The number of trials is as reported in Ref. [7] for the state $|\Psi^-\rangle$.

trial-by-trial biases in the settings choices need to be made. We do this by modifying the distribution constraints. Instead of having a fixed uniform settings distribution, for any bias b , we consider the set

$$\mathcal{B}_{V,b} = \{v : v \text{ is a distribution of } V \text{ and for all } v, |v(v) - \text{Unif}_V(v)| \leq b/|\text{Rng}(V)|\}. \quad (164)$$

For $\text{Rng}(V) = \{0, 1\}$, the extreme points of $\mathcal{B}_{V,b}$ are distributions v of the form $v(v) = (1 \pm (-1)^v b)/2$. For the (2,2,2) configuration, $Z = XY$ and we use $\mathcal{C}_{Z,b} = \text{Cvx}(\mathcal{B}_{X,b} \otimes \mathcal{B}_{Y,b})$ for the settings distribution. We then consider the free-for- Z distributions given by $\mathcal{N}_{C|Z} \times \mathcal{C}_{Z,b}$ and $\mathcal{Q}_{C|Z} \times \mathcal{C}_{Z,b}$. Thus, we allow any settings distribution that is in the convex span of

FIG. 3. Asymptotic gain rate (base 2) as a function of bias at ρ_{atoms} . This is the $\log_2$ -prob rate achieved for power $\beta \rightarrow 0_+$. The probability distribution of the settings $Z = XY$ at each trial is in the convex closure of the distributions where X and Y are independent and their probabilities lie in the interval $[(1-b)/2, (1+b)/2]$. Note that except for these convex constraints, the probabilities can vary arbitrarily from trial to trial.

FIG. 4. $\log_2$ -prob rates for representative biases at ρ_{atoms} . Here we show the $\log_2$ -prob rate as a function of power β at three representative biases. See also Fig. 3.

independent settings choices with limited bias. We do not consider the convex polytope of arbitrarily biased joint settings distributions given by $\mathcal{B}_{Z,b}$ but note that $\mathcal{C}_{Z,b} \subseteq \mathcal{B}_{Z,2b+b^2}$ and $\mathcal{B}_{Z,b^2} \subseteq \mathcal{C}_{Z,b}$. Another convex polytope that could be used is $\mathcal{B}'_{V,\epsilon} = \{v[V] : v(v) \geq \epsilon \text{ for all } v\}$. We have $\mathcal{B}_{Z,b} \subseteq \mathcal{B}'_{Z,(1-b)/4}$. It is known that there are Bell inequalities that can be violated for certain quantum states and settings-associated measurements when the settings distribution is arbitrary in $\mathcal{B}'_{Z,\epsilon}$ for all $1/4 \geq \epsilon > 0$ [55].

FIG. 5. Expected net entropy (in bits) at ρ_{atoms} . We optimized the power β and the probability r according to Eq. (165) given the number of trials n and the error bound ϵ_h as explained in the text. The break-even points are where the curves cross zero. The parameters for those points are in Table II.

TABLE II. Parameters for the break-even points for randomness expansion with ρ_{atoms} . These are the parameters for the points where the curves in Fig. 5 cross zero.

	ϵ_h	n_c	β	r	σ
$\mathcal{N}_{C Z}$	1×10^{-3}	2588821	1.3624×10^{-4}	3.7451×10^{-3}	0.060561
	1×10^{-6}	5177642			
	1×10^{-9}	7766462			
$\mathcal{Q}_{C Z}$	1×10^{-3}	1333781	1.4945×10^{-4}	7.3516×10^{-3}	0.108035
	1×10^{-6}	2667562			
	1×10^{-9}	4001343			

C. Exploration of parameters for a representative distribution of the outcomes

We switch to logarithms base 2 for all log-probs and entropies in the explorations to follow. Thus all explicitly given numerical values for such quantities can be directly related to bits without conversion. For clarity, we indicate the base as a subscript throughout.

For the first example, we use outcome table XII from the supplementary material of Ref. [7] to determine a representative distribution ρ_{atoms} for state-of-the-art loophole-free Bell tests involving entangled atoms. This experiment involved two neutral atoms in separate traps at a distance of 398 m. They were prepared in a heralded entangled state before choosing uniformly random measurement settings and determining the outcomes. The experiment was loophole-free and showed violation of LR at a significance level of 1.05×10^{-9} . The outcome table shows outcome counts for two different initial states labeled Ψ^+ and Ψ^- . We used the counts for the state Ψ^- , which has better violation. We determined the optimal conditional distribution satisfying $\mathcal{Q}_{C|Z}$ as explained in the previous section, which gives the settings-conditional distribution ρ_{atoms} shown in Table I.

We begin by exploring the relationship between power, $\log_2$ -prob rate and net $\log_2$ -prob rate. Figure 1 shows the optimal $\log_2$ -prob rates and net $\log_2$ -prob rates as a function of power β for a few choices of error bound ϵ_h . As expected, the $\log_2$ -prob rates decrease as the power β increases. The asymptotic gain rates can be inferred from the values approached as β goes to zero. The inferred asymptotic gain rates are approximately 0.088 and 0.191 for $\mathcal{N}_{C|Z}$ and $\mathcal{Q}_{C|Z}$, respectively. For small β , constraints $\mathcal{Q}_{C|Z}$ improve the $\log_2$ -prob rates by around a factor of 2 over $\mathcal{N}_{C|Z}$. The net $\log_2$ -prob rates show the trade-off between error bound and $\log_2$ -prob rate, with clear maxima at $\beta > 0$.

Another view of the trade-off is obtained from the expected net $\log_2$ -prob obtained after 27683 trials at different final error bounds, where we optimize the expected $\log_2$ -prob with respect to power β . This is shown in Fig. 2. The expected net $\log_2$ -prob is essentially the amount of min-entropy that we expect to certify by probability estimation in measurement outcomes from the 27683 trials in the experiment of Ref. [7]. The number of near-uniform bits that could have been extracted from this experiment's data is substantially larger than that shown in Fig. 8 for another Bell test using two ions.

We determined the robustness of the asymptotic gain rates and the $\log_2$ -prob rates against biases in the settings random variables X and Y , with the settings distribution constrained by $\mathcal{C}_{Z,b}$ as explained in the previous section. The bias is bounded but may vary within these bounds from trial to trial. Figure 3 shows how the asymptotic gain rate depends on bias. While the maximum bias tolerated in this case is less than 0.05, based on the results of Ref. [55], we expect that there are quantum states and measurement settings for which all biases strictly below 1 can be tolerated in the ideal case. See also Ref. [56]. Figure 4 shows the dependence of the $\log_2$ -prob rate on power β at three representative biases.

Finally, we consider the challenge of producing more random bits than are consumed. This requires using a strategy to minimize the entropy used for settings. We assume i.i.d. trials, with settings distributed according to

$$v_r(xy) = (1 - r)[xy = 11] + r/4, \quad (165)$$

which uses uniform settings distribution with probability r and a default setting 11 with probability $1 - r$. Let $S(r) = -(3r/4)\log_2(r/4) - (1 - 3r/4)\log_2(1 - 3r/4)$ be the base-2 entropy of v_r , $\sigma(\beta)$ the $\log_2$ -prob rate for a given PEF with power β at ρ_{atoms} , and $\epsilon_h = 2^{-\kappa}$ the desired error bound parameterized in terms of the κ and independent of n . If we

TABLE III. Settings-conditional distribution inferred for the XOR 3 training data. The training set consists of the first 50×10^6 trials. The table is organized as described in the caption of Table I.

ab	00	10	01	11
xy				
00	0.999596756631154	0.000106695746779	0.000100495174505	0.000196052447562
10	0.999039892488787	0.000663559889146	0.000086780398739	0.000209767223328
01	0.998967962694884	0.000089505202208	0.000729289110776	0.000213242992132
11	0.998187653168081	0.000869814729010	0.000939019719445	0.000003512383464

FIG. 6. Net $\log_2$ -probs achieved in the XOR 3 data set from Ref. [9]. This is the net $\log_2$ -prob achieved by probability estimation applied to the analysis set after determining the best PEF and power according to the conditional distribution inferred from the training set. The error bound is an input parameter. The curves show the achieved net $\log_2$ -probs for nonsignaling ($\mathcal{N}_{C|Z}$) and quantum ($\mathcal{Q}_{C|Z}$) constraints and three representative biases. Here, we performed probability estimation according to Eq. (53). The curve that is lowest on the right is the net $\log_2$ -prob reported in Ref. [13], where only nonsignaling constraints were exploited.

optimistically assume that all the available min-entropy is extractable, then the expected net number of random bits after n trials is given by

$$\sigma_{\text{net}}(n, \kappa, \beta, r) = n\sigma(\beta) - \kappa/\beta - nS(r). \quad (166)$$

We refer to σ_{net} as the expected *net entropy*. At given values of n and κ , we can maximize the expected net entropy over $\beta > 0$ and $r > 0$. Randomness expansion requires that $\sigma_{\text{net}} \geq 0$. Here we do not account for the seed requirements or the extractor constraints. Figure 5 shows the expected net entropy as a function of n at ρ_{atoms} . Of interest here are the *break-even* points, which are where the expected net entropy becomes positive. For this, consider the break-even equation $n\sigma(\beta) - \kappa/\beta - nS(r) = 0$. Equivalently, $\kappa = \beta(\sigma(\beta) - S(r))n$, and to minimize the number of trials n required for break-even, it suffices to maximize $\beta(\sigma(\beta) - S(r))$, independently of κ . (This independence motivates our choice of parametrization of ϵ_h .) Given the maximizing values of β and r , the critical

FIG. 7. Asymptotic gain rate (base 2) as a function of bias estimated from the training data in the XOR 3 data set in Ref. [9] for nonsignaling and quantum constraints.

value for n at the break-even point is expressed as

$$n_c = \frac{\kappa}{\beta(\sigma(\beta) - S(r))}. \quad (167)$$

One can see that n_c scales linearly with κ . Following this strategy, we obtained the values of the parameters at the break-even points. They are shown in Table II.

D. Randomness from an optical loophole-free Bell test

Reference [13] describes the generation of 256 bits of randomness within 0.001 of uniform from the optical loophole-free Bell test of Ref. [9], assuming only nonsignaling constraints for certification. The data set used is labeled XOR 3 (pulses 3-9) and exhibits a p -value against LR of 2.03×10^{-7} according to the analysis in Ref. [9]. It consists of about 1.82×10^8 trials with a low violation of LR per trial. The protocol of Ref. [13] was developed specifically to obtain randomness from Bell-test data with little violation per trial. It implicitly involves probability estimation. Table III shows the closest distribution in $\mathcal{Q}_{C|Z}$ inferred from the training set consisting of the first 50×10^6 trials of XOR 3.

The analysis in Ref. [13] left open the question of how much min-entropy one can certify in XOR 3 given bias in the settings choices. As reported in Ref. [9], small biases were present in the experiment, due to fluctuations in temperature

TABLE IV. Settings-conditional distribution inferred for the experiment reported in Ref. [10]. The distribution is determined from the training set consisting of the first 1000 trials. The table is organized as described in the caption of Table I.

ab	00	10	01	11
xy				
00	0.395306466091468	0.117610486235535	0.093816274721118	0.393266772951878
10	0.385009648861242	0.101263041214040	0.104113091951344	0.409614217973373
01	0.411397337408393	0.101519614918611	0.097960367844259	0.389122679828738
11	0.077378395334667	0.408894294740616	0.431979309917985	0.081748000006733

that affected the physical random sources used and issues with the high-speed synchronization electronics. A plot of the net $\log_2$ -prob achieved as a function of error bound at a few representative biases is shown in Fig. 6. We conclude that randomness certification from the XOR 3 data set is robust against biases substantially larger than the average ones inferred in Ref. [9]. The robustness is similar to that reported in Ref. [9] for the p -value bounds against LR. To get the results in Fig. 6, we determined the optimal PEFs and powers according to the conditional distribution in Table III. Then we applied the PEFs obtained to the remaining trials (the *analysis* trials). Here, we did not adapt the PEFs while processing the analysis trials.

We also determined the asymptotic gain rates from the conditional distribution in Table III. The dependence of the asymptotic gain rate on the bias parameter b is shown in Fig. 7. Notably, the robustness against bias is similar to that for ρ_{atoms} shown in Fig. 3. This may be due to both experiments being designed for the same family of Bell inequalities. For higher Bell-test bias tolerance, different inequalities and measurement settings need to be used [55].

E. Reanalysis of “random numbers certified by Bell’s theorem [10]”

Finally, we applied the procedure that we used in reanalyzing the XOR 3 data set in the previous section to the data from the experiment reported in Ref. [10]. This experiment involved two ions separated by about 1 meter in two different ion traps. As a Bell test, the experiment closed the detection loophole, but was subject to the locality loophole. The authors certified a min-entropy of 42 bits with respect to classical side information at an error bound of 0.01 with 3016 trials. They assumed quantum constraints (implemented by the NPA hierarchy) for this certification. For the reanalysis, we set aside a training set consisting of the first 1000 trials to estimate a constraints-satisfying conditional distribution. The distribution obtained is given in Table IV. We optimized PEFs and their powers as described in the previous section and applied them to the remaining 2016 trials, the analysis set. The

FIG. 8. Net $\log_2$ -probs achieved in the data from Ref. [10]. The analysis set consisted of the last 2016 trials. Results for nonsignaling and quantum constraints are shown. For these curves, we performed probability estimation according to Eq. (53). The reference reported 42 bits at $\epsilon_h = 10^{-2}$ for quantum constraints.

result is shown in Fig. 8. We expect that if it had been possible to optimize the PEFs based on the calibration experiments preceding the 3016 trials, then the $\log_2$ -probs obtained would have been approximately 50% larger, assuming that the trial statistics were sufficiently stable.

ACKNOWLEDGMENTS

We thank D. N. Matsukevich for providing the experimental data for Ref. [10], Yi-Kai Liu and Carl Miller for help with reviewing this paper, and Scott Glancy for discussions. P.B. acknowledges the support of NSF Grant No. 1839223. This work includes contributions of the National Institute of Standards and Technology, which are not subject to U.S. copyright.

-
- [1] R. Colbeck, Quantum and relativistic protocols for secure multi-party computation, Ph.D. thesis, University of Cambridge, 2007.
 - [2] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, Bell nonlocality, *Rev. Mod. Phys.* **86**, 419 (2014).
 - [3] A. Acín and L. Masanes, Certified randomness in quantum physics, *Nature (London)* **540**, 213 (2016).
 - [4] J. Bell, On the Einstein Podolsky Rosen paradox, *Physics* **1**, 195 (1964).
 - [5] M. Genovese, Research on hidden variable theories: A review of recent progresses, *Phys. Rep.* **413**, 319 (2005).
 - [6] B. Hensen *et al.*, Loophole-free Bell inequality violation using electron spins separated by 1.3 km, *Nature (London)* **526**, 682 (2015).
 - [7] W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortégel, M. Rau, and H. Weinfurter, Event-Ready Bell-Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes, *Phys. Rev. Lett.* **119**, 010402 (2017).
 - [8] M. Giustina, M. A. M. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. 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, *Phys. Rev. Lett.* **115**, 250401 (2015).
 - [9] 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, *Phys. Rev. Lett.* **115**, 250402 (2015).
- [10] S. Pironio, A. Acin, S. Massar, A. Boyer 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 (London)* **464**, 1021 (2010).
- [11] S. Fehr, R. Gelles, and C. Schaffner, Security and composability of randomness expansion from Bell inequalities, *Phys. Rev. A* **87**, 012335 (2013).
- [12] S. Pironio and S. Massar, Security of practical private randomness generation, *Phys. Rev. A* **87**, 012336 (2013).
- [13] P. Bierhorst, E. Knill, S. Glancy, A. Mink, S. Jordan, A. Rommal, Y.-K. Liu, B. Christensen, S. W. Nam, and L. K. Shalm, Experimentally generated random numbers certified by the impossibility of superluminal signaling (version 1), [arXiv:1702.05178v1](https://arxiv.org/abs/1702.05178v1).
- [14] Y. Zhang, S. Glancy, and E. Knill, Asymptotically optimal data analysis for rejecting local realism, *Phys. Rev. A* **84**, 062118 (2011).
- [15] Y. Zhang, S. Glancy, and E. Knill, Efficient quantification of experimental evidence against local realism, *Phys. Rev. A* **88**, 052119 (2013).
- [16] G. Shafer, A. Shen, N. Vereshchagin, and V. Vovk, Test martingales, Bayes factors and p -values, *Stat. Sci.* **26**, 84 (2011).
- [17] A. Acin, S. Massar, and S. Pironio, Randomness Versus Nonlocality and Entanglement, *Phys. Rev. Lett.* **108**, 100402 (2012).
- [18] O. Nieto-Silleras, S. Pironio, and J. Silman, Using complete measurement statistics for optimal device-independent randomness evaluation, *New J. Phys.* **16**, 013035 (2014).
- [19] J.-D. Bancal, L. Sheridan, and V. Scarani, More randomness from the same data, *New J. Phys.* **16**, 033011 (2014).
- [20] O. Nieto-Silleras, C. Bamps, J. Silman, and S. Pironio, Device-independent randomness generation from several Bell estimators, *New J. Phys.* **20**, 023049 (2018).
- [21] U. Vazirani and T. Vidick, Certifiable quantum dice: Or, true random number generation secure against quantum adversaries, in *Proceedings of the 44th Symposium on Theory of Computing, STOC 2012*, edited by H. J. Karloff and T. Pitassi (ACM, New York, 2012), pp. 61–76.
- [22] C. A. Miller and Y. Shi, Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices, *J. ACM* **63**, 33 (2016).
- [23] C. A. Miller and Y. Shi, Universal security for randomness expansion from the spot-checking protocol, *SIAM J. Comput.* **46**, 1304 (2017).
- [24] M. Coudron and H. Yuen, Infinite randomness expansion with a constant number of devices, in *STOC'14 Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (ACM, New York, 2014), pp. 427–436.
- [25] K.-M. Chung, Y. Shi, and X. Wu, Physical randomness extractors: Generating random numbers with minimal assumptions, [arXiv:1402.4797](https://arxiv.org/abs/1402.4797) [quant-ph].
- [26] R. Arnon-Friedman, F. Dupuis, O. Fawzi, R. Renner, and T. Vidick, Practical device-independent quantum cryptography via entropy accumulation, *Nat. Commun.* **9**, 459 (2018).
- [27] R. König and B. Terhal, The bounded-storage model in the presence of a quantum adversary, *IEEE Trans. Inf. Theory* **54**, 749 (2008).
- [28] Y. Zhang, E. Knill, and P. Bierhorst, Certifying quantum randomness by probability estimation, *Phys. Rev. A* **98**, 040304(R) (2018).
- [29] E. Knill, Y. Zhang, and H. Fu, Quantum probability estimation for randomness with quantum side information, [arXiv:1806.04553](https://arxiv.org/abs/1806.04553).
- [30] Y. Zhang, H. Fu, and E. Knill, Efficient randomness certification by quantum probability estimation, *Phys. Rev. Research* **2**, 013016 (2020).
- [31] M. Tomamichel, R. Colbeck, and R. Renner, A fully quantum asymptotic equipartition property, *IEEE Trans. Inf. Theory* **55**, 5840 (2009).
- [32] M. C. Pardo and I. Vajda, About distances of discrete distributions satisfying the data processing theorem of information theory, *IEEE Trans. Inf. Theory* **43**, 1288 (1997).
- [33] P. Hayden, R. Jozsa, D. Petz, and A. Winter, Structure of states which satisfy strong subadditivity of quantum entropy with equality, *Commun. Math. Phys.* **246**, 359 (2004).
- [34] M. Navascués, S. Pironio, and A. Acín, Bounding the Set of Quantum Correlations, *Phys. Rev. Lett.* **98**, 010401 (2007).
- [35] S. P. Vadhan, *Pseudorandomness*, Foundations and Trends in Theoretical Computer Science Vol. 7 (Now Publishers, Norwell, MA, 2012).
- [36] W. Maurer, C. Portmann, and V. B. Scholz, A modular framework for randomness extraction based on Trevisan's construction, [arXiv:1212.0520](https://arxiv.org/abs/1212.0520), code available on github (2012).
- [37] R. König, R. Renner, and C. Schaffner, The operational meaning of min- and max-entropy, *IEEE Trans. Inf. Theory* **55**, 4337 (2009).
- [38] J. Ville, *Etude Critique de la Notion de Collectif* (Gauthier-Villars, Paris, 1939).
- [39] B. G. Christensen, A. Hill, P. G. Kwiat, E. Knill, S. W. Nam, K. Coakley, S. Glancy, L. K. Shalm, and Y. Zhang, Analysis of coincidence-time loopholes in experimental Bell tests, *Phys. Rev. A* **92**, 032130 (2015).
- [40] C. Dellacherie and P. A. Meyer, *Probabilities and Potential, B: Theory of Martingales*, North-Holland Mathematics Studies (Elsevier Science, Amsterdam, Netherlands, 2011).
- [41] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Proposed Experiment to Test Local Hidden-Variable Theories, *Phys. Rev. Lett.* **23**, 880 (1969).
- [42] B. Schumacher, Information and quantum nonseparability, *Phys. Rev. A* **44**, 7047 (1991).
- [43] E. Knill, S. Glancy, S. W. Nam, K. Coakley, and Y. Zhang, Bell inequalities for continuously emitting sources, *Phys. Rev. A* **91**, 032105 (2015).
- [44] M. Müller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel, On quantum Rényi entropies: A new generalization and some properties, *J. Math. Phys.* **54**, 122203 (2013).
- [45] R. L. Frank and E. H. Lieb, Monotonicity of a relative Rényi entropy, *J. Math. Phys.* **54**, 122201 (2013).
- [46] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2001).
- [47] A. Uhlmann, Entropy and optimal decompositions of states relative to a maximal commutative subalgebra, *Open Systems Inf. Dynamics* **5**, 209 (1998).

- [48] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, UK, 2004).
- [49] R. V. Kadison and J. R. Ringrose, *Fundamentals of Theory of Operator Algebras. Vol. I: Elementary Theory*, Graduate Studies in Mathematics Vol. 15 (American Mathematical Society, Providence, RI, 1997).
- [50] L. K. Shalm, Y. Zhang, J. C. Bienfang, C. Schlager, M. J. Stevens, M. D. Mazurek, C. Abellán, W. Amaya, M. W. Mitchell, M. A. Alhejji, H. Fu, J. Ornstein, R. P. Mirin, S. W. Nam, and E. Knill, Device-independent randomness expansion with entangled photons, [arXiv:1912.11158](#).
- [51] B. S. Cirel'son, Quantum generalizations of Bell's inequality, *Lett. Math. Phys.* **4**, 93 (1980).
- [52] S. Popescu and D. Rohrlich, Quantum nonlocality as an axiom, *Found. Phys.* **24**, 379 (1994).
- [53] P. Bierhorst, Geometric decompositions of Bell polytopes with practical applications, *J. Phys. A: Math. Theor.* **49**, 215301 (2016).
- [54] P. Bierhorst and Y. Zhang, Tsirelson polytopes and randomness generation, *New J. Phys.* **22**, 083036 (2020).
- [55] G. Pütz, D. Rosset, T. J. Barnea, Y.-C. Liang, and N. Gisin, Arbitrarily Small Amount of Measurement Independence is Sufficient to Manifest Quantum Nonlocality, *Phys. Rev. Lett.* **113**, 190402 (2014).
- [56] M. Kessler and R. Arnon-friedman, Device-independent randomness amplification and privatization, [arXiv:1705.04148](#).