

The statistical strength of nonlocality proofs

Wim van Dam, Richard D. Gill, Peter D. Grünwald,

Abstract— There exist numerous proofs of Bell’s theorem, stating that quantum mechanics is incompatible with local realistic theories of nature. Here the strength of such nonlocality proofs is defined in terms of the amount of evidence against local realism provided by the corresponding experiments. Statistical considerations show that the amount of evidence should be measured by the *Kullback-Leibler* or *relative entropy* divergence.

The statistical strength of the following proofs is determined: Bell’s original proof and Peres’ optimized variant of it, and proofs by Clauser-Horne-Shimony-Holt, Hardy, Mermin, and Greenberger-Horne-Zeilinger. The GHZ proof is at least four and a half times stronger than all other proofs, while of the two-party proofs, the one of CHSH is the strongest.

Index Terms— nonlocality, Bell’s theorem, quantum correlations, Kullback-Leibler divergence

I. INTRODUCTION

A PLETHORA of proofs exist of Bell’s theorem (“quantum mechanics violates local realism”) encapsulated in inequalities and equalities of which the most celebrated are those of Bell [5], Clauser, Horne, Shimony and Holt (CHSH) [9], Greenberger, Horne and Zeilinger (GHZ) [16], Hardy [21], and Mermin [26]. Competing claims exist that one proof is stronger than another. For instance, a proof in which quantum predictions having probabilities 0 or 1 only are involved, is often said to be stronger than a proof that involves quantum predictions of probabilities between 0 and 1. Other researchers argue that one should compare the absolute differences between the probabilities that quantum mechanics predicts and those that are allowed by local theories, and so on. The main aim of this paper is to settle such questions once and for all: we formally define the strength of a nonlocality proof and we argue that our definition is the only one compatible with generally accepted notions in information theory and theoretical statistics.

To see the connection with statistics, note first that a *mathematical* nonlocality proof shows that the *predicted* probabilities of quantum theory are incompatible with local realism. Such a proof can be implemented as an *experimental proof* showing that *physical reality* conforms to those predictions and hence too is incompatible with local realism. We are interested in the strength of such experimental proofs, which should be

measured in statistical terms: how sure do we become that a certain theory is false, after observing a certain violation from that theory, in a certain number of experiments.

A. Our Game

We analyze the statistics of nonlocality proofs in terms of a two-player game. The two players are the pro-quantum theory experimenter QM, and a pro-local realism theoretician LR. The experimenter QM is armed with a specific proof of Bell’s theorem. A given proof—BELL, CHSH, HARDY, MERMIN, GHZ—involves a collection of equalities and inequalities between various experimentally accessible probabilities. The proof specifies a given quantum state (of a collection of entangled qubits, for instance) and experimental settings (orientations of polarization filters or Stern-Gerlach devices). All local realistic theories of LR will obey the (in)equalities, while the observations that QM will make when performing the experiment (assuming that quantum mechanics is true) will violate these (in)equalities. The experimenter QM still has a choice of the probabilities with which the different combinations of settings will be applied, in a long sequence of independent trials. In other words, he must still decide how to allocate his resources over the different combinations of settings. At the same time, the local realist can come up with all kinds of different local realistic theories, predicting different probabilities for the outcomes given the settings. She might put forward different theories in response to different specific experiments. Thus the quantum experimenter will choose that probability distribution over his settings, for which the *best* local realistic model explains the data *worst*, when compared with the true (quantum mechanical) description.

B. Quantifying Statistical Strength - Past Approaches

How should we measure the statistical strength of a given experimental setup? In the past it was often simply said that the largest deviation in the Bell inequality is attained with such and such filter settings, and hence the experiment which is done with these settings gives (potentially) the strongest proof of nonlocality. The argument is however not very convincing. One should take account of the statistical variability in finite samples. The experiment that might confirm the largest absolute deviation from local realistic theories, might be subject to the largest standard errors, and therefore be less convincing than an experiment where a much smaller deviation can be proportionally much more accurately determined.

Alternatively, the argument has just been that with a large enough sample size, even the smallest deviation between two theories can be made firm enough. For instance, [26] has said in the context of a particular example

A preliminary version [11] of this paper, containing extensive background material, appeared on the arXiv preprint server in the quantum physics archive as quant-ph/0307125.

Wim van Dam is with the Department of Computer Science, University of California, Santa Barbara, Santa Barbara, CA 93106-5110 United States of America. Email: vandam@cs.ucsb.edu

Richard Gill is with the Mathematical Institute, University Utrecht, Budapestlaan 6, NL-3584 CD Utrecht, The Netherlands. He is also at EURANDOM, Eindhoven, The Netherlands. Email: gill@math.uu.nl

Peter Grünwald is with the CWI, P.O. Box 94079, NL-1090 GB, The Netherlands. He is also at EURANDOM, Eindhoven, The Netherlands. Email: pdg@cwi.nl

“...to produce the conundrum it is necessary to run the experiment sufficiently many times to establish with overwhelming probability that the observed frequencies (which will be close to 25% and 75%) are not chance fluctuations away from expected frequencies of 33% and 66%. (A million runs is more than enough for this purpose)...”

We want to replace the words “sufficiently”, “overwhelming”, “more than enough” with something more scientific. (See Example 4 for our conclusion with respect to this.) And as experiments are carried out that are harder and harder to prepare, it becomes important to design them so that they give conclusive results with the smallest possible sample sizes. Initial work in this direction has been done by Peres [29]. Our approach is compatible with his, and extends it in a number of directions—see Section VII-B.

C. Quantifying Statistical Strength - Our Approach

We measure statistical strength using an information-theoretic quantification, namely the Kullback-Leibler (KL) divergence (also known as *information deficiency* or *relative entropy* [10]). We show (Appendix III) that for large samples, all reasonable definitions of statistical strength that can be found in the statistical and information-theoretic literature essentially coincide with our measure. For a given type of experiment, we consider the game in which the experimenter wants to maximize the divergence while the local theorist looks for theories that minimize it. The experimenter’s game space is the collection of probability distributions over joint settings, which we call in the sequel, for short, “setting distributions”. (More properly, these are “joint setting distributions”.) The local realist’s game space is the space of local realistic theories. This game defines an experiment, such that each trial (assuming quantum mechanics is true) provides on average, the maximal support for quantum theory against the best explanation that local realism can provide, at that setting distribution.

D. Our Results - Numerical

We determined the statistical strength of five two-party proofs: Bell’s original proof and Peres’ optimized variant of it, and the proofs of CHSH, Hardy, and Mermin. Among these, CHSH turns out to be the strongest by far. We also determined the strength of the three-party GHZ proof. Contrary to what has sometimes been claimed (see Section VII), even the GHZ experiment has to be repeated a fair number of times before a substantial violation of local realism is likely to be observed. Nevertheless, it is about 4.5 times stronger than the CHSH experiment, meaning that, in order to observe the same support for QM and against LR, the CHSH experiment has to be run about 4.5 times as often as the GHZ experiment — we provide precise numbers in Section VI, Example 4.

E. Our Results - Mathematical

To find the (joint) setting distribution that optimizes the strength of a nonlocality proof is a highly nontrivial computation. In the second part of this paper, we prove several

mathematical properties of our notion of statistical strength. These provide insights in the relation between local realist and quantum distributions that are interesting in their own right. They also imply that determining statistical strength of a given nonlocality proof may be viewed as a convex optimization problem which can be solved numerically. We also provide a game-theoretic analysis involving minimax and maximin KL divergences. This analysis allows us to shortcut the computations in some important special cases.

F. Organization of This Paper

Section II gives a formal definition of what we mean by a nonlocality proof and the corresponding experiment, as well as the notation that we will use throughout the article. The kinds of nonlocality proofs that this article analyzes are described in Section III, using the CHSH proof as a specific example; the other proofs are described in more detail in Appendices I and II. The definition of the ‘statistical strength of a nonlocality proof’ in terms of the Kullback-Leibler divergence is presented in Section IV, along with some standard facts about KL divergence and its role in hypothesis testing. The strength of a nonlocality proof is now well-defined, but it is not yet clear how to compute it. In Section V we develop the mathematical techniques needed to perform the computations efficiently: we establish topological, analytical and game-theoretic properties of statistical strength on which our (mostly numerical, but sometimes exact) calculations are based. This technical section may be skipped, since it is not crucial for understanding the remainder of the paper. Section VI lists the results of our calculations of statistical strength for six well-known proofs (with additional details again in Appendix II). The results are interpreted, discussed and compared in Section VII. In Section VIII we present five conjectures that are suggested by our results.

We defer all issues that require knowledge of the mathematical aspects of quantum mechanics to the appendices. There we provide more detailed information about the nonlocality proofs we analyzed, the relation of Kullback-Leibler divergence to hypothesis testing, and the proofs of the theorems we present in the main text.

II. FORMAL SETUP

A basic nonlocality proof (“quantum mechanics violates local realism”) has the following ingredients. There are two parties \mathcal{A} and \mathcal{B} , who can each do a measurement on one of two entangled qubits. They may each choose from two different measurement settings. In each trial of the experiment, \mathcal{A} and \mathcal{B} randomly sample from the four different joint settings and each observe one of two different binary outcomes, say “F” (false) and “T” (true). Quantum mechanics enables us to compute the joint probability distribution of the outcomes, as a function of the measurement settings and of the joint state of the two qubits. Thus possible design choices are: the state of the qubits, the values of the settings; and the probability distribution over the settings. More complicated experiments may involve more parties, more settings, and more outcomes. (We formalized the general set-up in [11]). In this text, we

focus primarily on the basic $2 \times 2 \times 2$ case, which stands for “2 parties \times 2 measurement settings per party \times 2 outcomes per measurement setting”. Below we introduce notation for all ingredients involved in nonlocality proofs.

A. Distribution of Settings

The random variable A denotes the measurement setting of party \mathcal{A} and the random variable B denotes the measurement setting of party \mathcal{B} . Both A and B take values in $\{1, 2\}$. The experimenter QM will decide on the distribution σ of (A, B) , giving the probabilities (and, after many trials of the experiment, the frequencies) with which each (joint) measurement setting is sampled. This *setting distribution* σ is identified with its probability vector $\sigma := (\sigma_{11}, \sigma_{12}, \sigma_{21}, \sigma_{22}) \in \Sigma$, and Σ is the unit simplex in \mathbb{R}^4 defined by

$$\Sigma := \left\{ (\sigma_{11}, \dots, \sigma_{22}) \mid \sum_{a,b \in \{1,2\}} \sigma_{ab} = 1, \text{ for all } a, b : \sigma_{ab} \geq 0 \right\}.$$

We use Σ^{UC} to denote the set of vectors representing *uncorrelated* distributions in Σ . Formally, $\sigma \in \Sigma^{\text{UC}}$ if and only if $\sigma_{ab} = (\sigma_{a1} + \sigma_{a2})(\sigma_{1b} + \sigma_{2b})$ for all $a, b \in \{1, 2\}$.

B. Measurement Outcomes

The random variable X denotes the measurement outcome of party \mathcal{A} and the random variable Y denotes that of party \mathcal{B} . Both X and Y take values in $\{\text{T}, \text{F}\}$, where F means ‘false’ and T means ‘true’. Thus, the statement ‘ $X = \text{F}, Y = \text{T}$ ’ and describes the event that party \mathcal{A} observed F and party \mathcal{B} observed T.

The distribution of (X, Y) depends on the chosen setting $(a, b) \in \{1, 2\}^2$. The state of the entangled qubits together with the measurement settings determines four conditional distributions $Q_{11}, Q_{12}, Q_{21}, Q_{22}$ for (X, Y) , one for each joint measurement setting, where Q_{ab} is the distribution of (X, Y) given that measurement setting (a, b) has been chosen. For example, $Q_{ab}(X = \text{F}, Y = \text{T})$, abbreviated to $Q_{ab}(\text{F}, \text{T})$, denotes the probability that party \mathcal{A} observes F and party \mathcal{B} observes T, given that the device of \mathcal{A} is in setting a and the device of \mathcal{B} is in setting b . According to QM, the total outcome (X, Y, A, B) of a single trial is then distributed as Q_σ , defined by $Q_\sigma(X = x, Y = y, A = a, B = b) := \sigma_{ab} Q_{ab}(X = x, Y = y)$.

C. Definition of a Nonlocality Proof and Corresponding Nonlocality Experiments

A *nonlocality proof* for 2 parties, 2 measurement settings per party, and 2 outcomes per measurement, is identified with an entangled quantum state of two qubits (realized, by, e.g., two photons) and two measurement devices (e.g., polarization filters) which each can be used in one of two different measurement settings (polarization angles). Everything about the quantum state, the measurement devices, and their settings that is relevant for the probability distribution of outcomes of the corresponding experiment can be summarized by the four distributions Q_{ab} of (X, Y) , one for each (joint) setting $(a, b) \in \{1, 2\}^2$. Henceforth, we will simply *identify* a $2 \times 2 \times 2$ nonlocality proof with the vector of distributions $Q := (Q_{11}, Q_{12}, Q_{21}, Q_{22})$.

This definition can easily be extended in an entirely straightforward manner to a situation with more than two parties, two settings per party, or two outcomes per setting [11].

We call a nonlocality proof $Q = (Q_{11}, Q_{12}, Q_{21}, Q_{22})$ *proper* if and only if it violates local realism, i.e. if there exists no local realist distribution π (as defined below) such that $P_{ab;\pi}(\cdot) = Q_{ab}(\cdot)$ for all $(a, b) \in \{1, 2\}^2$.

For the corresponding $2 \times 2 \times 2$ nonlocality *experiment* we have to specify the setting distribution σ with which the experimenter QM samples the different settings (a, b) . Thus, for a single nonlocality proof Q , QM can use different experiments (different in σ) to verify Nature’s nonlocality. Each experiment consists of a series of trials, where—per trial—the event (X, Y, A, B) occurs with probability $Q_\sigma(X = x, Y = y, A = a, B = b) = \sigma_{ab} Q_{ab}(X = x, Y = y)$.

D. Local Realist Theories

The local realist (LR) may provide any ‘local’ theory she likes to explain the results of the experiments.

Under such a theory it is possible to talk about “the outcome that \mathcal{A} would have observed, if she had used setting 1”, independently of which setting was used by \mathcal{B} and indeed of whether or not \mathcal{A} actually did use setting 1 or 2. Thus we have four binary random variables, which we will call X_1, X_2, Y_1 and Y_2 . As before, variables named X correspond to \mathcal{A} ’s observations, and variables named Y correspond to \mathcal{B} ’s observations. According to LR, each experiment determines values for the four random variables (X_1, X_2, Y_1, Y_2) . For $a \in \{1, 2\}$, $X_a \in \{\text{T}, \text{F}\}$ denotes the outcome that party \mathcal{A} would have observed if the measurement setting of \mathcal{A} had been a . Similarly, for $b \in \{1, 2\}$, $Y_b \in \{\text{T}, \text{F}\}$ denotes the outcome that party \mathcal{B} would have observed if the measurement setting of \mathcal{B} had been b .

A local theory π may be viewed as a probability distribution for (X_1, X_2, Y_1, Y_2) . Formally, we define π as a 16-dimensional probability vector with indices $(x_1, x_2, y_1, y_2) \in \{\text{T}, \text{F}\}^4$. By definition, $P_\pi(X_1 = x_1, X_2 = x_2, Y_1 = y_1, Y_2 = y_2) := \pi_{x_1 x_2 y_1 y_2}$. For example, π_{FFFF} denotes LR’s probability that, in all possible measurement settings, \mathcal{A} and \mathcal{B} would both have observed F. The set of local theories can thus be identified with the unit simplex in \mathbb{R}^{16} , which we will denote by Π .

As discussed above, the quantum state of the entangled qubits determines four distributions over measurement outcomes $Q_{ab}(X = \cdot, Y = \cdot)$, one for each joint setting $(a, b) \in \{1, 2\}^2$. Similarly, each LR theory $\pi \in \Pi$ determines four distributions $P_{ab;\pi}(X = \cdot, Y = \cdot)$. These are the distributions, according to the local realist theory π , of the random variables (X, Y) given that setting (a, b) has been chosen. Thus, the value $P_{ab;\pi}(X = \cdot, Y = \cdot)$ is defined as the sum of four terms:

$$P_{ab;\pi}(X = x, Y = y) := \sum_{\substack{x_1, x_2, y_1, y_2 \in \{\text{T}, \text{F}\} \\ x_a = x, y_b = y}} \pi_{x_1 x_2 y_1 y_2}.$$

We suppose that LR does not dispute the actual setting distribution σ which is used in the experiment, she only disputes the probability distributions of the measurement outcomes given the settings. According to LR therefore, the outcome of a single trial is distributed as $P_{\sigma;\pi}$ defined by $P_{\sigma;\pi}(X = x, Y = y, A = a, B = b) := \sigma_{ab} P_{ab;\pi}(X = x, Y = y)$.

III. THE NONLOCALITY PROOFS

In this paper we compute statistical strength for five (or six, since we have two versions of Bell’s proof) celebrated nonlocality proofs. In this section, we describe the general type of reasoning by which these nonlocality proofs are established, using CHSH as a concrete example. Details on the other proofs can be found in Appendices I and II.

Let us interpret the measurement outcomes F and T in terms of Boolean logic, i.e. F is “false” and T is “true”. We can then use Boolean expressions such as $X_2 \& Y_2$, which evaluates to true whenever both X_2 and Y_2 evaluate to ‘true’, i.e. when both $X_2 = T$ and $Y_2 = T$. We derive the proofs by applying the rule that if the event $X = T$ implies the event $Y = T$ (in short “ $X \implies Y$ ”), then $\Pr(X) \leq \Pr(Y)$. In similar vein, we will use rules like $\Pr(X \vee Y) \leq \Pr(X) + \Pr(Y)$ and $1 - \Pr(\neg X) - \Pr(\neg Y) \leq 1 - \Pr(\neg X \vee \neg Y) = \Pr(X \& Y)$.

As an aside we want to mention that the proofs of Bell, CHSH and Hardy all contain the following argument, which can be traced back to the nineteenth century logician George Boole (1815–1864) [8]. Consider four events such that $\neg B \cap \neg C \cap \neg D \implies \neg A$. Then it follows that $A \implies B \cup C \cup D$. And from this, it follows that $\Pr(A) \leq \Pr(B) + \Pr(C) + \Pr(D)$. In the CHSH argument and the Bell argument, the events concern the equality or inequality of one of the X_i with one of the Y_j . In the Hardy argument, the events concern the joint equality or inequality of one of the X_i , one of the Y_j , and a specific value F or T.

Example 1 (The CHSH Argument): For the CHSH argument one notes that the implication

$$[(X_1 = Y_1) \& (X_1 = Y_2) \& (X_2 = Y_1)] \implies (X_2 = Y_2)$$

is logically true, and hence $(X_2 \neq Y_2) \implies [(X_1 \neq Y_1) \vee (X_1 \neq Y_2) \vee (X_2 \neq Y_1)]$ holds. As a result, local realism implies the following “CHSH inequality”

$$\Pr(X_2 \neq Y_2) \leq \Pr(X_1 \neq Y_1) + \Pr(X_1 \neq Y_2) + \Pr(X_2 \neq Y_1), \quad (1)$$

which can be violated by many choices of settings and states under quantum theory. As a specific example, CHSH identified quantum states and settings such that the first probability equals (approximately) 0.8536 while the three probabilities on the right are each (approximately) 0.1464, thus clearly violating Equation (1). In full detail, the probability distribution that corresponds to CHSH’s proof is as follows

Pr	$X_1 = T$	$X_1 = F$	$X_2 = T$	$X_2 = F$
$Y_1 = T$	0.4268	0.0732	0.4268	0.0732
$Y_1 = F$	0.0732	0.4268	0.0732	0.4268
$Y_2 = T$	0.4268	0.0732	0.0732	0.4268
$Y_2 = F$	0.0732	0.4268	0.4268	0.0732

In Appendix II-C we explain how to arrive at this table. The table lists the four conditional distributions $Q = (Q_{11}, Q_{12}, Q_{21}, Q_{22})$ defined in Section II-C, and thus uniquely determines the nonlocality proof Q . As an example of how to read the table, note that $\Pr(X_2 \neq Y_2)$ is given by

$$\begin{aligned} \Pr(X_2 \neq Y_2) &= \Pr(X_2 = T \& Y_2 = F) + \Pr(X_2 = F \& Y_2 = T) \\ &\approx 0.4268 + 0.4268 = 0.8536, \end{aligned}$$

showing that the expression on the left in Equation (1) is approximately 0.8536. That on the right evaluates to approximately 0.4392.

The other nonlocality proofs are derived in a similar manner: one shows that according to any and all local realist theories, the random variables X_1, X_2, Y_1, Y_2 must satisfy certain logical constraints and hence probabilistic (in)equalities. One then shows that these constraints or (in)equalities can be violated by certain quantum mechanical states and settings corresponding to a table of probabilities of observations similar to (2). Details on the (in)equalities that must hold under local realism are given in Appendix I. Details about the entangled quantum states that give rise to the violations of the various (in)equalities are given in Appendix II.

IV. KULLBACK-LEIBLER DIVERGENCE AND STATISTICAL STRENGTH

A. Kullback-Leibler Divergence

In this section we formally define our notion of ‘statistical strength of a nonlocality proof’. The notion will be based on the KL divergence, an information theoretic quantity which we now introduce. Let \mathcal{Z} be an arbitrary finite set. For a distribution Q over \mathcal{Z} , $Q(z)$ denotes the probability of event $\{z\}$. For two (arbitrary) distributions Q and P defined over \mathcal{Z} , the Kullback-Leibler (KL) divergence from Q to P is defined as

$$D(Q\|P) := \sum_{z \in \mathcal{Z}} Q(z) \log \frac{Q(z)}{P(z)}$$

where the logarithm is taken here, as in the rest of the paper, to base 2. We use the conventions that, for $y > 0$, $y \log 0 := \infty$, and $0 \log 0 := \lim_{y \downarrow 0} y \log y = 0$.

The KL divergence is also known as relative entropy, cross-entropy, information deficiency or I -divergence. Introduced in [23], KL divergence has become a central notion in information theory, statistics and large deviation theory. A good reference is [10]. It is straightforward to show (using concavity of the logarithm and Jensen’s inequality) that $D(Q\|P) \geq 0$ with equality if and only if $P = Q$; in this sense, KL divergence behaves like a distance. However, in general $D(P\|Q) \neq D(Q\|P)$, so formally $D(\cdot\|\cdot)$ is not a distance. In Appendix III we explain at length why KL divergence can be interpreted as a measure of statistical closeness, and what exactly that means. Here we merely give a very short and informal explanation.

1) *KL Divergence and Statistical Strength in Simple Hypothesis Testing:* Let Z_1, Z_2, \dots be a sequence of random variables independently generated either by some distribution P or by some distribution Q with $Q \neq P$. Suppose we are given a sample (sequence of outcomes) z_1, \dots, z_n . We want to perform a statistical test in order to find out whether the sample is from P or Q . Suppose that the sample is, in fact, generated by Q (‘ Q is true’). Then, given enough data, the data will with very high (Q -) probability be overwhelmingly more likely according to Q than according to P . That is, the data strongly suggest that they were sampled from Q rather than P . The ‘statistical distance’ between P and Q indicates *how strongly* or, equivalently, *how convincingly* data that are

generated by Q will prove that they are from Q rather than P . It turns out that this notion of ‘statistical distance’ between two distributions is precisely captured by the KL divergence $D(Q\|P)$, which can be interpreted as *the average amount of support in favor of Q and against P per trial*. The larger the KL divergence, the larger the amount of support per trial. In Appendix III we explain at length what this means. For now we merely give an example: suppose $D(Q\|P) = d$ for some $d > 0$, and we sample from Q n times. Then we will observe a sample that strongly indicates that it is very unlikely that P generated the data. How unlikely? With $(Q-)$ probability very close to 1, our sample will make P about as unlikely as the hypothesis that a coin is fair if, after $n \cdot d$ throws, it has landed ‘heads’ all the time—see Example 4 below.

2) *KL Divergence and Statistical Strength in Composite Hypothesis Testing*: Trying to infer whether a sample was generated by Q or P is called *hypothesis testing* in the statistical literature. A hypothesis is *simple* if it consists of a single probability distribution. A hypothesis is called *composite* if it consists of a *set* of distributions. The composite hypothesis ‘ \mathcal{P} ’ should be interpreted as ‘there exists a $P \in \mathcal{P}$ that generated the data’. Above, we related the KL divergence to statistical strength when testing two simple hypotheses against each other. In this paper, the aim is to test two hypotheses, at least one of which is composite. For concreteness, suppose we want to test the distribution Q against the set of distributions \mathcal{P} . In this case, under some regularity conditions on \mathcal{P} and \mathcal{Z} , the element $P \in \mathcal{P}$ that is *closest* in statistical divergence to Q determines the statistical strength of the best test of Q against \mathcal{P} . Therefore, for a set of distributions \mathcal{P} on \mathcal{Z} we define (as is customary, [10])

$$D(Q\|\mathcal{P}) := \inf_{P \in \mathcal{P}} D(Q\|P).$$

Analogously to $D(Q\|P)$, $D(Q\|\mathcal{P})$ may be interpreted as *the average amount of support in favor of Q and against \mathcal{P} per trial*, if data are generated according to Q .

In our case, QM claims that data are generated by the distribution Q_σ . LR claims that data are generated by some $P \in \mathcal{P}_\sigma$, where $\mathcal{P}_\sigma := \{P_{\sigma,\pi} : \pi \in \Pi\}$. Here Q_σ corresponds to a nonlocality proof equipped with setting distribution σ , and \mathcal{P}_σ is the set of probability distributions of all possible local theories with the same σ — see Section II. QM and LR agree to test the hypothesis Q_σ against \mathcal{P}_σ . QM, who knows that data are really generated according to Q_σ , wants to select σ in such a way that the average amount of support in favor of Q and against \mathcal{P} is maximized. Let $\Sigma' \subseteq \Sigma$ denote the set of all settings σ that QM is allowed to choose from. The previous discussion suggests that QM should pick the $\sigma \in \Sigma'$ that *maximizes statistical strength* $D(Q_\sigma\|\mathcal{P}_\sigma)$. In Appendix III we show that this is (in some sense) also the optimal choice according to statistical theory.

B. Formal Definition of Statistical Strength

We define ‘the statistical strength of nonlocality proof Q ’ in three different manners, depending on the freedom that we allow QM in determining the sampling probabilities of the different measurement settings.

Definition 1 (Strength, Uniform Settings): When each measurement setting is sampled with equal probability, the resulting strength S_Q^{UNI} is defined by

$$S_Q^{\text{UNI}} := D(Q_{\sigma^\circ}\|\mathcal{P}_{\sigma^\circ}) = \inf_{\pi \in \Pi} D(Q_{\sigma^\circ}\|P_{\sigma^\circ,\pi}),$$

where σ° denotes the uniform distribution over the settings.

Definition 2 (Strength, Uncorrelated Settings): When the experimenter QM is allowed to choose any distribution on measurement settings, as long as the distribution for each party is uncorrelated with the distributions of the other parties, the resulting strength S_Q^{UC} is defined by

$$S_Q^{\text{UC}} := \sup_{\sigma \in \Sigma^{\text{UC}}} D(Q_\sigma\|\mathcal{P}_\sigma) = \sup_{\sigma \in \Sigma^{\text{UC}}} \inf_{\pi \in \Pi} D(Q_\sigma\|P_{\sigma,\pi}),$$

where $\sigma \in \Sigma^{\text{UC}}$ denotes the use of uncorrelated settings.

Definition 3 (Strength, Correlated Settings): When the experimenter QM is allowed to choose any distribution on measurement settings (including correlated distributions), the resulting strength S_Q^{COR} is defined by

$$S_Q^{\text{COR}} := \sup_{\sigma \in \Sigma} D(Q_\sigma\|\mathcal{P}_\sigma) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} D(Q_\sigma\|P_{\sigma,\pi}),$$

where $\sigma \in \Sigma$ denoted the use of correlated settings.

Throughout the remainder of the paper, we sometimes abbreviate the subscript $\sigma \in \Sigma^{\text{UC}}$ to Σ^{UC} , and $\pi \in \Pi$ to Π . As we explain in Section VII, we regard the definition S_Q^{UC} allowing maximization over uncorrelated distributions as the ‘right’ one. Henceforth, whenever we speak of ‘statistical strength’ without further qualification, we refer to S_Q^{UC} . Nevertheless, to facilitate comparisons, in Section VI we list our results also for the two alternative definitions of statistical strength.

We have now completed our formal definition of statistical strength. The paper now branches into two parts, which can be read separately: Section V is the mathematical part of this paper. Here we list some essential topological, analytical and game-theoretic properties of our three notions of strength, needed for computing statistical strength in practice. The other part consists of Sections VI and all sections thereafter. In Section VI we calculate statistical strength for various nonlocality proofs. The only mathematical result from Section V that is needed in Section VI and all sections thereafter is the following reassuring fact (Theorem 1, Section V-A, part 2(c)):

Fact 1: $S_Q^{\text{UNI}} \leq S_Q^{\text{UC}} \leq S_Q^{\text{COR}}$. Moreover, $S_Q^{\text{UNI}} > 0$ if and only if Q is a proper nonlocality proof.

V. MATHEMATICAL AND COMPUTATIONAL PROPERTIES OF STATISTICAL STRENGTH

In this section we prove several mathematical properties of our three variations of statistical strength. Some of these are interesting in their own right, giving new insights in the relation between distributions predicted by quantum theory and local realist approximations of it. But their main purpose is to help us compute S_Q^{UC} . We first establish some basic properties of our three notions of strength (Section V-A). Section V-B provides a game-theoretic analysis which will help to compute S_Q^{UC} very efficiently in certain special cases. Finally, in Section V-C we explicitly explain how to compute S_Q^{UC} in practice.

A. Basic Properties

We proceed to list some essential properties of $s_Q^{\text{UNI}}, s_Q^{\text{UC}}$ and s_Q^{COR} . We say that “nonlocality proof Q is *absolutely continuous with respect to local realist theory π* ” [13] if and only if for all $a, b \in \{1, 2\}, x, y \in \{\text{T}, \text{F}\}$, it holds that if $Q_{ab}(x, y) > 0$ then $P_{ab;\pi}(x, y) > 0$.

Theorem 1: Let Q be a given (not necessarily $2 \times 2 \times 2$) nonlocality proof and Π the (corresponding) set of local realist theories.

- 1) Let $U(\sigma, \pi) := D(Q_\sigma \| P_{\sigma;\pi})$, then:
 - a) For a $2 \times 2 \times 2$ proof, we have that

$$U(\sigma, \pi) = \sum_{a,b \in \{1,2\}} \sigma_{ab} D(Q_{ab} \| P_{ab;\pi}). \quad (3)$$

Hence, the KL divergence $D(Q_\sigma \| P_{\sigma;\pi})$ may alternatively be viewed as the average KL divergence between the conditional distributions of (X, Y) given the setting (A, B) , where the average is over the setting. For a generalized nonlocality proof, the analogous generalization of Equation (3) holds.

- b) For fixed σ , $U(\sigma, \pi)$ is convex and lower semicontinuous on Π , and continuous and differentiable on the interior of Π .
- c) If Q is absolutely continuous with respect to some fixed π , then $U(\sigma, \pi)$ is linear in σ .

- 2) Let

$$U(\sigma) := \inf_{\Pi} U(\sigma, \pi), \quad (4)$$

then

- a) For all $\sigma \in \Sigma$, the infimum in Equation (4) is achieved for some π^* .
- b) The function $U(\sigma)$ is nonnegative, bounded, concave and continuous on Σ .
- c) If Q is not a proper nonlocality proof, then for all $\sigma \in \Sigma, U(\sigma) = 0$. If Q is a proper nonlocality proof, then $U(\sigma) > 0$ for *all* σ in the interior of Σ .
- d) For a 2 party, 2 measurement settings per party nonlocality proof, we further have that, even if Q is proper, then still $U(\sigma) = 0$ for all σ on the boundary of Σ .

- 3) Suppose that σ is in the interior of Σ , then:
 - a) Let Q be a $2 \times 2 \times 2$ nonlocality proof. Suppose that Q is non-trivial in the sense that, for some a, b , Q_{ab} is not a point mass (i.e. $0 < Q_{ab}(x, y) < 1$ for some x, y). Then $\pi^* \in \Pi$ achieves the infimum in Equation (4) if and only if the following 16 (in)equalities hold:

$$\sum_{a,b \in \{1,2\}} \sigma_{ab} \frac{Q_{ab}(x_a, y_b)}{P_{ab;\pi^*}(x_a, y_b)} = 1 \quad (5)$$

for all $(x_1, x_2, y_1, y_2) \in \{\text{T}, \text{F}\}^4$ with $\pi_{x_1, x_2, y_1, y_2}^* > 0$ and

$$\sum_{a,b \in \{1,2\}} \sigma_{ab} \frac{Q_{ab}(x_a, y_b)}{P_{ab;\pi^*}(x_a, y_b)} \leq 1 \quad (6)$$

for all $(x_1, x_2, y_1, y_2) \in \{\text{T}, \text{F}\}^4$ with $\pi_{x_1, x_2, y_1, y_2}^* = 0$.

For generalized nonlocality proofs, the theory $\pi^* \in \Pi$ achieves Equation (4) if and only if the corresponding analogues of Equations (5) and (6) both hold.

- b) Suppose that π^* and π° both achieve the infimum in Equation (4). Then for all $x, y \in \{\text{T}, \text{F}\}$, $a, b \in \{1, 2\}$ with $Q_{ab}(x, y) > 0$, we have $P_{ab;\pi^*}(x, y) = P_{ab;\pi^\circ}(x, y) > 0$. In words, π^* and π° coincide in every measurement setting for every measurement outcome that has positive probability according to Q_σ , and Q is absolutely continuous with respect to π^* and π° .

The proof of this theorem is in Appendix IV-B.

In general, $\inf_{\Pi} U(\sigma, \pi)$ may be achieved for several, different π . By part 2 of the theorem, these must induce the same four marginal distributions $P_{ab;\pi}$. It also follows directly from part 2 of the theorem that, for $2 \times 2 \times 2$ proofs, $s_Q^{\text{UC}} := \sup_{\Sigma^{\text{UC}}} U(\sigma)$ is achieved for some $\sigma^* \in \Sigma^{\text{UC}}$, where $\sigma_{ab}^* > 0$ for all $a, b \in \{1, 2\}$.

B. Game-Theoretic Considerations

The following considerations will enable us to compute s_Q^{UC} very efficiently in some special cases, most notably the CHSH proof.

We consider the following variation of our basic scenario. Suppose that, before the experiments are actually conducted, LR has to decide on a *single* local theory π_0 (rather than the set Π) as an explanation of the outcomes that will be observed. QM then gets to see this π , and can choose σ depending on the π_0 that has been chosen. Since QM wants to maximize the strength of the experiment, he will pick the σ achieving $\sup_{\Sigma^{\text{UC}}} D(Q_\sigma \| P_{\sigma;\pi_0})$. In such a scenario, the ‘best’ LR theory, minimizing statistical strength, is the LR theory π_0 that minimizes, over $\pi \in \Pi$, $\sup_{\Sigma^{\text{UC}}} D(Q_\sigma \| P_{\sigma;\pi})$. Thus, in this slightly different setup, the statistical strength is determined by

$$\bar{s}_Q^{\text{UC}} := \inf_{\Pi} \sup_{\Sigma^{\text{UC}}} D(Q_\sigma \| P_{\sigma;\pi})$$

rather than $s_Q^{\text{UC}} := \sup_{\Sigma^{\text{UC}}} \inf_{\Pi} D(Q_\sigma \| P_{\sigma;\pi})$. Below we show that $\bar{s}_Q^{\text{UC}} \geq s_Q^{\text{UC}}$. As we already argued in Section VII, we consider the definition s_Q^{UC} to be preferable over \bar{s}_Q^{UC} . Nevertheless, it is useful to investigate under what conditions $s_Q^{\text{UC}} = \bar{s}_Q^{\text{UC}}$. Von Neumann’s famous minimax theorem of game theory [27] suggests that

$$\sup_{\Sigma^*} \inf_{\Pi} D(Q_\sigma \| P_{\sigma;\pi}) = \inf_{\Pi} \sup_{\Sigma^*} D(Q_\sigma \| P_{\sigma;\pi}), \quad (7)$$

if Σ^* is a convex subset of Σ . Indeed, Theorem 2 below shows that Equation (7) holds if we take $\Sigma^* = \Sigma$. Unfortunately, Σ^{UC} is *not* convex, and Equation (7) does not hold in general for $\Sigma^* = \Sigma^{\text{UC}}$, whence in general $s_Q^{\text{UC}} \neq \bar{s}_Q^{\text{UC}}$. Nevertheless, Theorem 3 provides some conditions under which Equation (7) does hold with $\Sigma^* = \Sigma^{\text{UC}}$. In Section V-C we put this fact to use in computing s_Q^{UC} for the CHSH nonlocality proof. But before presenting Theorems 2 and 3, we first need to introduce some game-theoretic terminology.

1) *Game-Theoretic Definitions:*

Definition 4 (Statistical Game [14]): A statistical game is a triplet (A, B, L) where A and B are arbitrary sets and $L : A \times B \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ is a loss function. If

$$\sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta) = \inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta),$$

we say that the game has value V with

$$V := \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta).$$

If for some $(\alpha^*, \beta^*) \in A \times B$ we have

$$\text{For all } \alpha \in A: L(\alpha, \beta^*) \leq L(\alpha^*, \beta^*)$$

$$\text{For all } \beta \in B: L(\alpha^*, \beta) \geq L(\alpha^*, \beta^*)$$

then we call (α^*, β^*) a *saddle point* of the game. It is easily seen (Proposition 1, Appendix IV) that, if α° achieves $\sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta)$ and β° achieves $\inf_{\beta \in B} L(\alpha, \beta)$ and the game has value V , then $(\alpha^\circ, \beta^\circ)$ is a saddle point and $L(\alpha^\circ, \beta^\circ) = V$.

Definition 5 (Correlated Game): With each nonlocality proof we associate a corresponding *correlated game*, which is just the statistical game defined by the triple (Σ, Π, U) , where $U : \Sigma \times \Pi \rightarrow \mathbb{R} \cup \{\infty\}$ is defined by

$$U(\sigma, \pi) := D(Q_\sigma \| P_{\sigma; \pi}).$$

By the definition above, if this game has a value then it is equal to V defined by

$$V := \inf_{\Pi} \sup_{\Sigma} U(\sigma, \pi) = \sup_{\Sigma} \inf_{\Pi} U(\sigma, \pi).$$

We call the game *correlated* because we allow distributions σ over measurement settings to be such that the probability that party \mathcal{A} is in setting a is correlated with (is dependent on) the setting b of party \mathcal{B} . The fact that each correlated game has a well defined value is made specific in Theorem 2 below.

Definition 6 (Uncorrelated Game): Recall that we use Σ^{UC} to denote the set of vectors representing uncorrelated distributions in Σ . With each nonlocality proof we can associate the game $(\Sigma^{\text{UC}}, \Pi, U)$ which we call the corresponding *uncorrelated game*.

2) *Game-Theoretic, Saddle Point Theorems:*

Theorem 2 (Saddle Point, Correlated Settings): For every (generalized) nonlocality proof, the correlated game (Π, Σ, U) corresponding to it has a finite value, i.e. there exist a $0 \leq V < \infty$ with $\inf_{\Pi} \sup_{\Sigma} U(\sigma, \pi) = V = \sup_{\Sigma} \inf_{\Pi} U(\sigma, \pi)$. The infimum on the left is achieved for some $\pi^* \in \Pi$; the supremum on the right is achieved for some $\sigma^* \in \Sigma$, so that (π^*, σ^*) is a saddle point.

The proof of this theorem is in Appendix IV-C.2.

In the information-theoretic literature, several well-known minimax and saddle point theorems involving the Kullback-Leibler divergence exist; we mention [22], [34]. However, all these deal with settings that are substantially different from ours.

In the case where there are two parties and two measurement settings per party, we can say a lot more.

Theorem 3 (Saddle Point, $2 \times 2 \times N$ Nonlocality Proofs):

Fix any proper nonlocality proof based on 2 parties with

2 measurement settings per party and let (Σ, Π, U) and $(\Sigma^{\text{UC}}, \Pi, U)$ be the corresponding correlated and uncorrelated games, then:

- 1) The correlated game has a saddle point with value $V > 0$. Moreover,

$$\sup_{\Sigma^{\text{UC}}} \inf_{\Pi} U(\sigma, \pi) \leq \sup_{\Sigma} \inf_{\Pi} U(\sigma, \pi) = V, \quad (8)$$

$$\inf_{\Pi} \sup_{\Sigma^{\text{UC}}} U(\sigma, \pi) = \inf_{\Pi} \sup_{\Sigma} U(\sigma, \pi) = V. \quad (9)$$

- 2) Let

$$\Pi^* := \{\pi : \pi \text{ achieves } \inf_{\Pi} \sup_{\Sigma} U(\sigma, \pi)\},$$

$$\Pi^{\text{UC}*} := \{\pi : \pi \text{ achieves } \inf_{\Pi} \sup_{\Sigma^{\text{UC}}} U(\sigma, \pi)\},$$

then

- a) Π^* is non-empty.
 - b) $\Pi^* = \Pi^{\text{UC}*}$.
 - c) All $\pi^* \in \Pi^*$ are ‘equalizer strategies’, i.e. for all $\sigma \in \Sigma$ we have the equality $U(\sigma, \pi^*) = V$.
- 3) The uncorrelated game has a saddle point if and only if there exists (π^*, σ^*) , with $\sigma^* \in \Sigma^{\text{UC}}$, such that
 - a) π^* achieves $\inf_{\Pi} U(\sigma^*, \pi)$.
 - b) π^* is an equalizer strategy.

If such (σ^*, π^*) exists, it is a saddle point.

The proof of this theorem is in Appendix IV-C.3.

C. Computing Statistical Strength

We are now armed with the mathematical tools needed to compute statistical strength. By convexity of $U(\sigma, \pi)$ in π , we see that for fixed σ , determining $D(Q_\sigma \| \mathcal{P}_\sigma) = \inf_{\Pi} U(\sigma, \pi)$ is a convex optimization problem, which suggests that numerical optimization is computationally feasible. Interestingly, it turns out that computing $\inf_{\Pi} U(\sigma, \pi)$ is formally equivalent to computing the maximum likelihood in a well-known statistical missing data problem. Indeed, we obtained our results by using a ‘vertex direction algorithm’ [17], a clever numerical optimization algorithm specifically designed for statistical missing data problems.

By concavity of $U(\sigma)$ as defined in Theorem 1, we see that determining S_Q^{COR} is a concave optimization problem. Thus, numerical optimization can again be performed. There are some difficulties in computing the measure S_Q^{UC} , since the set Σ^{UC} over which we maximize is not convex. Nevertheless, for the small problems (few parties, particles, measurement settings) we consider here it can be done.

In some special cases, including CHSH, we can do all the calculations *by hand* and do not have to resort to numerical optimization. We do this by making an educated guess of the σ^* achieving $\sup_{\Sigma^{\text{UC}}} D(Q_\sigma \| \mathcal{P}_\sigma)$, and then verify our guess using Theorem 1 and the game-theoretic tools developed in Theorem 3. This can best be illustrated using CHSH as an example.

Example 2 (CHSH, continued): Consider the CHSH nonlocality argument. The quantum distributions Q , given in the

table in Section III have traditionally been compared with the local theory $\tilde{\pi}$ defined by

$$\begin{aligned}\tilde{\pi}_{FFFF} &= \tilde{\pi}_{TTTT} = \tilde{\pi}_{FFTF} = \tilde{\pi}_{TTTF} = \\ \tilde{\pi}_{FFTF} &= \tilde{\pi}_{TTTF} = \tilde{\pi}_{TFTF} = \tilde{\pi}_{FTTF} = \frac{1}{8},\end{aligned}$$

and $\tilde{\pi}_{x_1x_2y_1y_2} = 0$ otherwise. This gives rise to the following probability table:

$P_{ab;\tilde{\pi}}$	$X_1 = T$	$X_1 = F$	$X_2 = T$	$X_2 = F$
$Y_1 = T$	0.375	0.125	0.375	0.125
$Y_1 = F$	0.125	0.375	0.125	0.375
$Y_2 = T$	0.375	0.125	0.125	0.375
$Y_2 = F$	0.125	0.375	0.375	0.125

(10)

There exists no local theory which has uniformly smaller absolute deviations from the quantum probabilities in all four tables. Even though, in general, absolute deviations are not a good indicator of statistical strength, based on the fact that all four tables ‘look the same’, we may still *guess* that, in this particular case, for uniform measurement settings $\tilde{\sigma}_{ab} = 1/4$, $a, b \in \{1, 2\}$, the optimal local realist theory is given by the $\tilde{\pi}$ defined above. We can now use Theorem 1, part 3(a) to check our guess. Checking the 16 equations (5) and (6) shows that our guess was correct: $\tilde{\pi}$ achieves $\inf U(\sigma, \pi)$ for the uniform measurement settings $\tilde{\sigma}$. It is clear that $\tilde{\pi}$ is an equalizer strategy and that $\tilde{\sigma}$ is uncorrelated. But now Theorem 3, part (3) tells us that $(\tilde{\sigma}, \tilde{\pi})$ is a saddle point in the uncorrelated game. This shows that $\tilde{\sigma}$ achieves $\sup_{\Sigma^{\text{uc}}} \inf_{\Pi} D(Q_{\sigma} \| P_{\sigma})$. Therefore, the statistical strength of the CHSH nonlocality proof must be given by

$$s_Q^{\text{uc}} = \sup_{\Sigma^{\text{uc}}} \inf_{\Pi} D(Q_{\sigma} \| P_{\sigma}) = D(Q_{\tilde{\sigma}} \| P_{\tilde{\sigma}; \tilde{\pi}}),$$

which is straightforward to evaluate.

VI. THE RESULTS

Table I summarizes the statistical strengths of the nonlocality proofs of Bell, CHSH, Hardy, Mermin and GHZ. For each proof we consider the three possibilities for the measurement settings, which can be uniform, uncorrelated or correlated as explained in Section IV-B. Note that the numbers in the *middle* column correspond to the ‘right’ definition s_Q^{uc} , which optimizes over measurement settings that are uncorrelated from each other. The data that underlies the strengths of the table are described in Appendix II; here we will highlight the most important cases.

Example 3 (The Strength of CHSH): To help interpret Table I, we continue our Examples 1 and 2 on CHSH. See the table of Equation (2) or, equivalently, Table XIV for the probabilities of this nonlocality proof that uses 2 parties, 2 measurement settings and 2 possible measurement outcomes (see Section II-C for the details). The entry in the first (‘uniform’) column for the CHSH proof in Table I was obtained as follows. The distribution of the measurement settings σ was set to the uniform distribution $\sigma^{\circ} = (0.25, 0.25, 0.25, 0.25)$. Together with the probabilities of the nonlocality proof this results in a joint distribution $Q_{\sigma^{\circ}}$ on measurement settings

and outcomes. This $Q_{\sigma^{\circ}}$ was used to determine the optimal local theory $\pi^* \in \Pi$ that obtains the minimum in

$$\inf_{\pi \in \Pi} D(Q_{\sigma^{\circ}} \| P_{\sigma^{\circ}, \pi}).$$

The resulting π^* can be found numerically and it, together with the corresponding $P_{ab;\pi^*}$ distributions, turns out to be the same as the local theory given in Example 2. The KL divergence between $Q_{\sigma^{\circ}}$ and $P_{\sigma^{\circ}, \pi^*}$ can now be calculated: It is equal to 0.0462738469, as can be seen from the left-most entry in Table I in the CHSH-row.

To get the rightmost entry in this row, we performed the same computation for *all* $\sigma \in \Sigma$ (we will explain later how to do this efficiently). We found that the resulting KL divergence $D(Q_{\sigma}, P_{\sigma, \pi^*})$ (where π^* depends on σ) was, in fact, maximized for $\sigma = \sigma^{\circ}$: there was no gain in trying any other value for σ . Thus, the rightmost column is equal to the leftmost column. Finally, Fact 1, listed at the end of Section IV, implies that the middle column entry must be in between the leftmost and the rightmost, explaining the entry in the middle column. In this highly symmetric case it is actually possible to calculate the strengths analytically, giving us $s_Q^{\text{uni}} = s_Q^{\text{cor}} =$

$$\begin{aligned}s_Q^{\text{uc}} &= \left(\frac{1}{2} + \frac{1}{\sqrt{8}}\right) \log \left(\frac{\frac{1}{2} + \frac{1}{\sqrt{8}}}{\frac{3}{4}}\right) + \left(\frac{1}{2} - \frac{1}{\sqrt{8}}\right) \log \left(\frac{\frac{1}{2} - \frac{1}{\sqrt{8}}}{\frac{1}{4}}\right) \\ &= \frac{1}{4} \sqrt{2} \log(1 + \frac{2}{3} \sqrt{2}) + \frac{1}{2} \log \frac{2}{3} \\ &\approx 0.046274\end{aligned}$$

for the CHSH proof.

The corresponding analyses for the other nonlocality proofs is done in Appendix II, which also gives the quantum states and measurements that each proof utilizes. Unlike the CHSH case, the other nonlocality proofs have different strengths depending on the allowed measurement settings (uniform, uncorrelated or correlated). All this results in a rather long list of tables of optimized distributions, which we will not duplicate in this section.

How to interpret and compare the values of the strengths of the nonlocality proofs is probably best explained in the following example.

Example 4 (Mermin’s “a million runs”): We recall Mermin’s quote from the Introduction where he says that “a million runs” of his experiment should be enough to convince us that “the observed frequencies ... are not chance fluctuations”. We now can put numbers to this.

Assuming that we perform Mermin’s experiment with the optimized, uncorrelated settings, we should get a strength of $1,000,000 \times 0.0191506613 \approx 19,150$. This means that after the million runs of the experiment, the likelihood of local realism still being true is comparable with the likelihood of a coin being fair after 19,150 tosses when the outcome was “tails” all the time.

From Table I we see that in the two-party setting, the nonlocality proof of CHSH is much stronger than those of Bell, Hardy or Mermin, and that this optimal strength is obtained for uniform measurement settings. Furthermore it is

TABLE I
STRENGTHS OF VARIOUS NONLOCALITY PROOFS

Strength	Uniform s_Q^{UNI}	Uncorrelated s_Q^{UC}	Correlated s_Q^{COR}
Original BELL	0.0141597409	0.0158003672	0.0169800305
Optimized BELL	0.0177632822	0.0191506613	0.0211293952
CHSH	0.0462738469	0.0462738469	0.0462738469
HARDY	0.0278585182	0.0279816333	0.0280347655
MERMIN	0.0157895843	0.0191506613	0.0211293952
GHZ	0.2075187496	0.2075187496	0.4150374993

clear that the three-party proof of GHZ is four and a half times stronger than all the two-party proofs.

We also note that the nonlocality proof of Mermin—in the case of non-uniform settings—is equally strong as the optimized version of Bell’s proof. The setting distributions tables in Appendix II-E shows why this is the case: the optimal setting distribution for Mermin exclude one setting on \mathcal{A} ’s side, and one setting on \mathcal{B} ’s side, thus reducing Mermin’s proof to that of Bell. One can view this is as an example of how a proof that is easier to understand (Mermin) is not necessarily stronger than one that has more subtle arguments (Bell).

We also see that in general, except for CHSH’s proof, uniform setting distributions do not give the optimal strength of a nonlocality proof. Rather, the experimenter obtains more evidence for the nonlocality of nature by employing sampling frequencies that are biased towards those settings that are more relevant for the nonlocality proof.

VII. INTERPRETATION AND DISCUSSION

A. Is our definition of statistical strength the right one?

We can think of two objections against our definition of statistical strength. First, we may wonder whether the KL divergence is really the right measure to use. Second, assuming that KL divergence is the right measure, is our game-theoretic set-up justified? We treat both issues in turn.

1) *Is Kullback-Leibler divergence justified?*: We can see two possible objections against KL divergence: (1) different statistical paradigms such as the ‘Bayesian’ and ‘frequentist’ paradigm define ‘amount of support’ in different manners (Appendix III); (2) ‘asymptopia’: KL divergence is an inherently asymptotic notion.

These two objections are inextricably intertwined: there exists no non-asymptotic measure which would (a) be acceptable to all statisticians; (b) would not depend on prior considerations, such as a ‘prior distribution’ for the distributions involved in the Bayesian framework, and a pre-set significance level in the frequentist framework. Thus, since we consider it most important to arrive at a generally acceptable and objective measure, we decided to opt for the KL divergence. We add here that even though this notion is asymptotic, it can be used to provide numerical bounds on the actual, non-asymptotic amount of support provided on each trial, both in Bayesian

and in frequentist terms. We have not pursued this option any further here.

2) *Game-Theoretic Justification*: There remains the question of whether to prefer s_Q^{UNI} , s_Q^{COR} or, as we do, s_Q^{UC} . The problem with s_Q^{UNI} is that, for any given combination of nonlocality proof Q and local theory π , some settings may provide, on average, more information about the nonlocality of nature than others. This is evident from Table I. We see no reason for the experimenter not to exploit this.

On the other hand, allowing QM to use *correlated* distributions makes QM’s case much weaker: LR might now argue that there is some hidden communication between the parties. Since QM’s goal is to provide an experiment that is as convincing as possible to LR, we do not allow for this situation. Thus, among the three definitions considered, s_Q^{UC} seems to be the most reasonable one. Nevertheless, one may still argue that *none* of the three definitions of strength are correct: they all seem unfavorable to QM, since we allow LR to adjust his theory to whatever frequency of measurement settings QM is going to use. In contrast, our definition does *not* allow QM to adjust his setting distribution to LR’s choice (which would lead to strength defined as inf sup rather than sup inf , Section V-B). The reason why we favor LR in this way is that the quantum experimenters QM should try to convince LR that nature is nonlocal *in a setting about which LR cannot complain*. Thus, if LR wants to entertain several local theories at the same time, or wants to have a look at the probabilities σ_{ab} before the experiment is conducted, QM should allow him to do so—he will still be able to convince LR, even though he may need to repeat the experiment a few more times. Nevertheless, in developing clever strategies for computing s_Q^{UC} , it turns out to be useful to investigate the inf sup scenario in more detail. This was done in Section V-B.

Summarizing, our approach is highly nonsymmetric between quantum mechanics and local realism. There is only one quantum theory, and QM believes in it, but he must arm

himself against any and all local realists.¹

B. Related Work by Peres

Earlier work in our direction has been done by Peres [29] who adopts a Bayesian approach. Peres implicitly uses the same definition of strength of nonlocality proofs as we do here, after merging equal probability joint outcomes of the experiment. Our work extends his in several ways; most importantly, we allow the experimentalist to optimize her experimental settings, whereas Peres assumes particular (usually uniform) distributions over the settings. Peres determines LR's best theory by an inspired guess. The proofs he considers have so many symmetries, that the best LR theory has the same equal probability joint outcomes as the QM experiment, the reduced experiment is binary, and his guess always gives the right answer. But his strategy would not work for, e.g., the Hardy proof, which is less symmetric.

Peres starts out with a nonlocality proof Q_σ to be tested against local theory $P_{\sigma;\pi}$, for some fixed distribution σ . Peres then defines the *confidence depressing factor for n trials*. In fact, Peres rediscovers the notion of KL divergence, since a straightforward calculation shows that for large n ,

$$D(Q_\sigma \| P_{\sigma;\pi}) = \frac{1}{n} \log(\text{confidence depressing factor}). \quad (11)$$

For any given large n , the larger the confidence depressing factor for n , the more evidence against $P_{\sigma;\pi}$ we are likely to get on the basis of n trials. Thus, when comparing a fixed quantum experiment (with fixed σ) Q_σ to a fixed local theory $P_{\sigma;\pi}$, Peres' notion of strength is equivalent to ours. Peres then goes on to say that, when comparing a fixed quantum experiment Q_σ to the corresponding set of *all* local theories \mathcal{P}_σ , we may expect that LR will choose the local theory with the least confidence depressing factor, i.e. the smallest KL divergence to Q_σ . Thus, whenever Peres chooses uniform σ , his notion of strength corresponds to our S_Q^{UNI} , represented in the first column of Table I. In practice, Peres chooses an intuitive σ , which is usually, but not always uniform in our sense. For example, in the GHZ scenario, Peres implicitly assumes that only those measurement settings are used that correspond to the probabilities (all 0 or 1) appearing in the GHZ-inequality (12), Appendix I-D. Thus, his experiment corresponds to a uniform distribution on those four settings. Interestingly, such a distribution on settings is *not* allowed under our definition of strength S_Q^{UC} , since it makes the probability of the setting at party A dependent on (correlated with) the other settings. This explains that Peres obtains a larger strength for GHZ than we do: he obtains $\log 0.75^{-n} = 0.4150 \dots n$, which corresponds

¹Some readers might wonder what would happen if one would replace the $D(Q\|P)$ in our analysis by $D(P\|Q)$. In short, $D(P\|Q)$ quantifies how strongly the predictions of quantum mechanics disagree with the outcomes of a classical system P . Hence such an analysis would be useful if one has to prove that the statistics of a local realistic experiment (say, a network of classically communicating computers) are not in correspondence with certain predictions of quantum mechanics. The minimax solution of the game based on $D(P\|Q)$ provides a value of Q which QM should specify as part of a challenge to LR to reproduce quantum predictions with LR's theory. With this challenge, the computer simulation using LR's theory can be run in as short as possible amount of time, before giving sufficient evidence that LR has failed.

to our S_Q^{COR} : the uniform distribution on the restricted set of settings appearing in the GHZ proof turns out to be the optimum over all distributions on measurement settings.

Our approach may be viewed as an extension of Peres' in several ways. First, we relate his confidence depressing factor to the Kullback-Leibler divergence and we argue that this is the right measure to use not just from a Bayesian point of view, but also from an information-theoretic point of view and the standard, 'orthodox' frequentist statistics point of view. Second, we extend his analysis to non-uniform distributions σ over measurement settings and show that in some cases, substantial statistical strength can be gained if QM uses non-uniform sampling distributions. Third, we give a game-theoretic treatment of the maximization of σ and develop the necessary mathematical tools to enable fast computations of statistical strength. Fourth, whereas Peres finds the best LR theory by cleverly guessing, we show the search for this theory can be performed automatically.

C. Which nonlocality proof is strongest and what does it mean?

1) *Caveat: statistical strength is not the whole story*: First of all, we stress that statistical strength is by no means the only factor in determining the 'goodness' of a nonlocality proof and its corresponding experiment. Various other aspects also come into play, such as: how easy is it to prepare certain types of particles in certain states? Can we arrange to have the time and spatial separations which are necessary to make the results convincing? Can we implement the necessary random changes in settings per trial, quickly enough? Our notion of strength neglects all these important practical aspects.

2) *Comparing GHZ and CHSH*: GHZ is the clear winner among all proofs that we investigated, being about 4.5 times stronger than CHSH, the strongest two-party proof that we found. This means that, to obtain a given level of support for QM and against LR, the optimal CHSH experiment has to be repeated about 4.5 times as often as the optimal GHZ experiment.

On the other hand, the GHZ proof is much harder to prepare experimentally. In light of the reasoning above, and assuming that both CHSH and GHZ can be given a convincing experimental implementation, it may be the case that repeating the CHSH experiment $4.5 \times n$ times is much cheaper than repeating GHZ n times.

3) *Nonlocality 'without inequality'?*: The GHZ proof was the first of a new class of proofs of Bell's theorem, "without inequalities". It specifies a state and collection of settings, such that all QM probabilities are zero or one, while this is impossible under LR. The QM probabilities involved are just the probabilities of the four events in Equation (12), Appendix I-D. The fact that all these must be either 0 or 1 has led some to claim that the corresponding experiment has to be performed only once in order to rule out local realism². As has been observed before [29], this is not the case. This can be seen

²Quoting [29], "The list of authors [claiming that a single experiment is sufficient to invalidate local realism] is too long to be given explicitly, and it would be unfair to give only a partial list."

immediately if we let LR adopt the uniform distribution on all possible observations. Then, although QM is correct, no matter how often the experiment is repeated, the resulting sequence of observations does not have zero probability under LR’s local theory — simply because *no* sequence of observations has probability 0 under LR’s theory. We can only decide that LR is wrong on a statistical basis: the observations are *much more likely* under QM than under LR. This happens even if, instead of using the uniform distribution, LR uses the local theory that is closest in KL divergence to the Q induced by the GHZ scenario. The reason is that there exists a positive ϵ such that any local realist theory which comes within ϵ of all the equalities but one, is forced to deviate by more than ϵ in the last. Thus, accompanying the GHZ style proof without inequalities, is an implied *inequality*, and it is this latter inequality that can be tested experimentally.

VIII. FUTURE EXTENSIONS AND CONJECTURES

The purpose of our paper has been to objectively compare the statistical strength of existing proofs of Bell’s theorem. The tools we have developed, can be used in many further ways.

Firstly, one can take a given quantum state, and ask the question, what is the best experiment which can be done with it. This leads to a measure of statistical nonlocality of a given joint state, whereby one is optimizing (in the outer optimization) not just over setting distributions, but also over the settings themselves, and even over the number of settings.

Secondly, one can take a given experimental type, for instance: the $2 \times 2 \times 2$ type, and ask what is the best state, settings, and setting distribution for that type of experiment? This comes down to replacing the outer optimization over setting distributions, with an optimization over states, settings, and setting distribution.

Using numerical optimization, we were able to analyze a number of situations, leading to the following conjectures.

Conjecture 1: Among all $2 \times 2 \times 2$ proofs, and allowing correlated setting distributions, CHSH is best.

Conjecture 2: Among all $3 \times 2 \times 2$ proofs, and allowing correlated setting distributions, GHZ is best.

Conjecture 3: The best experiment with the Bell singlet state is the CHSH experiment.

In [1] Acín et al. investigated the natural generalization of CHSH type experiments to qutrits. Their main interest was the resistance of a given experiment to noise, and to their surprise they discovered in the $2 \times 2 \times 3$ case, that a less entangled state was more resistant to noise than the maximally entangled state. After some preliminary investigations, we found that that a similar experiment with an *even less entangled* state gives a stronger nonlocality experiment.

Conjecture 4: The strongest possible $2 \times 2 \times 3$ nonlocality proof has statistical strength 0.077, and it uses the bipartite state $\approx 0.6475|1, 1\rangle + 0.6475|2, 2\rangle + 0.4019|3, 3\rangle$.

If true, this conjecture is in remarkable contrast with what appears to be the strongest possible $2 \times 2 \times 3$ nonlocality proof that uses the maximally entangled state $(|1, 1\rangle + |2, 2\rangle + |3, 3\rangle)/\sqrt{3}$, which has a statistical strength of only 0.058.

Conjecture 4 suggests that it is not always the case that a quantum state with more ‘entropy of entanglement’ [6] will always give a stronger nonlocality proof. Rather, it seems that entanglement and statistical nonlocality are different quantities. One possibility however is that the counterintuitive results just mentioned would disappear if one could do joint measurements on several pairs of entangled qubits, qutrits, or whatever. A regularized measure of nonlocality of a given state, would be the limit for $k \rightarrow \infty$, of the strength of the best experiment based on k copies of the state (where the parties are allowed to make joint measurements on k systems at the same time), divided by k . One may conjecture for instance that the best experiment based on two copies of the Bell singlet state is more than twice as good as the best experiment based on single states. That would be a form of “superadditivity of nonlocality”, quite in line with other forms of superadditivity which is known to follow from entanglement.

Conjecture 5: There is an experiment on pairs of Bell singlets, of the $2 \times 4 \times 4$ type, more than twice as strong as CHSH, and involving joint measurements on the pairs.

IX. ACKNOWLEDGMENTS

The authors are very grateful to Piet Groeneboom for providing us with the programs [17] needed to compute $\inf_{P \in \mathcal{P}} D(Q||P)$, and helping us to adjust them for our purposes.

The authors would like to thank EURANDOM for financial support. Part of this research was performed while Peter Grünwald was visiting the University of California at Santa Cruz (UCSC). The visit was funded by NWO (the Netherlands Organization for Scientific Research) and the Applied Math and Computer Science Departments of UCSC. Wim van Dam’s work was partly supported by funds provided by the U.S. Department of Energy (DOE) and cooperative research agreement DF-FC02-94ER40818, by a CMI postdoctoral fellowship, and by an earlier HP/MSRI fellowship. Richard Gill’s research was partially funded by project RESQ (IST-2001-37559) of the IST-FET programme of the European Union. Richard Gill is grateful for the hospitality of the Quantum Probability group at the department of mathematics of the University of Greifswald, Germany. His research there was supported by European Commission grant HPRN-CT-2002-00279, RTN QP-Applications. This work was also supported in part by the European Union’s PASCAL Network of Excellence, IST-2002-506778.

REFERENCES

- [1] A. Acín, T. Durt, N. Gisin, J.I. Latorre, “Quantum non-locality in two three-level systems”, *Physical Review A*, Volume 65, 052325, 2002.
- [2] R.R. Bahadur, “An optimal property of the likelihood ratio statistic”, In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1, pp. 13–26, 1967.
- [3] V. Balasubramanian, “A Geometric Formulation of Occam’s Razor for Inference of Parametric Distributions”, arXiv:adap-org/9601001, 1996.
- [4] A. Barron and T. Cover, “Minimum complexity density estimation”, *IEEE Transactions on Information Theory*, Volume 37(4), pp. 1034–1054, 1991.
- [5] J.S. Bell, “On the Einstein-Podolsky-Rosen paradox”, *Physics*, Volume 1, pp. 195–200, 1964.

- [6] C.H. Bennett, H.J. Bernstein, S. Popescu and B. Schumacher, “Concentrating partial entanglement by local operations”, *Physical Review A*, Volume 53, No. 4, pp. 2046–2052, 1996.
- [7] J.M. Bernardo and A.F.M. Smith, *Bayesian theory*, John Wiley, 1994.
- [8] G. Boole, *An Investigation of the Laws of Thought (on which are founded the mathematical theories of logic and probabilities)*, MacMillan and Co., Cambridge, 1854. Reprinted by Dover, 1951.
- [9] J.F. Clauser, M.A. Horne, A. Shimony, and R.A. Holt, “Proposed experiment to test local hidden-variable theories”, *Physical Review Letters*, Volume 23, pp. 880–884, 1969.
- [10] T.M. Cover and J.A. Thomas, *Elements of Information Theory*, Wiley Interscience, New York, 1991.
- [11] W. van Dam, R.D. Gill, and P.D. Grünwald, “The statistical strength of nonlocality proofs”, arXiv:quant-ph/0307125, 2003.
- [12] W. Feller, *An Introduction to Probability Theory and Its Applications*, Wiley, Volume 1, third edition, 1969.
- [13] W. Feller, *An Introduction to Probability Theory and Its Applications*, Wiley, Volume 2, third edition, 1969.
- [14] T.S. Ferguson, *Mathematical Statistics — a decision-theoretic approach*, Academic Press, 1967.
- [15] T.L. Fine, *Theories of Probability*, Academic Press, 1973.
- [16] D.M. Greenberger, M. Horne, and A. Zeilinger, “Going beyond Bell’s theorem”, In M. Kafatos, editor, *Bell’s Theorem, Quantum Theory, and Conceptions of the Universe*, pp. 73–76, Kluwer, Dordrecht, 1989.
- [17] P. Groeneboom, G. Jongbloed, and J.A. Wellner “Vertex direction algorithms for computing nonparametric function estimates in mixture models”, arXiv:math.ST/0405511, submitted, 2003.
- [18] P. Groeneboom and J. Oosterhoff, “Bahadur efficiency and probabilities of large deviations”, *Statistica Neerlandica*, Volume 31, pp. 1–24, 1977.
- [19] P.D. Grünwald, *The Minimum Description Length Principle and Reasoning under Uncertainty*, PhD thesis, University of Amsterdam, The Netherlands, October 1998. Available as ILLC Dissertation Series 1998-03.
- [20] P. Grünwald and A.P. Dawid, “Game theory, maximum entropy, minimum discrepancy, and robust Bayesian decision theory”, *Annals of Statistics*, Volume 32(4), pp. 1367–1433, August 2004.
- [21] L. Hardy, “Nonlocality for two particles without inequalities for almost all entangled states”, *Physical Review Letters*, Volume 71, pp. 1665–1668, 1993.
- [22] D. Haussler, “A general minimax result for relative entropy”, *IEEE Transactions on Information Theory*, Volume 43(4), pp. 1276–1280, 1997.
- [23] S. Kullback and R.A. Leibler, “On information and sufficiency”, *Annals of Mathematical Statistics*, Volume 22, pp. 76–86, 1951.
- [24] P.M. Lee, *Bayesian Statistics — an introduction*, Arnold & Oxford University Press, 1997.
- [25] M. Li and P.M.B. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, (revised and expanded Second edition), New York, Springer-Verlag, 1997.
- [26] N.D. Mermin, “Quantum mysteries for anyone”, *Journal of Philosophy*, Volume 78, pp. 397–408, 1981.
- [27] J. Von Neumann, “Zur Theorie der Gesellschaftsspiele”, *Mathematische Annalen*, Volume 100, pp. 295–320, 1928.
- [28] A. Peres, *Quantum Theory: Concepts and Methods*, Fundamental Theories of Physics, Volume 57, Kluwer Academic Publishers, 1995.
- [29] A. Peres, “Bayesian analysis of Bell inequalities”, *Fortschritte der Physik*, Volume 48, pp. 531–535, 2000.
- [30] E. Posner, “Random coding strategies for minimum entropy”, *IEEE Transactions on Information Theory*, Volume 21, pp. 388–391, 1975.
- [31] J.A. Rice, *Mathematical Statistics and Data Analysis*, Duxbury Press, 1995.
- [32] J. Rissanen, *Stochastic Complexity in Statistical Inquiry*, World Scientific Publishing Company, 1989.
- [33] R.T. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, New Jersey, 1970.
- [34] F. Topsøe, “Information-theoretical optimization techniques”, *Kybernetika*, Volume 15(1), 1989.
- [35] V. Vapnik, *Statistical Learning Theory*, John Wiley, 1998.

APPENDIX I THE NONLOCALITY ARGUMENTS

In this Appendix we present the inequalities and logical constraints that must hold under local realism yet can be violated under quantum mechanics. The specific quantum

states chosen to violate these inequalities, as well as the closest possible (in the KL divergence sense) local theories are listed in Appendix II.

A. Arguments of Bell and CHSH

CHSH’s argument was described in Example 1. By exactly the same line of reasoning as used in obtaining the CHSH inequality (1), one also obtains Bell’s inequality

$$\Pr(X_1 = Y_1) \leq \Pr(X_2 \neq Y_2) + \Pr(X_2 \neq Y_1) + \Pr(X_1 + Y_2).$$

See Sections II-A and II-B for how this inequality can be violated.

B. Hardy’s Argument

Hardy noted the following: if $(X_2 \& Y_2)$ is true, and $(X_2 \implies Y_1)$ is true, and $(Y_2 \implies X_1)$ is true, then $(X_1 \& Y_1)$ is true. Thus $(X_2 \& Y_2)$ implies: $\neg(X_2 \implies Y_1)$ or $\neg(Y_2 \implies X_1)$ or $(X_1 \& Y_1)$. Therefore

$$\Pr(X_2 \& Y_2) \leq \Pr(X_2 \& \neg Y_1) + \Pr(\neg X_1 \& Y_2) + \Pr(X_1 \& Y_1).$$

On the other hand, according to quantum mechanics it is possible that the first probability is positive, in particular, equals 0.09, while the three other probabilities here are all zero. See Section II-D for the precise probabilities.

C. Mermin’s Argument

Mermin’s argument uses three settings on both sides of the two parties, thus giving the set of six events $\{X_1, Y_1, X_2, Y_2, X_3, Y_3\}$. First, observe that the three equalities in $(X_1 = Y_1) \& (X_2 = Y_2) \& (X_3 = Y_3)$ implies at least one of the three statements in $((X_1 = Y_2) \& (X_2 = Y_1)) \vee ((X_1 = Y_3) \& (X_3 = Y_1)) \vee ((X_2 = Y_3) \& (X_3 = Y_2))$. By the standard arguments that we used before, we see that

$$\begin{aligned} 1 - \Pr(X_1 \neq Y_1) - \Pr(X_2 \neq Y_2) - \Pr(X_3 \neq Y_3) \\ \leq \Pr((X_1 = Y_1) \& (X_2 = Y_2) \& (X_3 = Y_3)), \end{aligned}$$

and that

$$\begin{aligned} \Pr \left(\begin{array}{c} ((X_1 = Y_2) \& (X_2 = Y_1)) \\ \vee \\ ((X_1 = Y_3) \& (X_3 = Y_1)) \\ \vee \\ ((X_2 = Y_3) \& (X_3 = Y_2)) \end{array} \right) \\ \leq \left(\begin{array}{c} \Pr((X_1 = Y_2) \& (X_2 = Y_1)) \\ + \\ \Pr((X_1 = Y_3) \& (X_3 = Y_1)) \\ + \\ \Pr((X_2 = Y_3) \& (X_3 = Y_2)) \end{array} \right) \\ \leq \frac{1}{2} \left(\begin{array}{c} \Pr(X_1 = Y_2) + \Pr(X_2 = Y_1) \\ + \\ \Pr(X_1 = Y_3) + \Pr(X_3 = Y_1) \\ + \\ \Pr(X_2 = Y_3) + \Pr(X_3 = Y_2) \end{array} \right). \end{aligned}$$

As a result we have the ‘Mermin inequality’

$$1 \leq \sum_{i=1}^3 \Pr(X_i \neq Y_i) + \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^3 \Pr(X_i = Y_j),$$

which gets violated by a state and measurement setting that has probabilities $\Pr(X_i \neq Y_i) = 0$ and $\Pr(X_i = Y_j) = \frac{1}{4}$ for $i \neq j$ (see Section II-E in the appendix).

D. GHZ’s Argument

Starting with [16], GHZ, proofs against local realism have been based on systems of three or more qubits, on systems of higher-dimensional quantum systems, and on larger sets of measurements (settings) per particle. Each time we are allowed to search over a wider space we may be able to obtain stronger nonlocality proofs, though each time the actual experiment may become harder to set up in the laboratory.

Let \oplus denote the exclusive or operation such that $X \oplus Y$ is true if and only if $X \neq Y$. Then the following implication must hold

$$\begin{aligned} & ((X_1 \oplus Y_2 = Z_2) \& (X_2 \oplus Y_1 = Z_2) \& (X_2 \oplus Y_2 = Z_1)) \\ \implies & (X_1 \oplus Y_1 = Z_1). \end{aligned}$$

Now, by considering the contrapositive, we get

$$\begin{aligned} & \Pr(X_1 \oplus Y_1 \neq Z_1) \\ & \leq \Pr((X_1 \oplus Y_2 \neq Z_2) \vee (X_2 \oplus Y_1 \neq Z_2) \vee (X_2 \oplus Y_2 \neq Z_1)). \end{aligned}$$

And because $\Pr(X \oplus Y \neq Z) = \Pr(X \oplus Y \oplus Z)$, this gives us GHZ’s inequality:

$$\begin{aligned} & \Pr(X_1 \oplus Y_1 \oplus Z_1) \tag{12} \\ & \leq \Pr(X_1 \oplus Y_2 \oplus Z_2) + \Pr(X_2 \oplus Y_1 \oplus Z_2) + \Pr(X_2 \oplus Y_2 \oplus Z_1). \end{aligned}$$

This inequality can be violated by a three way entangled state and measurement settings that give $\Pr(X_1 \oplus Y_1 \oplus Z_1) = 1$ and $\Pr(X_1 \oplus Y_2 \oplus Z_2) = \Pr(X_2 \oplus Y_1 \oplus Z_2) = \Pr(X_2 \oplus Y_2 \oplus Z_1) = 0$. The details of this proof are in Section II-F.

APPENDIX II

THE NONLOCALITY PROOFS, THEIR OPTIMAL SETTING DISTRIBUTIONS AND BEST CLASSICAL THEORIES

In this appendix we list the nonlocality proofs of Bell, an optimized version of Bell, CHSH, Hardy, Mermin and GHZ and their solutions. The proofs themselves are described by a multipartite quantum state and the measurement bases $|m_i\rangle$ of the parties. Because all bases are two dimensional in the proofs below, it is sufficient to only describe the vector $|m_i\rangle$, where it is understood that the other basis vector ($|\perp m_i\rangle$) is the orthogonal one. Because of its frequent use, we define for the whole appendix the rotated vector $|R(\phi)\rangle := \cos(\phi)|0\rangle + \sin(\phi)|1\rangle$. A *measurement setting* refers to the bases that parties use during a trial of the experiment. All proofs, except Mermin’s, have two different settings per party (in MERMIN they have three).

Given the state and the measurement bases, the proof is summarized in a table of probabilities of the possible measurement outcomes. Here we list these probabilities conditionally on the specific measurement settings. For example,

for Bell’s original nonlocality proof, which uses the state $|\Psi\rangle := \frac{1}{\sqrt{2}}(|0_A 0_B\rangle + |1_A 1_B\rangle)$ and the measurement vectors $|X = T\rangle_{a=1} := |R(0)\rangle$ and $|Y = T\rangle_{b=1} := |R(\frac{\pi}{8})\rangle$, we list the probability $Q_{11}(X = T, Y = T) = |\langle \Psi | X = T, Y = T \rangle_{a=1, b=1}|^2 \approx 0.4268$ in the table.

As discussed in the article (Section IV-B), the strength of a nonlocality proof will depend on the probabilities σ with which the parties use the different measurement settings. Recall that we defined three different notions of strength, depending on how these probabilities are determined: uniform settings (S_Q^{UNI}), uncorrelated settings (S_Q^{UC}) and correlated settings (S_Q^{COR}). For both the correlated and the uncorrelated settings, the parties can optimize their setting distributions to get the strongest possible statistics to prove the nonlocality of their measurement outcomes. We list these optimal distributions below where, for example, $\Pr(a = 1) = \sigma_{10} + \sigma_{11}$ stands for the probability that party \mathcal{A} uses the measurement basis $\{|(X = T|a = 1)\rangle, |(X = F|a = 1)\rangle\}$ and $\Pr(a = 1, b = 2) = \sigma_{ab}$ is the probability that \mathcal{A} uses the basis $\{|(X = T|a = 1)\rangle, |(X = F|a = 1)\rangle\}$ while \mathcal{B} uses the basis $\{|(Y = T|b = 2)\rangle, |(Y = F|b = 2)\rangle\}$, etc.

Associated with these optimal distributions there is an optimal local realist theory $\pi \in \Pi$ (see Section IV-B). We do not list the π -probabilities for such optimal classical theories as this would be too cumbersome and not very enlightening. Instead, we show the corresponding probabilities $P_{ab;\pi}$, which should be compared with the Q_{ab} -tables of the nonlocality proofs. Combining these data tables for each proof and each scenario we obtain the strengths that were listed in Section VI.

A. Original Bell

For Bell’s proof of nonlocality we have to make a distinction between the original version, which Bell described [5], and the optimized version, which is described by Peres in [28].

First we discuss Bell’s original proof. Take the bipartite state $\frac{1}{\sqrt{2}}|0_A 0_B\rangle + \frac{1}{\sqrt{2}}|1_A 1_B\rangle$, and the measurement settings

$$\begin{aligned} |X = T\rangle_{a=1} & := |R(0)\rangle \quad \text{and} \quad |X = T\rangle_{a=2} := |R(\frac{\pi}{8})\rangle \\ |Y = T\rangle_{b=1} & := |R(\frac{\pi}{8})\rangle \quad \text{and} \quad |Y = T\rangle_{b=2} := |R(\frac{\pi}{4})\rangle \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table II (where $\frac{1}{4} + \frac{1}{8}\sqrt{2} \approx 0.4267766953$ and $\frac{1}{4} - \frac{1}{8}\sqrt{2} \approx 0.0732233047$).

(1) *Uniform Settings, Original Bell:* When the two parties use uniform distributions for their settings, the optimal classical theory is the one described in Table III. The corresponding KL divergence is 0.0141597409.

(2) *Uncorrelated Settings, Original Bell:* The optimized, uncorrelated setting distribution is described in Table IV. The probabilities of the best classical theory for this uncorrelated setting distribution are those in Table V. The KL divergence for Bell’s original proof, with uncorrelated measurement settings is 0.0158003672.

(3) *Correlated Settings, Original Bell:* The optimized, correlated setting distribution is described in Table VI. The probabilities of the best classical theory for this distribution are described in Table VII. The corresponding KL divergence is 0.0169800305.

TABLE II
QUANTUM PREDICTIONS ORIGINAL BELL

$Q_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.4267766953	0.0732233047	0.5	0
	$y=F$	0.0732233047	0.4267766953	0	0.5
$b=2$	$y=T$	0.25	0.25	0.4267766953	0.0732233047
	$y=F$	0.25	0.25	0.0732233047	0.4267766953

TABLE III
BEST CLASSICAL THEORY FOR UNIFORM SETTINGS ORIGINAL BELL. KL DISTANCE: 0.0141597409.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.3970311357	0.1029688643	0.5000000000	0.0000000000
	$y=F$	0.1029688643	0.3970311357	0.0000000000	0.5000000000
$b=2$	$y=T$	0.2940622714	0.2059377286	0.3970311357	0.1029688643
	$y=F$	0.2059377286	0.2940622714	0.1029688643	0.3970311357

TABLE IV
OPTIMIZED UNCORRELATED SETTING DISTRIBUTION ORIGINAL BELL

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma^{uc}$	$a=1$	$a=2$	$\Pr(B=b)$
$b=1$	0.2316110419	0.1327830656	0.3643941076
$b=2$	0.4039948505	0.2316110419	0.6356058924
$\Pr(A=a)$	0.6356058924	0.3643941076	

TABLE V
BEST CLASSICAL THEORY FOR UNCORRELATED SETTINGS ORIGINAL BELL. KL DISTANCE: 0.0158003672.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.3901023259	0.1098976741	0.5000000000	0.0000000000
	$y=F$	0.1098976741	0.3901023259	0.0000000000	0.5000000000
$b=2$	$y=T$	0.2802046519	0.2197953481	0.3901023259	0.1098976741
	$y=F$	0.2197953481	0.2802046519	0.1098976741	0.3901023259

TABLE VI
OPTIMIZED CORRELATED SETTING DISTRIBUTION ORIGINAL BELL

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma$	$a=1$	$a=2$
$b=1$	0.2836084841	0.1020773549
$b=2$	0.3307056768	0.2836084841

TABLE VII
BEST CLASSICAL THEORY FOR CORRELATED SETTINGS ORIGINAL BELL. KL DISTANCE: 0.0169800305.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.3969913979	0.1030086021	0.4941498806	0.0058501194
	$y=F$	0.1030086021	0.3969913979	0.0058501194	0.4941498806
$b=2$	$y=T$	0.2881326764	0.2118673236	0.3969913979	0.1030086021
	$y=F$	0.2118673236	0.2881326764	0.1030086021	0.3969913979

B. Optimized Bell

Take the bipartite state $\frac{1}{\sqrt{2}}|0_A 0_B\rangle + \frac{1}{\sqrt{2}}|1_A 1_B\rangle$, and the measurement settings

$$\begin{aligned} |X = \mathsf{T}\rangle_{a=1} &:= |R(0)\rangle \text{ and } |X = \mathsf{T}\rangle_{a=2} := |R(\frac{\pi}{6})\rangle \\ |Y = \mathsf{T}\rangle_{b=1} &:= |R(0)\rangle \text{ and } |Y = \mathsf{T}\rangle_{b=2} := |R(\frac{\pi}{3})\rangle. \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table VIII.

(1) *Uniform Settings, Optimized Bell*: For the uniform setting distribution the best classical approximation is given in Table IX, which gives a KL divergence of 0.0177632822.

(2) *Uncorrelated Settings, Optimized Bell*: The optimal, uncorrelated setting distribution is given in Table X. The probabilities of the best classical theory for this distribution are those of Table XI. The corresponding KL divergence is 0.0191506613.

(3) *Correlated Settings, Optimized Bell*: The optimal correlated setting distribution is given in Table XII. The probabilities of the best classical theory for this distribution is given in Table XIII. The corresponding KL divergence is 0.0211293952.

C. CHSH

The bipartite state $\frac{1}{\sqrt{2}}|0_A 0_B\rangle + \frac{1}{\sqrt{2}}|1_A 1_B\rangle$. \mathcal{A} 's and \mathcal{B} 's measurement settings are:

$$\begin{aligned} |X = \mathsf{T}\rangle_{a=1} &:= |R(0)\rangle \text{ and } |X = \mathsf{T}\rangle_{a=2} := |R(\frac{\pi}{4})\rangle, \\ |Y = \mathsf{T}\rangle_{b=1} &:= |R(\frac{\pi}{8})\rangle \text{ and } |Y = \mathsf{T}\rangle_{b=2} := |R(-\frac{\pi}{8})\rangle. \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table XIV (with $\frac{1}{4} + \frac{1}{8}\sqrt{2} \approx 0.4267766953$ and $\frac{1}{4} - \frac{1}{8}\sqrt{2} \approx 0.0732233047$).

Uniform, Uncorrelated and Correlated Settings, CHSH: The optimal measurement settings is the uniform settings, where both \mathcal{A} and \mathcal{B} use one of the two measurements with probability 0.5 (that is $\sigma_{ab} = 0.25$)

The optimal classical theory in this scenario has the probabilities of Table XV.

D. Hardy

The bipartite state $\alpha|0_A 0_B\rangle - \beta|1_A 1_B\rangle$, with $\alpha := \frac{1}{2}\sqrt{2 + 2\sqrt{-13 + 6\sqrt{5}}} \approx 0.907$ and $\beta := \sqrt{1 - \alpha^2} \approx 0.421$ (such that indeed $\alpha^2 + \beta^2 = 1$). \mathcal{A} 's and \mathcal{B} 's measurement settings are now identical and given by:

$$\begin{aligned} |X = \mathsf{T}\rangle_{a=1} = |Y = \mathsf{T}\rangle_{b=1} &:= \sqrt{\frac{\beta}{\alpha+\beta}}|0\rangle + \sqrt{\frac{\alpha}{\alpha+\beta}}|1\rangle, \\ |X = \mathsf{T}\rangle_{a=2} = |Y = \mathsf{T}\rangle_{b=2} &:= -\sqrt{\frac{\beta^3}{\alpha^3+\beta^3}}|0\rangle + \sqrt{\frac{\alpha^3}{\alpha^3+\beta^3}}|1\rangle. \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table XVI.

(1) *Uniform Settings, Hardy*: For uniform measurement settings, the best classical theory to describe the quantum mechanical statistics is given in Table XVII, with KL divergence: 0.0278585182.

(2) *Uncorrelated Settings, Hardy*: The optimized uncorrelated setting distribution is given in Table XVIII. The probabilities of the best classical theory for this distribution are in Table XIX. The corresponding KL divergence is 0.0279816333.

(3) *Correlated Settings, Hardy*: The optimized correlated setting distribution is given in Table XX. The probabilities of the best classical theory for this distribution are described in Table XXI. The corresponding KL divergence is 0.0280347655.

E. Mermin

In [26], we find the following nonlocality proof with two parties, three measurement settings, and two possible outcomes. Let the entangled state be $\frac{1}{\sqrt{2}}(|0_A 0_B\rangle + |1_A 1_B\rangle)$, and the measurement settings:

$$\begin{aligned} |X = \mathsf{T}\rangle_{a=1} = |Y = \mathsf{T}\rangle_{b=1} &:= |0\rangle, \\ |X = \mathsf{T}\rangle_{a=2} = |Y = \mathsf{T}\rangle_{b=2} &:= |R(\frac{2}{3}\pi)\rangle, \\ |X = \mathsf{T}\rangle_{a=3} = |Y = \mathsf{T}\rangle_{b=3} &:= |R(\frac{4}{3}\pi)\rangle. \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table XXII.

(1) *Uniform Settings, Mermin*: The probabilities of the best classical theory for the uniform measurement settings is given in Table XXIII.

(2) *Uncorrelated Settings, Mermin*: The optimal uncorrelated setting distribution is given in Table XXIV. The probabilities of the best classical theory for this distribution is in Table XXV.

(3) *Correlated Settings, Mermin*: The optimal correlated setting distribution is given in Table XXVI (note that there are also other optimal distributions). The probabilities of the best classical theory for this specific distribution are described in Table XXVII. The corresponding KL divergence is 0.0211293952.

F. GHZ

The tripartite state $\frac{1}{\sqrt{2}}|0_A 0_B 0_C\rangle + \frac{1}{\sqrt{2}}|1_A 1_B 1_C\rangle$. The settings for all three parties are identical:

$$\begin{aligned} |X = \mathsf{T}\rangle_{a=1} = |Y = \mathsf{T}\rangle_{b=1} = |Z = \mathsf{T}\rangle_{c=1} &:= \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle, \\ |X = \mathsf{T}\rangle_{a=2} = |Y = \mathsf{T}\rangle_{b=2} = |Z = \mathsf{T}\rangle_{c=2} &:= \frac{1}{\sqrt{2}}|0\rangle + \frac{i}{\sqrt{2}}|1\rangle. \end{aligned}$$

With these settings, quantum mechanics predicts the conditional probabilities of Table XXVIII.

(1) *Uniform and Uncorrelated Settings, GHZ*: For all three settings, the best possible classical statistics that approximate the GHZ experiment is that of Table XXIX. The optimal uncorrelated setting is the uniform settings that samples all eight measurement settings with equal probability. The corresponding KL divergence is: 0.2075187496.

(2) *Correlated Settings, GHZ*: The optimal correlated setting samples with equal probability those four settings that yield the (0.125, 0) outcome probabilities (those are the settings where an even number of the measurements are measuring along the m_1 axis). The KL divergence in this setting is twice that of the previous uniform setting: 0.4150374993.

TABLE VIII
QUANTUM PREDICTIONS OPTIMIZED BELL

$Q_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.5	0	0.375	0.125
	$y=F$	0	0.5	0.125	0.375
$b=2$	$y=T$	0.125	0.375	0.375	0.125
	$y=F$	0.375	0.125	0.125	0.375

TABLE IX
BEST CLASSICAL THEORY FOR UNIFORM SETTINGS OPTIMIZED BELL. KL DISTANCE: 0.0177632822.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.5000000000	0.0000000000	0.3333333333	0.1666666667
	$y=F$	0.0000000000	0.5000000000	0.1666666667	0.3333333333
$b=2$	$y=T$	0.1666666667	0.3333333333	0.3333333333	0.1666666667
	$y=F$	0.3333333333	0.1666666667	0.1666666667	0.3333333333

TABLE X
OPTIMIZED UNCORRELATED SETTING DISTRIBUTION OPTIMIZED BELL

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma^{uc}$	$a=1$	$a=2$	$\Pr(B=b)$
$b=1$	0.1497077788	0.2372131160	0.3869208948
$b=2$	0.2372131160	0.3758659893	0.6130791052
$\Pr(A=a)$	0.3869208948	0.6130791052	

TABLE XI
BEST CLASSICAL THEORY FOR UNCORRELATED SETTINGS OPTIMIZED BELL. KL DISTANCE: 0.0191506613.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.5000000000	0.0000000000	0.3267978563	0.1732021436
	$y=F$	0.0000000000	0.5000000000	0.1732021436	0.3267978563
$b=2$	$y=T$	0.1732021436	0.3267978563	0.3464042873	0.1535957127
	$y=F$	0.3267978563	0.1732021436	0.1535957127	0.3464042873

TABLE XII
OPTIMIZED CORRELATED SETTING DISTRIBUTION OPTIMIZED BELL

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma$	$a=1$	$a=2$
$b=1$	0.1046493146	0.2984502285
$b=2$	0.2984502285	0.2984502285

TABLE XIII
BEST CLASSICAL THEORY FOR CORRELATED SETTINGS OPTIMIZED BELL. KL DISTANCE: 0.0211293952.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.4927305107	0.0072694892	0.3357564964	0.1642435036
	$y=F$	0.0072694892	0.4927305107	0.1642435036	0.3357564964
$b=2$	$y=T$	0.1642435036	0.3357564964	0.3357564964	0.1642435036
	$y=F$	0.3357564964	0.1642435036	0.1642435036	0.3357564964

TABLE XIV
QUANTUM PREDICTIONS CHSH

$Q_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.4267766953	0.0732233047	0.4267766953	0.0732233047
	$y=F$	0.0732233047	0.4267766953	0.0732233047	0.4267766953
$b=2$	$y=T$	0.4267766953	0.0732233047	0.0732233047	0.4267766953
	$y=F$	0.0732233047	0.4267766953	0.4267766953	0.0732233047

TABLE XV
BEST CLASSICAL THEORY FOR UNIFORM SETTINGS CHSH. KL DISTANCE: 0.0462738469.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.375	0.125	0.375	0.125
	$y=F$	0.125	0.375	0.125	0.375
$b=2$	$y=T$	0.375	0.125	0.125	0.375
	$y=F$	0.125	0.375	0.375	0.125

APPENDIX III

THE KULLBACK-LEIBLER DIVERGENCE AND STATISTICAL STRENGTH

This appendix provides in-depth information about the Kullback-Leiber divergence and its relation to statistical strength. The KL divergence was defined in Section IV as

$$D(Q\|P) := \sum_{z \in \mathcal{Z}} Q(z) \log \frac{Q(z)}{P(z)}.$$

We immediately see that the KL divergence expresses something like the average disbelief in P , when observing random outcomes Z from Q . Thus occasionally (with respect to Q) one observes an outcome Z that is more (log-) likely under P than Q , but on average (with respect to Q), the outcomes are more likely under Q than P , expressed by the fact that $D(Q\|P) \geq 0$. Apart from this intuitive meaning, KL divergence has several different concrete interpretations and applications. Here we focus on the interpretation we are concerned with in this paper: KL divergence as a measure of ‘statistical distance’ in the context of statistical hypothesis testing. Below we first (Section III-A) give an intuitive explanation of ‘statistical strength’ and ‘statistical distance’. Although there exist at least three different approaches to measure statistical distance, in Sections III-B, III-C, III-D, we show that for large samples, KL divergence is the appropriate measure according to all three of them. In [11] we provide a more extensive treatment, listing several properties and examples of the KL divergence. There we also explain why any reasonable notion of ‘statistical distance’ must be asymmetric, and, relatedly, why other common distance measures such as absolute deviations between probabilities are *not* well-suited for this purpose.

A. Interpreting ‘Statistical Strength’

Consider the scenario of Section IV-A.1: Z_1, Z_2, \dots are independently generated either by some distribution P or by some distribution Q with $Q \neq P$. We are given a sample (sequence of outcomes) z_1, \dots, z_n . We want to perform a statistical test in order to find out whether the sample is from

P or Q . Suppose that the sample is, in fact, generated by Q (‘ Q is true’). Then it turns out that

- 1) For a fixed sample size n , the larger $D(Q\|P)$, the more support there will be in the sample z_1, \dots, z_n for Q versus P (with high probability under Q).
- 2) For a pre-determined fixed level of support in favor of Q against P (equivalently, level of ‘confidence’ in Q , level of ‘convincingness’ of Q), we have that the larger $D(Q\|P)$, the smaller the sample size before this level of support is achieved (with high probability under Q).
- 3) If, based on observed data z_1, \dots, z_n , an experimenter decides that Q rather than P must have generated the data, then, the larger $D(Q\|P)$, the larger the *confidence* the experimenter should have in this decision (with high probability under Q).

What exactly do we mean by ‘level of support or convincingness’? Different approaches to statistical inference define this notion in a different manner. Nevertheless, *for large samples*, all definitions of support one finds in the literature become essentially equivalent, and are determined by the KL divergence up to lower order terms in the exponent. We consider three methods for statistical hypothesis testing: *frequentist* hypothesis testing [31], *Bayesian* hypothesis testing [24] and *information-theoretic* hypothesis testing [25], [32]. Nearly all state-of-the-art, theoretically motivated statistical methodology falls in either the Bayesian or the frequentist categories. Frequentist hypothesis testing is the most common, the most taught in statistics classes and is the standard method in, for example, the medical sciences. Bayesian hypothesis testing is becoming more and more popular in, for example, econometrics and biological applications. While theoretically important, the information-theoretic methods are less used in practice and are discussed mainly because they lead to a very concrete interpretation of statistical strength in terms of *bits* of information.

We illustrate below that in all three approaches the KL divergence indeed captures the notion of ‘statistical strength’. We consider the general situation with a sample Z_1, Z_2, \dots , with the Z_i independently and identically distributed according

TABLE XVI
QUANTUM PREDICTIONS HARDY

$Q_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0	0.38196601125	0.23606797750	0.14589803375
	$y=F$	0.38196601125	0.23606797750	0	0.61803398875
$b=2$	$y=T$	0.23606797750	0	0.09016994375	0.14589803375
	$y=F$	0.14589803375	0.61803398875	0.14589803375	0.61803398875

TABLE XVII
BEST CLASSICAL THEORY FOR UNIFORM SETTINGS HARDY. KL DISTANCE: 0.0278585182.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.0338829434	0.3543640363	0.2190090188	0.1692379609
	$y=F$	0.3543640363	0.2573889840	0.0075052045	0.6042478158
$b=2$	$y=T$	0.2190090188	0.0075052045	0.0488933524	0.1776208709
	$y=F$	0.1692379609	0.6042478158	0.1776208709	0.5958649058

TABLE XVIII
OPTIMIZED UNCORRELATED SETTING DISTRIBUTION HARDY

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma^{\text{UC}}$	$a=1$	$a=2$	$\Pr(B=b)$
$b=1$	0.2603092699	0.2498958554	0.5102051253
$b=2$	0.2498958554	0.2398990193	0.4897948747
$\Pr(A=a)$	0.5102051253	0.4897948747	

TABLE XIX
BEST CLASSICAL THEORY FOR UNCORRELATED SETTINGS HARDY. KL DISTANCE: 0.0279816333.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.0198831449	0.3612213769	0.2143180373	0.1667864844
	$y=F$	0.3612213769	0.2576741013	0.0141212511	0.6047742271
$b=2$	$y=T$	0.2143180373	0.0141212511	0.0481256471	0.1803136414
	$y=F$	0.1667864844	0.6047742271	0.1803136414	0.5912470702

TABLE XX
OPTIMIZED CORRELATED SETTING DISTRIBUTION HARDY

$\Pr(A=a, B=b) = \sigma_{ab} \in \Sigma$	$a=1$	$a=2$
$b=1$	0.2562288294	0.2431695652
$b=2$	0.2431695652	0.2574320402

TABLE XXI
BEST CLASSICAL THEORY FOR CORRELATED SETTINGS HARDY. KL DISTANCE: 0.0280347655.

$P_{ab}(X=x, Y=y)$		$a=1$		$a=2$	
		$x=T$	$x=F$	$x=T$	$x=F$
$b=1$	$y=T$	0.0173443545	0.3620376608	0.2123471649	0.1670348504
	$y=F$	0.3620376608	0.2585803238	0.0165954828	0.6040225019
$b=2$	$y=T$	0.2123471649	0.0165954828	0.0505353201	0.1784073276
	$y=F$	0.1670348504	0.6040225019	0.1784073276	0.5926500247

TABLE XXVIII
QUANTUM PREDICTIONS GHZ

$Q_{abc}(X=x, Y=y, Z=z)$			$a=1$		$a=2$	
			$x=T$	$x=F$	$x=T$	$x=F$
$c=1$	$b=1$	$z=T$ $y=T$	0.25	0	0.125	0.125
		$z=T$ $y=F$	0	0.25	0.125	0.125
		$z=F$ $y=T$	0	0.25	0.125	0.125
	$b=2$	$z=T$ $y=T$	0.125	0.125	0	0.25
		$z=T$ $y=F$	0.125	0.125	0.25	0
		$z=F$ $y=T$	0.125	0.125	0.25	0
$c=2$	$b=1$	$z=T$ $y=T$	0.125	0.125	0	0.25
		$z=T$ $y=F$	0.125	0.125	0.25	0
		$z=F$ $y=T$	0.125	0.125	0.25	0
	$b=2$	$z=T$ $y=T$	0	0.25	0.125	0.125
		$z=T$ $y=F$	0.25	0	0.125	0.125
		$z=F$ $y=T$	0.25	0	0.125	0.125
		$z=F$ $y=F$	0	0.25	0.125	0.125

TABLE XXIX
BEST CLASSICAL THEORY FOR UNIFORM SETTINGS GHZ. KL DISTANCE: 0.2075187496.

$P_{abc}(X=x, Y=y, Z=z)$			$a=1$		$a=2$	
			$x=T$	$x=F$	$x=T$	$x=F$
$c=1$	$b=1$	$z=T$ $y=T$	0.1875	0.0625	0.125	0.125
		$z=T$ $y=F$	0.0625	0.1875	0.125	0.125
		$z=F$ $y=T$	0.0625	0.1875	0.125	0.125
	$b=2$	$z=T$ $y=T$	0.125	0.125	0.0625	0.1875
		$z=T$ $y=F$	0.125	0.125	0.1875	0.0625
		$z=F$ $y=T$	0.125	0.125	0.1875	0.0625
$c=2$	$b=1$	$z=T$ $y=T$	0.125	0.125	0.0625	0.1875
		$z=T$ $y=F$	0.125	0.125	0.1875	0.0625
		$z=F$ $y=T$	0.125	0.125	0.1875	0.0625
	$b=2$	$z=T$ $y=T$	0.0625	0.1875	0.125	0.125
		$z=T$ $y=F$	0.1875	0.0625	0.125	0.125
		$z=F$ $y=T$	0.1875	0.0625	0.125	0.125
		$z=F$ $y=F$	0.0625	0.1875	0.125	0.125

to some Q_σ , Q_σ being some distribution over some finite set \mathcal{Z} . For each n , the first n outcomes are distributed according to the n -fold product distribution of Q_σ , which we shall also refer to as Q_σ . Hence $Q_\sigma(z_1, \dots, z_n) = \prod_{i=1}^n Q_\sigma(z_i)$. The independence assumption also induces a distribution over the set \mathcal{Z}^∞ of all infinite sequences³ which we shall *also* refer to as Q_σ .

We test Q_σ against a set of distributions \mathcal{P}_σ . Thus, Q_σ and \mathcal{P}_σ may, but do not necessarily refer to quantum and local realist theories — the statements below hold more generally.

B. Frequentist Justification

In frequentist hypothesis testing, \mathcal{P}_σ is called the *null-hypothesis* and Q_σ the alternative hypothesis. Frequentist hypothesis testing can be implemented in a number of different ways, depending on what *statistical test* one adopts. A statistical test is a procedure that, when input an arbitrary sample of arbitrary length, outputs a *decision*. The decision is either ‘ Q_σ generated the data’ or ‘ \mathcal{P}_σ generated the data’. Each test

is defined relative to some *test statistic* T and *critical value* c . A test statistic T is a function defined on samples of arbitrary length, that for each sample outputs a real number. Intuitively, large values of the test statistic indicate that something has happened which is much more unlikely under any of the distributions in the null-hypothesis than under the alternative hypothesis. A function that is often used as a test statistic is the *likelihood ratio*

$$T(z_1, \dots, z_n) := \frac{Q_\sigma(z_1, \dots, z_n)}{\sup_{P \in \mathcal{P}_\sigma} P(z_1, \dots, z_n)}, \quad (13)$$

but many other choices are possible as well.

The critical value c determines the threshold for the test’s decision: if, for the observed data z_1, \dots, z_n , it holds that $T(z_1, \dots, z_n) \geq c$, the test says ‘ Q_σ generates the data’; if $T(z_1, \dots, z_n) < c$, the test says ‘ \mathcal{P}_σ generated the data’.

The confidence in a given decision is determined by a quantity known as *the p-value*. This is a function of the data that was actually observed in the statistical experiment. It only depends on the observed value of the test statistic

³Readers familiar with measure theory should note that throughout this paper, we tacitly assume that \mathcal{Z}^∞ is endowed with a suitable σ -algebra such that all sets mentioned in this paper become measurable.

$t_{\text{observed}} := T(z_1, \dots, z_n)$. It is defined as

$$p\text{-value} := \sup_{P \in \mathcal{P}_\sigma} P(T(Z_1, \dots, Z_n) \geq t_{\text{observed}}). \quad (14)$$

Here the Z_1, \dots, Z_n are distributed according to P and thus do not refer to the data that was actually observed in the experiment. Thus, the p -value is the maximum probability, under any distribution in \mathcal{P}_σ , that the test statistic takes on a value that is at least as extreme as its actually observed outcome. Typically, the test is defined such that the critical value c depends on sample size n . It is set to the value c_0 such that the test outputs ‘ Q_σ ’ iff the p -value is smaller than some pre-defined *significance level*, typically 0.05.

Large p -values mean small confidence: for example, suppose the test outputs Q_σ whenever the p -value is smaller than 0.05. Suppose further that data are observed with a p -value of 0.04. Then the test says “ Q_σ ” but since the p -value is large, this is not that convincing to someone who considers the possibility that some $P \in \mathcal{P}_\sigma$ has generated the data: the large p -value indicates that the test may very well have given the wrong answer. On the other hand, if data are observed with a p -value of 0.001, this gives a lot more confidence in the decision output by the test.

We call a test statistic asymptotically optimal for identifying Q_σ if, under the assumption that Q_σ generated the data, the p -value goes to 0 at the fastest possible rate. Now let us assume that Q_σ generates the data, and an optimal test is used. A well-known result due to Bahadur [2, Theorem 1] says that, under some regularity conditions on Q_σ and \mathcal{P}_σ , with Q_σ -probability 1, for all large n ,

$$p\text{-value} = 2^{-nD(Q_\sigma \| \mathcal{P}_\sigma) + o(n)}. \quad (15)$$

where $\lim_{n \rightarrow \infty} o(n)/n = 0$. We say ‘the p -value is determined, to first order in the exponent, by $D(Q_\sigma \| \mathcal{P}_\sigma)$ ’. Note that what we called the ‘actually observed test statistic t_{observed} ’ in (14) has become a random variable in (15), distributed according to Q_σ . It turns out that the regularity conditions, needed for Equation (15) to hold, apply when Q_σ is instantiated to a quantum theory Q with measurement setting distributions σ , and \mathcal{P}_σ is instantiated to the corresponding set of LR theories as defined in Section II.

Now imagine that QM, who knows that Q_σ generates the data, wonders whether to use the experimental setup corresponding to σ_1 or σ_2 . Suppose that $D(Q_{\sigma_1} \| \mathcal{P}_{\sigma_1}) > D(Q_{\sigma_1} \| \mathcal{P}_{\sigma_2})$. It follows from Equation (15) that if the experiment corresponding to σ_1 is performed, the p -value will go to 0 exponentially faster (in the number of trials) than if the experiment corresponding to σ_2 is performed. It therefore makes sense to say that ‘the statistical strength of the experiment corresponding to σ_1 is larger than the strength of σ_2 ’. This provides a frequentist justification of adopting $D(Q_\sigma \| \mathcal{P}_\sigma)$ as an indicator of statistical strength.

Remark: Bahadur [2, Theorem 2] also provides a variation of Equation (15), which (roughly speaking) says the following: suppose Q_σ generates the data. For $\varepsilon > 0$, let N_ε be the *minimum number of observations* such that, for all $n \geq N_\varepsilon$, the test rejects \mathcal{P}_σ (if \mathcal{P}_σ is not rejected for infinitely many n , then N_ε is defined to be infinite). Suppose that an

optimal (in the sense we used previously) test is used. Then, for small ε , N_ε is inversely proportional to $D(Q_\sigma \| \mathcal{P}_\sigma)$: with Q_σ -probability 1, the smaller $D(Q_\sigma \| \mathcal{P}_\sigma)$, the larger N_ε . If a ‘non-optimal’ test is used, then N_ε can only be larger, never smaller.

The rate at which the p -value of a test converges to 0 is known in statistics as *Bahadur efficiency*. For an overview of the area, see [18]. For an easy introduction to the main ideas, focusing on ‘Stein’s lemma’ (a theorem related to Bahadur’s), see [4, Chapter 12, Section 8]. For an introduction to Stein’s lemma with a physicist audience in mind, see [3].

C. Bayesian Justification

In the Bayesian approach to hypothesis testing [7], [24], when testing Q_σ against \mathcal{P}_σ , we must first determine an *a priori probability distribution* over Q_σ and \mathcal{P}_σ . This distribution over distributions is usually just called ‘the prior’. It can be interpreted as indicating the prior (i.e., before seeing the data) ‘degree of belief’ in Q_σ vs. \mathcal{P}_σ . It is often used to incorporate prior knowledge into the statistical decision process. In order to set up the test as fairly as possible, QM and LR may agree to use the prior $\Pr(Q_\sigma) = \Pr(\mathcal{P}_\sigma) = 1/2$ (this should be read as ‘the prior probability that Q_σ obtains is equal to the prior probability that some $P \in \mathcal{P}_\sigma$ obtains’). Yet as long as $\Pr(Q_\sigma) > 0$ and there is a smooth and positive probability density for $P_\sigma \in \mathcal{P}_\sigma$, the specific values for the priors will be irrelevant for the result below.

For given prior probabilities and a given sample z_1, \dots, z_n , Bayesian statistics provides a method to compute the *posterior probabilities* of the two hypotheses, *conditioned* on the observed data: $\Pr(Q_\sigma)$ is transformed into $\Pr(Q_\sigma | z_1, \dots, z_n)$. Similarly, $\Pr(\mathcal{P}_\sigma)$ is transformed to $\Pr(\mathcal{P}_\sigma | z_1, \dots, z_n)$. One then adopts the hypothesis $H \in \{Q_\sigma, \mathcal{P}_\sigma\}$ with the larger posterior probability $\Pr(H | z_1, \dots, z_n)$. The confidence in this decision is given by the *posterior odds* of Q_σ against \mathcal{P}_σ , defined, for given sample z_1, \dots, z_n , as

$$\text{post-odds}(Q_\sigma, \mathcal{P}_\sigma) := \frac{\Pr(Q_\sigma | z_1, \dots, z_n)}{\Pr(\mathcal{P}_\sigma | z_1, \dots, z_n)}.$$

The larger *post-odds*, the larger the confidence. Now suppose that data are distributed according to Q_σ . It can be shown that, under some regularity conditions on Q_σ and \mathcal{P}_σ , with Q_σ -probability 1,

$$\text{post-odds} = 2^{nD(Q_\sigma \| \mathcal{P}_\sigma) + O(\log n)}, \quad (16)$$

In our previously introduced terminology, ‘the Bayesian confidence (posterior odds) is determined by $D(Q_\sigma \| \mathcal{P}_\sigma)$, up to first order in the exponent’. We may now reason exactly as in the frequentist case to conclude that it makes sense to adopt $D(Q_\sigma \| \mathcal{P}_\sigma)$ as an indicator of statistical strength, and that it makes sense for QM to choose the setting probabilities σ so as to maximize $D(Q_\sigma \| \mathcal{P}_\sigma)$.

Equation (16) is a ‘folklore result’ which ‘usually’ holds. In [11], we show that it does indeed hold with Q_σ and \mathcal{P}_σ defined as nonlocality proofs and local realist theories, respectively.

D. Information-Theoretic Justification

There exist several approaches to information-theoretic or compression-based hypothesis testing; see, for example, [4], [25]. The most influential of these is the so-called *Minimum Description Length Principle* [32]. The basic idea is always that the more one can compress a given sequence of data, the more regularity one has extracted from the data, and thus, the better one has captured the ‘underlying regularities in the data’. Thus, the hypothesis that allows for the maximum compression of the data should be adopted.

Let us first consider testing a simple hypothesis Q against another simple hypothesis P . Two basic facts of coding theory say that

- 1) There exists a uniquely decodable code with lengths L_Q that satisfy, for all $z_1, \dots, z_n \in \mathcal{Z}^n$,

$$L_Q(z_1, \dots, z_n) = \lceil -\log Q(z_1, \dots, z_n) \rceil.$$

The code with lengths L_Q is called the *Shannon-Fano code*, and its existence follows from the so-called *Kraft Inequality*, [10].

- 2) If data Z_1, \dots, Z_n are independently identically distributed $\sim Q$, then among all uniquely decodable codes, the code with length function L_Q has the shortest expected code-length. That is, let L be the length function of any uniquely decodable code over n outcomes, then

$$\mathbf{E}_Q[L(Z_1, \dots, Z_n)] \geq \mathbf{E}_Q[-\log Q(Z_1, \dots, Z_n)].$$

Thus, under the assumption that Q generated the data, the optimal (maximally compressing) code to use will be the Shannon-Fano code with lengths $-\log Q(Z^n)$ (here, as in the remainder of this section, we ignored the integer requirement for code lengths). Similarly, under the assumption that some P with $P \neq Q$ generated the data the optimal code will be the code with lengths $-\log P(Z^n)$. Thus, from the information-theoretic point of view, if one wants to find out whether P or Q better explains the data, one should check whether the optimal code under P or the optimal code under Q allows for more compression of the data. That is, one should look at the difference

$$\text{bit-diff} := -\log P(z_1, \dots, z_n) - [-\log Q(z_1, \dots, z_n)]. \quad (17)$$

If $\text{bit-diff} > 0$, then one decides that Q better explains the data. The confidence in this decision is given by the magnitude of bit-diff : the larger bit-diff , the more extra bits one needs to encode the data under P rather than Q , thus the larger the confidence in Q .

Now suppose that Q actually generates the data. The *expected code length difference*, measured in bits, between coding the data using the optimal code for Q and coding using the optimal code for P , is given by $\mathbf{E}_Q[-\log P(Z^n) - [-\log Q(Z^n)]] = nD(Q\|P)$. Thus, the KL divergence can be interpreted as *the expected additional number of bits needed to encode outcomes generated by Q , if outcomes are encoded using a code that is optimal for P rather than for Q* . Thus, the natural ‘unit’ of $D(\cdot\|\cdot)$ is the ‘bit’, and $D(Q\|P)$ may be viewed as ‘average amount of information about Z that is lost if Z is wrongfully regarded as being distributed by Q rather

than P ’. By the law of large numbers, Equation (17) implies that, with Q -probability 1, as $n \rightarrow \infty$,

$$\frac{1}{n}(\text{bit-diff}) \rightarrow D(Q\|P). \quad (18)$$

Thus, if Q generates the data, then the information-theoretic confidence *bit-diff* in decision “ Q explains the data better than P ” is, up to first order, determined by the KL divergence between Q and P : the larger $D(Q\|P)$, the larger the confidence. This gives an information-theoretic justification of the use of the KL divergence as an indicator of statistical strength for simple hypothesis testing. We now turn to composite hypothesis testing.

Composite Hypothesis Testing: If one compares Q_σ against a set of hypotheses \mathcal{P}_σ , then one has to associate \mathcal{P}_σ with a code that is ‘optimal under the assumption that some $P \in \mathcal{P}_\sigma$ generated the data’. It turns out that there exist codes with lengths $L_{\mathcal{P}_\sigma}$ satisfying, for all $z_1, \dots, z_n \in \mathcal{Z}^n$,

$$L_{\mathcal{P}_\sigma}(z_1, \dots, z_n) \leq \inf_{P \in \mathcal{P}_\sigma} -\log P(z_1, \dots, z_n) + O(\log n).$$

An example of such a code is given in [11]. The code $L_{\mathcal{P}_\sigma}$ is optimal, up to logarithmic terms, for whatever distribution $P \in \mathcal{P}_\sigma$ that might actually generate data. The information theoretic approach to hypothesis testing now tells us that, to test Q_σ against \mathcal{P}_σ , we should compute the difference in code lengths

$$\text{bit-diff} := L_{\mathcal{P}_\sigma}(z_1, \dots, z_n) - [-\log Q_\sigma(z_1, \dots, z_n)].$$

The larger this difference, the larger the confidence that Q_σ rather than \mathcal{P}_σ generated the data. The article [11] shows that, in analogy to Equation (18), as $n \rightarrow \infty$,

$$\frac{1}{n}(\text{bit-diff}) \rightarrow D(Q_\sigma\|\mathcal{P}_\sigma) \quad (19)$$

Thus, up to sublinear terms, the information-theoretic confidence in Q_σ is given by $nD(Q_\sigma\|\mathcal{P}_\sigma)$. This provides an information-theoretic justification of adopting $D(Q_\sigma\|\mathcal{P}_\sigma)$ as an indicator of statistical strength.

APPENDIX IV

PROOFS OF THEOREMS 1, 2 AND 3

A. Preparation

The proof of Theorem 1 uses the following lemma, which is of some independent interest.

Lemma 1: Let (Σ, Π, U) be the game corresponding to an arbitrary 2 party, 2 measurement settings per party nonlocality proof. For any $(a_0, b_0) \in \{1, 2\}^2$, there exists a $\pi \in \Pi$ such that for all $(a, b) \in \{1, 2\}^2 \setminus \{(a_0, b_0)\}$ we have $Q_{ab} = P_{ab;\pi}$. Thus, for any three of the four measurement settings, the probability distribution on outcomes can be perfectly explained by a local realist theory.

Proof: We give a detailed proof for the case that the measurement outcomes are two values $\{T, F\}$; the general case can be proved in a similar way.

Without loss of generality let $(a_0, b_0) = (2, 2)$. Now we must prove that the equation $Q_{ab} = P_{ab;\pi}$ holds for the three settings $(a, b) \in \{(1, 1), (1, 2), (2, 1)\}$. Every triple of distributions $P_{ab;\pi}$

for these three settings may be represented by a table of the form:

Pr	$X_1 = T$	$X_1 = F$	$X_2 = T$	$X_2 = F$
$Y_1 = T$	p_1	p_2	p_5	p_6
$Y_1 = F$	p_3	p_4	p_7	p_8
$Y_2 = T$	p_9	p_{10}		
$Y_2 = F$	p_{11}	p_{12}		

with $p_1, \dots, p_{12} \geq 0$ and the normalization restrictions $p_1 + \dots + p_4 = p_5 + \dots + p_8 = p_9 + \dots + p_{12} = 1$. Given any table of this form, we say that the LR distribution P_π corresponds to the p -table if $P_{11;\pi}(T, T) = p_1$, $P_{12;\pi}(F, T) = p_{10}$ etc., for all p_i .

The no-signaling restriction implies that the realized measurement setting on \mathcal{A} 's side should not influence the probability on \mathcal{B} 's side and vice versa. Hence, for example, $\Pr_{00}(Y = T) = \Pr_{10}(Y = T)$, which gives $p_1 + p_2 = p_5 + p_6$. In total there are four such no-signaling restrictions:

$$\begin{cases} p_1 + p_2 = p_5 + p_6 \\ p_3 + p_4 = p_7 + p_8 \\ p_1 + p_3 = p_9 + p_{11} \\ p_2 + p_4 = p_{10} + p_{12}. \end{cases}$$

We call a table with $p_1, \dots, p_{12} \geq 0$, that obeys the normalization restriction on the sub-tables and that satisfies Equations (20) a Γ -table. We already showed that each triple of conditional LR distributions may be represented as a Γ -table. In exactly the same way one shows that each triple of conditional quantum experimentalist distributions Q_{00}, Q_{01}, Q_{10} can be represented as a Γ -table. It therefore suffices if we can show that every Γ -table corresponds to some LR theory P_π . We show this by considering the 16 possible deterministic theories T_{x_1, x_2, y_1, y_2} . Here T_{x_1, x_2, y_1, y_2} is defined as the theory with $P_\pi(X_1 = x_1, X_2 = x_2, Y_1 = y_1, Y_2 = y_2) = \pi_{x_1, x_2, y_1, y_2} = 1$. Each deterministic theory T_{x_1, x_2, y_1, y_2} corresponds to a specific Γ -table denoted by $\Gamma_{x_1, x_2, y_1, y_2}$. For example, the theory T_{FFTF} gives the following Γ_{FFTF} -table:

Pr	$X_1 = T$	$X_1 = F$	$X_2 = T$	$X_2 = F$
$Y_1 = T$	0	1	0	1
$Y_1 = F$	0	0	0	1
$Y_2 = T$	0	0		
$Y_2 = F$	0	1		

We will prove that the set of Γ -tables is in fact the convex hull of the 16 tables $\Gamma_{x_1, x_2, y_1, y_2}$ corresponding to deterministic theories. This shows that any Γ -table can be reproduced by a mixture of deterministic theories. Since every LR theory $\pi \in \Pi$ can be written as such a mixture, this proves the lemma.

First we observe that a Γ -table with all entries 0 or 1 has to be one of the 16 deterministic theories. Given a Γ -table that is not a deterministic theory, we focus on its smallest nonzero entry $\Gamma_{ab} = \varepsilon > 0$. By the restrictions of Equations (20) there exists a deterministic theory T_k such that the table $(\Gamma - \varepsilon \Gamma_k)/(1 - \varepsilon)$ has no negative entries. For example, suppose that the smallest element in Γ corresponds to $P_\pi(X_1 = F, Y_1 = T)$ (denoted as p_2 in the first table above). By the restrictions of Equation (20), either the table $(\Gamma - p_2 \Gamma_{FFTF})/(1 - p_2)$ (where Γ_{FFTF} is shown above) or one of the

three tables $(\Gamma - p_2 \Gamma_{FFTT})/(1 - p_2)$, $(\Gamma - p_2 \Gamma_{FTTF})/(1 - p_2)$, $(\Gamma - p_2 \Gamma_{FTTT})/(1 - p_2)$ has only nonnegative entries.

Let $\Gamma' := (\Gamma - \varepsilon \Gamma_k)/(1 - \varepsilon)$ where k is chosen such that Γ' has no negative entries. Clearly, either Γ' describes a deterministic theory with entries 0 and 1, or Γ' is a Γ -table with number of nonzero entries one less than that of Γ . Hence by applying the above procedure at most 16 times, we obtain a decomposition $\Gamma = \varepsilon_1 \Gamma_{k_1} + \dots + \varepsilon_{16} \Gamma_{k_{16}}$, which shows that Γ lies in the convex hull of the Γ -tables corresponding to deterministic theories. Hence, any such Γ can be described as a LR theory.

For measurement settings with more than two outcomes, the proof can be generalized in a straightforward manner. \square

B. Proof of Theorem 1

We only give proofs for the $2 \times 2 \times 2$ case; extension to the general case is entirely straightforward. We define

$$\begin{aligned} U((a, b), \pi) & \\ &:= D(Q_{ab}(\cdot) \| P_{ab;\pi}(\cdot)) \\ &= \sum_{\substack{x, y \in \{T, F\} \\ Q_{ab}(x, y) > 0}} Q_{ab}(x, y) [\log Q_{ab}(x, y) - \log P_{ab;\pi}(x, y)]. \end{aligned} \quad (20)$$

Note that $U(\sigma, \pi)$ can be written as

$$U(\sigma, \pi) = \sum_{a, b \in \{1, 2\}} \sigma_{ab} U((a, b), \pi)$$

Part 1: Equation (3) follows directly from the additivity property of KL divergence, see [11, Appendix IV-B] or [10]. Convexity is immediate by Jensen's inequality applied to the logarithm in Equation (20) and the fact that $P_{ab;\pi}(x, y)$ is linear in π_{x_1, x_2, y_1, y_2} for each $(x_1, x_2, y_1, y_2) \in \{T, F\}^4$. If π lies in the interior of Π , then $P_{ab;\pi}(x, y) > 0$ for $a, b \in \{1, 2\}$ so that $U(\sigma, \pi)$ is finite. Continuity and differentiability are then immediate by continuity and differentiability of $\log x$ for $x > 0$. Lower semicontinuity of $U(\sigma, \pi)$ on Π is implied by the fact that, on general spaces, $D(Q \| P)$ is jointly lower semicontinuous in Q and P in the weak topology, as proved by Posner [30, Theorem 2]. Part 1(c) is immediate.

Part 2: We have already shown that for fixed σ , $U(\sigma, \pi)$ is lower semicontinuous on Π . Lower semicontinuous functions achieve their infimum on a compact domain (see for example [14, page 84]), so that for each σ , Equation (4) is achieved for some π^* . This proves (a). To prove (b), note that nonnegativity of $U(\sigma)$ is immediate by nonnegativity of the KL divergence. Boundedness of $U(\sigma)$ follows by considering the uniform distribution π° , with, for all x_1, x_2, y_1, y_2 , $\pi_{x_1, x_2, y_1, y_2}^\circ = 1/16$. π° is in Π , so that

$$\begin{aligned} U(\sigma) &\leq U(\sigma, \pi^\circ) \\ &= \sum_{a, b \in \{1, 2\}} \sigma_{ab} \left(\sum_{\substack{x, y \in \{T, F\} \\ Q_{ab}(x, y) > 0}} Q_{ab}(x, y) [\log Q_{ab}(x, y) + 2] \right) \\ &\leq - \sum_{a, b \in \{1, 2\}} \sigma_{ab} H(Q_{ab}) + 8, \end{aligned}$$

where $H(Q_{ab})$ is the Shannon-entropy of the distribution Q_{ab} . Boundedness of $U(\sigma)$ now follows from the fact that $H(Q) \geq$

0 for every distribution Q , which is a standard result (see, e.g. [10]).

Let σ be in the interior of Σ and let $\pi^* \in \Pi$ achieve $\inf_{\Pi} U(\sigma, \pi)$. Since $U(\sigma)$ is bounded, Q is absolutely continuous with respect to π^* (otherwise $U(\sigma) = \infty$, a contradiction). Thus, $U(\sigma)$ satisfies

$$U(\sigma) = \inf_{\pi: Q \ll \pi} U(\sigma, \pi), \quad (21)$$

where $Q \ll \pi$ means that Q is absolutely continuous with respect to $\pi \in \Pi$. We already proved that if Q is absolutely continuous with respect to π^* , then $U(\sigma, \pi^*)$ is linear in σ . Thus, by Equation (21), $U(\sigma)$ is an infimum of linear functions, which (by a standard result of convex analysis, see e.g. [33]) is concave. A concave and bounded function with a convex domain must be continuous on the interior of this domain (see, e.g., [33]). It remains to show that $U(\sigma)$ is continuous at boundary points of Σ . Showing this is straightforward by taking limits (but tedious). We omit the details.

Now for part (c). If Q is not a proper nonlocality proof, then by definition there exists a $\pi_0 \in \Pi$ such that, for $a, b \in \{1, 2\}$, we have $Q_{ab} = P_{ab; \pi_0}$ and hence $U(\sigma, \pi_0) = 0$ for all $\sigma \in \Sigma$.

Now suppose Q is a proper nonlocality proof. Let σ be in the interior of Σ . $\inf_{\Pi} U(\sigma, \pi)$ is achieved for some π^* . Suppose, by means of contradiction, that $U(\sigma, \pi^*) = 0$. Since $\sigma_{ab} > 0$ for $a, b \in \{1, 2\}$, we must have $Q_{ab} = P_{ab; \pi^*}$ for $a, b \in \{1, 2\}$. But then Q is not a proper nonlocality proof; contradiction. For part (d), if σ is on the boundary of Σ , then for some a, b , $\sigma_{ab} = 0$. It then follows from Lemma 1 and the fact that, for all P , $D(P||P) = 0$ that $U(\sigma, \pi^*) = 0$.

Part 3: Part (a) The condition that Q_{ab} is not a point mass for some a, b , implies that all π^* that achieve the infimum must have $\pi_{x_1, x_2, y_1, y_2}^* < 1$ for all x_1, x_2, y_1, y_2 , (otherwise $U(\sigma, \pi^*) = \infty$, which is a contradiction). Thus, we assume that $\pi^* \in \Pi_0$, with Π_0 the set of π s that satisfy this “ < 1 ” restriction.

For $\rho \in [0, \infty)^{16}$, let

$$\bar{\rho}_{x_1, x_2, y_1, y_2}(\rho) := \frac{\rho_{x_1, x_2, y_1, y_2}}{\sum_{x'_1, x'_2, y'_1, y'_2 \in \{T, F\}} \rho_{x'_1, x'_2, y'_1, y'_2}}.$$

In this way, each vector ρ with at least one non-zero component uniquely defines a local theory $\bar{\rho} \in \Pi$, and

$$\left\{ \bar{\rho} : \rho \in [0, \infty)^{16} \text{ and } \sum_{x_1, x_2, y_1, y_2 \in \{T, F\}} \rho_{x_1, x_2, y_1, y_2} > 0 \right\} = \Pi_0.$$

Let ρ^* be such that $\bar{\rho}^*$ achieves the infimum in Equation (4). Then Q is absolutely continuous with respect to $\bar{\rho}^*$. One can now show that for each $(x_1, x_2, y_1, y_2) \in \{T, F\}^4$, the partial derivative $\partial U(\sigma, \bar{\rho}) / \partial \rho_{x_1, x_2, y_1, y_2}$ evaluated at $\rho = \rho^*$ exists (even if $\rho_{x_1, x_2, y_1, y_2}^* = 0$). Since $\bar{\rho}^*$ achieves the infimum, it follows that, for each $(x_1, x_2, y_1, y_2) \in \{T, F\}^4$, we must have that $(\partial / \partial \rho_{x_1, x_2, y_1, y_2}) U(\sigma, \bar{\rho})$ evaluated at ρ^* is not less than 0, or, equivalently,

$$\left\{ \frac{\partial U(\sigma, \bar{\rho})}{\partial \rho_{x_1, x_2, y_1, y_2}} \right\}_{\rho = \rho^*} \cdot \left(\sum_{x_1, x_2, y_1, y_2} \rho_{x_1, x_2, y_1, y_2} \right) \geq 0 \quad (22)$$

with equality if $\rho_{x_1, x_2, y_1, y_2}^* > 0$. Straightforward evaluation of Equation (22) gives Equations (5) and (6). This shows that

each π^* achieving Equation (4) satisfies Equations (5) and (6). On the other hand, each π^* corresponding to a ρ^* with $\bar{\rho}^* = \pi^*$ such that Equation (22) holds for each $(x_1, x_2, y_1, y_2) \in \{T, F\}^4$ must achieve a local minimum of $U(\sigma, \pi)$ (viewed as a function of π). Since $U(\sigma, \pi)$ is convex, π^* must achieve the infimum of Equation (4).

For part (b), suppose, by way of contradiction, that for at least one $(x_1, y_1) \in \{T, F\}^2$, $a, b_0 \in \{1, 2\}$ with $Q_{a_0 b_0}(x_1, y_1) > 0$, we have $P_{a_0 b_0; \pi^*}(x_1, y_1) \neq P_{a_0 b_0; \pi^\circ}(x_1, y_1)$. For each $x, y \in \{T, F\}$, $a, b \in \{1, 2\}$, we can write

$$\begin{aligned} P_{ab; \pi^*}(x, y) &= \pi_{k_1}^* + \pi_{k_2}^* + \pi_{k_3}^* + \pi_{k_4}^*, \\ P_{ab; \pi^\circ}(x, y) &= \pi_{k_1}^\circ + \pi_{k_2}^\circ + \pi_{k_3}^\circ + \pi_{k_4}^\circ, \end{aligned} \quad (23)$$

for some k_1, \dots, k_4 depending on x, y, a, b . Here each k_j is of the form $x_1 x_2 y_1 y_2$ with $x_i, y_i \in \{T, F\}$. Now consider $\pi^+ := (\pi^* + \pi^\circ) / 2$. Clearly $\pi^+ \in \Pi$. By Jensen's inequality applied to the logarithm and using Equation (23), we have for $a, b \in \{1, 2\}$: $Q_{ab}(x, y) [\log Q_{ab}(x, y) - \log P_{ab; \pi^+}(x, y)] \leq Q_{ab}(x, y) [\log Q_{ab}(x, y) - \frac{1}{2} \log P_{ab; \pi^*}(x, y) - \frac{1}{2} \log P_{ab; \pi^\circ}(x, y)]$, where the inequality is strict if $x = x_1, y = y_1, a = a_0$ and $b = b_0$. But then for $a, b \in \{1, 2\}$: $U((a, b), \pi^+) \leq \frac{1}{2} U((a, b), \pi^*) + \frac{1}{2} U((a, b), \pi^\circ)$, which for $(a, b) = (a_0, b_0)$ must be strict. By assumption, $\sigma_{a_0 b_0} > 0$. But that implies $U(\sigma, \pi^+) < U(\sigma, \pi^*) = \inf_{\Pi} U(\sigma, \pi)$ and we have arrived at the desired contradiction.

C. Proofs of Game-Theoretic Theorems

1) Game-Theoretic Preliminaries: Proposition 1 gives a few standard game-theoretic results (partially copied from [14]). We will use these results at several stages in later proofs.

Proposition 1: Let A and B be arbitrary sets and let $L : A \times B \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ be an arbitrary function on $A \times B$. We have

- 1) $\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) \geq \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta)$.
- 2) Suppose the following conditions hold:
 - a) The game (A, B, L) has a value $V \in \mathbb{R} \cup \{-\infty, \infty\}$, that is

$$\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) = V = \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta).$$

- b) There is α^* that achieves $\sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta)$.
- c) There is β^* that achieves $\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta)$.

Then (α^*, β^*) is a saddle point and $L(\alpha^*, \beta^*) = V$.

- 3) Suppose there exists a pair (α^*, β^*) such that
 - a) β^* achieves $\inf_{\beta \in B} L(\alpha^*, \beta)$ and
 - b) β^* is an *equalizer strategy*, that is, there exists a $K \in \mathbb{R} \cup \{-\infty, \infty\}$ with for all $\alpha \in A$, $L(\alpha, \beta^*) = K$.

Then the game (A, B, L) has value K , i.e.

$$\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) = \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta) = K,$$

and (α^*, β^*) is a saddle point.

Proof: (1) For all $\alpha' \in A$,

$$\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) \geq \inf_{\beta \in B} L(\alpha', \beta).$$

Therefore, $\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) \geq \sup_{\alpha' \in A} \inf_{\beta \in B} L(\alpha', \beta)$.

(2) Under our assumptions,

$$\begin{aligned}
L(\alpha^*, \beta^*) &\leq \sup_{\alpha \in A} L(\alpha, \beta^*) \\
&= \inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) \\
&= V \\
&= \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta) \\
&= \inf_{\beta \in B} L(\alpha^*, \beta) \leq L(\alpha^*, \beta^*),
\end{aligned}$$

so $L(\alpha^*, \beta^*) = V = \inf_{\beta \in B} L(\alpha^*, \beta)$ and $L(\alpha^*, \beta^*) = V = \sup_{\alpha \in A} L(\alpha, \beta^*)$.

(3) To show that the game has a value, by (1) it is sufficient to show that

$$\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) \leq \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta).$$

But this is indeed the case:

$$\begin{aligned}
\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta) &\leq \sup_{\alpha \in A} L(\alpha, \beta^*) \\
&= L(\alpha^*, \beta^*) = K = \inf_{\beta \in B} L(\alpha^*, \beta) \\
&\leq \sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta),
\end{aligned}$$

where the first equalities follow because β^* is an equalizer strategy. Thus, the game has a value equal to K . Since $\sup_{\alpha \in A} L(\alpha, \beta^*) = K$, β^* achieves $\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta)$. Since $\inf_{\beta \in B} L(\alpha^*, \beta) = K$, α^* achieves $\sup_{\alpha \in A} \inf_{\beta \in B} L(\alpha, \beta)$. Therefore, (α^*, β^*) is a saddle point. \square

2) *Proof of Theorem 2, the Saddle Point Theorem for Correlated Settings and Generalized Nonlocality Proofs:* We use the following well-known minimax theorem due to Ferguson. The form in which we state it is a straightforward combination of Ferguson's [14] Theorem 1, page 78 and Theorem 2.1, page 85, specialized to the Euclidean topology.

Theorem 4 (Ferguson 1967): Let (A, B, L) be a statistical game where A is a finite set, B is a convex compact subset of \mathbb{R}^k for some $k > 0$ and L is such that for all $\alpha \in A$,

- 1) $L(\alpha, \beta)$ is a convex function of $\beta \in B$.
- 2) $L(\alpha, \beta)$ is lower semicontinuous in $\beta \in B$.

Let \mathcal{A} be the set of distributions on A and define, for $P \in \mathcal{A}$, $L(P, \beta) = \mathbf{E}_P L(\alpha, \beta) = \sum_{\alpha \in A} P_\alpha L(\alpha, \beta)$. Then the game (\mathcal{A}, B, L) has a value V , i.e.

$$\sup_{P \in \mathcal{A}} \inf_{\beta \in B} L(P, \beta) = \inf_{\beta \in B} \sup_{P \in \mathcal{A}} L(P, \beta),$$

and a minimax $\beta^* \in B$ achieving $\inf_{\beta \in B} \sup_{\alpha \in A} L(\alpha, \beta)$ exists. By Theorem 1, part (1), $U(\sigma, \pi) = D(Q_\sigma \| P_{\sigma; \pi})$ is lower semicontinuous in π for all $\sigma \in \Sigma$. Let us now focus on the case of a $2 \times 2 \times 2$ game. We can apply Theorem 4 with $A = \{11, 12, 21, 22\}$, $\mathcal{A} = \Sigma$ and $B = \Pi$. It follows that the game (Σ, Π, U) has a value V , and $\inf_{\Pi} \sup_{\Sigma} U(\sigma, \pi) = V$ is achieved for some $\pi^* \in \Pi$. By Theorem 1, part (2), $0 \leq V < \infty$, and, since $U(\sigma)$ is continuous in σ , there exists some σ^* achieving $\sup_{\Sigma} \inf_{\Pi} U(\sigma, \pi)$.

The proof for generalized nonlocality proofs is completely analogous; we omit details.

3) *Proof of Theorem 3, Saddle Points and Equalizer Strategies for $2 \times 2 \times 2$ Nonlocality Proofs:* The correlated game has a value V by Theorem 2 and $V > 0$ by Theorem 1. Inequality (8) is immediate. Let $U((a, b), \pi)$ be defined as in the proof of Theorem 1, Equation (20). To prove Equation (9), note that for every $\pi \in \Pi$,

$$\begin{aligned}
\sup_{\Sigma \cup C} U(\sigma, \pi) &= \sup_{\Sigma} U(\sigma, \pi) \\
&= \max_{a, b \in \{1, 2\}} U((a, b), \pi).
\end{aligned}$$

Thus, Equation (9) and part 2(b) of the theorem follow. Part 2(a) is immediate from Theorem 2. To prove part 2(c), suppose, by way of contradiction, that there exists a $\pi^* \in \Pi^*$ that is not an equalizer strategy. Then the set $\{(a, b) \mid U((a, b), \pi^*) = \max_{a, b \in \{1, 2\}} U((a, b), \pi^*)\}$ has less than four elements. By Theorem 2, there exists a $\sigma^* \in \Sigma$ such that (σ^*, π^*) is a saddle point in the correlated game. Since $\sigma^* \in \Sigma$ achieves $\sup_{\Sigma} U(\sigma, \pi^*)$, it follows that for some $a_0, b_0 \in \{1, 2\}$, $\sigma_{a_0 b_0}^* = 0$. But then σ^* lies on the boundary of Σ . By Theorem 1, part 2(d), this is impossible, and we have arrived at the desired contradiction.

It remains to prove part (3). Part (3), 'if' follows directly from Proposition 1. To prove part (3), 'only if', suppose the uncorrelated game has saddle point (σ^*, π^*) . It is clear that π^* achieves $\inf_{\Pi} U(\sigma^*, \pi)$. We have already shown above that π^* is an equalizer strategy.