Certified randomness from quantum speed limits

Caroline L. Jones,[1,2,*] Albert Aloy,[1,2,†] Gerard Higgins,[1,3] and Markus P. Müller[1,2,4]

[1] *Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, A-1090 Vienna, Austria*
[2] *Vienna Center for Quantum Science and Technology (VCQ), Faculty of Physics, University of Vienna, Vienna, Austria*
[3] *Institute of High Energy Physics, Austrian Academy of Sciences, Dominikanerbastei 16, A-1010 Vienna, Austria*
[4] *Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada*

(Dated: June 17, 2025)

Quantum speed limits are usually regarded as fundamental restrictions, constraining the amount of computation that can be achieved within some given time and energy. Complementary to this intuition, here we show that these limitations are also of operational value: they enable the secure generation of certified randomness. We consider a prepare-and-measure scenario with some (experimentally determined or promised) upper bound on the energy uncertainty ΔE of the average prepared quantum state, but without any further assumptions on the devices, Hilbert space or Hamiltonian. Given that we can freely choose the time at which to apply the untrusted preparation procedure, we show that this scenario admits the generation of randomness that is secure against adversaries with additional classical information. We show how to determine the amount of certified randomness given the observed correlations, discuss how interactions with the environment are taken into account, and sketch a conceivable experimental implementation. Remarkably, even single-mode coherent states admit this kind of certification of non-zero randomness in some parameter regimes, reinforcing ongoing approaches to demonstrate versions of nonclassicality in the simple harmonic oscillator. Our results extend existing efforts to devise semi-device-independent protocols grounded in reasonable physical assumptions, and they contribute to the understanding of time-energy uncertainty relations via their operational consequences.

I. INTRODUCTION

Historically, the probabilistic nature of quantum mechanics has been the root of scientific disquiet – with Einstein famously remarking that "[Nature] does not throw dice" [1]. Nevertheless, converse to this intuition, the role of indeterminacy became further entrenched in quantum theory with Heisenberg's posited uncertainty principles [2, 3], which claimed that pairs of canonical variables could only be known up to some jointly bounded precision:

$$\begin{aligned} \Delta x \Delta p &\geq \hbar/2, \\ \Delta E \Delta t &\geq \hbar/2. \end{aligned} \tag{1}$$

Whilst the position-momentum uncertainty relation was soon established mathematically [4], its time-energy analogue proved more cumbersome. In particular, without a formulation of time as a quantum observable, it was initially unclear how to interpret Δt as "time uncertainty". An alternative paradigm was proposed by Mandelstam and Tamm [5], reformulating the relation as a *quantum speed limit* (QSL):

$$\Delta t \geq \tau_{\mathrm{QSL}} := \frac{\pi\hbar}{2\Delta E}, \tag{2}$$

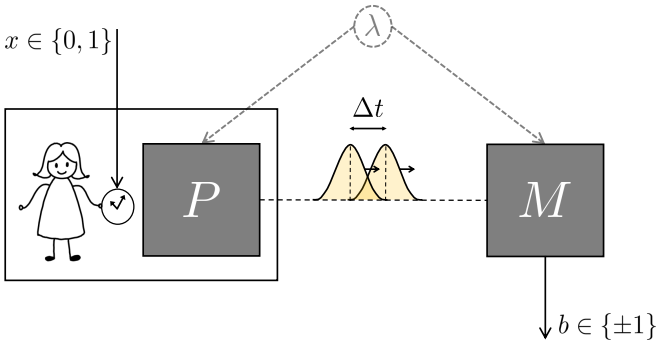

FIG. 1. The preparation device P prepares a fixed quantum state which is subsequently sent to a measurement device M. We treat both P and M as black boxes, but we assume that we have full control of the time t_x at which P is triggered: at t_0 or $t_1 = t_0 + \Delta t$, where $x \in \{0, 1\}$ labels the input to P. In contrast, M is implemented at a predetermined time. A classical variable λ is allowed to influence and possibly correlate P and M. It may be known to an adversary, but it is unknown to the user of the randomness generator.

specifying a minimum time τ_{QSL} that a system takes to evolve to an orthogonal state. The Mandelstam-Tamm bound was further generalised to arbitrary states [6]:

$$\Delta t \geq \frac{\hbar \arccos |\langle\psi_1|\psi_0\rangle|}{\Delta E}, \tag{3}$$

thereby bounding how fast a system can evolve between any two states $|\psi_0\rangle$ and $|\psi_1\rangle$, where $\Delta E = \Delta E_{\psi_0} = \Delta E_{\psi_1}$ is the standard deviation of the energy in the initial and

* CarolineLouise.Jones@oeaw.ac.at;
† Albert.Aloy@oeaw.ac.at;
AA and CLJ contributed equally to this work.

final states. Note that when the states are fully distinguishable (i.e. $\arccos|\langle\psi_1|\psi_0\rangle| = \pi/2$), Eq. (2) is recovered from Eq. (3).

Whilst typically viewed as a *limitation* [7–11], in this paper we demonstrate the utility of the QSL for quantum information protocols. In particular, we show its application for randomness generation. At first sight, it seems almost trivial to go from uncertainty relations to randomness: why not just prepare a particle in some state with very small Δp, ensuring large Δx, and measure its position? Or why do we not simply send a photon on a half-silvered mirror and record the outcome? A reason to be careful is that we may not trust our devices (the mirror, or the state preparation procedure): perhaps these device are producing outcomes that look random to us, but are in fact predictable by an adversary who holds additional information about the state of the system.

These restrictions can be overcome with device-independent or *semi-device-independent* protocols. In the latter case, we have a prepare-and-measure scenario (such as the one in Figure 1) where we do not trust the preparation and measurement devices P and M and treat them as black boxes, but we trust the validity of an assumption about the physical system that is sent from P to M. For example, we may assume that the quantum system to be sent is a qubit, or, more generally, that it is described by a Hilbert space of some fixed dimension d [12–18]. Observing some correlations between the inputs of P (and perhaps M), and the output of M, we can then ensure that this output must be at least partially random, even relative to every adversary holding an arbitrary amount of additional classical information about the devices P and M and all variables that are physically relevant for these procedures.

However, the assumption of Hilbert space dimension d is arguably not very well physically motivated. Therefore, various protocols have been proposed which replace the dimensionality assumption by assumptions that have a more direct physical or information-theoretic meaning. Several works have analysed semi-device-independent randomness generation schemes based on assumptions about the overlaps between the prepared states [19–23], or about the information content of the transmitted system [24, 25], with important conceptual and formal relations between the different proposals [26]. In particular, van Himbeeck et al. [27] have shown that assuming a bound on the expected energy $\langle E\rangle$ of the prepared state admits the certification of randomness [28]. However, this assumes that the corresponding Hamiltonian $\hat{H}$ has a unique ground state, and assumes knowledge of the relation between the energy expectation value and the gap above the ground state. Here, we show how randomness can be certified by assuming an upper bound on the energy uncertainty ΔE of the prepared state, without knowing or assuming anything about the Hamiltonian.

Conceptually, our scheme builds on the idea that we can perform *trusted operations on untrusted devices*: Even if our preparation procedure is treated as a black box, we typically assume that we have an evident, pretheoretic, "macroscopic" notion of some operations that we can apply to it, regardless of the "microscopic" details of the device. For example, we think that we know what it means to supply an input to a black box, or to rewire the outputs of several boxes, or to place several boxes far apart in spacelike separation (as in Bell experiments). In previous work [29], we have assumed that we know how to rotate preparation devices around a fixed axis in space, and we have analyzed the resulting prepare-and-measure correlations within and beyond quantum physics [30]. Similarly, here we imagine that we can freely decide the time at which we operate a given box (say, by pressing the button that triggers the untrusted preparation procedure), which yields a semi-device-independent randomness generation scheme based solely on an upper bound to the system's energy standard deviation ΔE.

We emphasize that our results build to a significant extent on earlier works by van Himbeeck et al. [27, 28], but with some important differences and novelties. For example, the fact that the standard deviation ΔE is not an observable implies that the "classical max average" set of correlations $\bar{\mathcal{C}}_{\mathfrak{E},\Delta t}$ is not a subset of the quantum set $\mathcal{Q}_{\mathfrak{E},\Delta t}$, and that their algorithm must be adapted to determine the amount of certified randomness. The use of quantum speed limits leads to the necessity of further physical and conceptual considerations. See Section V for a more detailed comparison to existing work.

Our article is organized as follows. In Section II, we give a theoretical description of the prepare-and-measure scenario. We characterize the set of correlations that are consistent with quantum theory for any given ΔE and Δt, and use concavity of the variance to determine the set of correlations that are "classical", i.e. that do not admit the generation of certified randomness. We show how the results generalize from closed system evolution to interactions with an environment. In Section III, we show how the amount of certifiable randomness can be determined, and give an example plot of the result. In Section IV, we describe a possible experimental implementation involving the quantum harmonic oscillator, and we conclude in Section V.

II. THEORETICAL DESCRIPTION

We follow the semi-device-independent (semi-DI) protocol of [27], consisting of a simple prepare-and-measure scenario as depicted in Figure 1. We begin by describing the scenario without taking the variable λ into account. A preparation box P takes an input $x \in \{0, 1\}$, and sends some system to a measurement box M, which yields one of two outputs $b \in \{\pm 1\}$. By construction, the only effect of the input is to control the time at which the preparation procedure is implemented: in the case of input $x = 0$, the state is prepared at time t_0, or, in the case $x = 1$, the state is prepared at time $t_1 = t_0 + \Delta t$. While we do not

trust the device P, we assume that we can fully control the time at which we perform P, similarly as we believe that we can control the choice of input or whether we implement any operation at all. Thus, the two possible states ρ_x that may be sent to M are time-displaced with respect to one another by some delay Δt.

The measurement device then produces outcome b, and is described by some POVM $\{M_b\}$. Minimal assumptions are made about the devices, therefore ρ_x and M_b are treated as unknown and may fluctuate according to some classical random variable(s) λ, unbeknownst to the experimenter. This variable λ can be thought of as containing additional information about the world, which could in principle constitute some predictive advantage for the outcome b. Our goal is then to quantify over all possible assignments of $\lambda \in \Lambda$ such that security is guaranteed independently of knowledge of these variables. This means that we allow for shared randomness that is able to correlate the workings of the devices P and M. The joint behaviour of the devices is therefore characterised by the probabilities

$$P(b|x) = \sum_{\lambda} p(\lambda) \mathrm{tr}\big[M_b^{\lambda} U_{t_x} \rho_0^{\lambda} U_{t_x}^{\dagger}\big], \tag{4}$$

with $t_x \in \{0, \Delta t\}$, and $U_{t_x} = e^{-iHt_x/\hbar}$ describes the unitary evolution of the system under some fixed Hamiltonian $\hat{H}$. For now, we assume that the prepared quantum state ρ_x undergoes closed-system evolution according to a fixed Hamiltonian $\hat{H}$ which defines its energy's standard deviation

$$\Delta E_{\rho_x} = \sqrt{\mathrm{tr}(\rho_x \hat{H}^2) - (\mathrm{tr}(\rho_x \hat{H}))^2},$$

which gives identical values for $x = 0$ and $x = 1$. We will relax this assumption further below.

The presence of useful correlations is expressed by deviation from the line $C_0 = C_1$, where C_x characterises the bias of the output towards ± 1 for the input x:

$$C_x = P(+1|x) - P(-1|x) \quad (x \in \{0, 1\}). \tag{5}$$

In particular, $C_0 \neq C_1$ indicates that the outcome b is influenced by the choice of input x. We will now analyse the possible correlations given some value of Δt and upper bound on ΔE.

Quantum correlations

Let us begin by considering the prepare-and-measure scenario without the shared randomness λ, and discuss how to reintroduce it at the end of this section. In this case, we have probabilities $P(b|x) = \mathrm{tr}(M_b \rho_x)$ with $\rho_x = U_{t_x} \rho_x U_{t_x}^{\dagger}$, such that $C_x = \mathrm{tr}(\rho_x M)$ with $M = M_{+1} - M_{-1}$. Let us initially assume that the preparation device prepares a pure state $|\psi_0\rangle$ at time t_0. Then, at time $t_1 = t_0 + \Delta t$, the state is given by

$$|\psi_1\rangle = U_{\Delta t} |\psi_0\rangle = \sum_n e^{-iE_n \Delta t/\hbar} c_n |E_n\rangle, \tag{6}$$

where $|\psi_0\rangle = \sum_n c_n |E_n\rangle$ is decomposed into its energy eigenbasis (we absorb all further time evolution into the definition of the measurement procedure M).

We define the quantum set of pure state correlations for time delay $\Delta t \geq 0$ and maximal energy uncertainty $\mathfrak{E} \geq 0$ as follows:

$$\mathcal{Q}_{\mathfrak{E},\Delta t} := \Big\{ (C_0, C_1) \,\Big|\, C_x = \langle \psi_x | M | \psi_x \rangle, -\mathbb{1} \leq M \leq \mathbb{1}, \\ \exists \hat{H} \text{ s. t. } |\psi_1\rangle = U_{\Delta t} |\psi_0\rangle, \Delta E_{\psi_0} \leq \mathfrak{E} \Big\}. \tag{7}$$

The penultimate condition imposes the existence of some fixed Hamiltonian $\hat{H}$ that relates states $|\psi_0\rangle$ to $|\psi_1\rangle$ via unitary time evolution, where the energy uncertainty of the initial state ΔE_{ψ_0} is less than some maximal fixed value $\mathfrak{E}$. This condition $\Delta E_{\psi_0} \leq \mathfrak{E}$ implements the semi-DI assumption for our scenario, expressing our belief of an upper bound on the energy uncertainty.

It has already been proven in [27] that, in a prepare-and-measure scenario communicating pure states, the possible correlations are characterised by the inequality

$$\frac{1}{2}\left(\sqrt{1+C_0}\sqrt{1+C_1} + \sqrt{1-C_0}\sqrt{1-C_1}\right) \geq \gamma, \tag{8}$$

where γ is defined as the smallest possible overlap between the two states $|\psi_0\rangle$ and $|\psi_1\rangle$. An overlap $\gamma = 1$ corresponds to the maximally restricted set of correlations $C_0 = C_1$, whilst as $\gamma \to 1$ the range of accessible correlations increases. For our scenario, the quantum speed limit constrains the overlap as follows:

$$\gamma := \min |\langle \psi_1 | \psi_0 \rangle| = \begin{cases} \cos(\mathfrak{E}\Delta t) & \text{if } \mathfrak{E}\Delta t < \frac{\pi}{2}, \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

(Note that from here, we use natural units $\hbar = 1$ for simplicity.) This allows us to evaluate the set of pure-state quantum correlations, given the energy uncertainty of the state prepared by P and the time delay between possible preparation times. In Appendix B, we show that all correlations that satisfy Eq. (8) with equality (corresponding to the two curves that bound the blue set (for any fixed γ) in Figure 2 have a quantum model, i.e. are contained in $\mathcal{Q}_{\mathfrak{E},\Delta t}$. Furthermore, by considering purifications with an ancilla system, we prove in Appendix A that the set of correlations is unchanged when one allows for mixed states, related according to $\rho_1 = U_{\Delta t} \rho_0 U_{\Delta t}^{\dagger}$ and with energy constraint $\Delta E_{\rho_0} \leq \mathfrak{E}$. This shows that $\mathcal{Q}_{\mathfrak{E},\Delta t}$ is convex, and thus, that $\mathcal{Q}_{\mathfrak{E},\Delta t}$ is exactly the set of (C_0, C_1) that satisfy (8).

Now we will consider the case that there is shared randomness λ. In this case, the observed correlation $\mathbf{C} = (C_0, C_1)$ arises from probabilities of the form of Eq. (4), i.e. $C_x = \sum_{\lambda} p(\lambda) \mathrm{tr}[M^{\lambda} \rho_x^{\lambda}]$, where $M^{\lambda} = M_{+1}^{\lambda} - M_{-1}^{\lambda}$. Our goal is to show that $\mathbf{C} \in \mathcal{Q}_{\mathfrak{E},\Delta t}$ for all $\mathfrak{E} \geq \Delta E_{\rho_0}$ for

$$\rho_x = \sum_{\lambda} p(\lambda) \rho_x^{\lambda}. \tag{10}$$

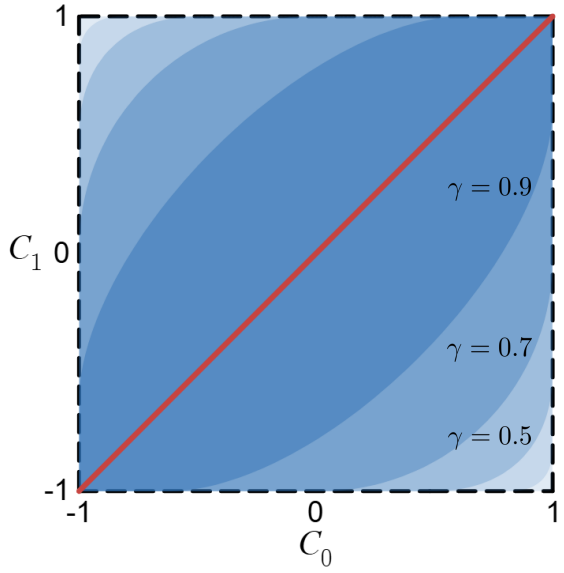

FIG. 2. Set of quantum (blue) and classical (red) correlations for different state overlaps $\gamma \in \{0.9, 0.7, 0.5, 0.3, 0.1\}$ (from darkest to lightest).

On a larger Hilbert space, consider the block matrices $\bar{\rho}_0 := \bigoplus_\lambda p(\lambda)\rho_0^\lambda$ and $\bar{M} := \bigoplus_\lambda M^\lambda$, then $\bar{\rho}_0$ is a density operator and $-\mathbb{1} \leq \bar{M} \leq \mathbb{1}$ because $\bar{M} = \bar{M}^\dagger$ and all its eigenvalues are in the interval $[-1, 1]$. Furthermore, $\bar{H} := \bigoplus \hat{H}$ defines a Hamiltonian that evolves each λ-subspace independently via $\bar{U}_t = e^{-i\bar{H}t} = \bigoplus_\lambda e^{-i\hat{H}t} = \bigoplus_\lambda U_t$. For $\bar{\rho}_x := \bar{U}_{t_x}\bar{\rho}_0\bar{U}_{t_x}^\dagger$, we get $\mathrm{tr}[\bar{M}\bar{\rho}_x] = C_x$, i.e. we have reproduced the correlation $\mathbf{C}$ in a scenario with a *fixed* measurement M. Furthermore, $\mathrm{tr}(\bar{\rho}_0\bar{H}^k) = \mathrm{tr}(\rho_0\hat{H}^k)$ for $k \in \{1, 2\}$, and so the construction also preserves the energy's standard deviation $\Delta E_{\rho_0} = \Delta E_{\bar{\rho}_0}$. This shows that $\mathbf{C} \in \mathcal{Q}_{\mathfrak{E},\Delta t}$ for all $\mathfrak{E} \geq \Delta E_{\rho_0}$: that is, the sets $\mathcal{Q}_{\mathfrak{E},\Delta t}$ describe the correlations in the scenario of Fig. 1 with or without shared randomness λ, obtainable from average prepared states ρ_0 with energy uncertainty $\Delta E_{\rho_0} \leq \mathfrak{E}$.

For $\mathfrak{E}\Delta t \to 0$, the states have overlap $\gamma \to 1$, and so are indistinguishable by the measurement device. In this case, the outputs b must be independent of x; i.e. the correlations are on the line $C_0 = C_1$. However, for increasing $\mathfrak{E}\Delta t$, the overlap decreases such that the measurement device can at least partially distinguish the two states – therefore a larger set of correlations are available, as shown in Figure 2. For $\mathfrak{E}\Delta t \geq \pi/2$, the states may be perfectly distinguishable, and therefore all correlations are possible.

Classical correlations

Suppose that an adversary has complete knowledge of the variable λ in the decomposition (10), and knows the value of x. Under what conditions does this allow the adversary to predict the measurement outcome b perfectly, for both $x \in \{0, 1\}$? To answer this question, we have to consider the energy uncertainty of the states ρ_0^λ if we know that $\Delta E_{\rho_0} \leq \mathfrak{E}$. In Appendix C, we show that the standard deviation ΔE is *concave*, and so

$$\Delta E_{\rho_0} \geq \sum_\lambda p(\lambda)\Delta E_{\rho_0^\lambda}.$$

That is, while some ρ_0^λ may have energy uncertainty strictly larger than $\mathfrak{E}$, the average energy uncertainty is still upper-bounded by $\mathfrak{E}$.

Now suppose that the eavesdropper, knowing λ, can perfectly predict the output of M conditioned on x. This means that the associated correlations $\mathbf{C}^\lambda = (C_0^\lambda, C_1^\lambda)$ must be in the set of deterministic correlations, $\{\pm 1\} \times \{\pm 1\}$, and that the energy bound must be satisfied on average. Thus, the average correlation $\mathbf{C} = \sum_\lambda p(\lambda)\mathbf{C}^\lambda$ must be contained in the following set, the *classical max-average set*, defined similarly as in [27]:

$$\overline{\mathcal{C}}_{\mathfrak{E},\Delta t} = \Bigg\{ \mathbf{C} = \sum_\lambda p(\lambda)\mathbf{C}^\lambda \;\Bigg|\; \mathbf{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}^\lambda,\Delta t}, \; \mathbf{C}^\lambda \in \{\pm 1\} \times \{\pm 1\}, \sum_\lambda p(\lambda)\mathfrak{E}^\lambda \leq \mathfrak{E} \Bigg\}. \quad (11)$$

To characterise this set, we consider two possibilities for each λ; either $\mathfrak{E}^\lambda < \frac{\pi}{2\Delta t}$, in which case the only deterministic correlations are those represented by the line $|C_0^\lambda - C_1^\lambda| = 0$; or $\mathfrak{E}^\lambda \geq \frac{\pi}{2\Delta t}$, in which case all correlations are possible, i.e. $|C_0^\lambda - C_1^\lambda| \leq 2$. We collect the variables according to $\Lambda_1 = \{\lambda_1 : \mathfrak{E}^{\lambda_1} < \frac{\pi}{2\Delta t}\}$, which occurs with probability $\sum_{\lambda_1 \in \Lambda_1} p(\lambda_1) =: p_1$ and $\Lambda_2 = \{\lambda_2 : \mathfrak{E}^{\lambda_2} \geq \frac{\pi}{2\Delta t}\}$, which occurs with probability $\sum_{\lambda_2 \in \Lambda_2} p(\lambda_2) =: p_2 = 1 - p_1$, and where $\Lambda = \Lambda_1 \dot{\cup} \Lambda_2$. First, consider the case that both p_1 and p_2 are non-zero. The correlations are bounded according to

$$|C_0 - C_1| \leq \sum_{\lambda \in \Lambda} p(\lambda)|C_0^\lambda - C_1^\lambda| \leq 2p_2.$$

We have

$$\begin{aligned} &\sum_{\lambda_1 \in \Lambda_1} p(\lambda_1)\mathfrak{E}^{\lambda_1} + \sum_{\lambda_2 \in \Lambda_2} p(\lambda_2)\mathfrak{E}^{\lambda_2} \leq \mathfrak{E} \\ \Rightarrow\; & p_1 \sum_{\lambda_1 \in \Lambda_1} \underbrace{\frac{p(\lambda_1)}{p_1}}_{:=q_{\lambda_1}} \mathfrak{E}^{\lambda_1} + p_2 \sum_{\lambda_2 \in \Lambda_2} \underbrace{\frac{p(\lambda_2)}{p_2}}_{:=q_{\lambda_2}} \mathfrak{E}^{\lambda_2} \leq \mathfrak{E} \\ \Rightarrow\; & p_1\mathfrak{E}_1 + p_2\mathfrak{E}_2 \leq \mathfrak{E}. \end{aligned}$$

In the last line, we have defined $\mathfrak{E}_1 := \sum_{\lambda_1 \in \Lambda_1} q_{\lambda_1}\mathfrak{E}^{\lambda_1}$ and $\mathfrak{E}_2 := \sum_{\lambda_2 \in \Lambda_2} q_{\lambda_2}\mathfrak{E}^{\lambda_2}$ (the averages of the energy uncertainties below and above $\frac{\pi}{2\Delta t}$ respectively), for which $\sum_{\lambda_1 \in \Lambda_1} q_{\lambda_1} = 1$ and $\sum_{\lambda_2 \in \Lambda_2} q_{\lambda_2} = 1$. From this, we can then calculate a bound on p_2:

$$p_2 \leq \frac{\mathfrak{E} - \mathfrak{E}_1}{\mathfrak{E}_2 - \mathfrak{E}_1} = \frac{\delta}{\epsilon + \delta}$$

where, by construction, $\mathfrak{E}_1 < \mathfrak{E} < \mathfrak{E}_2$, so we can write $\mathfrak{E}_1 = \mathfrak{E} - \delta$ and $\mathfrak{E}_2 = \mathfrak{E} + \epsilon$, with $0 < \delta \leq \mathfrak{E}$ and $\epsilon \geq \frac{\pi}{2\Delta t} - \mathfrak{E} \geq 0$. This is clearly maximised by taking the lower bound $\epsilon = \frac{\pi}{2\Delta t} - \mathfrak{E}$ (i.e. $\mathfrak{E}_2 = \frac{\pi}{2\Delta t}$, the minimal $\mathfrak{E}_2$ for which all correlations are possible). Then, the remaining fraction is maximised by taking the upper bound $\delta = \mathfrak{E}$ (i.e. $\mathfrak{E}_1 = 0$). Therefore:

$$p_2 \leq \frac{2\mathfrak{E}\Delta t}{\pi}. \tag{12}$$

So far, we have assumed that $p_1, p_2 \neq 0$. But if $p_2 = 0$, then this bound is trivially true. Furthermore, if $p_1 = 0$, then $\mathfrak{E} = \mathfrak{E}_2 \geq \frac{\pi}{2\Delta t}$, and so $\frac{2\mathfrak{E}\Delta t}{\pi} \geq 1 = p_2$. Hence (12) is true in all cases.

Therefore we have the following bound on the correlations, where the lower case is the one of interest for the max-average assumption:

$$|C_0 - C_1| \leq \begin{cases} 2 & \text{if } \mathfrak{E} \geq \frac{\pi}{2\Delta t}, \\ \frac{4\mathfrak{E}\Delta t}{\pi} & \text{if } \mathfrak{E} < \frac{\pi}{2\Delta t}. \end{cases} \tag{13}$$

Conversely, suppose some correlation $\boldsymbol{k} = (k_0, k_1)$ satisfies inequality (13). It is geometrically clear that this must be in $\overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$, following the same line of reasoning as in [29, Appendix E]. Suppose $k_0 < k_1$, then one can draw a line through $\boldsymbol{k}$, from one corner $(-1, 1)$ to the diagonal line $|C_0 - C_1| = 0$, intersecting at some $\boldsymbol{C} = (C, C)$. Therefore, this correlation can be written as $\boldsymbol{k} = \kappa(-1, 1) + (1-\kappa)\boldsymbol{C}$, for some $\kappa \in [0, 1]$. From this, we know $\kappa = \frac{1}{2}(k_1 - k_0) \leq \frac{2\mathfrak{E}\Delta t}{\pi}$. Define $\Lambda = \{1, 2, 3\}$, $\mathbf{C}^1 = (-1, 1)$, $\mathbf{C}^2 = (-1, -1)$, and $\mathbf{C}^3 = (+1, +1)$. We have just shown that $\mathbf{k}$ can be written as a convex combination of these three correlations, where the weight $p(1)$ of $\mathbf{C}^1$ is at most $\frac{2\mathfrak{E}\Delta t}{\pi}$. Moreover, $\mathbf{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}^\lambda,\Delta t}$ for $\mathfrak{E}^1 = \frac{\pi}{2\Delta t}$ and $\mathfrak{E}^2 = \mathfrak{E}^3 = 0$. Thus, the energy constraint in definition (11) is satisfied:

$$\sum_\lambda p(\lambda)\mathfrak{E}^\lambda = p(1)\mathfrak{E}^1 \leq \frac{2\mathfrak{E}\Delta t}{\pi} \cdot \frac{\pi}{2\Delta t} = \mathfrak{E}.$$

This shows that $\mathbf{k} \in \overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$. The case $k_0 \geq k_1$ can be treated analogously swapping the extremal point $(1, -1)$ for $(-1, 1)$. Thus, Eq. (13) characterises the $\overline{\mathcal{C}}_{E,\Delta t}$ precisely.

For an illustration of the classical max-average set $\bar{\mathcal{C}}_{\mathfrak{E},\Delta t}$ see Fig. 3. Interestingly, it is *not* in general a strict subset of the quantum set, since ΔE is concave under mixtures rather than linear, as for observables. In particular, for small $\mathfrak{E}\Delta t$, there are correlations that can be modelled classically under the max-average assumption that are not predicted by quantum theory under the stricter $\mathfrak{E}$ constraint.

Hence, if we have an upper bound $\mathfrak{E}$ on the energy uncertainty of the initial state, and if our scenario generates some correlation $\mathbf{C} \in \mathcal{Q}_{\mathfrak{E},\Delta t} \cap \bar{\mathcal{C}}_{\mathfrak{E},\Delta t}$, then it is possible that this correlation comes from an ensemble of deterministic correlations, which makes the outcome b potentially predictable by an adversary. But importantly for our protocol, for all $0 < \mathfrak{E}\Delta t < \frac{\pi}{2}$, there are correlations consistent with $\mathcal{Q}_{\mathfrak{E},\Delta t}$ that are outside of $\overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$. That is, such correlations are inconsistent with a deterministic model, even when we can only measure the *average* value of the energy uncertainty, rather than assuming its validity for every value of λ separately. This allows us to certify randomness, the quantification of which we discuss in Section III.

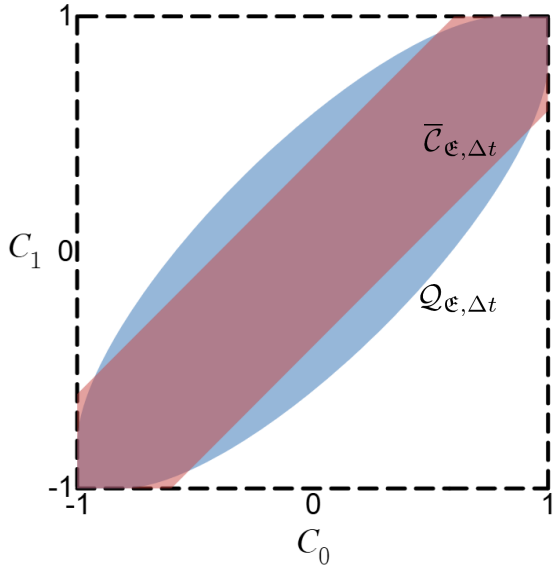

FIG. 3. Set of quantum correlations $\mathcal{Q}_{\mathfrak{E},\Delta t}$ (blue) and classical max-average correlations $\bar{\mathcal{C}}_{\mathfrak{E},\Delta t}$ (red) for $\mathfrak{E}\Delta t = 0.314$ (and thus $\gamma = 0.951$ according to Eq. (9)).

Open system evolution

So far we have formulated our semi-DI assumption for closed system communication, under unitary evolution. However, we can consider a natural extension to open systems, in which the inevitable environmental effects are taken into account. In particular, [31] derive a speed limit for open systems S coupled to general environments E via an arbitrary coupling Hamiltonian

$$\hat{H} = \hat{H}_S + \hat{H}_{SE} + \hat{H}_E = \hat{H}'_S + \hat{H}_E,$$

where $\hat{H}'_S = \hat{H}_S + \hat{H}_{SE}$ describes the Hamiltonian of S together with its coupling to the environment. If ρ_0 is the quantum state of S at time 0 and ρ_1 at time Δt, then they show that

$$\Delta t \geq \frac{\arccos \mathcal{F}(\rho_0, \rho_1)}{\langle \Delta E \rangle_{\Delta t}}, \tag{14}$$

where $\mathcal{F}(\rho_0, \rho_1) = \mathrm{tr}\sqrt{\rho_0^{1/2}\rho_1\rho_0^{1/2}}$ is the Uhlmann fidelity, and $\langle \Delta E \rangle_{\Delta t} = \frac{1}{\Delta t}\int_0^{\Delta t} dt\, \Delta E(t)$ is the time average of the energy uncertainty $\Delta E(t) = \sqrt{\langle \hat{H}'^2_S \rangle_{\rho(t)} - (\langle \hat{H}'_S \rangle_{\rho(t)})^2}$,

with $\rho(t)$ the reduced state of S at time t. For pure states $\rho_k = |\psi_k\rangle\langle\psi_k|$, the fidelity is $\mathcal{F}(\rho_0, \rho_1) = |\langle\psi_0|\psi_1\rangle|$, and for $\hat{H} = \hat{H}'_S$, this bound reduces to the Mandelstam-Tamm bound (3).

In cases where S is "small" compared to its environment E, the authors of [31] show that their bound (14) is a significant improvement of the original Mandelstam-Tamm bound (which, indeed, becomes trivial if the environment is very large). Moreover, as we will now show, our results above continue to hold unchanged if the assumption $\Delta E_{\rho_0} \leq \mathfrak{E}$ is replaced by an assumption $\langle\Delta E\rangle_{\Delta t} \leq \mathfrak{E}$ for the time-averaged open-system energy uncertainty. In some sense, rather than having to assume something about the system's energy uncertainty *within the whole universe*, we now only have to assume a bound on its energy uncertainty *within its vicinity of influence*.

To see this, we will first show that $\mathcal{F}(\rho_0, \rho_1) \geq \gamma$ implies that Eq. (9) holds for all measurement procedures, i.e. all $-\mathbb{1} \leq M \leq \mathbb{1}$ that generate the correlations $C_x = \mathrm{tr}(\rho_x M)$. This can be seen as follows. According to Uhlmann's theorem, there exist purifications $|\psi_i\rangle_{SA}$ of ρ_i $(i = 0, 1)$, where A is some ancillary system, such that $\mathcal{F}(\rho_0, \rho_1) = |\langle\psi_0|\psi_1\rangle_{SA}|$. But then, Eq. (9) applies whenever $C_x = \langle\psi_x|M'_{SA}|\psi_x\rangle_{SA}$, where $-\mathbb{1} \leq M'_{SA} \leq \mathbb{1}$. But this is true in particular for $M'_{SA} = M_S \otimes \mathbf{1}_A$, where $C_x = \langle\psi_x|M_S \otimes \mathbf{1}_A|\psi_x\rangle_{SA} = \mathrm{tr}(\rho_x M)$.

From Eq. (14), we have $\mathcal{F}(\rho_0, \rho_1) \geq \gamma$, where

$$\gamma = \begin{cases} \cos\left(\langle\Delta E\rangle_{\Delta t}\Delta t\right) & \text{if } |\langle\Delta E\rangle_{\Delta t}\Delta t| < \frac{\pi}{2} \\ 0 & \text{otherwise.} \end{cases}$$

Thus the possible quantum correlations obtainable under open-system evolution as described above under the assumption $\langle\Delta E\rangle_{\Delta t} \leq \mathfrak{E}$ are exactly given by $\mathcal{Q}_{\mathfrak{E},\Delta t}$. All further arguments, including the randomness certification results below, continue to hold without changes.

For a possibly even tighter quantum bound, Ref. [32] proves a general constraint for arbitrary physical processes in terms of the quantum Fisher information $\mathcal{F}_Q(t)$. Whilst their inequality serves as a tighter bound in general on the state transformation, the quantum Fisher information is less easy to directly determine for time-dependent evolution [31, 32]. Therefore the formulation of [31] in terms of the energy uncertainty $\langle E_S\rangle_\tau$ is more readily suitable for our operationally motivated scenario.

III. RANDOMNESS CERTIFICATION

For correlations $\boldsymbol{C} \in \mathcal{Q}_{\mathfrak{E},\Delta t} \backslash \overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$, there is no classical model that could reproduce correlations predicted by quantum theory, even assuming the energy constraint is only respected on average. Therefore, observing such correlations certifies randomness in that they have no deterministic description, and can be used for the generation of random numbers. The amount of certified randomness can be quantified adapting a method of [28].

Imagine that an adversary (Eve) wishes to guess the output of b. Like the experimenter (Alice), she may have knowledge of the inputs x, but (unlike Alice) she may also have access to additional information about the scenario. We may even consider that she knows all additional physical parameters, characterised by λ, that determine the behaviour of the device. We could imagine in some situations that some physical parameter, of which Eve knows the value, is correlated with the outcome b, allowing her a predictive advantage over the experimenter. We wish to prove for our scenario that, irrespective of knowing any such parameters, the outcome of the experiment is random, even to Eve.

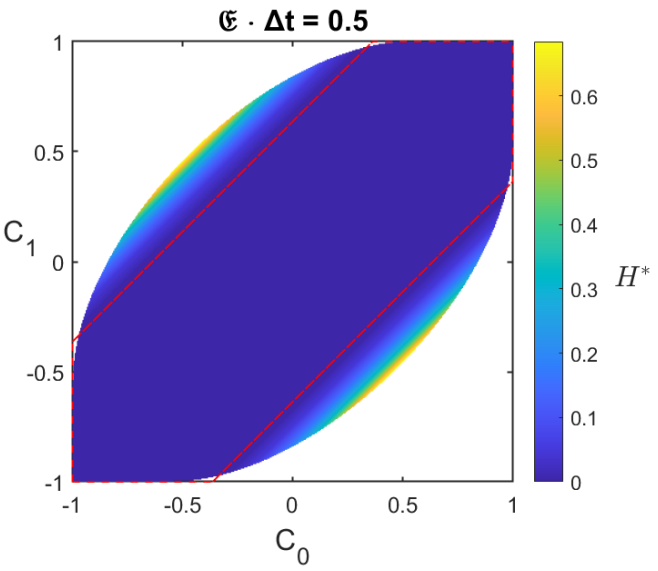

FIG. 4. Numerical estimates of the certifiable entropy H^* under the assumption $\Delta E \leq \mathfrak{E}$ with $\mathfrak{E}\Delta t = 0.5$ in units where $\hbar = 1$. Points with $H^* > 0$ correspond to correlations that allow for certified randomness extraction. As expected, $H^* = 0$ (only) on the red-rimmed classical max-average set $\bar{\mathcal{C}}_{\mathfrak{E},\Delta t}$. The maximal numerical value of the entropy we observe is $H^* = 0.68$. See Appendix E for details on the methodology we use to evaluate (15).

For Eve, the scenario is characterised by the ensemble $\{p(\lambda), (\boldsymbol{C}^\lambda, \mathfrak{E}^\lambda)\}$. The correlations as viewed by the experimenter can be more precisely characterised according to Eve's knowledge, by $\boldsymbol{C} = \sum_\lambda p(\lambda)\boldsymbol{C}^\lambda$. Nevertheless, the experimenter has already checked that the energy uncertainty is on average bounded by $\mathfrak{E}$, therefore the constraint $\sum_\lambda p(\lambda)\mathfrak{E}^\lambda \leq \mathfrak{E}$ applies, which follows from the concavity of variances (see Appendix C). For Alice to be sure that Eve cannot predict the outcome b reliably, she can quantify the randomness via e.g. the conditional Shannon entropy $H(B|X,\Lambda) = -\sum_{b,x,\lambda} p(b,x,\lambda)\log_2 p(b|x,\lambda)$. This quantifies the difficulty for Eve to predict b, given x and λ. Provided the inputs are independent of λ, the conditional entropy can be written as

$$H(B|X,\Lambda) = \sum_\lambda p(\lambda) H(\mathbf{C}^\lambda),$$

where

$$H(\boldsymbol{C}) := -\frac{1}{2}\sum_{b,x}\frac{1+bC_x}{2}\log\frac{1+bC_x}{2}.$$

This can be determined by the optimisation problem $H(B|X,\Lambda) \geq H^\star$, where

$$\begin{aligned} H^\star \;=\; & \min_{\{p(\lambda),\boldsymbol{C}^\lambda,\mathfrak{E}^\lambda\}} \sum_\lambda p(\lambda) H(\boldsymbol{C}^\lambda) && \text{(15a)}\\ & \text{s.t. } \sum_\lambda p(\lambda)\mathbf{C}^\lambda = \boldsymbol{C}, && \text{(15b)}\\ & \sum_\lambda p(\lambda)\mathfrak{E}^\lambda = \mathfrak{E}, && \text{(15c)}\\ & \mathbf{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}^\lambda,\Delta t}. && \text{(15d)} \end{aligned}$$

As in [28, Eq. 26], we take the condition in (15c) to be equality without loss of generality; see Appendix D. In contrast to their setting, however, our feasible set defined by the constraint (15d) is non-convex when ranging over all possible $\mathfrak{E}^\lambda$. Despite this, we adopt a similar strategy by formulating the Lagrange dual of the problem, which yields certified lower bounds on the entropy H^*. This dual approach enables us to work directly with our non-convex and non-linear constraints while still ensuring that any feasible dual solution provides a valid bound for randomness certification. See Appendix E for full details of such dual formulation and the numerical method used to evaluate it. In Figure 4, we plot the resulting estimates of H^* over the space of correlation pairs, for a fixed average energy uncertainty constraint $\mathfrak{E} = 1/2$. Any value $H^* > 0$ can be used to extract a semi-device-independent certificate of randomness. As expected, points lying within the classical average set $\bar{\mathcal{C}}$ yield $H^* = 0$ (and vice versa), indicating that no randomness can be certified against an adversary in that region.

IV. IMPLEMENTATION WITH COHERENT STATES

In this section, we show that our protocol can in principle be performed with coherent states of a single harmonic oscillator, and we sketch a quantum-optical implementation.

Consider a situation in which the initial state is given by $|\psi_0\rangle = |\alpha\rangle$ a single-mode coherent state, described in its Fock basis by

$$|\alpha\rangle = e^{-\frac{|\alpha|^2}{2}}\sum_{n=0}^{\infty}\frac{\alpha^n}{\sqrt{n!}}|n\rangle\,. \tag{16}$$

The state evolves in time as a harmonic oscillator $\hat{H} = \hbar\omega(\hat{a}^\dagger\hat{a} + \frac{1}{2})$, and we use the quantum optics convention where $\hat{a} = \frac{1}{\sqrt{2}}(\hat{q} + i\hat{p})$ and $[\hat{a}, \hat{a}^\dagger] = \mathbf{1}$ such that $\hat{H} = \frac{\hbar\omega}{2}(\hat{q}^2+\hat{p}^2)$ (an analogous but more cumbersome calculation for the mechanical oscillator $\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega^2\hat{q}^2$ leads to the same result for the correlations $\mathbf{C} = (C_0, C_1)$ as below). The coherent state evolves in time by rotating in phase space,

$$e^{-i\hat{H}t}|\alpha\rangle = e^{-i\omega t/2}|\alpha(t)\rangle, \tag{17}$$

where $\alpha(t) = e^{-i\omega t}\alpha$. Up to an irrelevant global phase $\theta \in \mathbb{R}$, its representation in the quadrature q basis is given by

$$\langle q|\alpha\rangle = e^{i\theta}\pi^{-1/4}\exp\left[-\frac{1}{2}\left(q - \sqrt{2}\,\mathrm{Re}(\alpha)\right)^2\right],$$

where $|q\rangle$ is an eigenstate of the quadrature operator $\hat{q} = \frac{1}{\sqrt{2}}(\hat{a} + \hat{a}^\dagger)$. This expression shows that the coherent state is a Gaussian wavepacket centered at $\sqrt{2}\,\mathrm{Re}(\alpha)$ in the quadrature q.

At the measurement device M, we follow a protocol similar to the BPSK example of [27, Section 2.3.1]. In particular, M implements a quadrature measurement and assigns the binary output $b = \mathrm{sign}(q)$, i.e.

$$b = \begin{cases} +1 & \text{if } q > 0, \\ -1 & \text{if } q < 0. \end{cases} \tag{18}$$

The measurement outcome is sampled from

$$|\langle q|\alpha(t)\rangle|^2 = \pi^{-1/2}\exp\left[-\left(q - \sqrt{2}\,\mathrm{Re}(\alpha e^{-i\omega t})\right)^2\right].$$

That is, $b = b(t)$ is time-dependent. This leaves room for there to be interesting correlations from our scenario. In particular, the correlations are given by

$$\begin{aligned} C(t) \;\equiv\; \mathbb{E}[b(t)] &= \int_{-\infty}^{\infty}\mathrm{sign}(q)|\langle q|\alpha(t)\rangle|^2 dq \\ &= \mathrm{erf}(\sqrt{2}\mathrm{Re}(\alpha e^{-i\omega t})), \end{aligned} \tag{19}$$

where the final line uses the well-known result $\mathbb{E}[\mathrm{sign}(x)] = \mathrm{erf}(\mu/\sqrt{2}\sigma)$ for a Gaussian distribution centred at μ with standard deviation σ (here, we have $\sigma = 1/\sqrt{2}$). Then the correlations (C_0, C_1) are given by

$$\begin{aligned} C_0 &= C(t=0) = \mathrm{erf}(\sqrt{2}\mathrm{Re}(\alpha)), \\ C_1 &= C(t=\Delta t) = \mathrm{erf}(\sqrt{2}\mathrm{Re}(\alpha e^{-i\omega\Delta t})). \end{aligned} \tag{20}$$

The energy expectation is given by

$$\langle\hat{H}\rangle_{\alpha(t)} = \hbar\omega\left(|\alpha(t)|^2 + \frac{1}{2}\right) = \hbar\omega\left(|\alpha|^2 + \frac{1}{2}\right),$$

and the squared energy:

$$\langle\hat{H}^2\rangle_{\alpha(t)} = \hbar^2\omega^2\left(|\alpha(t)|^4 + 2|\alpha(t)|^2 + \frac{1}{4}\right).$$

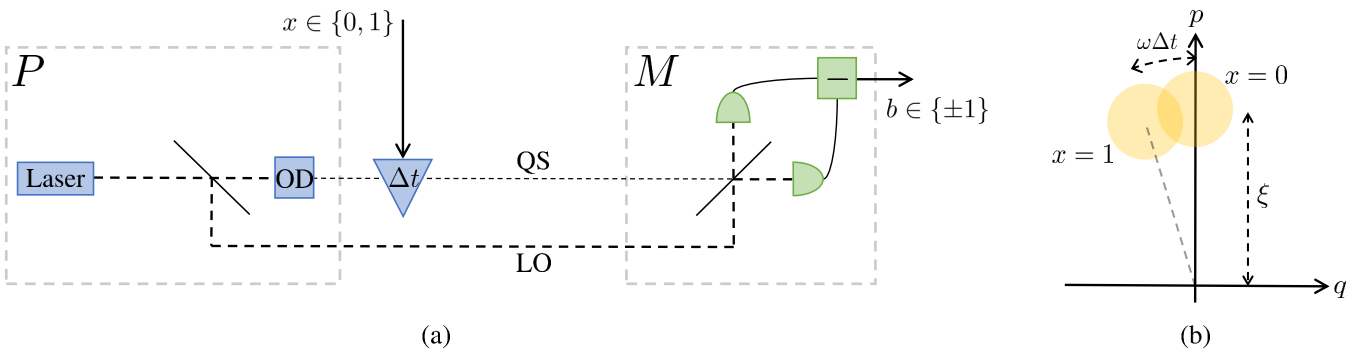

FIG. 5. (a) Diagram of proposed experimental setup; technically similar to, but conceptually distinct from the BPSK implementation of [28]. The preparation device P consists of a monochromatic laser emitting a coherent pulse, which is sent through a beam splitter for homodyne detection. The transmitted beam (the quantum signal, QS) is attenuated by an optical density filter (OD) and then, for the input $x = 1$, delayed by some time Δt. This can be achieved using an optical wedge, whose position and thickness depends on the input x. By causing a time delay, the wedge introduces a corresponding phase shift. At the measurement device M, the signal is interfered with the reflected signal (the local oscillator, LO), which acts as a phase reference, and is measured. The quadrature q is then measured, the sign of which determines the output b. (b) Diagram of the phase space of the coherent state, showing the relative phase $\omega\Delta t$ between the two states that may arrive at M. For $x = 0$, the state $|\alpha\rangle = |i\xi\rangle$ is centred on the y-axis (in the rotating frame of reference of the local oscillator), evolving counter-clockwise with period $2\pi/\omega$. For $x = 1$, the state is given by $|\alpha(t = \Delta t)\rangle = \left|i\xi e^{-i\omega\Delta t}\right\rangle$.

The energy uncertainty is thus given by

$$\Delta E_{\alpha(t)} = \sqrt{\langle \hat{H}^2\rangle - \langle \hat{H}\rangle^2} = \hbar\omega|\alpha(t)| = \hbar\omega|\alpha|.$$

In order to realise interesting quantum correlations that certify randomness, suppose we begin with an initial state $|\alpha\rangle = |i\xi\rangle$ where $\xi \in \mathbb{R}$. In this case, the energy uncertainty is given by $\Delta E_\alpha = \hbar\omega\xi$. For the correlation function, from Eq. (20), we get

$$C_0 = 0, \tag{21}$$

$$C_1 = \mathrm{erf}(\sqrt{2}\xi \sin(\omega\Delta t)). \tag{22}$$

To summarise, in order to obtain interesting correlations that certify randomness from coherent states, we want values of $\boldsymbol{C} = (0, \mathrm{erf}(\sqrt{2}\xi \sin\omega\Delta t))$ that are in the set $\mathcal{Q}_{\mathfrak{E},\Delta t}$ but outside $\overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$, where $\hbar\omega\xi \leq \mathfrak{E}$. From the characterisation of the max-average classical set of equation (13), with units $\hbar$ put back in, we get the following condition for non-zero entropy:

$$\left|\mathrm{erf}\left(\sqrt{2}\xi \sin(\omega\Delta t)\right)\right| > \frac{4\mathfrak{E}\Delta t}{\pi\hbar}, \tag{23}$$

for $\mathfrak{E}\Delta t/\hbar < \pi/2$ and $\hbar\omega\xi \leq \mathfrak{E}$.

As Fig. 6 shows, there are indeed choices of frequency ω, time delay Δt and of assumed upper bounds $\mathfrak{E}$ to the energy uncertainty such that a non-zero amount of randomness can be certified with a standard coherent state of a simple harmonic oscillator. An example is for $\hbar\omega\xi = \mathfrak{E} = 0.5$, $\omega\Delta t = 0.4$ (i.e. the point $(0.4, 0.2)$ in Fig. 6), for which we obtain $C_1 = 0.303$. This places $\boldsymbol{C} \in \mathcal{Q}_{\mathfrak{E},\Delta t}\backslash\overline{\mathcal{C}}_{\mathfrak{E},\Delta t}$, and yields a certified entropy $H^* = 0.0242$.

In Fig. 5, we sketch a possible quantum-optical implementation, see the figure caption for a description. Note that the local oscillator can in principle be replaced by a second laser, and then only the upper path corresponds to the physical system that is prepared and measured. Since the implementation of the time delay relies on the specific physics of light propagation in materials, it does not accomplish the device-independence of the anticipated "trusted operation on an untrusted device" depicted in Fig. 1. However, assuming that a photon state has been prepared, and that the optical element works as desired, it admits the certification of randomness independently of the actually prepared photon state (and of the actually performed measurement).

There is an interesting conceptual relation to work by Tsirelson [33], which has recently received renewed attention (see e.g. [34]). Tsirelson has shown that a simple quantum harmonic oscillator can behave probabilistically in a way that a classical harmonic oscillator cannot. In more detail, asking whether the position q of the oscillator is positive (which is exactly the measurement in our protocol above) after zero, one-third or two-thirds of a period yields probabilities between 1/3 and 2/3 for classical distributions in phase space, whereas some states of the quantum harmonic oscillator (necessarily with negativity in the Wigner function) give a probability of 0.71. Analogously, our result shows a nonclassical effect of the quantum harmonic oscillator in some sense, but for standard coherent states with a positive Wigner function.

V. DISCUSSION & OUTLOOK

In this work we have demonstrated the applicability of the quantum speed limit for performing tasks in quantum information – namely, for generating certifiably secure random numbers. We have described a simple prepare-

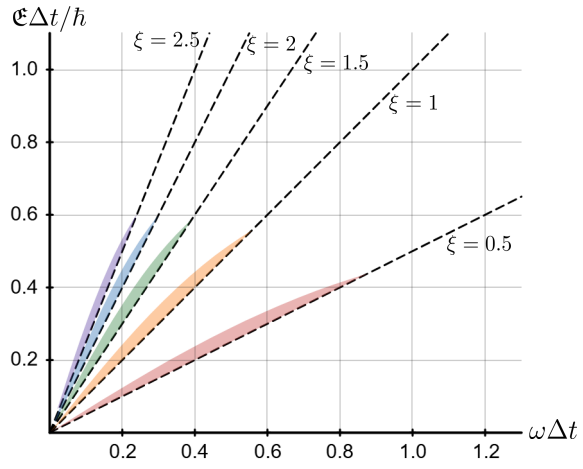

FIG. 6. Regions for which non-zero randomness can be certified according to Eq. (23), for $\xi \in \{0.5, 1, 1.5, 2, 2.5\}$ (red, yellow, green, blue, purple), where $\xi = |\alpha|$ is related to the average photon number $\langle n \rangle = \xi^2$ in the coherent state $|\alpha\rangle$. Dashed lines represent the lines $\hbar\omega\xi = \mathfrak{E}$.

and-measure scenario, where the inputs correspond to two time-displaced preparations of a transmitted system. This communicated system is constrained only in terms of some upper bound on its energy uncertainty. We have characterised the possible quantum correlations of the scenario, for pure state, mixed state and open systems, and shown that quantum theory predicts statistics that are incompatible with any deterministic explanation – even for an adversary who has knowledge of the inputs and who has access to extra information about the communication between devices. Moreover, we have provided a numerical estimate for the amount of certifiable entropy $H^\star$ under our energy and time assumptions, and sketched an experimental implementation involving coherent states that yields non-zero certified entropy. It is interesting particularly to note that some single-mode coherent states admit the certification of genuine randomness, which bears a close resemblance to an observation by Tsirelson [33, 34] indicating some form of nonclassical behavior of the simple harmonic oscillator.

Our protocol contributes to ongoing efforts to replace the traditional dimensionality assumption by physically better motivated alternatives, building on the work of e.g. [19–28]. The security of our protocol is grounded in reliable upper bounds on the energy uncertainty and the time delay between preparations. The former could be verified "from the outside", in a similar spirit as proposed by [27]; by performing many tests on the state ρ_x emitted by P, one could do a statistical analysis to determine the energy uncertainty. The latter is imposed manifestly by the nature of the experiment; the device time-displaces the two possible states by Δt, by virtue of implementing a time delay during one of the two preparations. Ideally, we would consider that the experimenter themselves waits some time between preparation times, but the nano-timescale prohibits a fully operational implementation of this assumption. Nevertheless, we can conceive of this as a trusted super-operation on an otherwise uncharacterised preparation box.

The consideration of semi-DI protocols via trusted operations in spacetime on untrusted devices was initiated in [29]. There, the input for a prepare-and-measure scenario is given by some rotation of the preparation device around a fixed axis by some angle α, while assuming an upper bound to the spin J of the transmitted physical system. For small enough angles, the gap between quantum and classical (deterministic) sets of correlations can be used to generate secure random numbers. Moreover, it has been shown that the set of quantum correlations in this setup can be recovered even without assuming quantum theory. In a similar spirit to the "spacetimes boxes" framework of [29, 30], there may be scope to extend the results of this paper to a theory-independent setting. There has been some interesting work in this direction by other authors [35], who derive a speed limit from purely information-theoretic principles. This could motivate the generalisation of the results of this paper beyond quantum theory, or even within a framework guided only by spatiotemporal considerations. If this were possible, one could look to ground quantum information protocols, such as the certification of randomness, on assumptions about time translation symmetry alone.

ACKNOWLEDGMENTS

We acknowledge support from the Austrian Science Fund (FWF) via project P 33730-N. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science, and Economic Development, and by the Province of Ontario through the Ministry of Colleges and Universities.

[1] I. Born (ed.), Letter to Max Born (4 December 1926), in *The Born-Einstein Letters*, Walker and Company, New York, 1971.

[2] W. Heisenberg, *Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik*, Z. Phys. **43**, 172 (1927).

[3] W. Heisenberg, *Die physikalischen Prinzipien der Quantumtheorie*, Hirzel, Stuttgart, Germany, 5th edition (1944).

[4] N. Bohr, *The quantum postulate and recent developments of atomic theory*, Nature **121**, 580 (1928).

[5] L. Mandelstam, I. Tamm, *The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics*, J. Phys. (USSR) **9**, 249 (1945).

[6] J. Anandan, and Y. Aharonov, *Geometry of quantum evolution*, Phys. Rev. Lett. **65**(14), 1697 (1990).

[7] J. D. Bekenstein, *Energy Cost of Information Transfer*, Phys. Rev. Lett. **46**, 623 (1981).

[8] S. Lloyd, *Ultimate physical limits to computation*, Nature **406**, 1047 (2000).

[9] S. Deffner and E. Lutz, *Generalized Clausius inequality for nonequilibrium quantum processes*, Phys. Rev. Lett. **105**, 170402 (2010).

[10] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, *Optimal control at the quantum speed limit*, Phys. Rev. Lett. **103**, 240501 (2009).

[11] V. Giovannetti, S. Lloyd, and L. Maccone, *Advances in quantum metrology*, Nat. Photonics **5**, 222–229 (2011).

[12] R. Gallego, N. Brunner, C. Hadley, and A. Acín, *Device-independent tests of classical and quantum dimensions*, Phys. Rev. Lett. **105**, 230501 (2010).

[13] J. Bowles, M. T. Quintino, and N. Brunner, *Certifying the dimension of classical and quantum systems in a prepare-and-measure scenario with independent devices*, Phys. Rev. Lett. **112**, 140407 (2014).

[14] N. Brunner, S. Pironio, A. Acín, N. Gisin, A. A. Méthot, and V. Scarani, *Testing the dimension of Hilbert spaces*, Phys. Rev. Lett. **100**, 210503 (2008).

[15] M. Pawłowski and N. Brunner, *Semi-device-independent securuty of one-way quantum key distribution*, Phys. Rev. A **84**, 010302 (2011).

[16] H.-W. Li, Y. Yhen-Qiang, Y.-C. Wu, X.-B. Zou, S. Wang, W. Chen, G.-C. Guo, and Z.-F. Han, *Semi-device-independent random-number expansion without entanglement*, Phys. Rev. A **84**, 034301 (2011).

[17] H.-W. Li, M. Pawłowski, Z.-Q. Yin, G.-C. Guo, and Z.-F. Han, *Semi-device-independent randomness certification using $n \to 1$ quantum random access codes*, Phys. Rev. A **85**, 052308 (2012).

[18] P. Mironowicz, A. Tavakoli, A. Hameedi, B. Marques, M. Pawłowski, and M. Bourennane, *Increased certification of semi-device independent random numbers using many inputs and more post-processing*, New J. Phys. **18**, 065004 (2016).

[19] J. Bohr Brask, A. Martin, W. Esposito, R. Houlmann, J. Bowles, H. Zbinden, and N. Brunner, *Megahertz-Rate Semi-Device-Independent Quantum Random Number Generators Based on Unambiguous State Discrimination*, Phys. Rev. Appl. **7**, 054018 (2017).

[20] H. Tebyanian, M. Zahidy, M. Avesani, A. Stanco, P. Villoresi, and G. Vallone, *Semi-device independent randomness generation based on quantum state's indistinguishability*, Quantum Sci. Technol. **6**, 045026 (2021).

[21] C. Roch i Carceller, L. Nunes Faria, Z.-H. Liu, N. Sguerso, U. L. Andersen, J. S. Nergaard-Nielsen, and J. B. Brask, *Improving semi-device-independent randomness certification by entropy accumulation*, arXiv:2405.04244.

[22] M. Ioannou, P. Sekatski, A. A. Abbott, D. Rosset, J.-D. Bancal, and N. Brunner, *Receiver-device-independent quantum key distribution protocols*, New J. Phys. **24**, 063006 (2022).

[23] A. Tavakoli, *Semi-Device-Independent Framework Based on Restricted Distrust in Prepare-and-Measure Experiments*, Phys. Rev. Lett. **126**, 210503 (2021).

[24] A. Tavakoli, E. Z. Cruzeiro, E. Woodhead, and S. Pironio, *Informationally restricted correlations: A general framework for classical and quantum systems*, Quantum **6**, 620 (2022).

[25] A. Tavakoli, E. Z. Cruzeiro, J. B. Brask, N. Gisin, and N. Brunner, *Informationally restricted correlations*, Quantum **4**, 332 (2020).

[26] J. Pauwels, S. Pironio, and A. Tavakoli, *Information capacity of quantum communication under natural physical assumptions*, Quantum **9**, 1637 (2025).

[27] T. Van Himbeeck, E. Woodhead, N. J. Cerf, R. García-Patrón, and S. Pironio, *Semi-device-independent framework based on natural physical assumptions*, Quantum **1**, 33 (2017).

[28] T. van Himbeeck and S. Pironio, *Correlations and randomness generation based on energy constraints*, arXiv:1905.09117.

[29] C. L. Jones, S. L. Ludescher, A. Aloy, and M. P. Müller, *Theory-independent randomness generation with spacetime symmetries*, arXiv:2210.14811.

[30] A. Aloy, T. D. Galley, C. L. Jones, S. L. Ludescher, and M. P. Müller, *Spin-bounded correlations: rotation boxes within and beyond quantum theory*, Commun. Math. Phys. **405**(12), 292 (2024).

[31] N. Shiraishi, and K. Saito, *Speed limit for open systems coupled to general environments*. Phys. Rev. Research **3**(2), 023074 (2021).

[32] M. M. Taddei, B. M. Escher, L. Davidovich, and R. L. de Matos Filho, *Quantum Speed Limit for Physical Processes*, Phys. Rev. Lett. **110**, 050402 (2013).

[33] B. Tsirelson, *How often is the coordinate of a harmonic oscillator positive?*, arXiv:quant-ph/0611147.

[34] L. H. Zaw, M. Weilenmann, and V. Scarani, *Tsirelson's Inequality for the Precession Protocol Is Maximally Violated by Quantum Theory*, Phys. Rev. Lett. **134**, 190201 (2025).

[35] L. Giannelli and G. Chiribella, *Information-theoretic derivation of energy and speed bounds*, arXiv:2403.13223.

APPENDIX

A. MIXED STATES

We now consider the set of quantum correlations for mixed states, which we prove is equal to that of pure states. For $\mathfrak{E} \geq 0$ and $\Delta t \geq 0$, we define the following:

$$\mathcal{Q}'_{\mathfrak{E},\Delta t} := \Big\{ (C_0, C_1) \,\Big|\, C_x = \mathrm{tr}[M\rho_x], -\mathbb{1} \leq M \leq \mathbb{1}, \\ \exists \hat{H} \text{ s. t. } \rho_1 = U_{\Delta t}^{\dagger} \rho_0 U_{\Delta t},\, \Delta E_{\rho_0} \leq \mathfrak{E} \Big\}, \quad (24)$$

which we show to be equal to $\mathcal{Q}_{\mathfrak{E},\Delta t}$. To do so, we consider the purification of mixed states, using ancilla systems, and show that the purification procedure does not change the energy variance. In particular, consider a mixed state ρ_0^A acting on $\mathcal{H}_A$, which is diagonalised as $\rho_0^A = \sum_{k^A} \tilde{c}_k \left|k^A\right\rangle \left\langle k^A\right|$. To characterise its purification, we introduce a fictitious system B with Hilbert space $\mathcal{H}_B$ such that $\dim(\mathcal{H}_A) = \dim(\mathcal{H}_B)$. This defines a pure state

$|\psi_{AB}\rangle = \sum_k \sqrt{\tilde{c}_k}\,|k^A k^B\rangle$, evolving under the Hamiltonian $\hat{H}_{AB} := \hat{H}_A + \hat{0}_B$ via $U_{\Delta t} \otimes \mathbb{1}_B$. This embedding does not affect the energy variance. In particular,

$$\langle \hat{H}^2_{AB}\rangle_{\psi_{AB}} = \langle\psi_{AB}|\, \hat{H}^2_A \otimes \mathbb{1}_B\, |\psi_{AB}\rangle = \mathrm{tr}\Big[\rho_A \hat{H}^2_A\Big] = \langle \hat{H}^2_A\rangle_{\rho_A},$$

and likewise for $\langle \hat{H}_{AB}\rangle_{\psi_{AB}}$. Therefore, $\Delta E_{\psi_{AB}} = \Delta E_{\rho_A}$. We can think of the communication of a mixed state ρ_A as being realised by the communication of a pure state $|\psi\rangle_{AB}$ (which is therefore already contained within the characterisation $\mathcal{Q}_{\mathfrak{E},\Delta t}$), but such that the measurement device does not "listen" to the subsystem B. This means that we embed the POVM $\{M_+, M_-\}$ on A via $\{M_+ \otimes \mathbb{1}_B, M_- \otimes \mathbb{1}_B\}$ on AB, reproducing the original correlation via a pure state on the larger system AB.

Accordingly, we have proven the equality

$$\mathcal{Q}'_{\mathfrak{E},\Delta t} = \mathcal{Q}_{\mathfrak{E},\Delta t}. \tag{25}$$

That is, the set of quantum correlations is unchanged by including mixed states. This also implies that $\mathcal{Q}_{\mathfrak{E},\Delta t}$ is convex.

B. QUANTUM MODEL FOR ALL CORRELATIONS OF (8)

Having shown that all correlations in $\mathcal{Q}_{\mathfrak{E},\Delta t}$ are constrained by inequality (8), we now prove the tightness of this bound by providing a quantum model for all correlations that satisfy (8), that is:

$$\frac{1}{2}\left(\sqrt{1+C_0}\sqrt{1+C_1} + \sqrt{1-C_0}\sqrt{1-C_1}\right) \geq \gamma.$$

To do so, we follow a similar approach to that of [29], in which we have already provided an analogous analysis for the sets called $\mathcal{Q}_{J,\alpha}$. For these purposes, it is useful to switch description from correlations $\boldsymbol{C} = (C_0, C_1)$ to probabilities $\boldsymbol{P}^+ = (P^+_0, P^+_1)$, where probabilities $P^+_x := P(+1|x)$ are related to the correlations by the bijective affine transformation $\boldsymbol{P}^+ = (\boldsymbol{C}+1)/2$. Following [29], we can write down a description for the inequality above in terms of the extremal points $(0,0)$, $(1,1)$, along with the following curves p_1 and p_2 parametrised by τ:

$$p_1(\tau) = \left(\cos^2(\mathfrak{E}\tau), \cos^2(\mathfrak{E}(\tau+\Delta t))\right),\ \tau \in \mathcal{I}_1, \tag{26}$$

$$p_2(\tau) = \left(\cos^2(\mathfrak{E}\tau), \cos^2(\mathfrak{E}(\tau-\Delta t))\right),\ \tau \in \mathcal{I}_2, \tag{27}$$

where $\mathcal{I}_1 = [0, \frac{\pi}{2\mathfrak{E}} - \Delta t]$ and $\mathcal{I}_2 = [\Delta t, \frac{\pi}{2\mathfrak{E}}]$.

We will now provide a quantum model for all of these extremal points. First, the points $(0,0)$ and $(1,1)$ are given by constant probability distributions, and so can be trivially modelled. For example, the latter can be reproduced by the state $|\psi_0\rangle = |\psi_1\rangle = |E\rangle$, for which $\Delta E = 0$, and the measurement operators $M_+ = |E\rangle\langle E|$ and $M_- = \mathbb{1} - M_+$.

Next we take the curve given by $p_1(\tau)$, which can be modelled (for instance) by the state

$$|\psi_0\rangle = \frac{|0\rangle + |2\mathfrak{E}\rangle}{\sqrt{2}}$$

for some Hamiltonian that has 0 and $2\mathfrak{E}$ among its energy eigenvalues, which has an energy uncertainty of $\Delta E = \mathfrak{E}$. This state evolves in a time Δt to the following:

$$\begin{aligned}|\psi_1\rangle &= U_{\Delta t}\left(\frac{|0\rangle + |2\mathfrak{E}\rangle}{\sqrt{2}}\right)\\ &= \frac{1}{\sqrt{2}}\left(|0\rangle + e^{-2i\mathfrak{E}\Delta t}\,|2\mathfrak{E}\rangle\right)\\ &= \frac{1}{\sqrt{2}} e^{-i\mathfrak{E}\Delta t}\left(e^{i\mathfrak{E}\Delta t}\,|0\rangle + e^{-i\mathfrak{E}\Delta t}\,|2\mathfrak{E}\rangle\right).\end{aligned}$$

Then we define the measurement operators $M_+ = U^\dagger_\tau\,|\psi_0\rangle\langle\psi_0|\,U_\tau$ and $M_- = \mathbb{1} - M_+$, for which a simple calculation shows that we achieve the required probabilities: $P^+_0 = \langle\psi_0|M_+|\psi_0\rangle = \cos^2(\mathfrak{E}\tau)$ and $P^+_1 = \langle\psi_1|M_+|\psi_1\rangle = \cos^2(\mathfrak{E}(\tau+\Delta t))$. The curve $p_2(\tau)$ can be modelled similarly, using the same state but the measurement operators $M_+ = U^\dagger_{-\tau}\,|\psi_0\rangle\langle\psi_0|\,U_{-\tau}$ and $M_- = \mathbb{1} - M_+$.

Having modelled all extreme points, we then use that $\mathcal{Q}_{\mathfrak{E},\Delta t}$ is convex to conclude that all correlations that obey inequality (8) have a quantum model, and thus it precisely characterises the quantum set $\mathcal{Q}_{E,\Delta t}$.

C. CONCAVITY OF ΔE

To show concavity of ΔE_ρ in $\rho = \lambda\rho_1 + (1-\lambda)\rho_2$, first we note convexity of $\langle\hat{H}\rangle^2_\rho$:

$$\begin{aligned}\langle\hat{H}\rangle^2_\rho &= (\lambda\langle\hat{H}\rangle_{\rho_1} + (1-\lambda)\langle\hat{H}\rangle_{\rho_2})^2\\ &\leq \lambda\langle\hat{H}\rangle^2_{\rho_1} + (1-\lambda)\langle\hat{H}\rangle^2_{\rho_2},\end{aligned}$$

due to the convexity of $(\cdot)^2$. We also have the equality

$$\lambda\langle\hat{H}^2\rangle_{\rho_1} + (1-\lambda)\langle\hat{H}^2\rangle_{\rho_2} = \langle\hat{H}^2\rangle_\rho.$$

Together with the concavity of the square root function, these show concavity of ΔE_ρ in ρ:

$$\begin{aligned}\Delta E_\rho &= \sqrt{\langle\hat{H}^2\rangle_\rho - \langle\hat{H}\rangle^2_\rho}\\ &\geq \sqrt{\lambda\langle\hat{H}^2\rangle_{\rho_1} + (1-\lambda)\langle\hat{H}^2\rangle_{\rho_2} - \lambda\langle\hat{H}\rangle^2_{\rho_1} - (1-\lambda)\langle\hat{H}\rangle^2_{\rho_2}}\\ &= \sqrt{\lambda(\langle\hat{H}^2\rangle_{\rho_1} - \langle\hat{H}\rangle^2_{\rho_1}) + (1-\lambda)(\langle\hat{H}^2\rangle_{\rho_2} - \langle\hat{H}\rangle^2_{\rho_2})}\\ &= \sqrt{\lambda(\Delta E_{\rho_1})^2 + (1-\lambda)(\Delta E_{\rho_2})^2}\\ &\geq \lambda\Delta E_{\rho_1} + (1-\lambda)\Delta E_{\rho_2}\,.\end{aligned}$$

This extends in an obvious way to more general convex combinations $\rho = \sum_\lambda p(\lambda)\rho_\lambda$.

D. EQUALITY (15c) OF THE OPTIMISATION PROBLEM

We need to show the following: if $\{p(\lambda), \mathbf{C}^\lambda, \mathfrak{E}^\lambda\}$ satisfies (15b), (15c') and (15d), where

$$\sum_\lambda p(\lambda)\mathfrak{E}^\lambda \leq \mathfrak{E}, \tag{15c'}$$

then there are $\mathfrak{E}'^\lambda \geq \mathfrak{E}^\lambda$ such that $\{p(\lambda), \mathbf{C}^\lambda, E'^\lambda\}$ satisfies (15b), (15c) and (15d). Consequently, we can replace (15c') by (15c) in the optimisation problem.

This can be seen as follows. Clearly, $\sum_\lambda p(\lambda)\mathfrak{E}'^\lambda$ is increasing in $\mathfrak{E}'^\lambda$, and it can in fact be made arbitrarily large. Thus, we can pick the $\mathfrak{E}'^\lambda$ such that $\sum_\lambda p(\lambda)\mathfrak{E}'^\lambda = \mathfrak{E}$, i.e. (15c) is satisfied. Since (15b) depends only on the correlations $\mathbf{C}^\lambda$, it is unchanged. Finally, we have

$$\mathfrak{E} \leq \mathfrak{E}' \Rightarrow \mathcal{Q}_{\mathfrak{E},\Delta t} \subseteq \mathcal{Q}_{\mathfrak{E}',\Delta t},$$

and so (15d) implies that $\mathbf{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}'^\lambda,\Delta t}$.

E. OPTIMIZATION FRAMEWORK FOR BOUNDING THE ENTROPY H^*

We are interested in bounding the amount of randomness (entropy H^*) that can be certified in our prepare-and-measure scenario from given observed correlations $\boldsymbol{C} = (C_0, C_1)$, under an average energy uncertainty constraint $\Delta E_{\rho_0} \leq \mathfrak{E}$ and evolution time Δt. The goal is to minimise the conditional entropy $H(B|X,\Lambda)$ over all compatible hidden-variable decompositions. Formally, the optimisation problem can be expressed as:

$$\begin{aligned} H^* := \min_{\{p(\lambda),\, \boldsymbol{C}^\lambda,\, \mathfrak{E}^\lambda\}} \quad & \sum_\lambda p(\lambda) H(\boldsymbol{C}^\lambda) \\ \text{subject to} \quad & \sum_\lambda p(\lambda)\boldsymbol{C}^\lambda = \boldsymbol{C}, \\ & \sum_\lambda p(\lambda)\mathfrak{E}^\lambda = \mathfrak{E}, \\ & \boldsymbol{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}^\lambda,\Delta t}, \end{aligned} \tag{28}$$

where $H(\boldsymbol{C}) = \sum_x p(x) h_{\text{bin}}(C_x)$, using the binary entropy function $h_{\text{bin}}(C_x) = -\sum_b \frac{1+bC_x}{2}\log_2\frac{1+bC_x}{2}$. Here we assume $p(x) = 1/2$ for both x, i.e. there is equal a priori probability for both inputs.

Finding an optimal solution requires, in principle, optimising over all possible ensembles $\{p(\lambda), \mathfrak{E}^\lambda, \boldsymbol{C}^\lambda\}$. Since the number of hidden variables $\lambda \in \Lambda$ is unbounded, this problem is generally intractable in its full form. To address this challenge, we follow a strategy similar to the one presented in [28] by formulating the dual problem. This reformulation allows us to obtain safe lower bounds on the entropy $H^*_{\text{dual}} \leq H^*$, even when the primal problem is not directly solvable. Moreover, we demonstrate that our problem exhibits the strong duality property, ensuring that the optimal dual solution coincides with the optimal primal solution, *i.e.*, $H^*_{\text{dual}} = H^*$.

Before deriving the dual, it is convenient to first express the primal problem as part of a broader class of non-linear constrained optimisation problems. Let f be a continuous function defined over the set of feasible solutions $\mathcal{Q} \subset \mathbb{R}^{\dim(\mathcal{Q})}$. The optimization problem can then be expressed as:

$$\begin{aligned} f^*(\boldsymbol{x}_0) = \min_{\{\boldsymbol{x}^\lambda, p(\lambda)\}} \quad & \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda) \\ \text{subject to} \quad & \sum_\lambda p(\lambda)\boldsymbol{x}^\lambda = \boldsymbol{x}_0, \\ & \boldsymbol{x}^\lambda \in \mathcal{Q}, \\ & p(\lambda) \in \mathcal{P}(\Lambda), \end{aligned} \tag{29}$$

where $\mathcal{P}(\lambda)$ denotes the set of all probability distributions over the λ. For our specific problem of Eq. (28), the feasible set $\mathcal{Q} \subset \mathbb{R}^3$ contains all possible correlations $\boldsymbol{C}^\lambda \in \mathcal{Q}_{\mathfrak{E}^\lambda,\Delta t}$ for all energy uncertainties within $0 \leq \mathfrak{E}^\lambda \leq \pi/(2\Delta t)$. Therefore, the optimization variables can be identified as $\boldsymbol{x}^\lambda = (\boldsymbol{C}^\lambda, \mathfrak{E}^\lambda) \in \mathcal{Q}$, while the given observed data and constraint are $\boldsymbol{x}_0 = (\boldsymbol{C}, \mathfrak{E})$.

In contrast to the case in [28], our feasible set $\mathcal{Q}$ is non-convex, as illustrated in Figure 7. Despite this non-convexity, it can be demonstrated that the dual formulation of the problem maintains strong duality, ensuring that the optimal solution obtained from the dual problem coincides with the optimal primal solution.

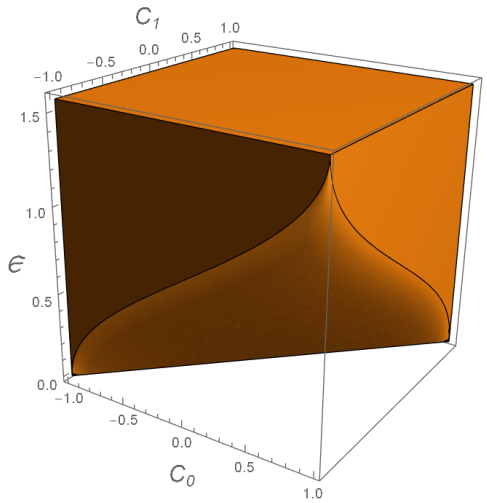

FIG. 7. The non-convex feasible set $\mathcal{Q}$ for the optimization problem 28, representing the set of allowed correlations and energy uncertainties used as optimization variables. Specifically, the elements of the set are of the form $(C_0^\lambda, C_1^\lambda, \mathfrak{E}^\lambda)$.

Dual formulation

To derive the dual formulation of the optimisation problem, we begin by defining the Lagrangian:

$$\mathcal{L}(\boldsymbol{x}_0, \{p(\lambda), \boldsymbol{x}^\lambda\}, \boldsymbol{t}) = \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda) + \boldsymbol{t} \cdot \left(\boldsymbol{x}_0 - \sum_\lambda p(\lambda) \boldsymbol{x}^\lambda \right),$$

where $\boldsymbol{t} \in \mathbb{R}^m$ are the Lagrange multipliers (or *dual variables*) associated to each constraint in the primal problem. In our case (28), we have $m = 3$ corresponding to the constraints on the observed values C_0, C_1 and $\mathfrak{E}$.

The dual function $g : \mathbb{R}^m \to \mathbb{R}$ is then defined as:

$$\begin{aligned} g(\boldsymbol{x}_0, \boldsymbol{t}) &= \inf_{\{p(\lambda), \boldsymbol{x}^\lambda\}} \mathcal{L}(\boldsymbol{x}_0, \{p(\lambda), \boldsymbol{x}^\lambda\}, \boldsymbol{t}) \\ &= \inf_{\{p(\lambda), \boldsymbol{x}^\lambda\}} \left\{ \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda) + \boldsymbol{t} \cdot \left(\boldsymbol{x}_0 - \sum_\lambda p(\lambda) \boldsymbol{x}^\lambda \right) \right\} \\ &= \boldsymbol{t} \cdot \boldsymbol{x}_0 + \inf_{\{p(\lambda), \boldsymbol{x}^\lambda\}} \left\{ \sum_\lambda p(\lambda) \left(f(\boldsymbol{x}^\lambda) - \boldsymbol{t} \cdot \boldsymbol{x}^\lambda \right) \right\}. \end{aligned} \tag{30}$$

The dual function g is concave by construction, as it is a pointwise infimum over affine functions. Next, let $t \in \mathbb{R}$ be the infimum $t := \inf_{\{p(\lambda), \boldsymbol{x}^\lambda\}} \left\{ \sum_\lambda p(\lambda) \left(f(\boldsymbol{x}^\lambda) - \boldsymbol{t} \cdot \boldsymbol{x}^\lambda \right) \right\}$, so that we can express the dual function simply as $g(\boldsymbol{x}_0, \boldsymbol{t}) = t + \boldsymbol{t} \cdot \boldsymbol{x}_0$.

This dual Lagrange function satisfies the *weak duality* property, yielding lower bounds on the optimal value:

$$g(\boldsymbol{x}_0, \boldsymbol{t}) \leq f^*(\boldsymbol{x}_0),$$

where $f^*(\boldsymbol{x}_0) = \min_{\{p(\lambda), \boldsymbol{x}^\lambda\}} \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda)$ is the optimal value of the primal problem. To see that weak duality is satisfied, consider any feasible ensemble $\{p(\lambda), \boldsymbol{x}^\lambda\}$ such that $\sum_\lambda p(\lambda) \boldsymbol{x}^\lambda = \boldsymbol{x}_0$. Then:

$$\begin{aligned} g(\boldsymbol{x}_0, \boldsymbol{t}) &\leq \boldsymbol{t} \cdot \boldsymbol{x}_0 + \sum_\lambda p(\lambda) (f(\boldsymbol{x}^\lambda) - \boldsymbol{t} \cdot \boldsymbol{x}^\lambda) \\ &= \boldsymbol{t} \cdot \boldsymbol{x}_0 + \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda) - \boldsymbol{t} \cdot \sum_\lambda p(\lambda) \boldsymbol{x}^\lambda \\ &= \sum_\lambda p(\lambda) f(\boldsymbol{x}^\lambda). \end{aligned} \tag{31}$$

Since this holds for any feasible ensemble, we conclude $g(\boldsymbol{x}_0, \boldsymbol{t}) \leq f^*(\boldsymbol{x}_0)$. Furthermore, note that for all $\boldsymbol{x} \in \mathcal{Q}$ one has $t + \boldsymbol{t} \cdot \boldsymbol{x} \leq f(\boldsymbol{x})$.

The inequality

$$t + \boldsymbol{t} \cdot \boldsymbol{x} \leq f(\boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{Q}$$

then defines a global affine underestimator of the function f. Therefore, the dual problem corresponding to the primal (29) can be expressed as

$$\begin{aligned} &\sup_{t, \boldsymbol{t}} \quad t + \boldsymbol{t} \cdot \boldsymbol{x}_0 \\ &\text{subject to} \quad t + \boldsymbol{t} \cdot \boldsymbol{x} \leq f(\boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{Q}. \end{aligned} \tag{32}$$

In our prepare-and-measure scenario, each variable $\boldsymbol{x} = (C_0, C_1, \mathfrak{E}) \in \mathbb{R}^3$ represents a possible correlation pair and corresponding energy uncertainty. The observed quantities $\boldsymbol{x}_0 = (C_0^{\text{obs}}, C_1^{\text{obs}}, \mathfrak{E}^{\text{obs}})$ are given and fixed. The feasible set is:

$$\mathcal{Q} := \{(C_0, C_1, \mathfrak{E}) \mid (C_0, C_1) \in \mathcal{Q}_{\mathfrak{E}, \Delta t}, 0 \leq \mathfrak{E} \leq \frac{\pi}{2\Delta t}\},$$

where recall that $\mathcal{Q}_{\mathfrak{E}, \Delta t} \subset [-1, 1]^2$ denotes the set of physically allowed correlations for a given energy uncertainty $\mathfrak{E}$ and evolution time Δt. Finally, the primal objective function $f(\boldsymbol{x})$ is the entropy function $f(\boldsymbol{x}) = H(\boldsymbol{C})$ defined above.

Thus, the dual formulation of our problem becomes:

$$\begin{aligned} H^*_{\text{dual}} = \sup_{t, \boldsymbol{t}} \quad & t + \boldsymbol{t} \cdot \begin{pmatrix} \boldsymbol{C} \\ \mathfrak{E} \end{pmatrix} \\ \text{subject to} \quad & t + \boldsymbol{t} \cdot \begin{pmatrix} \boldsymbol{C}' \\ \mathfrak{E}' \end{pmatrix} \leq H(\boldsymbol{C}') \quad \forall \begin{pmatrix} \boldsymbol{C}' \\ \mathfrak{E}' \end{pmatrix} \in \mathcal{Q}. \end{aligned} \tag{33}$$

The objective function H is concave. It is easy to see that this implies that both in the primal and in the dual problem, we can replace the domain of optimisation $\mathcal{Q}$ by its convex hull $\text{conv}(Q)$. Since this is a convex set, Proposition 18 in [28] implies strong duality, i.e. $H^* = H^*_{\text{dual}}$ for all $\mathbf{C}$ and $\mathfrak{E}$. Yet, as long as we are only aiming for lower bounds to H^*, it is not necessary to appeal to this result.

Numerical lower bound via brute-force dual optimization

Solving the dual problem in (33) exactly is beyond the scope of this manuscript. However, since any feasible solution to the dual provides a valid lower bound on the optimal entropy H^*, we adopt a brute-force numerical strategy. While this method does not guarantee optimality, it reliably yields a lower bound to H^*, sufficient for certifying randomness under the semi-device-independent assumptions from the main text.

The brute-force method proceeds as follows:

- While the Lagrange multipliers $\boldsymbol{t} \in \mathbb{R}^3$ are in principle unbounded, for computational purposes we restrict them to lie within a cube of side length $2L$, centered at the origin, where L is chosen to be finite at the cost of losing optimality. This cube is uniformly discretized with N finite points per axis, leading to N^3 candidate dual vectors.

- For each candidate dual vector $\boldsymbol{t}$, we evaluate the dual constraint:

  $$t + \boldsymbol{t} \cdot \boldsymbol{x} \leq H(\boldsymbol{C}) \quad \text{for all } \boldsymbol{x} = (\boldsymbol{C}, \mathfrak{E}) \in \mathcal{Q},$$

  by computing $t = \min_{\boldsymbol{x} \in \mathcal{Q}} [H(\boldsymbol{C}) - \boldsymbol{t} \cdot \boldsymbol{x}]$.

- We then evaluate the dual objective at the target point $\boldsymbol{x}_0 = (C_0^{\text{obs}}, C_1^{\text{obs}}, \mathfrak{E}^{\text{obs}})$ as:

$$H_{\text{dual}}(\boldsymbol{t}) = t + \boldsymbol{t} \cdot \boldsymbol{x}_0,$$

and retain the maximum value found over all $\boldsymbol{t}$.

To ensure that the membership constraint $\boldsymbol{x} \in \mathcal{Q}$ is satisfied, we discretize the domain as follows:

- The correlation coordinates $(C_0, C_1) \in [-1, 1]^2$ and the energy uncertainty interval $\mathfrak{E} \in [0, \pi/2]$ are discretized into N equally spaced steps each.

- We retain only those tuples $(C_0, C_1, \mathfrak{E})$ that satisfy the physical constraint:

$$\frac{1}{2}\left(\sqrt{1+C_0}\sqrt{1+C_1} + \sqrt{1-C_0}\sqrt{1-C_1}\right) \geq \cos(\mathfrak{E}\Delta t).$$

This defines the discretized feasible set $\mathcal{Q}_{\text{discretized}} \subset \mathcal{Q}$. Note that, for efficiency, it suffices to take only the extremal points of $\mathcal{Q}_{\text{discretized}}$ (*i.e.,* solving the physical constraint for strict equality).

Heuristically, we observe that setting $L \sim 20$ and using $N \sim 200$ discretization steps provides a good trade-off between computational cost and approximation accuracy. Increasing these values further yields only marginal improvements in the estimated entropy, while significantly increasing memory usage and runtime.

Algorithm Summary

In practice, the algorithm proceeds as follows:

1. Pre-compute all tuples $\boldsymbol{x} = (C_0, C_1, \mathfrak{E}) \in \mathcal{Q}_{\text{discretized}} \subset \mathcal{Q}$ satisfying the above inequality.

2. For each $\boldsymbol{t} \in [-L, L]^3$, compute:

$$\tau(\boldsymbol{t}) = \min_{(C_0, C_1, \mathfrak{E}) \in \mathcal{Q}_{\text{discretized}}} \left[H(C_0, C_1) - \boldsymbol{t} \cdot (C_0, C_1, \mathfrak{E})\right],$$

where τ gathers all the evaluated $\boldsymbol{t}$ values in the search space.

3. Evaluate the dual value at the target point $\boldsymbol{x}_0 = (C_0^{\text{obs}}, C_1^{\text{obs}}, \mathfrak{E}^{\text{obs}})$:

$$H_{\text{dual}}(\boldsymbol{t}) = \tau(\boldsymbol{t}) + \boldsymbol{t} \cdot \boldsymbol{x}_0,$$

where $H_{\text{dual}}(\boldsymbol{t})$ gathers all the evaluated $\boldsymbol{t}$ values in the search space.

4. Return:

$$H^*_{\text{dual}} = \max_{\boldsymbol{t}} H_{\text{dual}}(\boldsymbol{t}),$$

which serves as a certified lower bound to the true entropy H^*.

This procedure to obtain H^*_{dual} provides a valid lower bound on the entropy $H^*_{\text{dual}} \leq H^*$, thus enabling randomness certification under the semi-device-independent assumptions of the main text.