

Bell inequalities for continuously emitting sourcesEmanuel Knill,^{1,*} Scott Glancy,¹ Sae Woo Nam,¹ Kevin Coakley,¹ and Yanbao Zhang^{2,1,†}¹*National Institute of Standards and Technology, Boulder, Colorado, 80305, USA*²*Department of Physics, University of Colorado Boulder, Boulder, Colorado, 80309, USA*

(Received 23 October 2014; published 4 March 2015)

A common experimental strategy for demonstrating nonclassical correlations is to show violation of a Bell inequality by measuring a continuously emitted stream of entangled photon pairs. The measurements involve the detection of photons by two spatially separated parties. The detection times are recorded and compared to quantify the violation. The violation critically depends on determining which detections are coincident. Because the recorded detection times have “jitter,” coincidences cannot be inferred perfectly. In the presence of settings-dependent timing errors, this can allow a local-realistic system to show apparent violation, the so-called “coincidence loophole.” Here, we introduce a family of Bell inequalities based on signed, directed distances between the parties’ sequences of recorded time tags. Given that the time tags are recorded for synchronized, fixed observation periods and that the settings choices are random and independent of the source, violation of these inequalities unambiguously shows nonclassical correlations violating local realism. Distance-based Bell inequalities are generally useful for two-party configurations where the effective size of the measurement outcome space is large or infinite. We show how to systematically modify the underlying Bell functions to improve the signal-to-noise ratio and to quantify the significance of the violation.

DOI: [10.1103/PhysRevA.91.032105](https://doi.org/10.1103/PhysRevA.91.032105)

PACS number(s): 03.65.Ud, 42.50.Xa, 02.50.Cw

I. INTRODUCTION

Quantum mechanical systems can give rise to measurement correlations that local realistic (LR) systems are unable to produce. Physical theories that satisfy the principle of local realism (LR) posit a set of hidden variables associated with the physical systems. The hidden variables cannot be influenced by spacelike-separated events. They are not observable, but they determine the outcomes of all measurements. The existence of hidden variables constrains the probability distributions that can describe LR systems. In 1964, Bell constructed an inequality that is satisfied by all correlations accessible by LR and showed that correlations between spacelike-separated measurements on two quantum systems can violate this inequality [1]. The realization that quantum mechanics allows more general probability distributions than LR has motivated many experimental tests that have shown Bell-inequality violations (see Ref. [2] for a review). In addition to the fundamental importance of tests of LR, systems violating LR can be used for quantum information tasks such as quantum key distribution [3–5] and secure randomness generation [6–8]. For these cryptographic applications, one must demonstrate violation of LR with high statistical significance in the presence of adversarial effects, such as a hacker who has tampered with the system in an attempt to learn a secret key.

So far, all tests of LR have invoked additional assumptions about the types of LR theories governing their experiments. Examples include the assumption that photon detection probabilities are not correlated with the measurement choices or photon polarizations (the “fair sampling” assumption), the assumption that measurement choices at one location cannot influence events at another location even when they are not spacelike separated, and the assumption that the sequence

of measurement choices and outcomes in an experiment are independent and identically distributed (i.i.d.). For a review, see [9]. Various experiments have been able to relax some of these assumptions, but no single experiment has been able to reject the most general LR theories. Due to recent advances in entangled photon generation and photon detection, we anticipate that an optical experiment that is free of additional assumptions will be accomplished in the near future.

To test a Bell inequality in an experiment, one repeats the preparation and joint measurement of spatially separated systems some finite number N times. We call each such repetition a “trial” of the experiment. During a trial, entangled systems are sent to two or more measurement locations. At each location, a random choice is made that determines which property will be measured. Ideally, each trial is clearly identifiable so that measurement choices and outcomes at the various locations can be matched with one another. However, in many experiments this trial identification cannot be achieved with perfect certainty. A popular experiment design involves the continuous pumping of a nonlinear crystal that produces photon pairs through spontaneous parametric down-conversion. In these experiments, the entangled pairs are produced randomly in time. Furthermore, the detectors used at the two measurement locations (“Alice” and “Bob”) have nonzero timing jitter. These effects can create confusion about which events at Alice correspond to which events at Bob.

To resolve this confusion, one typically defines a “coincidence window” in time, so that if the period of time between photon detections at Alice and Bob is less than the coincidence window width, the two events are considered to be part of the same trial. If Alice or Bob observe multiple detections within one coincidence window, more sophisticated algorithms can be used to attempt to match Alice’s and Bob’s detections. The choice of the coincidence window width depends on balancing the expected interarrival time between entangled pair creations and the detector timing jitter, so that the probability for a trial to contain multiple photon pairs is small, and the probability

*emanuel.knill@nist.gov

†Current address: Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.

for a photon to be lost by falling outside of the coincidence window is also small.

Unfortunately, local realistic theories or hackers can use this uncertainty about trial identification to produce apparent violation of Bell inequalities. The photons being measured could have correlations between their measurement outcomes and arrival times. Correlations could also exist between multiple photon pairs detected during the same coincidence window. Previous tests of LR that use continuously pumped spontaneous parametric down-conversion have (often implicitly) assumed that such correlations do not exist. Therefore, these experiments do not reject the most general LR theories; they only reject LR theories that do not allow timing or cross-pair correlations. Experiments requiring these assumptions are said to suffer from the “coincidence loophole.” Larsson and Gill described local realistic theories that exploit the coincidence loophole and methods to defeat them in [10]. Their closure of this loophole requires a bound on the probability that a true coincidence is missed. This bound cannot be measured without additional assumptions and must be trusted.

In this paper, we describe a different method for closing the coincidence loophole that does not require additional assumptions about the systems being measured. Instead, we define a “trial” as all events occurring within a predetermined time interval. The data produced by the trial are the measurement choices at Alice and Bob and the lists of times at which Alice and Bob detected photons (the “time tags”). The trial may be many times longer than the expected time between photon-pair creations, and Alice and Bob may observe many photon detection events during one trial. (In practice, the length of a trial will be constrained if one desires to achieve spacelike separation between measurement choices on one side and detections on the other.)

To test LR with such trial data, we develop distance-based Bell inequalities. The use of distances and triangle inequalities was suggested by Schumacher [11] in the case of the conventional correlation-based Bell inequalities. Signed, directed distances between trial outcomes that measure the dissimilarity between Alice’s and Bob’s lists of time tags are defined such that they obey a directed triangle inequality. These distances are closely related to edit distances used to compare words in spell checking or to align DNA sequences in computational biology. Using LR to compute the expectation value of sums of these distances yields an inequality satisfied by all LR theories. These distance-based Bell inequalities provide a rigorous analysis of tests of LR based on continuously emitting sources; they enable the rejection of a larger class of LR theories than previously known methods. Also, the triangle inequality can be a powerful tool for finding new Bell inequalities in other contexts. It is explored in the works of Dzhamalov and Kujala [12] and of Kurzynski and Kaslikowski [13].

In Sec. II, we introduce basic notions and define relevant mathematical notation. In Sec. III, we explain the experimental setup of “time-tag Bell tests” in more detail and describe a simple LR model that exploits the coincidence loophole for an apparent violation of a Bell inequality. In Sec. IV, we define “distance” functions between measurement outcomes that obey the directed triangle inequality. We then use the triangle inequality to derive Bell inequalities satisfied by any LR theory regardless of the choice of distance function. In Sec. V,

we describe functions for computing the distance between time-tag sequences and obtain the associated Bell inequalities. In Sec. VI, we introduce nonsignaling equalities that constrain all theories that prohibit Alice from sending information to Bob by use of her measurement choice (and vice versa). Although both quantum and LR theories obey the nonsignaling equalities, these equalities can be used to transform Bell inequalities and improve the signal-to-noise ratio (SNR) of the inequalities’ violation in an experiment. In Sec. VII, we provide a protocol that sets aside an initial segment of the data as a training set to determine a good distance function. In Sec. VIII, we discuss the relationship between the SNR for the violation of a Bell inequality and p -value bounds for rejecting LR. Bounds on p values can be computed with Gill’s martingale-based protocol [14,15] or the prediction-based ratio (PBR) protocol [16,17]. The main result of this section is a method for truncating distance functions to enable application of these protocols. The technique is general and can be used on any Bell function derived from a triangle inequality. Here, a useful step is to balance the violation between the measurement settings by means of the nonsignaling equalities. In Sec. IX, we apply time-tag Bell inequalities to simulated data. We discuss the effects of detector inefficiency and detector jitter on the violation of time-tag Bell inequalities and on the p -value bounds computed with the PBR protocol. We quantify the violation and p -value bounds as functions of the jitter distribution’s width and quote lower bounds on the maximum jitter width at which violation can be observed for photon detection efficiencies ranging from 0.74 to 0.95. The simulations include a LR model that exploits the coincidence loophole while closely mimicking the measurement statistics of a Poisson source of entangled photons measured with jittery detectors. In the Appendixes, we describe numerical methods to optimize parameters of the distance functions to give high inequality violation, to compute distances for time-tag Bell inequalities, and to compute the SNR of the violation of an inequality. The Appendixes also contain further details of the coincidence-loophole-exploiting LR model.

II. PRELIMINARIES

We consider experiments to test LR, where an experiment consists of a sequence of trials. The trials’ measurement outcomes need not be independent from one trial to the next, but before the next trial, there is a probabilistic description of the next trial’s outcome, where the probabilities may depend on the past and current conditions. The class of LR models of interest is defined by specifying constraints on these probabilities. We consider the case where a trial consists of observations by two parties, A and B , each of whom can choose one of two measurement settings for their observation. We leave extensions to more parties and settings for future work. The full trial outcome includes the settings chosen as well as the measurement outcomes. In many cases, the measurement outcomes are two valued. For example, the outcome may indicate whether a photodetector “clicked” or not. Here, we consider arbitrary outcome spaces, but focus on the case where a party’s measurement outcome is an ordered sequence of time tags of events, for example, detection events. Thus, there is no bound on the size of the outcome space.

We denote the random variable for a trial's outcome including the settings by T . This random variable is a tuple of four random variables $T = (O^A, S^A, O^B, S^B)$ where O^X is X 's measurement outcome and S^X is X 's chosen setting. The two possible settings are denoted by $\bar{1}$ and $\bar{2}$. We also use the notation $T^X = (O^X, S^X)$. We follow the notational convention that random variables (O, S, T, \dots) are denoted by roman uppercase letters. This is also true of party labels, but the distinction should be clear from context. The range of random variable R is denoted by \mathcal{R} . Observed values of random variables are denoted by their corresponding lowercase letters; for example, r denotes an observed value of R . Superscripts and subscripts serve to identify members of a family of conceptually related random variables or to select out parts of tuple-valued random variables. Formally, the random variables for a trial are functions on an underlying probability space that includes any "hidden" variables that may play a role, but we do not need to explicitly refer to this space here.

A deterministic LR model must, before a trial and independent of the settings, commit to a specific measurement outcome d_c^X for each party $X = A, B$ and each setting $c = \bar{1}, \bar{2}$. A general LR model is a probabilistic mixture of deterministic models. (One can imagine that a hidden random variable selects which deterministic model controls a trial's measurement outcomes.) Thus, a LR model is described by a random variable $D_{LR} = (D_1^A, D_2^A, D_1^B, D_2^B)$, where T relates to D_{LR} according to $T = (D_{S^A}^A, S^A, D_{S^B}^B, S^B)$. Although the parties cannot simultaneously measure both settings $\bar{1}$ and $\bar{2}$ in a single trial, the LR model allows for that possibility by preassigning measurement outcomes to both settings. That such a preassignment exists is the essential claim of realism. Quantum theory does not preassign outcomes and disallows the possibility that the two settings can be measured simultaneously. Quantum theory can thereby achieve a larger set of trial probability distributions.

In an idealized test of LR, the settings choices are made randomly and independently of D_{LR} according to a probability distribution that is under experimenter control. In this case, LR models satisfy that $S = (S^A, S^B)$ is independent of D , and the probability distribution of S is known before the trial. This defines LR models satisfying the free-choice assumption. For the remainder of the paper, LR models are assumed to satisfy free choice, and by default the settings' distribution is uniform.

From a mathematical and statistical point of view, a successful test of LR shows that probabilistic LR models are statistically inconsistent with the data. Interpretation of the inconsistency requires additional analysis and can depend on the experimental context. In fact, such interpretations could attribute the inconsistency to the presence of so-called "loopholes" rather than to the falsity of LR. An experimental goal is to convincingly exclude the presence of such loopholes.

III. TIME-TAG BELL TESTS

A common method for performing Bell tests is to use a source that continuously emits pairs of polarization-entangled photons. The photons are delivered to two measurement setups. A trial consists of choosing the settings and then recording photodetection events for a fixed observation window. We focus on the simplest case, where the measurement setups

involve polarizers whose angles determine the settings. Each setup has one photodetector that records photons that passed through the polarizer. Thus, the record of a trial includes two time-tag sequences recording the times at which photons were detected. The experiments reported in [18,19] used this setup.

One way to think about such an experiment is that fundamentally, each photon pair's emission and detection constitutes a trial. In this case, the first step in an analysis is to identify the detection pattern for each emitted pair. The record does not identify when neither photon was detected, but the Bell inequalities used can be chosen so that the total Bell-inequality violation is insensitive to the number of photon-pair trials where neither photon was detected. Thus, the analysis requires identifying coincidences, that is, pairs of detections that are due to one photon pair. Identifying coincidences is complicated by the fact that the recorded time tags have "jitter," that is, the difference between the time tag t and the "true" time of arrival of a photon t_0 is a random variable with non-negligible width j (to be defined in Sec. IX for specific jitter distributions). Furthermore, since photon pairs are continuously emitted, their creation times and their times of arrival are also random. Pair emission can usually be modeled as a Poisson process. Denote the mean interarrival time between successive photon pairs as τ . It is necessary to determine which pairs of close time tags t and r of A and B are due to the same photon pair. This cannot be done without error, as there is always the possibility that photons detected by A and B around the same time are from two different photon pairs that were created with small time separation. The probability of this event grows with j/τ .

Given that coincidences cannot be identified exactly, it is necessary to determine how this affects the interpretation of a Bell-inequality violation. In cases where the nominal mean violation per photon pair is small and j/τ is relatively large, the evidence against LR may be weakened substantially. An example of this situation is the experiment reported in Ref. [18], which aimed to close the fair sampling loophole with photons. The violation was limited by the overall detection efficiency realized in the experiment. To interpret the violation, one can analyze the effects of coincidence identification error by making the assumptions that the source is idealized Poisson, the jitter is settings independent, the photon pairs' states are independent and identical, and the method for recording detections is memoryless. Parts of such an analysis are in [20]. But, these are highly idealizing assumptions unlikely to be satisfied in a real experiment. Relying on them precludes making strong claims on having demonstrated non-LR effects. Of particular concern is that the presence of jitter in combination with a conventional coincidence analysis requires a fair coincidence sampling assumption [10]. For conventional analyses, this assumption can be avoided by using "pulsed" trials, which can significantly reduce the rate of detections. Such an experiment was reported in Ref. [19].

Unfair coincidence sampling can arise from local, settings-dependent properties of the detectors, including the associated settings-related apparatus such as polarization filters. A simple LR model that exploits unfair coincidence sampling to show violation of a Bell inequality is illustrated in Fig. 1. Suppose that for A , the difference $t - t_0$ between the recorded time tag and the photon arrival time is 0 on setting $\bar{1}$ and Δ on setting

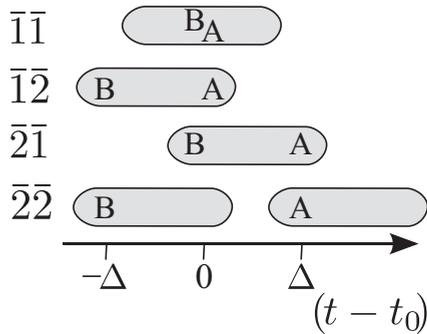


FIG. 1. Illustration of a LR model that uses settings-dependent detection times to violate a Bell inequality. Detection time $t - t_0$ proceeds along the horizontal axis. Each row above the time axis corresponds to one measurement setting combination used by A and B . Photons are detected by A and B at the times labeled “A” and “B.” (At the $\bar{1}\bar{1}$ setting, both A and B detect their respective photons near 0.) The shaded regions indicate coincidence windows of width 1.5Δ . At the $\bar{2}\bar{2}$ setting, the photons detection times are separated by 2Δ , so A and B can never observe coincidences with $\bar{2}\bar{2}$.

$\bar{2}$. For B , suppose that the difference is 0 on setting $\bar{1}$ and $-\Delta$ on setting $\bar{2}$. To identify coincidences, one can choose a coincidence window width w and declare that time tags of A and B whose separation is less than w are coincident. It is necessary to have a method for resolving coincidence conflicts such as when a time tag of A is within w of more than one time tag of B . Here, we just assume that j/τ is small enough for this not to be considered an issue. Suppose that the coincidence window width is chosen to be $w = 1.5\Delta$. If the LR model for the photon pair is to “detect” no matter what the setting is, then A and B record coincidences on all settings except $\bar{2}\bar{2}$, where the two detections are inferred as being noncoincidence detections. This LR model strongly violates commonly used Bell inequalities, such as the one introduced in Sec. IV, Eq. (2), whose violation is increased by anticorrelation between A ’s and B ’s detections on the $\bar{2}\bar{2}$ setting. The choice of w may seem arbitrary, but a natural way to choose w is to optimize the violation on a training set or on preliminary data. In this case, $w = 1.5\Delta$ is an optimal choice. Note that locality assumptions and the assumption that settings’ choices are uniformly random and independent of hidden variables affecting the measurement outcomes are satisfied in this example.

A detailed theoretical treatment of unfair coincidence sampling is given in Ref. [10], including a more sophisticated example that can respond to the continuous angular settings choices available when measuring photons. Given the assumptions in Ref. [10], there are valid adjustments to a Bell inequality based on knowledge of the probability of missing a coincidence. Our approach based on time-tag Bell inequalities does not require such additional knowledge.

It may seem like the presence of unfair coincidence sampling due to a dependence of detector timing on settings can be excluded by checking that there is no widening of the time-separation distribution of the nearest time tags of A and B for the $\bar{2}\bar{2}$ settings compared to the others. Further, one can attempt to choose w after studying this distribution

to ensure that the fraction of missed coincidences for the $\bar{2}\bar{2}$ setting is sufficiently small. Any such attempt would require an hypothesis test (or some other way to quantify evidence) for the claim that no unfair coincidence sampling is present in the experiment. Depending on how the data for these tests are acquired, additional assumptions on consistency of detector behavior may be needed. In pursuing this approach, one must then decide at what significance level one wishes to exclude excessive unfair coincidence sampling. This significance level should be similar to the claimed significance of the violation. This may not be feasible in an experiment without greatly weakening the significance of the result.

Coincidence sampling effects can also be exploited directly by a LR source. The simplest example just simulates the detector timing issue. There are two operationally different versions of this example. In the first, the local hidden variable of a photon also identifies the time at which the detector records it in a setting-dependent way. One could imagine that once the photon arrives at the detector and senses the setting, it “pauses” a variable amount of time. It now suffices to emit photon pairs where photon A ’s local variable assigns “detect” at time t on setting $\bar{1}$ and “detect” at time $t + \Delta$ on setting $\bar{2}$, whereas photon B ’s variable assigns “detect” at time t on setting $\bar{1}$ and “detect” at time $t - \Delta$ on setting $\bar{2}$. Such photons seem impossible to realize, but the following version of the example may be realizable. The source sends a photon to B at time $t - \Delta$, a pair to A and B at time t , and another to A at time $t + \Delta$. The first photon is prepared so as to be detected by B only on setting $\bar{2}$. The two middle photons are detected only on setting $\bar{1}$. The last photon is prepared so as to be detected by A only on setting $\bar{2}$. When A and B compare their time tags and use $w = 1.5\Delta$ for their coincidence analysis, they again see coincidences on all settings except for $\bar{2}\bar{2}$. A LR source that wishes to hide having manipulated emission times can intersperse a small number of photons with the emission pattern above with regular LR pairs of photons that ensure equality for the Bell inequality of interest. Other opportunities to hide the presence of unfair coincidence sampling from the experimenters exist. The LR source can systematically introduce LR photons with varying timing features, and it can conditionally omit “normal” photon pairs to make it more difficult to see excess numbers of close detections outside the coincidence window. While these possibilities may seem physically unrealistic, they are of concern in cryptographic applications of experimental configurations for violating Bell inequalities. To show that these concerns are justified, in Sec. IX we show simulation results for a LR source whose statistics would be difficult to distinguish from a quantum source in an experimental setting. The conventional coincidence analysis shows a false violation of a Bell inequality for this source.

IV. DISTANCE-BASED BELL FUNCTIONS

Distance-based Bell functions and associated Bell inequalities generalize the conventional two-party, two-setting Bell inequalities such as the CHSH [21] and CH inequalities [22] (the abbreviations stand for the authors’ initials). Consider two parties A and B , where each can choose from two settings labeled $\bar{1}$ and $\bar{2}$ and the physical meaning of the setting

labels depends on the party. Suppose that the measurement outcomes at a given setting are -1 or 1 . Let O^X and S^X be the measurement outcome and setting of party X , respectively. One of the CHSH inequalities for this configuration is

$$\begin{aligned} & \langle O^A O^B | S^A = \bar{2}, S^B = \bar{2} \rangle - \langle O^A O^B | S^A = \bar{2}, S^B = \bar{1} \rangle \\ & - \langle O^A O^B | S^A = \bar{1}, S^B = \bar{1} \rangle \\ & - \langle O^A O^B | S^A = \bar{1}, S^B = \bar{2} \rangle \geq -2, \end{aligned} \quad (1)$$

where $\langle U \rangle$ denotes the expectation of U . This inequality is satisfied by all LR models. Let $l(x, y) = |x - y|$. Since $O^A O^B = 1 - l(O^A, O^B)$, Eq. (1) can be rewritten as

$$\begin{aligned} & \langle l(O^A, O^B) | S^A = \bar{2}, S^B = \bar{1} \rangle + \langle l(O^B, O^A) | S^A = \bar{1}, S^B = \bar{1} \rangle \\ & + \langle l(O^A, O^B) | S^A = \bar{1}, S^B = \bar{2} \rangle \\ & - \langle l(O^A, O^B) | S^A = \bar{2}, S^B = \bar{2} \rangle \geq 0. \end{aligned} \quad (2)$$

A deterministic LR model assigns a specific value d_c^X for the measurement outcome of each party X and setting c before the experiment. In this case, the left-hand side of Eq. (2) is given by $l(d_2^A, d_1^B) + l(d_1^B, d_1^A) + l(d_1^A, d_2^B) - l(d_2^A, d_2^B)$. This is at least 0 because l satisfies the triangle inequality as illustrated in Fig. 2 and observed by Schumacher in Ref. [11]. The inequality of Eq. (2) follows because general LR models are probabilistic mixtures of deterministic ones. We have switched the A and B arguments of l in the contribution for the $\bar{1}\bar{1}$ setting in preparation for applying asymmetric functions l .

The Bell inequalities of Eqs. (1) and (2) are expressed in terms of quantities that are conditional on settings. Since measurement settings need to be chosen randomly, experimentally estimating these quantities requires dividing by the actual number of times the relevant settings are chosen. This complicates the estimation of experimental uncertainty. To avoid this complication, recall that the probability distribution of the settings is known beforehand and is independent of the LR model. Let p_{ab} be the probability that A and B choose settings a and b , respectively. Then, the left-hand side of Eq. (2) is equal to

$$\langle (-1)^{[S^A=\bar{2} \text{ and } S^B=\bar{2}]} l(O^A, O^B) / p_{S^A S^B} \rangle, \quad (3)$$

where the expression $[\phi]$ in the exponent of -1 evaluates to 1 when the logical formula ϕ is true and to 0 when it is false.

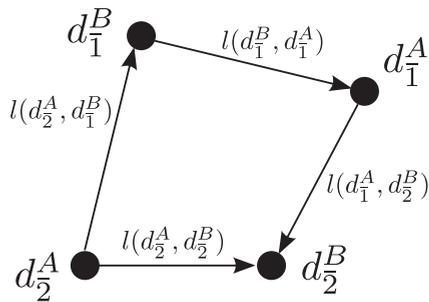


FIG. 2. Illustration of the twice-iterated triangle inequality. Each node represents one of the four potential measurement outcomes in d_{LR} , and each directed edge represents one of the lengths l in Eq. (4). The edges on the path correspond to “compatible” measurements pairs that are experimentally measured. The locations of the nodes are arbitrary in this illustration.

Equation (3) is the expectation of a function of the settings and measurement outcomes. This function is the Bell function for the inequality of Eq. (2). For our default assumption of a uniform settings probability distribution, $p_{S^A S^B} = \frac{1}{4}$.

In general, we define a Bell function to be a function of trial outcomes whose expectation with respect to every LR model is non-negative, where the probability distribution of the settings is fixed. We now obtain such Bell functions from functions $l : (x, y) \in \mathcal{O} \times \mathcal{O} \mapsto l(x, y) \in \mathbb{R}$ that satisfy the twice-iterated triangle inequality $l(x_1, x_4) \leq l(x_1, x_2) + l(x_2, x_3) + l(x_3, x_4)$, where \mathcal{O} is a common measurement outcome space for all parties and settings. Unlike a true distance, we do not require l to be non-negative or symmetric. Note that the twice-iterated triangle inequality follows from the usual triangle inequality $d(u, w) \leq d(u, v) + d(v, w)$ by substituting the right-hand side of $d(x_1, x_3) \leq d(x_1, x_2) + d(x_2, x_3)$ for $d(x_1, x_3)$ in the inequality $d(x_1, x_4) \leq d(x_1, x_3) + d(x_3, x_4)$. While it is convenient to construct functions satisfying the twice-iterated triangle inequality from functions satisfying the usual triangle inequality, we do not require that l satisfy the usual triangle inequality. Note that in the context of metrics and quantum information, the twice-iterated triangle inequality is occasionally referred to as a “quadrilateral inequality” [11] or a “quadrangle inequality” (these terms have other meanings depending on context).

Consider a deterministic LR model given as above by its outcome assignments $d_{LR} = (d_1^A, d_2^A, d_1^B, d_2^B)$. The elements of d_{LR} obey the inequality

$$0 \leq l(d_2^A, d_1^B) + l(d_1^B, d_1^A) + l(d_1^A, d_2^B) - l(d_2^A, d_2^B), \quad (4)$$

obtained by arranging the twice-iterated triangle inequality as illustrated in Fig. 2. The reversal of the parties in the term $l(d_1^B, d_1^A)$ (the middle edge of the indirect path) requires argument reversals in the expressions following.

We can define a Bell function based on l as follows:

$$B_l(t) = \begin{cases} 4l(o^B, o^A) & \text{if } s^A = \bar{1} \text{ and } s^B = \bar{1}, \\ 4l(o^A, o^B) & \text{if } s^A \neq s^B, \\ -4l(o^A, o^B) & \text{if } s^A = \bar{2} \text{ and } s^B = \bar{2}, \end{cases} \quad (5)$$

where $t = (o^A, s^A, o^B, s^B)$ is a trial outcome. The factor of 4 originates as the inverse of the probability $\frac{1}{4}$ of each of the settings. Consider a LR model with measurement outcome random variables D_c^X . In the trial outcome random variable T , $O^A = D_{S^A}^A$ and $O^B = D_{S^B}^B$. Because the settings are independent of the D_c^X , the expectation of B_l can be computed as

$$\begin{aligned} \langle B_l(T) \rangle_{LR} &= \langle l(D_2^A, D_1^B) \rangle_{LR} + \langle l(D_1^B, D_1^A) \rangle_{LR} \\ &+ \langle l(D_1^A, D_2^B) \rangle_{LR} - \langle l(D_2^A, D_2^B) \rangle_{LR} \geq 0, \end{aligned} \quad (6)$$

where the factors of 4 were canceled by the probabilities of the settings. The inequality can be checked directly for deterministic LR models, by replacing the random variable D_c^X with the constant d_c^X . For general LR models, both sides can be integrated with respect to the appropriate distributions over deterministic LR models. For a general settings distribution, the values of l in Eq. (5) are multiplied by the inverse of the applicable settings’ probability so that the expression for the expectation in Eq. (6) is unchanged.

Equation (6) is the Bell inequality associated with the Bell function B_l . To test it in an experiment involving a sequence T_k of trial outcomes, one computes the values $B_l(t_k)$ on the actual trial outcomes t_k . The sum $v = \sum_k B_l(t_k)$ is then an estimate of $\sum_k \langle B_l(T_k) \rangle$, that can be compared to 0. The test is considered successful if $v < 0$, and the difference between v and 0 is statistically significant. According to the conventional approach, this involves determining an uncertainty for v . See Appendix A for an effective method for obtaining such an uncertainty that takes into account the known probability distribution of the settings. While this prescription is seemingly straightforward, care must be taken when interpreting the results for trials that may not be independent and identical [14,23,24]. Statements on the strength of evidence against LR require additional analysis based on statistical hypothesis testing (see Sec. VIII).

For our applications to time-tag sequences, we further generalize the Bell functions by allowing l to depend on the settings. Consider functions l_{ab} of two measurement outcomes that satisfy the following version of the iterated triangle inequality:

$$0 \leq l_{\bar{2}\bar{1}}(d_2^A, d_1^B) + l_{\bar{1}\bar{1}}(d_1^B, d_1^A) + l_{\bar{1}\bar{2}}(d_1^A, d_2^B) - l_{\bar{2}\bar{2}}(d_2^A, d_2^B). \quad (7)$$

We call such an l a CH function. For the first and shorter expression in the next definition, we use the notation $\tilde{l}_{\bar{1}\bar{1}}(o_1, o_2) = l_{\bar{1}\bar{1}}(o_2, o_1)$ and $\tilde{l}_{ab}(o_1, o_2) = l_{ab}(o_1, o_2)$ for $ab \neq \bar{1}\bar{1}$. A Bell function B_l can now be defined by generalizing Eq. (5) according to

$$B_l(t) = 4(-1)^{[s^A = \bar{2} \text{ and } s^B = \bar{2}]} \tilde{l}_{s^A s^B}(o^A, o^B) \\ = \begin{cases} 4l_{s^A s^B}(o^B, o^A) & \text{if } s^A = \bar{1} \text{ and } s^B = \bar{1}, \\ 4l_{s^A s^B}(o^A, o^B) & \text{if } s^A \neq s^B, \\ -4l_{s^A s^B}(o^A, o^B) & \text{if } s^A = \bar{2} \text{ and } s^B = \bar{2}. \end{cases} \quad (8)$$

We call B_l a CH Bell function. It satisfies Eq. (6) after adding the appropriate indices to the occurrences of l .

We remark that every Bell function for a two-party, two-settings configuration can be put in the form of a CH Bell function for some choice of CH function l . Thus, many of the techniques discussed in the remainder of the paper are generally applicable. Consider an arbitrary Bell function $B(o^A, s^A, o^B, s^B)$ standardized as above so that $\langle B(T) \rangle_{\text{LR}} \geq 0$ for the settings probability distribution p_{ab} . We can define

$$l_{B,ab}(o_1, o_2) = \begin{cases} B(o_2, a, o_1, b)p_{ab} & \text{if } ab = \bar{1}\bar{1}, \\ B(o_1, a, o_2, b)p_{ab} & \text{if } a \neq b, \\ -B(o_1, a, o_2, b)p_{ab} & \text{if } ab = \bar{2}\bar{2}. \end{cases} \quad (9)$$

Since there are no constraints on the measurement outcomes in deterministic LR models, the Bell inequality implies the iterated triangle inequality of Eq. (7). To show this, fix the LR model so that it is deterministic according to d_{LR} , where d_{LR} is arbitrary. Then,

$$0 \leq \langle B(T) \rangle_{\text{LR}} = \sum_{ab} p_{ab} \langle B(o^A, a, o^B, b) | ab \rangle_{\text{LR}} \\ = \sum_{ab} p_{ab} B(d_a^A, a, d_b^B, b)$$

$$= l_{B,\bar{2}\bar{1}}(d_2^A, d_1^B) + l_{B,\bar{1}\bar{1}}(d_1^B, d_1^A) + l_{B,\bar{1}\bar{2}}(d_1^A, d_2^B) \\ - l_{B,\bar{2}\bar{2}}(d_2^A, d_2^B).$$

V. BELL FUNCTIONS FOR TIME-TAG SEQUENCES

For our construction of CH functions for general time-tag sequences, we require functions l_{ab} whose domains are two lists of real numbers representing the time-tag sequences at settings ab , whose ranges are the real numbers, and that satisfy the iterated triangle inequality. Our general strategy is to construct l_{ab} so that it computes a quantity similar to an edit distance, which quantifies the degree of dissimilarity between the time-tag sequences obtained by A and B . (However, our CH functions may be negative and are not symmetric in their arguments, so they are not strictly distances.) As mentioned in the Introduction, edit distances are commonly used to compare strings and are defined by the minimum number (or cost) of elementary edits required to convert one string to another. Here, the edits consist of deleting time tags or moving them, where the cost of the move is related to the distance of the move. In particular, our implementation matches time tags of A with time tags of B and assigns a cost to the difference between the matched time tags and a cost to unmatched time tags.

To compute the cost for matched time tags we use function-tuples $(f_{ab})_{a,b \in \{\bar{1}, \bar{2}\}}$, where the $f_{ab} : x \in \mathbb{R} \mapsto f_{ab}(x) \in \mathbb{R}$ satisfy that for all $x, y, z \in \mathbb{R}$, $f_{\bar{2}\bar{2}}(x + y + z) \leq f_{\bar{2}\bar{1}}(x) + f_{\bar{1}\bar{1}}(y) + f_{\bar{1}\bar{2}}(z)$. We denote the set of such function-tuples by \mathcal{T}_4 . The goal is to compare time tags r and t obtained at settings ab by computing $f_{ab}(t - r)$. To construct a function-tuple in \mathcal{T}_4 from any given $f_{\bar{2}\bar{1}}$, $f_{\bar{1}\bar{1}}$, and $f_{\bar{1}\bar{2}}$, we can choose $f_{\bar{2}\bar{2}}$ such that $f_{\bar{2}\bar{2}}(u) \leq \inf_{x,y \in \mathbb{R}} (f_{\bar{2}\bar{1}}(x) + f_{\bar{1}\bar{1}}(y) + f_{\bar{1}\bar{2}}(u - x - y))$, provided the expression on the right-hand side is bounded from below. This condition is satisfied if the given f_{ab} are lower bounded. Three immediate examples of members of \mathcal{T}_4 are the linear tuples with $f_{ab}(x) = \lambda x$, the constant tuples with $f_{ab}(x) = c_{ab}$ where $c_{\bar{2}\bar{1}} + c_{\bar{1}\bar{1}} + c_{\bar{1}\bar{2}} = c_{\bar{2}\bar{2}}$ (an *exact* and *constant* function-tuple), and $f_{ab}(x) = [x \geq 0]$. Here, we again used the convention that for a logical formula ϕ , $[\phi] = 1$ if ϕ is true and 0 otherwise. To construct other members of \mathcal{T}_4 , it helps to apply closure properties of \mathcal{T}_4 .

Theorem 1. The set of function-tuples \mathcal{T}_4 is closed under the following operations.

- Componentwise addition and multiplication by a positive real number (\mathcal{T}_4 is a convex cone).
- The reflection defined by $f'_{ab}(x) = f_{ab}(-x)$.
- Componentwise maximum of two function-tuples.
- For real numbers t_{ab} satisfying $t_{\bar{2}\bar{2}} = \sum_{ab \neq \bar{2}\bar{2}} t_{ab}$, the shift transforming the components according to $f'_{ab}(x) = f_{ab}(x + t_{ab})$.
- For non-negative function-tuples and $c \geq 0$, the transformation defined by $f'_{ab}(x) = \min(f_{ab}(x), c)$.
- Let f and f' be function-tuples with $f_{\bar{2}\bar{2}}$ monotone nondecreasing. Then, the function-tuple defined by $f''_{ab} = f'_{ab} \circ f_{ab}$ is a function-tuple. If $f_{\bar{2}\bar{2}}$ is also monotone nondecreasing, then so is $f''_{\bar{2}\bar{2}}$.

Proof. (a) can be checked by direct application of the definitions. (c) follows from the observation that the maximum is monotone in each argument and the maximum of two

matched sums is at most the sum of the maximums of the terms. (b) and (d) follow from invariance of the defining inequalities under reflection and under the shift specified in (d). To check (e) note that $f'_{22}(x + y + z) \leq f_{22}(x + y + z) \leq f_{21}(x) + f_{11}(y) + f_{12}(z)$. If a term on the right-hand side is greater than c , non-negativity implies that the right-hand side is at least c , an upper bound on the left-hand side by definition of f'_{22} . If not, the right-hand side is equal to $f'_{21}(x) + f'_{11}(y) + f'_{12}(z)$. The following inequalities show (f):

$$\begin{aligned} f''_{22}(x + y + z) &= f''_{22}(f_{22}(x + y + z)) \\ &\leq f'_{22}(f_{21}(x) + f_{11}(y) + f_{12}(z)) \\ &\leq f'_{21}(f_{21}(x)) + f'_{11}(f_{11}(y)) + f'_{12}(f_{12}(z)) \\ &= f''_{21}(x) + f''_{11}(y) + f''_{12}(z), \end{aligned} \tag{10}$$

and the composition of monotone nondecreasing functions is monotone nondecreasing. ■

Here are a few more examples of function-tuples in \mathcal{T}_4 .

(1) $f_{ab}(x) = |x| = \max(x, -x)$. In this case, the condition is just the twice-iterated triangle inequality for the reals.

(2) One-sided threshold functions. Let real numbers $(t_{ab})_{a,b}$ satisfy $t_{22} = \sum_{ab \neq 22} t_{ab}$ and define $f_{ab}(x) = [x \geq t_{ab}]$. Note that this is an application of the shift in Theorem 1(d) with parameters $(-t_{ab})_{a,b}$ to the function-tuple $g_{ab}(x) = [x \geq 0]$.

(3) Half-linear functions. For t_{ab} as in (2), a positive slope m and an exact constant function-tuple c_{ab} , define $f_{ab}(x) = \max(m(x - t_{ab}), c_{ab})$. That this tuple is in \mathcal{T}_4 follows from the closure properties applied to the generating examples given above.

(4) Linear-edge window functions. Choose thresholds $t_{l,ab} \leq t_{h,ab}$ such that for $d = l$ and for $d = h$, $t_{d,22} = \sum_{ab \neq 22} t_{d,ab}$, and positive slopes m_l and m_h . Define

$$f_{ab}(x) = \min(1, \max(0, m_h(x - t_{h,ab}), m_l(t_{l,ab} - x))). \tag{11}$$

As illustrated in Fig. 3, these functions are 0 between $t_{l,ab}$ and $t_{h,ab}$, and rise linearly away from these thresholds up to a value of 1. That they form tuples in \mathcal{T}_4 follows from the closure properties.

In our applications, we use the linear-edge window functions. In spot checks using linear programming, they appear to optimize the sought for violations among non-negative tuples in \mathcal{T}_4 that are 1 outside a fixed interval.

We now construct a CH function l_f for pairs of time-tag sequences from an arbitrary function-tuple $f = (f_{ab})_{a,b \in \{1,2\}}$ in \mathcal{T}_4 . Let \mathbf{r}, \mathbf{t} be two ordered time-tag sequences, $\mathbf{r} = (r_1 \leq \dots \leq r_m)$ and $\mathbf{t} = (t_1 \leq \dots \leq t_n)$. Let \mathcal{M} be the family

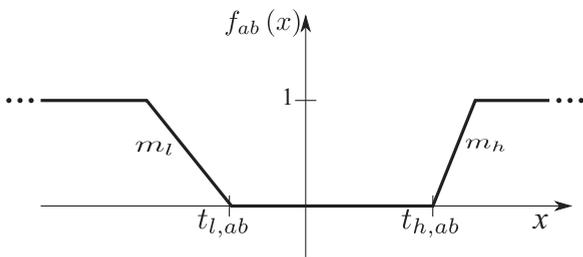


FIG. 3. Illustration of the linear-edge window functions. See the text for the definitions.

of partial, noncrossing matchings between \mathbf{r} and \mathbf{t} . Such matchings M can be identified with one-to-one, partial, monotone functions $M : k \in \text{dom}(M) \subseteq [m] \mapsto M(k) \in [n]$. Monotonicity implies that if $k, l \in \text{dom}(M)$ and $k < l$, then $M(k) < M(l)$. (We use the notation $[j] = \{1, \dots, j\}$.) Let $l_{f,ab}(\mathbf{r}, \mathbf{t})$ be the minimum over all $M \in \mathcal{M}$ of the “cost”

$$l(f_{ab}, M, \mathbf{r}, \mathbf{t}) = m - |\text{dom}(M)| + \sum_{k \in \text{dom}(M)} f_{ab}(t_{M(k)} - r_k). \tag{12}$$

One way to think of this is as the minimum total cost of editing \mathbf{r} into \mathbf{t} by deleting time tags in \mathbf{r} at a cost of 1 (assessed by $m - |\text{dom}(M)|$), deleting time tags of \mathbf{t} at no cost, and by shifting the remaining time tags of \mathbf{r} by x at a cost of $f_{ab}(x)$ (assessed by the sum over $k \in \text{dom}(M)$), where each time tag can be shifted at most once and the final time ordering is the same as the initial one. We can also view this as a maximum weighted bipartite noncrossing matching problem, where the matching is between indices of \mathbf{r} and indices of \mathbf{t} with the weight of (k, l) being $(1 - f_{ab}(t_l - r_k))$. The cost is then given by m minus the weight of the maximum-weight matching. An example of the cost computation is in Fig. 4.

Theorem 2. Suppose that the measurement outcome space \mathcal{O} consists of time-tag sequences. Let f be a function-tuple in \mathcal{T}_4 . Then, l_f [as defined before Eq. (12)] is a CH function.

Proof. Consider deterministic LR outcomes d_c^X as introduced previously, but with outcomes consisting of time-tag sequences. Let M_{ab} be the cost-minimizing matchings for which $l(f_{ab}, M_{ab}, d_a^A, d_b^B) = l_{f,ab}(d_a^A, d_b^B)$ if $ab \neq \bar{1}\bar{1}$, and $l(f_{\bar{1}\bar{1}}, M_{\bar{1}\bar{1}}, d_1^B, d_1^A) = l_{f,\bar{1}\bar{1}}(d_1^B, d_1^A)$. We can construct a matching M' from d_2^A to d_2^B by composing $M' = M_{\bar{1}\bar{2}} \circ M_{\bar{1}\bar{1}} \circ M_{2\bar{1}}$, with domain consisting of those indices for which the composition is defined. Then, M' is monotone and $l(f_{22}, M', d_2^A, d_2^B) \geq l_{f,22}(d_2^A, d_2^B)$. Therefore, it suffices to show that

$$\begin{aligned} l(f_{22}, M', d_2^A, d_2^B) &\leq l(f_{2\bar{1}}, M_{2\bar{1}}, d_2^A, d_1^B) + l(f_{\bar{1}\bar{1}}, M_{\bar{1}\bar{1}}, d_1^B, d_1^A) \\ &\quad + l(f_{\bar{1}\bar{2}}, M_{\bar{1}\bar{2}}, d_1^A, d_2^B). \end{aligned} \tag{13}$$

The composition of functions defining M' terminates at the first step where the mapped-to index fails to be in the domain of the next matching. This allows us to associate with each index not in the domain of M' a unique index along the way that is “deleted” in the next step. Indices in the domain pass through each matching and accumulate separate distances that bound

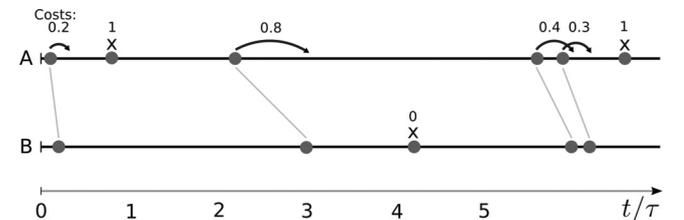


FIG. 4. Example of a time-tag–distance calculation. The beginnings of two time-tag sequences are shown. The gray circles are positioned at the recorded times of the time tags. The settings may be assumed to be $\bar{1}\bar{1}$. The cost function is $f_{\bar{1}\bar{1}}(x) = \min(1, |x/\tau|)$. The edit action is indicated, with “x” marking a deleted time tag and arrows showing the order-preserving moves. The costs are shown above the edit action. The total cost of the edits shown here is 3.7.

the corresponding distance in M' . To formalize this idea, for each index k of d_2^A , we define $N(k)$ as a pair consisting of an index and a party and setting label as follows: If $k \notin \text{dom}(M_{2\bar{1}})$, then let $N(k) = (k, A\bar{2})$. Else, if $M_{2\bar{1}}(k) \notin \text{dom}(M_{1\bar{1}})$, then let $N(k) = (M_{2\bar{1}}(k), B\bar{1})$. Else, if $M_{1\bar{1}}(M_{2\bar{1}}(k)) \notin \text{dom}(M_{1\bar{2}})$, then let $N(k) = (M_{1\bar{1}}(M_{2\bar{1}}(k)), A\bar{1})$. If none of these conditions apply, k is in the domain of M' and we let $N(k) = (M'(k), B\bar{2})$. The definition implies that N is one to one, $k \in \text{dom}(M')$ iff the second component of $N(k)$ is $B\bar{2}$, and if $k \notin \text{dom}(M')$, then the first component of $N(k)$ is not in the domain of one of $M_{2\bar{1}}$, $M_{1\bar{1}}$, or $M_{1\bar{2}}$. This ensures that all members of d_2^A deleted according to M' are matched to deleted members of one of the time-tag sequences in the composition. In particular, the deletion cost on the left-hand side of Eq. (13) is at most that on the right-hand side.

We now focus on the shift costs contributing to Eq. (13). Let $(d_c^X)_k$ be the k th time tag of d_c^X . It remains to be shown that

$$\begin{aligned} & \sum_{k \in \text{dom}(M')} f_{2\bar{2}}[(d_2^B)_{M'(k)} - (d_2^A)_k] \\ & \leq \sum_{k \in \text{dom}(M_{2\bar{1}})} f_{2\bar{1}}[(d_1^B)_{M_{2\bar{1}}(k)} - (d_2^A)_k] \\ & \quad + \sum_{k \in \text{dom}(M_{1\bar{1}})} f_{1\bar{1}}[(d_1^A)_{M_{1\bar{1}}(k)} - (d_1^B)_k] \\ & \quad + \sum_{k \in \text{dom}(M_{1\bar{2}})} f_{1\bar{2}}[(d_1^B)_{M_{2\bar{1}}(k)} - (d_2^A)_k]. \end{aligned} \quad (14)$$

Because all shift costs are positive, it suffices to show that

$$\begin{aligned} & \sum_{k \in \text{dom}(M')} f_{2\bar{2}}[(d_2^B)_{M'(k)} - (d_2^A)_k] \\ & \leq \sum_{k \in \text{dom}(M')} \{f_{1\bar{2}}[(d_2^B)_{M'(k)} - (d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)}] \\ & \quad + f_{1\bar{1}}[(d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)} - (d_1^B)_{M_{2\bar{1}}(k)}] \\ & \quad + f_{2\bar{1}}[(d_1^B)_{M_{2\bar{1}}(k)} - (d_2^A)_k]\}, \end{aligned} \quad (15)$$

in which terms on the right-hand side that are not included in the composed matchings have been neglected. For each $k \in \text{dom}(M')$,

$$\begin{aligned} & f_{2\bar{2}}[(d_2^B)_{M'(k)} - (d_2^A)_k] \\ & = f_{2\bar{2}}[(d_2^B)_{M'(k)} - (d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)} + (d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)} \\ & \quad - (d_1^B)_{M_{2\bar{1}}(k)} + (d_1^B)_{M_{2\bar{1}}(k)} - (d_2^A)_k] \\ & \leq f_{1\bar{2}}[(d_2^B)_{M'(k)} - (d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)}] \\ & \quad + f_{1\bar{1}}[(d_1^A)_{M_{1\bar{1}} \circ M_{2\bar{1}}(k)} - (d_1^B)_{M_{2\bar{1}}(k)}] \\ & \quad + f_{2\bar{1}}[(d_1^B)_{M_{2\bar{1}}(k)} - (d_2^A)_k], \end{aligned} \quad (16)$$

according to the defining inequality of function-tuples. These summands are separate contributions [for distinct $k \in \text{dom}(M')$] to the right-hand side of Eq. (13), completing the proof. \blacksquare

An algorithm to compute the minimum costs $l_{f,ab}(\mathbf{r}, \mathbf{t})$ is described in Appendix B.

VI. NONSIGNALING ADJUSTMENTS TO CH FUNCTIONS

The nonsignaling conditions are a set of equalities that constrain probability distributions describing A 's and B 's measurement outcomes. They ensure that each party's outcome distribution is independent of the other party's setting. Otherwise, one party could signal their setting to the other party. Formally, given nonsignaling and any real-valued function h on \mathcal{O} , $\langle h(O^A) | S^A = a, S^B = \bar{1} \rangle = \langle h(O^A) | S^A = a, S^B = \bar{2} \rangle$ and similarly for reversing the roles of A and B . Although both quantum and LR theories obey the nonsignaling equalities, using these equalities to transform Bell inequalities can increase the signal-to-noise ratio (SNR) of a Bell-inequality violation observed in an experimental test of LR.

Consider the general iterated triangle inequality for a CH function l . We can modify l by defining

$$l'_{ab}(x, y) = l_{ab}(x, y) + \begin{cases} -f_{\bar{1}}(y) - g_{\bar{1}}(x) & \text{if } a = \bar{1} \text{ and } b = \bar{1}, \\ f_a(x) + g_b(y) & \text{otherwise,} \end{cases} \quad (18)$$

where f_a and g_b are arbitrary real-valued functions. Replacing l by l' leaves the right-hand side of Eq. (7) unchanged, so l' is also a CH function. Furthermore, we have $\langle B_l(T) \rangle = \langle B_{l'}(T) \rangle$ for any model satisfying the nonsignaling conditions. We call l' a nonsignaling adjustment of l . Note that the functions

$$l_{A,ab}(x, y) = \begin{cases} -f_{\bar{1}}(y) & \text{if } a = \bar{1} \text{ and } b = \bar{1}, \\ f_a(x) & \text{otherwise} \end{cases} \quad (19)$$

and

$$l_{B,ab}(x, y) = \begin{cases} -g_{\bar{1}}(x) & \text{if } a = \bar{1} \text{ and } b = \bar{1}, \\ g_b(y) & \text{otherwise} \end{cases} \quad (20)$$

are CH functions, and $l'_{ab} = l_{ab} + l_{A,ab} + l_{B,ab}$.

Nonsignaling adjustments can be used to improve the SNR of the empirical estimate of a Bell function B_l obtained from a sequence of trials. As a simple example with a two-point outcome space, consider N trials to test a version of the inequality of Eq. (6). The experiment is configured so that each trial involves emission of exactly one photon pair with some probability (and emission of nothing otherwise), and the measurement outcomes 1 and 0 correspond to whether a photon was detected or not. We start with $l_{ab}(x, y) = \max(x - y, 0)$, which is a CH function. Let p_c^X be the probability that X detects a photon at setting c in a trial. Let c_{ab} be the probability of coincident detections when the parties use settings ab . The expected value of the Bell function of Eq. (5) is then

$$\begin{aligned} & \left\{ \begin{aligned} & \langle 4[S^A = \bar{2} \text{ and } S^B = \bar{1}]l_{2\bar{1}}(O^A, O^B) \rangle \\ & + \langle 4[S^A = \bar{1} \text{ and } S^B = \bar{1}]l_{1\bar{1}}(O^B, O^A) \rangle \\ & + \langle 4[S^A = \bar{1} \text{ and } S^B = \bar{2}]l_{1\bar{2}}(O^A, O^B) \rangle \\ & - \langle 4[S^A = \bar{2} \text{ and } S^B = \bar{2}]l_{2\bar{2}}(O^A, O^B) \rangle \end{aligned} \right\} \\ & = \left\{ \begin{aligned} & \langle 4[S^A = \bar{2} \text{ and } S^B = \bar{1}][O^A = 1 \& O^B = 0] \rangle \\ & + \langle 4[S^A = \bar{1} \text{ and } S^B = \bar{1}][O^B = 1 \& O^A = 0] \rangle \\ & + \langle 4[S^A = \bar{1} \text{ and } S^B = \bar{2}][O^A = 1 \& O^B = 0] \rangle \\ & - \langle 4[S^A = \bar{2} \text{ and } S^B = \bar{2}][O^A = 1 \& O^B = 0] \rangle \end{aligned} \right\} \end{aligned}$$

$$= \begin{cases} p_2^A - c_{2\bar{1}} \\ + p_1^B - c_{1\bar{1}} \\ + p_1^A - c_{1\bar{2}} \\ - (p_2^A - c_{2\bar{2}}). \end{cases} \quad (21)$$

Note that the terms p_2^A cancel in this expression. (The last expression matches half of the original CH inequality in the paragraph of Eq. (B8) of Ref. [22] after relabeling and a change of sign.) However, when estimating the expectation by evaluating the Bell function on the trials, the two p_2^A 's are contributed by values at different settings, namely $\bar{2}\bar{1}$ and $\bar{2}\bar{2}$. Consequently, two sources of counting statistics variability associated with detections of A at setting $\bar{2}$ affect the the SNR of the Bell function. We can eliminate this problem by defining $f_{\bar{2}}(x) = -x$ in Eq. (18). (Here, f_a and g_b are set to zero if not explicitly assigned.) By the nonsignaling constraints and construction, the modified Bell function has the same expectation as the original Bell function. This Bell function was used to show a Bell-inequality violation in the experiment reported in Ref. [18]. A further improvement of the SNR is obtained by ‘‘distributing’’ the terms whose expectations are p_1^B and p_1^A over the different settings. It is equivalent to averaging them over the other party’s setting choices and involves setting $f_{\bar{1}}(x) = -x/2$ and $g_{\bar{1}}(x) = x/2$ in Eq. (18), in addition to setting $f_{\bar{2}}(x) = -x$. This modification was introduced for the explicit purpose of improving the SNR of the violation in Ref. [19].

VII. PROTOCOL FOR DATA ANALYSIS

When analyzing a set of time-tag–sequence pairs from trial measurement outcomes, it is necessary to choose a Bell function B_l that, before the experiment, can be expected to show good results. It may not be feasible to make a good choice of B_l on purely theoretical grounds. One can instead acquire a statistically useful training set consisting of the outcomes from the first N_l trials and set aside the remainder in the analysis set. The notion of ‘‘statistically useful’’ is not formalized here. The training set is used to choose the parameters required for analyzing the rest of the data. The training set is excluded from the final analysis. The main task is to determine a function-tuple in \mathcal{T}_4 and nonsignaling adjustments to use for defining B_l . In principle, one can optimize the function-tuple on the training set. That is, one can compute B_l for all such function-tuples and pick the one that minimizes the empirically computed value of B_l . This optimization is difficult, but one can use nonlinear optimization on the eight independent parameters of the linear-edge window functions. We note that on spot checks, this subset of \mathcal{T}_4 appears to contain the optimal solution among function-tuples of \mathcal{T}_4 whose members are constrained to be 1 outside a fixed interval $[-u, u]$. In the simulations following, the number of independent parameters was reduced to 2 by taking advantage of symmetries. The suggested nonlinear optimization can still be too resource intensive. In our simulations, we used an effective approximation; see Appendix C. We do not discuss methods for optimizing the nonsignaling adjustments here. In the simulations, we just use the ones described in the last paragraph of Sec. VI. An even more ambitious optimization could seek to optimize the SNR or the statistical significance

of the violation of the inequality rather than the expected value of B_l .

VIII. BELL FUNCTION TRUNCATION FOR P -VALUE BOUNDS

Consider N trials whose trial outcomes are t_k and a Bell function B . A direct way to analyze the trial outcomes is to compute $b_k = B(t_k)$, let $f = \sum_k b_k$, and determine the sample standard error for f as $\sigma_e = \sqrt{N \sum_k (b_k - f/N)^2 / (N - 1)}$. The violation can then be quantified by the ‘‘number of standard deviations of violation’’, $-f/\sigma_e$, which is the SNR of the total violating signal. This number needs to be interpreted with care, particularly if it is very large, a desirable outcome of an experiment. If the trials are i.i.d., then $f/N \pm \sigma_e/N$ is an approximate 68% confidence interval for the expectation of B . (A better method that takes advantage of the fact that the settings probability distribution is known is given in Appendix A.) However, even in the case of i.i.d. trials, $-f/\sigma_e$ does not quantify how strongly the experiment ‘‘rejects’’ LR models. This is because the central limit theorem cannot be reliably used to estimate extreme tail probabilities. Furthermore, the assumption that the trials are identical rarely holds to high precision, and independence cannot be assumed in applications to cryptographic protocols. For more details on these issues, see [16,17].

To determine the statistical significance of a Bell-inequality violation, one can compute a bound on the largest probability with which any LR model could produce a violation at least as large as that observed. This upper bounds a p value according to the theory of statistical hypothesis testing with respect to the composite null hypothesis consisting of all possible LR models. Typical Bell-inequality experiments thus aim for extremely small p values. Given a p -value bound p , it is convenient to quantify the violation in terms of the (negative) log- p -value bound, formally defined as $-\log_2(p)$. In many cases, for example when reporting a discovery in particle physics, p values are converted to equivalent standard deviations with respect to the one-sided tail probabilities of the standard normal distribution. For comparison, the log- p -values corresponding to 1, 2, 3, 4, and 5 standard deviations are 2.7, 5, 9.5, 14.9, and 21.7.

The first rigorous method for computing such a bound was given by Gill [14,15] and is based on martingale theory. This ‘‘martingale-based protocol’’ does not require that trials be independent from one trial to the next, or that they have identically distributed measurement outcomes, an important feature for its application to quantum randomness expansion [6]. The bound obtained by the martingale-based protocol is suboptimal, but there is a protocol, the PBR protocol [16], that is optimal in an asymptotic sense. Like the martingale-based protocol, the PBR protocol also does not require independent or identical outcomes. The PBR protocol has the advantage that it does not require a predetermined Bell inequality, nor does it require that the number of trials be decided in advance: it gives valid p -value bounds for any stopping rule (see the last paragraph of this section). The full PBR protocol is computationally infeasible when the measurement outcome spaces or the number of settings are large, but there is a simplified and efficient version of the protocol

that still outperforms the martingale-based protocol [17] while retaining the advantages of the full protocol.

The PBR protocol is based on the following observation. Suppose that before the k th trial, we can determine a “test factor” P_k on the space of possible trial outcomes such that $P_k \geq 0$ and for all LR models, $\langle P_k(T_k) \rangle_{\text{LR}} \leq 1$, where the bound holds regardless of what happened before the k th trial in the experiment. Then, $P = \langle \prod_{k=1}^N P_k(T_k) \rangle_{\text{LR}} \leq 1$. Thus, given LR, according to the Markov inequality, the probability that $P > 1/p$ is bounded above by p , and therefore $1/P$ is a p -value bound. In general, candidate test factors R can be obtained from Bell functions B bounded above by z according to $R = (z - B)/z$. More generally, given a collection of candidate test factors $R_i \geq 0$ satisfying $\langle R_i(T) \rangle_{\text{LR}} \leq 1$, any convex combination of the R_i can be used as P_k . (The collection may depend on k .) The simplified PBR protocol takes as input such a collection $(R_i)_{i=0}^m$ and chooses P_k before the k th trial by optimizing the convex combination on the outcomes of the previous trials. (R_0 is always chosen to be the “trivial” test factor 1.) Specifically, it maximizes the empirical estimate of the log- p -value increase per trial given by $\sum_{i=1}^{k-1} -\log_2[P_k(t_i)]/(k-1)$. For more details, see Ref. [17]. The possibility of adjusting the test factors for the upcoming trial after each trial makes it possible to avoid making predetermined choices for the Bell function parameters.

In practice, there is little gained by reoptimizing the test factors before every trial, and instead the factor is reused until sufficiently many new trials have been obtained. In principle, the number of trials before P_k can be productively updated may be determined from statistical considerations. (Our implementations so far are largely based on heuristic considerations.) The simplest method is to determine the optimal test factor from the training set and use it uniformly on the analysis trials. For example, the starting Bell functions can include multiple choices of parameters. The optimal convex combination of the corresponding test factors can then be determined empirically on the training set and can be directly applied to the trials in the analysis set.

To apply the PBR protocol to time-tag data, we make two modifications to the CH Bell functions. The first modification ensures that we have a set of Bell functions that are bounded from above as required by the simplified PBR protocol. Finite bounds are normally not available for the time-tag-sequence Bell functions discussed so far, or the bounds are too high to be useful, so we describe a truncation strategy below. Our second modification increases the expected log- p -value bound produced by the PBR protocol by shifting the CH functions so that the contributions to the violation of the CH Bell inequality are equalized across measurement settings.

Consider a CH function l . To obtain a bounded CH Bell function, it suffices to modify l by composition with a function-tuple f_{ab} in \mathcal{T}_4 for which $f_{\bar{2}\bar{2}}$ is monotone nondecreasing and the f_{ab} are bounded. We call such an f_{ab} a monotone and bounded function-tuple. To see that this preserves the desired inequalities, consider Eq. (7) and define $l'_{ab}(d_a^X, d_b^Y) = f_{ab}(l_{ab}(d_a^X, d_b^Y))$. Because of monotonicity of $f_{\bar{2}\bar{2}}$, Theorem 1(f) ensures that l' satisfies the inequality in Eq. (7), thus l' is also a CH function.

A convenient family of monotone and bounded function-tuples in \mathcal{T}_4 is provided by

$$g_{ab}(x; b, u, c) = \min(\max(x + b_{ab}, 0), c) - u_{ab}, \quad (22)$$

where b_{ab} and u_{ab} are exact constant tuples in \mathcal{T}_4 and $c \geq 0$. That g is a function-tuple follows from the closure properties of \mathcal{T}_4 (Theorem 1). If l is modified to l' by means of a function-tuple of the form g_{ab} , the Bell function $B_{l'}$ is guaranteed to be bounded above by the maximum of $c - u_{\bar{1}\bar{1}}, c - u_{\bar{1}\bar{2}}, c - u_{\bar{2}\bar{1}}, u_{\bar{2}\bar{2}}$. The following is a method for systematically choosing the truncation parameters.

Consider a collection of training trials. Let $\mathbf{l}_{ab} = (l_{ab,k})_{k=1}^{N_{ab}}$ consist of the observed values of $l_{s^A, s^B}(o^A, o^B)$ for the trials where $s^A = a$ and $s^B = b$. Let $\bar{l}_{ab} = \sum_k l_{ab,k}/N_{ab}$. The bounds b_{ab} can be chosen as a compromise between having small bounds on the Bell function and preserving the variation in the values of $l_{ab,k}$. The violating signal is reduced if we truncate the values of $l_{\bar{2}\bar{2}}$ so that the maximum value is too close to the mean $\bar{l}_{\bar{2}\bar{2}}$. This truncation point is determined by solving $x_{\bar{2}\bar{2}} + b_{\bar{2}\bar{2}} = c$, that is, $x_{\bar{2}\bar{2}} = c - b_{\bar{2}\bar{2}}$. (The upper or lower truncation point is the value of x for which g_{ab} reaches its upper or lower bound.) Similarly, we should not truncate the other l_{ab} so that their minimum values are too close to their means \bar{l}_{ab} . These truncation points are at $x_{ab} = -b_{ab}$. Let w_{ab} be a “safe” separation between the truncation points and \bar{l}_{ab} , to be determined from the distributions of the $l_{ab,k}$ (see the end of the next paragraph). For $ab \neq \bar{2}\bar{2}$, we can set b_{ab} by solving $-b_{ab} = \bar{l}_{ab} - w_{ab}$. This determines $b_{\bar{2}\bar{2}} = \sum_{ab \neq \bar{2}\bar{2}} b_{ab}$. We can then choose c so that $c - b_{\bar{2}\bar{2}} = \bar{l}_{\bar{2}\bar{2}} + w_{\bar{2}\bar{2}}$.

The next step is to choose u_{ab} so as to ensure that the contributions to the violation conditional on the settings are equalized. This is done to improve the expectation of the log- p -value bound by exploiting the concavity of the logarithm. Let $l'_{ab,k} = g_{ab}(l_{ab,k}; b, 0, c)$ with b and c as obtained so far. Define \bar{l}'_{ab} accordingly. We choose u_{ab} so as to balance the average violation for the different settings. This is accomplished by defining

$$u_{ab} = \bar{l}'_{ab} - (-1)^{\{a=\bar{2}\&b=\bar{2}\}}(\bar{l}'_{\bar{1}\bar{1}} + \bar{l}'_{\bar{1}\bar{2}} + \bar{l}'_{\bar{2}\bar{1}} - \bar{l}'_{\bar{2}\bar{2}})/4. \quad (23)$$

If we then use g_{ab} as defined in Eq. (22) and modify l as described there with $f_{ab}(x) = g_{ab}(x; b_{ab}, u_{ab}, c)$, this ensures that each trial contributes the same estimated violation $-v = (\bar{l}'_{\bar{1}\bar{1}} + \bar{l}'_{\bar{1}\bar{2}} + \bar{l}'_{\bar{2}\bar{1}} - \bar{l}'_{\bar{2}\bar{2}})/4$ on average. If this violation is not negative, then this truncation is not helpful for use in the PBR protocol. Assuming there is an empirical violation according to the original \bar{l}_{ab} , the separations w_{ab} need to be increased until a violation remains visible in the \bar{l}'_{ab} . In this case, if the violation persists in future trials, the PBR protocol can take advantage of the truncation. In our implementation, rather than attempting to find the optimal choice for w_{ab} , we consider a small set of good candidates, obtain the associated Bell functions, and convert them to the nontrivial test factors $R_i, i \geq 1$ that are then convexly combined with the trivial test factor by the simplified PBR protocol as described above. We found that the convex combination chosen by the protocol normally involved more than one choice of w_{ab} , suggesting that a single choice is not optimal.

We remark that the shift u_{ab} can be expressed as trivial nonsignaling adjustments by constant functions. In addition to

increasing the SNR overall, a goal of nonsignaling adjustments might be to equalize the SNR conditional on the settings, but we did not attempt to achieve this.

The truncation and shifting strategy above results in factors P_k (and candidate test factors R_i) whose predicted expectations are upper bounded by $\frac{4}{3}$, according to the training set. To see this, we show that the upper bound z of the modified Bell function that determines a candidate test factor is at least $3v$, while by construction, the predicted expectation is $-v$. Here, v depends on the training data and choice of w_{ab} for this test factor. Since the test factor is given by $R = (z - B)/z$, we find that the predicted expectation of R is at most $(z + v)/z \leq \frac{4}{3}$. The upper bound is given by $z = \max_{ab,x} (-1)^{[a=2 \text{ and } b=2]} g_{ab}(x; b_{ab}, u_{ab}, c)$. The expression for g_{ab} shows that $\max_x g_{ab}(x; b_{ab}, u_{ab}, c) = c - u_{ab}$ and $\max_x -g_{22}(x; b_{22}, u_{22}, c) = u_{22}$. Consider u_{22} , which is the sum of the other u_{ab} (by definition of exact constant tuples and by construction). For $ab \neq 22$, the lower bound on $g_{ab}(x; b_{ab}, u_{ab}, c)$ is $-u_{ab}$. But, the sample mean at setting ab of the truncated Bell function is $-v$, which requires that its lower bound satisfies $-u_{ab} \leq -v$. It follows that $u_{22} \geq 3v$, which completes the argument.

A consequence of this observation is that the expected increase in the log- p -value per trial is bounded by $\log_2(\frac{4}{3})$, and it is not possible to take advantage of seemingly strong violating signals per trial. To some extent, this is unavoidable: We are making no assumptions on the probability distribution of the time-tag sequences, and an extremely adversarial LR model could take advantage of this in future trials given our choice of parameters for the PBR protocol. For the purpose of making the most of the PBR protocol, it is therefore advantageous to design the individual trials to have statistically small violating signals. In particular, it is favorable to have the one-trial SNR be well below 1. A simple way to accomplish this and get a better overall log- p -value bound is to shorten the durations of the trials and increase the number of trials proportionally.

We finish this section by explaining our comment that the PBR protocol can be used with any stopping rule, such as one according to which one collects trials until a desired p -value bound is observed. To see this, virtually replace the experiment with the stopping rule by one that performs a fixed, large number of trials, larger than the maximum number of trials that could be performed by the original experiment. When the original experiment's stopping rule says "stop," the new experiment sets all future test factors to 1. This is justified because the experimenter's choice of the test factors for trial k is only constrained by $P_k \geq 0$ and $\langle P_k(T_k) \rangle_{\text{LR}} \leq 1$, which are satisfied by $P_k = 1$. The two experiments have the same statistics for the p -value bounds obtained and the virtual experiment's p -value bounds are valid according to the theory of the PBR protocol. We remark that the PBR protocol can be viewed as an application of the theory of test supermartingales as reviewed in Ref. [25].

IX. DEMONSTRATIONS ON SIMULATED DATA

To illustrate time-tag-sequence analysis, we simulated experiments intended to test inequalities such as Eq. (5) and its nonsignaling variations. The situation is as described in Sec. III

with a Poisson source of polarization-entangled photon pairs and high overall efficiency. We assume that both arms of the experiment have identical efficiency η . In principle, such tests can succeed if $\eta > \frac{2}{3}$ [26]. We explored the effects of uniform jitter (defined in the next paragraph) on the performance of such experiments at an efficiency of $\eta = 0.8$. We also found lower bounds on the maximum uniform and exponential jitter (defined in the next paragraph) for which our techniques can show LR violation for $\eta \geq 0.74$. For each efficiency being considered, we first optimized the violation of the CHSH inequality in Eq. (1) by varying the settings' choices and the parameter θ in the family of unbalanced Bell states defined by $\cos(\theta)|00\rangle + \sin(\theta)|11\rangle$. Given the efficiency, an optimal state, and settings, we computed the probabilities of the measurement outcomes conditional on the settings.

We considered two families of time-tag jitter distributions for the difference between the recorded time tag and the true arrival time of a photon. The first is the uniform distribution on an interval of width j_u . That is, given the true arrival time t , the recorded time tag t' is uniformly distributed in the interval $[t, t + j_u]$. The second family is an exponential distribution with density $\gamma e^{-\gamma(t'-t)}$ for $t' \geq t$. The two families of distributions were chosen for ease of calculation and to illustrate the effect of no tail versus long-tail behavior, with long tails leading to greater loss of violating signal.

The procedures for simulating and analyzing an experiment were automated. We generated simulated photon pairs at a normalized rate of 1 per (arbitrary) unit of time. Thus, the numerical value of τ (the mean photon-pair interarrival time) is 1 in these units. From here on, time quantities such as j_u are given as numerical values with respect to these units. The procedure is based on a choice of observation window T for the time-tag sequences, number of training trials N_t , and number of analysis trials N_a . In principle, these can be chosen before the simulation is started to ensure sufficient data for determining the needed analysis parameters from the training set. We recall that for the PBR analysis to take full advantage of the SNR, it is a good idea to choose T so that the SNR for one trial is below 1. On the other hand, if T is too small, loss of coincidences near the boundary due to jitter leads to an additional reduction of the violating signal. To reflect the conditions of experiments with always-on pump lasers, we start generating photon pairs 2 units of time before the beginning of the observation window. The data presented below use $T = 1000$, $N_t = 10\,000$, $N_a = 200\,000$. With these parameters and given the jitter distribution, the procedure for generating and analyzing data is as follows:

- (1) Generate the trials for the training and analysis sets. For each trial, first produce the sequence of times at which photon pairs arrive at the detectors according to a Poisson process of rate 1 as described above. The jitter distribution is used to delay the recorded time of detection independently for the two parties. The time tags inside the observation window are saved to the parties' time-tag sequences.

- (2) Determine the analysis parameters from the training set. Three studies are performed, each of which requires optimized parameters. The first is a conventional coincidence analysis, for which the coincidence window width is determined by optimizing the resulting violation on the training set. The second computes the time-tag-sequence

distances based on the linear-edge window function-tuples. The function-tuple parameters are optimized on the training set. The tuples are restricted to be reflection symmetric around 0, and the thresholds for the settings other than $\bar{2}\bar{2}$ are taken to be identical. Thus, only two parameters need to be optimized. (The optimization algorithm is described in Appendix C.) The third is the PBR analysis on the truncated Bell function. The truncation method and optimization of the anticipated $\log -p$ -value increase per trial described in Sec. VIII are used.

(3) Perform the three analyses on the analysis set using the parameters determined from the training set.

The conventional coincidence analysis used in this procedure is equivalent to using identical window functions with infinite edge slopes to compute the distances between each pair of time-tag sequences according to the prescription for using function-tuples. Because the function-tuple effectively used is not in \mathcal{T}_4 , there is no guarantee that the targeted Bell inequality is strictly satisfied by LR models. The coincidence loophole examples demonstrate that such an inequality requires additional assumptions on the nature of the detection events. We remark that the analysis described in [27] is related to using window functions with infinite edge slopes, but the width of the window function for setting $ab = \bar{2}\bar{2}$ is three times the width of that used for the other settings. This set of window functions is a function-tuple in \mathcal{T}_4 and produces an analysis free of the coincidence loophole.

A. Uniform jitter distributions at efficiency = 0.8

We simulated experiments where the jitter has a uniform distribution on the interval $[0, j_u]$, where $j_u \in [0.001, 0.225]$. This covers most of the range for which the conventional coincidence analysis shows a violating signal. Figure 5 shows the nominal SNRs for the conventional and the time-tag analyses. The nominal SNR is the ratio of the violating signal to the sample standard deviation. Positive values are violating, negative ones are nonviolating. We refer to this SNR as “nominal” because it cannot be interpreted in terms of Gaussian tail distributions. See [17] for a discussion of this issue. Our method for determining the SNRs is given in Appendix A. The figure also shows the $\log -p$ -value bounds from the PBR analysis on a matched scale. The $\log -p$ -value bound and the time-tag analysis’ SNR both drop to zero around $j_u = 0.06$.

We next considered the following question: For what j_u does there exist an LR source that has the same statistics as our simulated source? We do not have a definite answer to this question. However, we constructed a LR source whose one- and two-point statistics closely match those of the ideal Poisson source of quantum photon pairs for $j_u \geq 0.11$. More generally, we tweaked the LR source so that for all positive values of j_u it looks like the Poisson source that we simulated for the data in Fig. 5, except that it may have more coincidences at the $\bar{2}\bar{2}$ setting depending on the jitter. Further details are in Appendix D. Figure 6 shows the results from applying the conventional and time-tag analysis methods to data generated by a simulation of the LR source. As expected, the time-tag analysis shows no violation. But, the conventional analysis falsely shows violation. For $j_u \geq 0.11$ the violation is similar

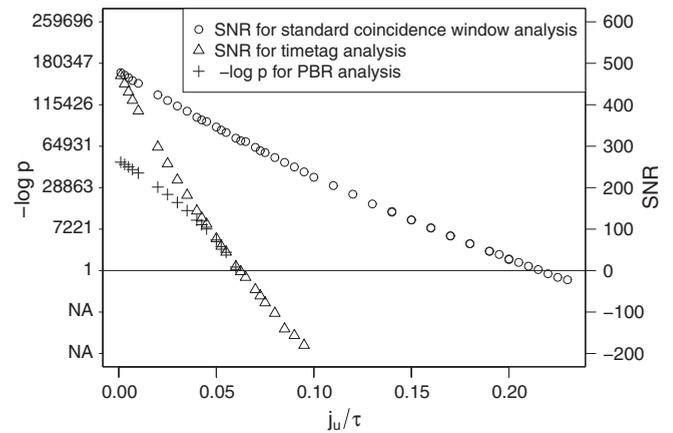


FIG. 5. Comparison of methods for simulated time-tag data from a quantum source and detectors with efficiency $\eta = 0.8$ and uniform jitter. The nominal SNRs for the standard and the time-tag analyses are shown on the right axis. Negative SNRs mean that the signal is positive and therefore not violating. The $\log -p$ -value bound for the PBR analysis is shown on the left axis. The horizontal axis shows the relative jitter width j_u/τ . To match the two vertical axes with one another, we converted $\log -p$ -values to equivalent Gaussian SNRs by computing the value for which the one-sided tail probability for the standard normal distribution matches the p -value bound. The computed $\log -p$ -value bounds are 0 for j_u above approximately 0.06 and not shown on the plot.

if somewhat lower than that for the corresponding quantum photon pairs in Fig. 5.

B. Jitter thresholds for efficiency ≥ 0.74

The simulations discussed above show that when there is too much jitter, a nominally violating source of entangled photon pairs produces time-tag data that become indistin-

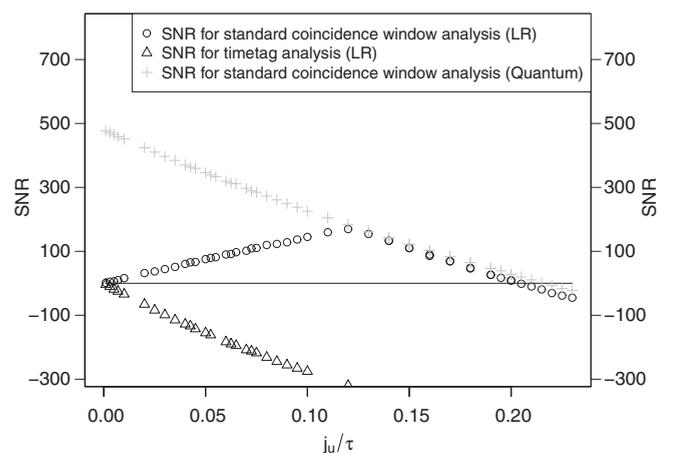


FIG. 6. Comparison of methods on LR-generated time-tag data. The conventional coincidence analysis shows a false violation over almost the entire range. The time-tag analysis shows no violation. The $\log -p$ -value bounds from the PBR analysis are 0 everywhere and are not shown. The SNR of the Poisson quantum source whose one- and two-point statistics are approximated by the LR source for jitter larger than $j_u/\tau \geq 0.11$ is also shown.

TABLE I. Lower bounds on maximum jitter at which the time-tag analysis can still detect violation of LR. The rows give the maximum median jitter at which our simulations showed a violation as determined by the log- p -value being strictly positive. The simulation parameters other than jitter are the same as for the other simulations. The jitter bounds are shown for the uniform and for the exponential distributions. For ease of comparison, they are parametrized in terms of the median delays of the recorded time tags. For low violation, unfeasibly large training and analysis sets are required to make the violation apparent in a simulation. Thus, the entries in the table are lower bounds on the maximum jitter at which the time-tag analysis can still detect violation. The missing entry was not computed due to excessive computational resource requirements for our implementation.

Efficiency	Uniform jitter median ($j_u/2$)	Exponential jitter median
0.74	0.013	0.0033
0.76	0.018	0.0049
0.78	0.024	0.0070
0.80	0.031	0.0095
0.85	0.052	0.017
0.90	0.07	0.029
0.95		0.051

guishable from that produced by an LR model. While we cannot determine the minimum jitter at which this happens, it is possible to lower bound the maximum jitter at which the time-tag analysis methods can see LR violation. We simulated the photon-pair source and measurement configuration introduced above at various efficiencies and varied the jitter distribution. We considered the uniform jitter model and the exponential jitter model and determined the maximum jitter widths at which the time-tag analysis found violation. The results are shown in Table I in terms of the median of the jitter delay of the recorded time tags.

X. CONCLUDING REMARKS

We have given a strategy based on the twice-iterated triangle inequality for constructing arbitrary two-party, two-settings Bell inequalities and applied it to the problem of analyzing data from Bell trials with time-tag-sequence measurement outcomes. We believe that the strategy can be used to analyze any Bell experiment where settings are changed slowly compared to the rate of detections. The benefit of considering all the data while the settings are held fixed as contributing to one trial is that loopholes associated with independence and stability assumptions can be closed. For applications involving continuously emitting sources, our strategy also closes any timing loopholes such as the coincidence loophole. We explored the behavior of the time-tag Bell functions in simulated data where jitter makes it difficult to assign coincidences reliably. The simulations demonstrate the practicality of the method and also demonstrate that the coincidence loophole can be exploited surreptitiously, with little sign of the exploit in the statistics of the detections.

ACKNOWLEDGMENTS

We thank Y.-K. Liu and K. Shalm for their help in revising this paper. This work is a contribution of the National Institute of Standards and Technology and is not subject to U. S. copyright.

APPENDIX A: DETERMINING THE VIOLATION'S NOMINAL SNR

The trials of an experiment result in a sequence of Bell function values $(b_i)_{i=1}^N$, one for each trial. Let s_i be the settings of trial i and N_{ab} the number of trials at settings ab . A standard approach to estimating the violation is to compute the sample mean $\hat{B}_{ab} = \sum_{i:s_i=ab} b_i / N_{ab}$ and define the estimated total violation as $N \sum_{ab} \hat{B}_{ab} / 4$. The variance of this value can then be estimated to first order with respect to the variance of N_{ab} from the sample variances of the subsequences $(b_i)_{i:s_i=ab}$. Instead of this procedure, we used a method that makes no first-order approximations and can be meaningfully applied even if the trials' Bell function values B_i are not independent. We consider this method better motivated, and it gives results that are statistically close to those obtained by the standard approach. Here, we describe the method for the case of i.i.d. trials.

The goal is to estimate the expectation of the sum of the trials' Bell function values and obtain a nearly tight bound on the variance of the estimate. The method is adaptive and applied to the analysis data given some initial training data. (Calibration data or theoretical predictions can be used instead.) The desired expectation is $\bar{B}_{\text{tot}} = \sum_i \langle B_i \rangle$. (Since we are considering i.i.d. trials, the distribution of B_i is the same for all i .) Before considering the i th trial, we use the previous trials and the training set to obtain four unbiased estimates $\hat{B}_{i,ab}$ of $\langle B_i | S_i = ab \rangle$, where S_i is the settings' random variable for the i th trial. This estimate can be obtained by any means desired. We used the formula for \hat{B}_{ab} given above and applied it to the training set and the first $i - 1$ trials. For the i th trial, we then define the random variables $F_i = \hat{B}_{i,S_i}$ and $\Delta_i = B_i - F_i$. Because the settings' probability distribution is known, so is the expectation of F_i : $\langle F_i \rangle = \sum_{ab} \hat{B}_{i,ab} / 4$. We then record $\delta_i = b_i - f_i$ for the i th trial and continue. Note that

$$\begin{aligned} \sum_i \langle B_i \rangle &= \sum_i \langle B_i - F_i \rangle + \sum_i \langle F_i \rangle \\ &= \sum_i \langle \Delta_i \rangle + \sum_{i,ab} \hat{B}_{i,ab} / 4. \end{aligned} \quad (\text{A1})$$

We can therefore empirically estimate \bar{B}_{tot} as $\hat{B}_{\text{tot}} = \sum_i \delta_i + \sum_{i,ab} \hat{B}_{i,ab} / 4$, which ensures that $\langle \hat{B}_{\text{tot}} \rangle = \bar{B}_{\text{tot}}$. We are considering the case of i.i.d. trials only, so $\hat{v} = \sum_i \delta_i^2$ is an estimate of the variance of \hat{B}_{tot} that is biased high. This is because the variance of a random variable W is the minimum of $\langle (W - a)^2 \rangle$ over a . Since Δ_i is designed to asymptotically converge to a zero-mean random variable, the variance estimate is asymptotically unbiased. The SNRs for the conventional coincidence analysis and the time-tag-sequence analysis shown in Figs. 5 and 6 are defined as $\hat{B}_{\text{tot}} / \sqrt{\hat{v}}$ multiplied by the sign of the violation.

APPENDIX B: DETERMINING FUNCTION-TUPLE-BASED DISTANCES

For general time-tag–sequence pairs, computing a CH function l_f [see Eq. (12)] can be accomplished by dynamic programming. The simplest implementation of this technique has a quadratic time cost. For sequences such as those produced by our simulations, which are associated with detections and coincidences generated uniformly randomly in time, this cost can be reduced to average linear time.

Let $d = (r_1 \leq \dots \leq r_m)$ and $e = (t_1 \leq \dots \leq t_n)$ be time-tag sequences. We wish to determine the distance $l_{f,ab}(d,e)$. For this purpose, let $d_k = (r_1, \dots, r_k)$ and $e_l = (t_1, \dots, t_l)$ be their initial segments. Let $c(k,l) = l_{f,ab}(d_k, e_l)$. We can determine $c(k,l)$ inductively. We have $c(k,0) = k$ and $c(0,l) = 0$. Suppose we have determined $c(k',l')$ for $k' < k$ and $l' \leq l$. Let $M_{k,l}$ be the (not-yet-known) matching minimizing $l(f_{ab}, M, d_k, e_l)$. There are two possibilities to consider. Either $k \notin \text{dom}(M_{k,l})$, in which case $c(k,l) = c(k-1,l) + 1$; or $M_{k,l}(k) = l' \leq l$, in which case $c(k,l) = c(k-1, l' - 1) + f_{ab}(t_{l'} - r_k)$. This reduction works because $M_{k,l}$ is monotone. Thus, $c(k,l)$ can be determined as the minimum of these possibilities. In an algorithm, one can store the $c(k,l)$ in an $(m+1) \times (n+1)$ matrix and fill its entries in the order suggested by this inductive construction. It is possible to reduce memory requirements by filling the matrix row by row and discarding the rows no longer needed. However, if it is desirable to extract the cost-minimizing matching after $c(m,n)$ has been determined, it helps to keep the full matrix, and work backward from the (m,n) entry to determine which of the cases above was used when the entry encountered was computed. The next entry as we work backwards is determined by the case used, and the relationships in the matchings arise from the encountered entries where the second case was used.

The construction as given above can be too resource intensive for long time-tag sequences. For pairs of sequences with sufficiently low rates of time tags, a simple way to speed up the algorithm is to break up the sequences at sufficiently large gaps and apply the algorithm to each resulting pair of subsequences. The subsequence costs are added. To formalize this idea, let $u > 0$ be such that for all $x \notin [-u, u]$ and for all ab , $f_{ab}(x) \geq 1$. In this case, no pair of indices k, l with $|t_l - r_k| \geq u$ needs to be considered for matching. Break up the two sequences into pairs of subsequences $d(k_{i,1}, k_{i,2}) = (r_{k_{i,1}} \leq \dots \leq r_{k_{i,2}})$ and $e(l_{i,1}, l_{i,2}) = (t_{l_{i,1}} \leq \dots \leq t_{l_{i,2}})$ such that $\max(d(k_{i,1}, k_{i,2}), e(l_{i,1}, l_{i,2})) + u \leq \min(d(k_{i+1,1}, k_{i+1,2}), e(l_{i+1,1}, l_{i+1,2}))$ for every i for which both sides of the inequality are defined. (The maximum of a collection of sequences given as arguments to \max is defined as the maximum of the set of all time tags occurring in the arguments, and similarly for \min .) The algorithm is then applied to each corresponding pair of subsequences and the costs computed are added to determine the cost of the original sequence.

APPENDIX C: OPTIMIZING FUNCTION-TUPLE PARAMETERS

When performing the time-tag–sequence analysis on simulated data, we optimized the parameters of the linear-

edge window function-tuple using a set of training data. To simplify the optimization, we restricted the linear-edge window functions in Eq. (11) to be reflection symmetric around 0, so that $m_l = m_h := m$ and $t_{l,ab} = t_{h,ab} := t_{ab}$. We also fixed $t_{\bar{1}\bar{1}} = t_{\bar{1}\bar{2}} = t_{\bar{2}\bar{1}} := t$. By definition of the linear-edge window function-tuple, $t_{\bar{2}\bar{2}} = 3t$, so only the slope m and truncation point t remain to be optimized. Instead of direct optimization, we used a simpler approximate optimization strategy based on compressing the relevant information in the time-tag sequences. This is accomplished as follows: For each pair of training-set time-tag sequences \mathbf{r} and \mathbf{t} at setting ab contributing to B_l , we first determine the optimum matching M for the “compression” function-tuple $f_{ab}(x) = \min(\lambda|x|, 1)$. In the simulations, $\lambda = 1$. In general, λ needs to be chosen so that the time-tag differences Δ for which the analysis function-tuple may be less than 1 satisfy $f_{ab}(\Delta) < 1$. Our choice of λ is given by the photon-pair creation rate. But, it should suffice to choose λ so that few entangled photon pairs have recorded time tags that are separated by more than $1/\lambda$. For each index k of \mathbf{r} in the domain of M , we determine the differences $x_k = (t_{M(k)} - r_k)$, where the notation for time tags of \mathbf{r} and \mathbf{t} is as before. We then collect all such time-tag differences for all pairs \mathbf{r} and \mathbf{t} at setting ab in a single sequence y_{ab} . For the optimization, we assume that the matchings M obtained in the construction of y_{ab} are close in cost to the optimal matchings on the training set that would be obtained according to function-tuples with the parameters we are optimizing. Since we are working on the training, not the analysis set, our computations can be approximate. Given this assumption, we can use y_{ab} to compute an approximation of the ab contribution to the Bell function B_{l_g} for a given function-tuple g . Let X_{ab} be the maximum number of deletions that can contribute to the ab cost. This is given by the sum over all time-tag–sequence pairs at settings ab in the training set of the number of time tags in the first (when $ab \neq \bar{1}\bar{1}$) or second (when $ab = \bar{1}\bar{1}$) time-tag sequence in each pair of sequences. Then, the estimated difference between the ab cost and $X_{ab} - |y_{ab}|$ is given by the sum of the values of g_{ab} applied to the time-tag differences in y_{ab} . (Here, $|y_{ab}|$ is the number of time-tag differences in y_{ab} .) These sums can be computed much faster than one can compute the exact minimum ab cost for g_{ab} . The approximate costs are used to optimize g , and the resulting g is subsequently used for the time-tag analysis of the analysis set. We note that the objective function in the optimization has irregularities that can result in high sensitivity of the parameters of g to statistical noise.

APPENDIX D: A LR SOURCE THAT EXPLOITS THE COINCIDENCE LOOPHOLE

In this section, we describe the LR source whose false violation of a Bell inequality under a standard coincidence window analysis is shown in Fig. 6. This LR source can closely mimic the one- and two-point statistics of a Poisson source of entangled photon pairs detected by detectors with uniform jitter distribution of width $j_u \geq 0.11$.

A LR source generates four time-tag sequences \mathbf{t}_c^X , where $X \in \{A, B\}$ and $c \in \{\bar{1}, \bar{2}\}$, for each trial. These sequences are the sequences that may be recorded by A and B depending on their settings. After the experiment, only the two sequences

corresponding to the actually chosen settings are visible to the parties. The goal is to match the visible statistics of the LR source to those of a Poisson source of quantum photon pairs with jitter. We consider the uniform jitter distribution of width j_u and a Poisson source whose detection statistics are determined by the single-pair settings-conditional outcome probabilities $p(o^A o^B | ab)$ and the uniform distribution for settings. We ensure that the LR source exhibits the same marginal detection rates $p(o^X | ab)$. But, since we are interested in how readily the conventional coincidence analysis can be deceived, we allow the LR source to adjust the apparent coincidence rates. Note that the inferred coincidence rates depend on the method used to determine coincidences. The construction of the LR source uses probabilities $p'(o^A o^B | ab)$ as a template in an attempt to deceive the experimenter into believing p' . The template satisfies $p'(o^X | ab) = p(o^X | ab)$. The apparent coincidence rates are the same except at the $\bar{2}\bar{2}$ setting, where we set $p'(11|\bar{2}\bar{2}) = p(11|\bar{2}\bar{2}) + \delta_c$, $p'(00|\bar{2}\bar{2}) = p(00|\bar{2}\bar{2}) + \delta_c$, $p'(01|\bar{2}\bar{2}) = p(01|\bar{2}\bar{2}) - \delta_c$, $p'(10|\bar{2}\bar{2}) = p(10|\bar{2}\bar{2}) - \delta_c$, and $p' = p$ otherwise. The coincidence rate adjustment δ_c is chosen to maximize the rate at which the LR source can successfully introduce an apparent violating signal. For $j_u \gtrsim 0.11$, we found that it is possible to match the coincidence rates ($\delta_c = 0$).

For successful deception, the LR source's time-tag-sequence statistics should match that of a Poisson source with template frequencies given by p' . (These frequencies account for the photon states, source statistics, and the jitter distribution.) We aimed for matching detection rates and correlation functions. While our source does not match the correlation functions exactly (see following), the residual correlation mismatches are sufficiently small to either escape detection or to be hidden by the typically much larger correlation artifacts of the same order as the jitter introduced by the detection apparatus. We also note that in an experiment, the source and jitter distributions are not known beforehand, making a statistical test for correlation artifacts introduced by the LR source difficult. Comparisons of the correlation functions for $j_u = 0.11$ are in Fig. 7. Note that the mean time separation between $\bar{2}\bar{2}$ coincidences seems to match that for other settings. That is, there is no tell-tale broadening that might be expected from the earlier coincidence-loop-hole examples in Sec. III.

Here is a sketch of the LR source construction. To generate LR time-tag sequences according to p' and j_u , we first decompose $p' = \lambda_{\text{LR}} p_{\text{LR}} + \lambda_{\text{PR}} p_{\text{PR}}$, where λ_{LR} is maximized subject to p_{LR} being an LR probability distribution and p_{PR} a Popescu-Rohrlich (PR) box's [28] probability distribution. Here, the PR box is perfectly correlated on all settings except at $\bar{2}\bar{2}$, where it is perfectly anticorrelated. To create an apparent PR box signal, we proceed as follows. Let $J(x, j_u)$ be the distribution of the time separation x between pairs of initially coincident time tags both of which are jittered uniformly with width j_u . The density $J(x, j_u)$ has a triangle shape centered at 0, supported on $[-j_u, j_u]$, with $J(0, j_u) = 1/j_u$. Let p_c^X be the detection rate of X at setting c . We begin by generating time tags for A at setting $\bar{2}$ at a uniform rate $p_{\bar{2}}^A$. This gives the time-tag sequence \mathbf{t}_2^A . For each time tag t thus generated, with probability $\lambda_{\text{PR}}/2$, we wish to generate a corresponding timetag of \mathbf{t}_2^B that is generally far enough from t to be

missed as a coincidence by the conventional analysis and is indistinguishable from background. The trick is to do this while preserving the ability to generate B 's detections at setting $\bar{2}$ according to a uniform process of rate $p_{\bar{2}}^B$. A first attempt is to generate time tags for \mathbf{t}_2^B so that the rate of detections at s is $\lambda_{\text{PR}} J(s - t, 3j_u)$ (three times the apparent jitter width). For fixed t , the probability of at least one detection for \mathbf{t}_2^B thus generated is $1 - e^{-\lambda_{\text{PR}}}$. If there is at least one such detection, we allocate it to a "hidden" coincidence for purposes of keeping track of the detection statistics. The hidden coincidence will be attributed to a PR-type anticorrelation if this detection is not recognized as a coincidence.

Note that detections generated at a uniform rate r can be realized by independently generating detections according to rate distributions ρ_i where $\sum_i \rho_i = r$. Thus, if $\rho_1(s) = \sum_t \lambda_{\text{PR}} J(s - t, 3j_u) \leq p_{\bar{2}}^B$, we can generate further detections to get the desired marginal statistics for B at setting $\bar{2}$. Actually, we need to exclude the rate of coincidences $p'(11|\bar{2}\bar{2})$ from $p_{\bar{2}}^B$ in this inequality because these "original" coincidences are to be added separately later. In addition, to ensure that ρ_1 is below the corrected bound $p_{\bar{2}}^B - p'(11|\bar{2}\bar{2})$, we have to deal with the problem that for nearby time tags of \mathbf{t}_2^A , the distributions $J(s - t, 3j_u)$ overlap. For this purpose, we made $\lambda_{\text{PR}} = \lambda_{\text{PR}}(t)$ depend on the time tag t and used a linear programming technique to maximize $\sum_{t \in \mathbf{t}_2^A} \lambda_{\text{PR}}(t)$ subject to the bound. The actual rate of apparent PR boxes is affected by the result. These rates were determined by a Monte Carlo method as a function of the jitter and matched to $\lambda_{\text{PR}}/2$ by adjusting δ_c as needed.

The PR boxes inserted into the time tags at setting $\bar{2}\bar{2}$ by hiding coincidences need to be extended to recognizable coincidences at the other settings. We first fill in these coincidences by spreading them across the $\bar{2}\bar{1}$, $\bar{1}\bar{1}$, and $\bar{1}\bar{2}$ settings at the regular jitter width j_u by dividing the hidden coincidences' $\bar{2}\bar{2}$ separations equally into the separations involving the time tags at the non- $\bar{2}$ settings. Extra coincidences with normal jitter width are then added to get the desired coincidence rates at the non- $\bar{2}\bar{2}$ settings. The resulting marginal detection rates are subtracted from the uniform rates $p_{\bar{1}}^X$ and extra detections are then filled in accordingly.

While our implementation required substantial elaboration of the above outline, the success of the strategy is witnessed by the false violation discovered by the conventional coincidence analysis visible in Fig. 6 and the empirical autocorrelation and cross-correlation functions shown in Fig. 7. The figure shows the correlation functions for a LR source that simulates the quantum sources used for the analyses shown in Fig. 5 with $j_u = 0.11$. In this case, the LR source's apparent coincidence rate at the $\bar{2}\bar{2}$ setting matches the corresponding quantum source's rate. Small deviations from the quantum source's correlation functions are visible in the plots. In practice, such deviations are common and therefore hard to distinguish from normal experimental artifacts. Thus, we expect that it would be difficult to find true LR violation in the quantum data, even if its statistics cannot be exactly simulated by a LR source. For comparison, the maximum jitter for which the time-tag analysis finds a violation is about $j_u \approx 0.06$.

The estimated correlation functions shown in Fig. 7 are determined from binned data as follows: Let \mathbf{r} and \mathbf{t} be two

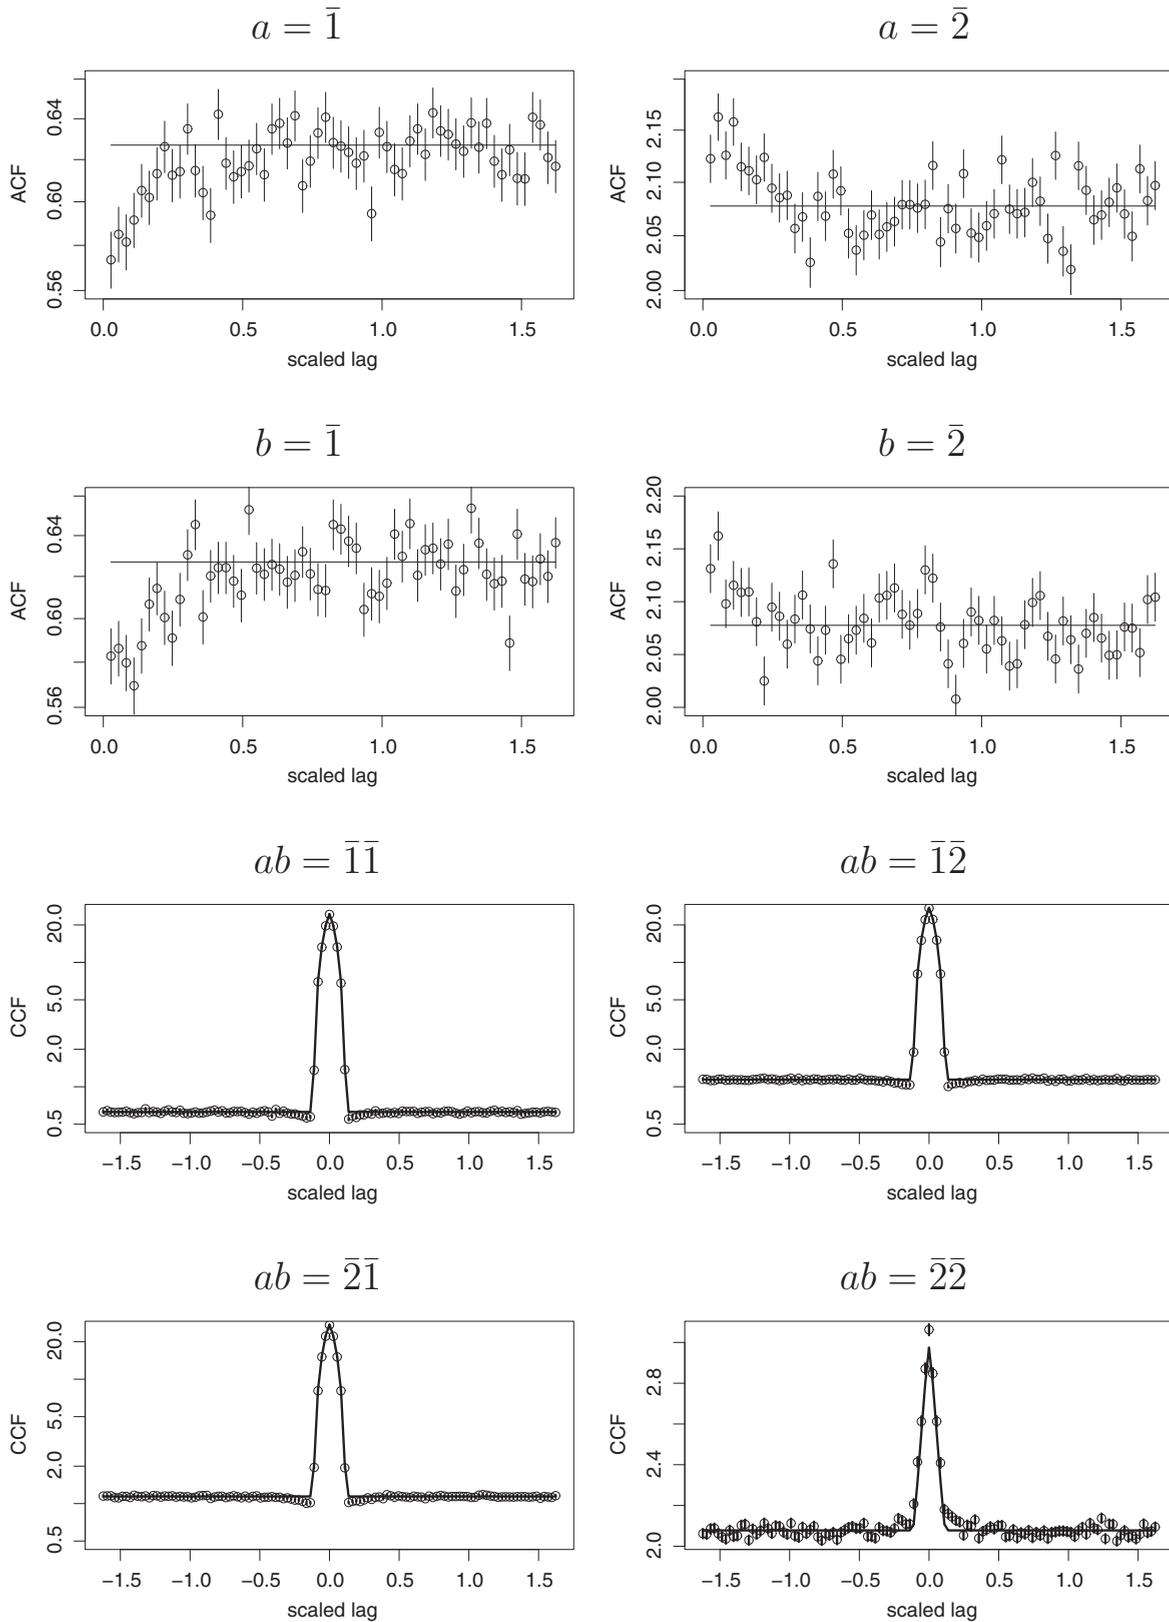


FIG. 7. Autocorrelation (ACF) and cross-correlation functions (CCF) (unnormalized) for LR-generated time-tag data. Each subfigure is labeled with the measurement setting(s) used to record the time-tag data. The points with error bars show estimated correlation function values and their associated approximate 68% confidence intervals for the LR time-tag data determined from 200 000 trials. The LR source was designed to match time-tag data from the quantum source with an efficiency $\eta = 0.8$ and uniform jitter $j_u = 0.11$ that was used for the results shown in Fig. 5. The continuous curves are the corresponding theoretical correlation functions for the quantum source. See the text for the definition of correlation functions used here. The bin width is $w_b = j_u/4$. The “lag” is defined as $d * w_b$.

time-tag sequences and w_b a bin width. Let T be the total observation period of the time-tag sequences and assume that the initial time is 0. We construct functions $b_s : \{0, 1, \dots, N = \lceil T/w_b \rceil\} \rightarrow \mathbb{N}$ for $\mathbf{s} = \mathbf{r}$ and $\mathbf{s} = \mathbf{t}$ by defining $b_s(k)$ to be the number of time tags s in \mathbf{s} with $kw_b \leq s < (k+1)w_b$. The estimated correlation function for \mathbf{r} and \mathbf{t} is defined by $c(d, \mathbf{r}, \mathbf{t}) = \sum_{i=0}^{i+d=N} b_{\mathbf{r}}(i)b_{\mathbf{t}}(i+d)$. (Note that we do not normalize the correlation function.) The empirical values shown in the figure are the sample means of the $c(d, \mathbf{r}, \mathbf{t})$ over the appropriate pairs of time-tag sequences from 200 000 trials. Note that the autocorrelation functions are of the form $c(d, \mathbf{r}, \mathbf{r})$ and are symmetric about $d = 0$, so only the half with $d \geq 0$ is shown.

We end with a brief note on how an adversary might realize the LR model of this section. Suppose that the adversary can surreptitiously manipulate the measurement instruments of each party so that the polarization angles of the two settings are approximately orthogonal. Also, suppose that the true jitter of

the detectors is small but the experimenter does not realize this. In this case, for each trial, the adversary can randomly generate four time-tag sequences according to the LR model and, for each event in the time-tag sequences, send a photon at the event's time to the appropriate party with the polarization that ensures it is only detected at the intended setting. Because of the way that the LR model's time-tag sequences are generated, the detections will appear to be detections from photon pairs with detector jitter. The manufacturer of the measurement instruments can build in the features that the adversary needs to exploit in this scenario. It is therefore unlikely that a protocol relying on a conventional coincidence analysis can achieve device-independent security. We leave open the question of whether a LR model based on our techniques can be realized by an adversary who can control only the source, when the characteristics of the measurement instruments are fixed and known to the experimenter.

-
- [1] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Physics* **1**, 195 (1964).
- [2] M. Genovese, Research on hidden variable theories: A review of recent progresses, *Phys. Rep.* **413**, 319 (2005).
- [3] J. Barrett, L. Hardy, and A. Kent, No signaling and quantum key distribution, *Phys. Rev. Lett.* **95**, 010503 (2005).
- [4] L. Masanes, Universally composable privacy amplification from causality constraints, *Phys. Rev. Lett.* **102**, 140501 (2009).
- [5] L. Masanes, S. Pironio, and A. Acin, Secure device-independent quantum key distribution with causally independent measurement devices, *Nat. Commun.* **2**, 238 (2011).
- [6] S. Pironio, A. Acin, S. Massar, A. B. de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning *et al.*, Random numbers certified by Bell's theorem, *Nature (London)* **464**, 1021 (2010).
- [7] R. Colbeck and A. Kent, Private randomness expansion with untrusted devices, *J. Phys. A: Math. Theor.* **44**, 095305 (2011).
- [8] U. Vazirani and T. Vidick, Certifiable quantum dice: Or, true random number generation secure against quantum adversaries, in *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2012), pp. 61–76.
- [9] J.-Å. Larsson, Loopholes in Bell inequality tests of local realism, *J. Phys. A: Math. Theor.* **47**, 424003 (2014).
- [10] J.-Å. Larsson and R. D. Gill, Bell's inequality and the coincidence-time loophole, *Europhys. Lett.* **67**, 707 (2004).
- [11] B. W. Schumacher, Information and quantum nonseparability, *Phys. Rev. A* **44**, 7047 (1991).
- [12] E. N. Dzhafarov and J. V. Kujala, Order-distance and other metric-like functions on jointly distributed random variables, *Proc. Am. Math. Soc.* **141**, 3291 (2013).
- [13] P. Kurzynski and D. Kaszlikowski, Information-theoretic metric as a tool to investigate nonclassical correlations, *Phys. Rev. A* **89**, 012103 (2014).
- [14] R. D. Gill, Accardi contra Bell (cum mundi): The impossible coupling, [arXiv:quant-ph/0110137](https://arxiv.org/abs/quant-ph/0110137).
- [15] R. D. Gill, Time, finite statistics, and Bell's fifth position, [arXiv:quant-ph/0301059](https://arxiv.org/abs/quant-ph/0301059).
- [16] Y. Zhang, S. Glancy, and E. Knill, Asymptotically optimal data analysis for rejecting local realism, *Phys. Rev. A* **84**, 062118 (2011).
- [17] Y. Zhang, S. Glancy, and E. Knill, Efficient quantification of experimental evidence against local realism, *Phys. Rev. A* **88**, 052119 (2013).
- [18] M. Giustina, A. Mech, S. Ramelow, B. Wittmann, J. Kofler, J. Beyer, A. Lita, B. Calkins, T. Gerrits, S. W. Nam *et al.*, Bell violation using entangled photons without the fair-sampling assumption, *Nature* **497**, 227 (2013).
- [19] B. G. Christensen, K. T. McCusker, J. Altepeter, B. Calkins, T. Gerrits, A. Lita, A. Miller, L. K. Shalm, Y. Zhang, S. W. Nam, and N. Brunner, Detection-loophole-free test of quantum nonlocality, and applications, *Phys. Rev. Lett.* **111**, 130406 (2013).
- [20] J. Kofler, S. Ramelow, M. Giustina, and A. Zeilinger, On 'Bell violation using entangled photons without the fair-sampling assumption', [arXiv:1307.6475](https://arxiv.org/abs/1307.6475).
- [21] 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).
- [22] J. F. Clauser and M. A. Horne, Experimental consequences of objective local theories, *Phys. Rev. D* **10**, 526 (1974).
- [23] J. Barrett, D. Collins, L. Hardy, A. Kent, and S. Popescu, Quantum nonlocality, Bell inequalities, and the memory loophole, *Phys. Rev. A* **66**, 042111 (2002).
- [24] R. D. Gill, G. Weihs, A. Zeilinger, and M. Żukowski, No Time Loophole in Bell's theorem: The Hess-Philipp model is nonlocal, *Proc. Natl. Acad. Sci. USA* **99**, 14632 (2002).
- [25] G. Shafer, A. Shen, N. Vereshchagin, and V. Vovk, Test martingales, Bayes factors and p -values, *Stat. Sci.* **26**, 84 (2011).
- [26] P. H. Eberhard, Background level and counter efficiencies required for a loophole-free Einstein-Podolsky-Rosen experiment, *Phys. Rev. A* **47**, R747 (1993).
- [27] J.-Å. Larsson, M. Giustina, J. Kofler, B. Wittmann, R. Ursin, and S. Ramelow, Bell violation with entangled photons, free of the coincidence-time loophole, *Phys. Rev. A* **90**, 032107 (2014).
- [28] S. Popescu and D. Rohrlich, Causality and nonlocality as axioms for quantum mechanics, *Found. Phys.* **24**, 379 (1997).