

Article

No-Signalling Constraints on Exponential Tilts in CHSH Scenarios

Camilla Maria Kyllikki Josephson 

Strategic Research Support, Aarhus University, 8000 Aarhus, Denmark; cmkjosephson@gmail.com

Abstract

We characterize when exponential reweightings of a no-signalling CHSH probability box preserve no-signalling. While such tilts are automatically positive and normalized within each measurement setting, they can modify cross-setting marginals and thereby introduce signalling into the probability table. We identify a four-dimensional setting-only redundancy in the residual parametrization, derive the exact nonlinear compatibility conditions for no-signalling preservation, and obtain the linearized no-signalling constraint around a no-signalling reference box. For the unbiased Tsirelson CHSH box, we compute the linearized constraint in closed form and show that the admissible tangent space has dimension twelve before quotienting and dimension eight after quotienting by the setting-only redundancy, matching the standard dimension of binary no-signalling boxes. Exact-probability calculations confirm the predicted scaling: generic residual directions produce first-order no-signalling leakage, while admissible tangent directions suppress the leakage to second order. We further show that local-additive residuals, despite their algebraic locality, are not generically no-signalling safe. These results give a sharp first-order admissibility criterion for exponential tilts of Bell probability boxes.

Keywords: bell inequalities; CHSH scenario; no-signalling; exponential tilts; probability reweighting; quantum correlations; Tsirelson bound; information geometry

1. Introduction

Exponential reweighting of a reference probability law,

$$P_i = \frac{P_i^{(0)} e^{-F_i}}{\sum_j P_j^{(0)} e^{-F_j}}, \quad (1)$$

is a standard operation. It appears in Gibbs reweighting, likelihood tilting, and maximum-entropy updating relative to a prior distribution [1]. In a single probability space, this construction is unproblematic: for any real array F , the reweighted distribution is positive and normalized.

The situation is different in a Bell scenario. A CHSH table is not a single probability distribution but a family of conditional distributions, one for each pair of measurement settings. Positivity and normalization inside each setting block are therefore not the only requirements. The four setting-conditioned distributions must also be compatible with each other through the no-signalling constraints.



Academic Editor: Aurelian Isar

Received: 28 April 2026

Revised: 21 June 2026

Accepted: 22 June 2026

Published: 29 June 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

A CHSH box is a family of distributions

$$P(a, b|x, y),$$

indexed by Alice's setting x , Bob's setting y , and outcomes a, b . Although each setting pair (x, y) can be normalized separately, the different setting pairs are linked by no-signalling constraints. These constraints require Alice's marginal distribution to be independent of Bob's setting, and Bob's marginal distribution to be independent of Alice's setting [2–5]. The Bell-local bound 2, the Tsirelson bound $2\sqrt{2}$, and the no-signalling constraints describe different levels of admissibility in the CHSH scenario. The Bell bound concerns the local hidden-variable polytope. The Tsirelson bound constrains quantum-realizable CHSH correlations. No-signalling is a broader compatibility condition on the family of setting-wise distributions: it requires that the local marginal readout of one party is indistinguishable under changes of the other party's setting. The present paper studies this no-signalling compatibility condition, not the characterization of the local or quantum sets themselves.

This paper studies the following question: if a no-signalling CHSH box is exponentially tilted setting by setting, when does the resulting box remain no-signalling? Let $P_0(a, b|x, y) > 0$ be a strictly positive reference box and let $F = (F_{abxy}) \in \mathbb{R}^{16}$ be a residual array. We define

$$P_F(a, b|x, y) = \frac{P_0(a, b|x, y) e^{-F_{abxy}}}{Z_{xy}(F)}, \quad Z_{xy}(F) = \sum_{a', b'} P_0(a', b'|x, y) e^{-F_{a'b'xy}}. \quad (2)$$

For each fixed (x, y) , this is a valid probability distribution. The nontrivial question is whether the four setting-wise distributions still fit together into a no-signalling box.

This is a probability-level admissibility problem. A setting-wise exponential tilt can be perfectly well defined inside each measurement context while still shifting Alice's or Bob's marginal distribution across the remote setting. In this operational sense, the tilt may introduce signalling into the probability table. We do not mean that a physical faster-than-light communication process has been generated. Rather, "signalling" means that the local marginal distribution becomes dependent on the remote measurement setting.

This question is relevant whenever Bell-type probability tables are reweighted, biased, post-processed, or used as diagnostic objects. Such situations arise naturally in robustness analyses, self-testing, device-independent reasoning, resource-theoretic treatments of nonlocality, and quantum-cryptographic contexts. In all such settings, one must distinguish genuine joint-correlation structure from transformations that inadvertently introduce remote-setting dependence in local marginals.

Recent pedagogical and experimental discussions of Bell tests, including optical emulations of quantum-state tomography and Bell-test scenarios, also illustrate the continuing role of CHSH-type correlations as a finite-dimensional testbed for quantum and no-signalling constraints [6].

We focus on the binary CHSH scenario, with $x, y, a, b \in \{0, 1\}$. The main reference point is the unbiased Tsirelson CHSH box, whose CHSH value is $2\sqrt{2}$ and whose local marginals are unbiased [7]. This box is useful because it is exactly no-signalling and lies on the quantum boundary.

The use of the Tsirelson box should not be read as a claim that this paper is primarily about Tsirelson's theorem. The unbiased Tsirelson box is used as a symmetric, strictly positive, no-signalling reference point. Its role is computational and geometric: it allows the linearized no-signalling Jacobian and its Gram matrix to be computed explicitly.

Contribution and Relation to Prior Work

The eight-dimensional structure of the binary no-signalling polytope is well known; it goes back to Barrett et al. [8] and is reviewed in [5]. The dimension count itself is therefore not new. The contribution of this paper is to recover and analyze the local no-signalling structure in exponential-tilt coordinates.

The present question is related to, but distinct from, the literature on no-signalling boxes as information-theoretic resources and on causal or locality-preserving maps. Work on no-signalling boxes, quantum no-signalling boxes, and causal quantum operations asks when a global operation preserves the operational no-signalling structure. In particular, Piani, M. Horodecki, P. Horodecki and R. Horodecki studied causal quantum operations and showed that strong restrictions arise when one requires global transformations to preserve no-signalling structure [9]. The object studied here is narrower: we do not consider general completely positive maps or arbitrary resource transformations. Instead, we study setting-wise exponential reweightings of classical CHSH probability tables.

The analysis has two stages. First, we write the exact nonlinear compatibility equations for an exponentially tilted box. Second, we linearize these equations around a no-signalling reference box in order to identify the residual directions that do not create first-order signalling leakage. Thus the main result is a first-order admissibility criterion, not a complete nonlinear classification of all exactly no-signalling residual tilts.

Specifically, we

- (i) Identify a four-dimensional setting-only redundancy $\mathcal{G}$ in the residual representation: residual arrays of the form $F_{abxy} = G_{xy}$ leave the box unchanged;
- (ii) Derive the exact no-signalling compatibility equations for the tilted box P_F ;
- (iii) Compute the linearized no-signalling map $D NS_{P_0} : \mathbb{R}^{16} \rightarrow \mathbb{R}^4$ at a no-signalling reference box;
- (iv) Specialize to the unbiased Tsirelson box and compute the Gram matrix $D NS_{P_0} (D NS_{P_0})^\top$ explicitly, obtaining singular values

$$\left\{ \frac{\sqrt{5}}{4}, \frac{\sqrt{5}}{4}, \frac{1}{4}, \frac{1}{4} \right\};$$

- (v) Conclude that the linearized admissible kernel has dimension 12, and that quotienting by the four-dimensional redundancy gives an eight-dimensional gauge-free admissible tangent space;
- (vi) Prove the scaling law

$$NS(P_{\epsilon f}) = \epsilon D NS_{P_0}[f] + O(\epsilon^2),$$

so that generic residual directions give first-order no-signalling leakage, while tangent-admissible directions give only second-order leakage;

- (vii) Exhibit an explicit local-additive residual $f_{abxy} = A_{ax} + B_{by}$ which is local in algebraic form but still produces first-order signalling leakage.

Thus the contribution is not the eight-dimensional no-signalling description itself, which is standard. The contribution is to recover and analyze this admissibility structure in exponential residual coordinates, to compute the corresponding Jacobian explicitly at the Tsirelson reference box, and to show that a residual which is local in algebraic form need not be no-signalling safe after setting-wise normalization of a correlated reference box.

The last point is important. A residual may be written as an Alice-only term plus a Bob-only term, yet still fail to preserve no-signalling after setting-wise reweighting of a correlated reference box. Thus locality of the residual expression is not the same as causal admissibility of the resulting probability table.

The scope of the paper is deliberately limited. No-signalling is a necessary condition for quantum realizability, but it is not sufficient. A tilted box may remain no-signalling while leaving the quantum set, and may even have a CHSH value above $2\sqrt{2}$. Such values should not be interpreted as physical violations of Tsirelson's theorem. They only show that the probability-level tilt need not correspond to measurements on a quantum state. The present paper studies first-order no-signalling admissibility, not full quantum admissibility.

This paper is organized as follows. Section 2 recalls CHSH boxes and the no-signalling constraints. Section 3 defines exponential tilts and the setting-only redundancy. Section 4 gives the exact compatibility equations. Section 5 derives the linearized constraint. Section 6 computes the Jacobian spectrum at the unbiased Tsirelson box. Section 7 gives the dimension count. Section 8 studies the residual classes, including the local-additive counterexample. Section 9 reports the numerical verification. Section 10 discusses finite-shot scaling, and Section 11 clarifies the relation to quantum admissibility.

2. CHSH Boxes and No-Signalling

We work in the standard binary Bell scenario. Alice chooses a setting $x \in \{0, 1\}$, Bob chooses a setting $y \in \{0, 1\}$, and the corresponding outcomes are $a, b \in \{0, 1\}$. A CHSH box is a family of conditional probabilities $P(a, b|x, y)$ satisfying

$$P(a, b|x, y) \geq 0, \quad \sum_{a,b} P(a, b|x, y) = 1 \quad (3)$$

for every setting pair (x, y) .

There are four setting pairs and four outcomes per setting pair. Since each setting pair distribution has one normalization constraint, the space of normalized binary boxes has dimension

$$4(4 - 1) = 12. \quad (4)$$

For later use, define the sign variables

$$s_0 = t_0 = +1, \quad s_1 = t_1 = -1. \quad (5)$$

The correlator associated with setting pair (x, y) is

$$E_{xy} = \sum_{a,b} s_a t_b P(a, b|x, y). \quad (6)$$

The CHSH value of a box P is the scalar

$$\text{CHSH}(P) = E_{00} + E_{01} + E_{10} - E_{11}. \quad (7)$$

The CHSH scenario is often used both as the simplest test case for Bell nonlocality and as a finite-dimensional model in which the local, quantum, and no-signalling sets can be compared geometrically. The present paper does not alter this standard geometry; rather, it studies how the no-signalling constraints appear when conditional probability tables are parametrized by setting-wise exponential residual tilts.

The no-signalling condition says that one party's marginal distribution cannot depend on the other party's choice of setting. Thus Alice's marginal

$$P_A(a|x, y) = \sum_b P(a, b|x, y) \quad (8)$$

must be independent of y , and Bob's marginal

$$P_B(b|x, y) = \sum_a P(a, b|x, y) \quad (9)$$

must be independent of x .

Because the outcomes are binary, it is enough to impose these conditions for outcome 0 on each side; the outcome 1 conditions then follow from normalization. We therefore define

$$N_{A,x}(P) = \sum_b P(0, b|x, 0) - \sum_b P(0, b|x, 1), \quad x \in \{0, 1\}, \quad (10)$$

and

$$N_{B,y}(P) = \sum_a P(a, 0|0, y) - \sum_a P(a, 0|1, y), \quad y \in \{0, 1\}. \quad (11)$$

The no-signalling map is

$$NS(P) = (N_{A,0}(P), N_{A,1}(P), N_{B,0}(P), N_{B,1}(P))^T \in \mathbb{R}^4. \quad (12)$$

Thus

$$P \text{ is no-signalling} \iff NS(P) = 0. \quad (13)$$

Throughout this paper, $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^4$ or $\mathbb{R}^{16}$ unless otherwise stated. Since all norms are equivalent in finite dimension, the asymptotic scaling statements are unchanged if another norm is used.

3. Exponential Tilts of CHSH Boxes

Definition 1 (Exponential tilt of a Bell box). Let $P_0(a, b|x, y) > 0$ be a strictly positive reference box. A residual array is $F = (F_{abxy}) \in \mathbb{R}^{16}$. The associated exponential tilt is

$$P_F(a, b|x, y) = \frac{P_0(a, b|x, y) e^{-F_{abxy}}}{Z_{xy}(F)}, \quad Z_{xy}(F) = \sum_{a', b'} P_0(a', b'|x, y) e^{-F_{a'b'xy}}. \quad (14)$$

Strict positivity of P_0 ensures that the tilt is globally well defined and that F can be recovered from P_F up to a coordinate redundancy described below.

For each fixed (x, y) , Equation (14) is a single-context exponential tilt of four outcome probabilities, hence positive and normalized. What is not automatic is compatibility between different setting pairs: setting-wise normalization may move Alice's or Bob's marginals and so take a no-signalling box outside the no-signalling set.

Thus the issue is not normalization inside a single measurement context, but compatibility across contexts. A residual tilt may be a perfectly valid probability update for each fixed setting pair while failing the joint no-signalling consistency test for the whole CHSH box.

Setting-Only Redundancy

The residual representation has an immediate redundancy. If F depends only on the settings,

$$F_{abxy} = G_{xy}, \quad (15)$$

the factor $e^{-G_{xy}}$ is common to all four outcomes and cancels in the following normalization:

$$P_F(a, b|x, y) = \frac{P_0(a, b|x, y) e^{-G_{xy}}}{\sum_{a', b'} P_0(a', b'|x, y) e^{-G_{xy}}} = \frac{P_0(a, b|x, y) e^{-G_{xy}}}{e^{-G_{xy}} \sum_{a', b'} P_0(a', b'|x, y)} = P_0(a, b|x, y). \quad (16)$$

Define the redundancy subspace

$$\mathcal{G} = \{F \in \mathbb{R}^{16} : F_{abxy} = G_{xy} \text{ for some } G \in \mathbb{R}^4\}, \quad \dim \mathcal{G} = 4. \quad (17)$$

The four free parameters are the constants $G_{00}, G_{01}, G_{10}, G_{11}$. We use the term “redundancy” (or “residual gauge” in the coordinate sense) only to mean that different residual arrays describe the same probability box; this is not a Yang–Mills-type gauge symmetry, just a consequence of setting-wise normalization. The physically relevant residual information is the equivalence class $[F] \in \mathbb{R}^{16}/\mathcal{G}$.

It is important to distinguish this setting-only residual redundancy from a marginal-zero perturbation of the probability table. A residual array of the form $F_{abxy} = G_{xy}$ does not leave the box unchanged because it has vanishing marginal probability. It leaves the box unchanged because the corresponding exponential factor is constant within each setting block and is removed before any marginal comparison is made. Consequently,

$$P_{F,A}(a|x, y) = P_{0,A}(a|x, y), \quad P_{F,B}(b|x, y) = P_{0,B}(b|x, y), \quad (18)$$

and therefore

$$NS(P_F) = NS(P_0). \quad (19)$$

Thus these directions are not physical deformations of the probability table; they are coordinate redundancies of the exponential residual parametrization.

4. Exact No-Signalling Compatibility

A residual F is no-signalling admissible relative to P_0 if $NS(P_F) = 0$. This condition is nonlinear in F because the normalization factor $Z_{xy}(F)$ depends on the setting pair.

Although the equations below are written for a representative residual array F , the condition $NS(P_F) = 0$ depends only on the tilted box P_F . Hence it is invariant under the setting-only redundancy $F_{abxy} \mapsto F_{abxy} + G_{xy}$. Equivalently, no-signalling admissibility is a property of the residual equivalence class $[F] \in \mathbb{R}^{16}/\mathcal{G}$, not of a particular representative.

The Alice marginal of the tilted box is

$$P_F(a|x, y) = \frac{\sum_b P_0(a, b|x, y) e^{-F_{abxy}}}{Z_{xy}(F)}, \quad (20)$$

so independence from y requires

$$\frac{\sum_b P_0(a, b|x, y) e^{-F_{abxy}}}{Z_{xy}(F)} = \frac{\sum_b P_0(a, b|x, y') e^{-F_{abxy'}}}{Z_{xy'}(F)} \quad (21)$$

for all a, x, y, y' . Similarly, Bob’s marginal independence from x requires

$$\frac{\sum_a P_0(a, b|x, y) e^{-F_{abxy}}}{Z_{xy}(F)} = \frac{\sum_a P_0(a, b|x', y) e^{-F_{abx'y}}}{Z_{x'y}(F)} \quad (22)$$

for all b, y, x, x' . Equations (21) and (22) are the exact admissibility conditions; the rest of this paper analyses their local form near $F = 0$.

These equations are necessary and sufficient for exact no-signalling of the tilted box P_F . They are nonlinear because each setting block has its own normalization factor. The linearized analysis below studies the tangent version of these exact conditions at $F = 0$.

5. Linearized No-Signalling Constraint and Analytic Jacobian

The exact compatibility equations are nonlinear in F . To extract the local admissible directions we expand around $F = 0$. Write $F = \epsilon f$ with $0 < \epsilon \ll 1$, and define the P_0 expectation at setting (x, y) using

$$\langle f \rangle_{xy} = \sum_{a,b} P_0(a, b|x, y) f_{abxy}. \quad (23)$$

From $e^{-\epsilon f_{abxy}} = 1 - \epsilon f_{abxy} + O(\epsilon^2)$ we have

$$Z_{xy}(\epsilon f) = 1 - \epsilon \langle f \rangle_{xy} + O(\epsilon^2), \quad (24)$$

hence

$$P_{\epsilon f}(a, b|x, y) = P_0(a, b|x, y) [1 - \epsilon (f_{abxy} - \langle f \rangle_{xy})] + O(\epsilon^2). \quad (25)$$

Only the centered residual $f_{abxy} - \langle f \rangle_{xy}$ matters at first order. The setting-wise average is removed by the normalization; this is the infinitesimal form of the redundancy of Section 3.

This already shows why the linearized problem is a tangent problem on residual equivalence classes. Adding a setting-only residual changes $\langle f \rangle_{xy}$ by exactly the same setting-only amount and therefore does not change the centered first-order deformation. Thus the Jacobian below acts trivially on the infinitesimal version of the setting-only redundancy $\mathcal{G}$.

Theorem 1 (First-order no-signalling constraint). *Assume P_0 is no-signalling. Then $P_{\epsilon f}$ is no-signalling at first order in ϵ if and only if*

$$\sum_b P_0(a, b|x, y) (f_{abxy} - \langle f \rangle_{xy}) \quad \text{is independent of } y \text{ for every fixed } a, x, \quad (26)$$

and

$$\sum_a P_0(a, b|x, y) (f_{abxy} - \langle f \rangle_{xy}) \quad \text{is independent of } x \text{ for every fixed } b, y. \quad (27)$$

Proof. According to (25), Alice's marginal is

$$P_{\epsilon f}(a|x, y) = P_0(a|x, y) - \epsilon \sum_b P_0(a, b|x, y) (f_{abxy} - \langle f \rangle_{xy}) + O(\epsilon^2). \quad (28)$$

Since P_0 is no-signalling, $P_0(a|x, y)$ is already independent of y . The tilted box is no-signalling at first order on Alice's side precisely when the first-order correction is independent of y , which is (26). Bob's condition (27) follows using the symmetric argument. $\square$

We denote the corresponding linear map as $DNS_{P_0} : \mathbb{R}^{16} \rightarrow \mathbb{R}^4$. The linearly admissible residual tangent space is

$$T_{\text{adm}}(P_0) = \ker DNS_{P_0}. \quad (29)$$

Membership in $T_{\text{adm}}(P_0)$ is a first-order condition. It does not imply that $P_{\epsilon f}$ is exactly no-signalling for finite ϵ ; it means that the first-order signalling leakage vanishes. Exact finite- ϵ no-signalling is governed by the nonlinear compatibility Equations (21) and (22).

Closed form of the Jacobian.

Differentiating (25) entry-wise gives

$$\left. \frac{\partial P_F(a, b|x, y)}{\partial F_{a'b'x'y'}} \right|_{F=0} = \delta_{xx'} \delta_{yy'} P_0(a, b|x, y) [P_0(a', b'|x, y) - \delta_{aa'} \delta_{bb'}]. \quad (30)$$

The derivative is setting-local: variables $F_{a'b'x'y'}$ affect only the probability block with the same setting pair (x', y') . Consequently each row of DNS_{P_0} has support only on the two setting blocks compared by the corresponding no-signalling component: $N_{A,x}$ couples blocks $(x, 0)$ and $(x, 1)$; $N_{B,y}$ couples blocks $(0, y)$ and $(1, y)$. This setting-local block structure is the key to the rank computation in the next section.

In this sense the Jacobian is not an arbitrary 4×16 matrix: its sparsity encodes the operational structure of no-signalling. Each row compares two setting blocks that differ only in the remote setting.

6. The Tsirelson CHSH Box

We now specialize to the unbiased Tsirelson CHSH box. Let $s_0 = t_0 = +1$ and $s_1 = t_1 = -1$, and define

$$P_0(a, b|x, y) = \frac{1}{4} [1 + E_{xy} s_a t_b], \quad E_{00} = E_{01} = E_{10} = \frac{1}{\sqrt{2}}, \quad E_{11} = -\frac{1}{\sqrt{2}}. \quad (31)$$

The local marginals are unbiased, $P_0(a|x, y) = P_0(b|x, y) = \frac{1}{2}$, so $NS(P_0) = 0$. The CHSH value is

$$E_{00} + E_{01} + E_{10} - E_{11} = 2\sqrt{2}, \quad (32)$$

the Tsirelson value [7].

The role of this reference point is local and geometric. We use the unbiased Tsirelson box because it is strictly positive, exactly no-signalling, has unbiased local marginals, and lies on the quantum boundary. The calculation below should therefore be read as a closed-form linearization at a symmetric reference box, not as a claim about all quantum correlations.

Proposition 1 (Closed-form Gram matrix and singular values). *At the unbiased Tsirelson box, the Gram matrix of DNS_{P_0} is*

$$DNS_{P_0} (DNS_{P_0})^\top = \frac{1}{32} \begin{pmatrix} 6 & 0 & 2\sqrt{2} & -2\sqrt{2} \\ 0 & 6 & -2\sqrt{2} & -2\sqrt{2} \\ 2\sqrt{2} & -2\sqrt{2} & 6 & 0 \\ -2\sqrt{2} & -2\sqrt{2} & 0 & 6 \end{pmatrix}. \quad (33)$$

Its eigenvalues are $\{5/16, 5/16, 1/16, 1/16\}$, hence the singular values of DNS_{P_0} are

$$\sigma_1 = \sigma_2 = \frac{\sqrt{5}}{4}, \quad \sigma_3 = \sigma_4 = \frac{1}{4}. \quad (34)$$

In particular $\text{rank}(DNS_{P_0}) = 4$.

Geometrically, these singular values measure the first-order sensitivity of the no-signalling constraints to residual-coordinate perturbations at the Tsirelson reference box. The two singular-value scales show that the linearized leakage is anisotropic: two orthogonal combinations of marginal-compatibility constraints respond with strength $\sqrt{5}/4$, while two respond with strength $1/4$. This should not be read as a new physical scale; it is a local geometric property of the exponential-tilt parametrization at the chosen reference box.

Proof. We compute the row vectors of $DNS_{P_0} \in \mathbb{R}^{4 \times 16}$ from (30). Using Theorem 1 and (30), the entry of the row corresponding to $N_{A,x}$ in column (a', b', x', y') is

$$(DNS_{P_0})_{N_{A,x}(a'b'x'y')} = \sum_b \left[\delta_{x'x} \delta_{y'0} P_0(0, b|x, 0) (P_0(a', b'|x, 0) - \delta_{0a'} \delta_{bb'}) - \delta_{x'x} \delta_{y'1} P_0(0, b|x, 1) (P_0(a', b'|x, 1) - \delta_{0a'} \delta_{bb'}) \right], \quad (35)$$

and similarly for the $N_{B,y}$ rows by exchanging the roles of the two parties.

For the unbiased Tsirelson box, $P_0(a, b|x, y) = \frac{1}{4}[1 + E_{xy}s_a t_b]$ with $\sum_a s_a = \sum_b t_b = 0$, so the marginals are $P_0(a|x, y) = P_0(b|x, y) = \frac{1}{2}$ for all settings. We will use the identities

$$\sum_{a', b'} P_0(a', b'|x, y)^2 = \frac{1+E_{xy}^2}{4}, \quad \sum_{b'} P_0(0, b'|x, y)^2 = \sum_{a'} P_0(a', 0|x, y)^2 = \frac{1+E_{xy}^2}{8}, \quad (36)$$

both of which follow directly from $P_0(a, b|x, y) = \frac{1}{4}(1 + E_{xy}s_a t_b)$ and the orthogonality relations $\sum_a s_a = \sum_b t_b = 0$, $s_a^2 = t_b^2 = 1$.

Diagonal entries. Fix x , and consider the $N_{A,x}$ row inner product with itself. By the setting-local block structure, this row has support only on the two blocks $(x, 0)$ and $(x, 1)$, with sign $\sigma_0 = +1$ and $\sigma_1 = -1$ respectively. The entry of the row at column (a', b', x, y) is

$$J_{a'b'}^{(x,y)} = \sigma_y \sum_b P_0(0, b|x, y) [P_0(a', b'|x, y) - \delta_{0a'} \delta_{bb'}] = \sigma_y \left[\frac{1}{2} P_0(a', b'|x, y) - \delta_{0a'} P_0(0, b'|x, y) \right], \quad (37)$$

where the second equality uses $\sum_b P_0(0, b|x, y) = P_0(0|x, y) = \frac{1}{2}$. Squaring and summing,

$$\begin{aligned} \sum_{a', b'} (J_{a'b'}^{(x,y)})^2 &= \sum_{a', b'} \left[\frac{1}{4} P_0(a', b'|x, y)^2 - \delta_{0a'} P_0(a', b'|x, y) P_0(0, b'|x, y) + \delta_{0a'} P_0(0, b'|x, y)^2 \right] \\ &= \frac{1}{4} \sum_{a', b'} P_0(a', b'|x, y)^2 - \sum_{b'} P_0(0, b'|x, y)^2 + \sum_{b'} P_0(0, b'|x, y)^2 \\ &= \frac{1}{4} \cdot \frac{1+E_{xy}^2}{4} = \frac{1+E_{xy}^2}{16}, \end{aligned} \quad (38)$$

using (36) and the fact that, when $a' = 0$, $P_0(a', b'|x, y) = P_0(0, b'|x, y)$, the cross and Y^2 terms cancel. At the unbiased Tsirelson value $E_{xy}^2 = \frac{1}{2}$, this gives $\frac{3}{32}$ per block; summing over the two blocks $(x, 0)$ and $(x, 1)$ that the $N_{A,x}$ row touches gives the diagonal entry $2 \cdot \frac{3}{32} = \frac{6}{32}$. The same value arises for each of the four diagonal entries according to the symmetric argument on $N_{B,y}$.

Off-diagonal Alice–Alice and Bob–Bob entries. The rows $N_{A,0}$ and $N_{A,1}$ are supported on disjoint setting blocks (settings $(0, y)$ versus $(1, y)$ for all y), so their inner product vanishes. The same argument shows the $(N_{B,0}, N_{B,1})$ entry vanishes.

Cross Alice–Bob entries. The rows $N_{A,x}$ and $N_{B,y}$ touch the single common block (x, y) . On that block, the A -row entries are

$$J_{a'b'}^{(x,y),A} = \sigma_y^A \left[\frac{1}{2} P_0(a', b'|x, y) - \delta_{0a'} P_0(0, b'|x, y) \right]$$

and the B -row entries are

$$J_{a'b'}^{(x,y),B} = \sigma_x^B \left[\frac{1}{2} P_0(a', b'|x, y) - \delta_{0b'} P_0(a', 0|x, y) \right],$$

with $\sigma_0^A = \sigma_0^B = +1$, $\sigma_1^A = \sigma_1^B = -1$. Expanding the product and using (36),

$$\begin{aligned} \sum_{a',b'} J_{a'b'}^{(x,y),A} J_{a'b'}^{(x,y),B} &= \sigma_y^A \sigma_x^B \left[\frac{1}{4} \sum_{a',b'} P_0(a', b' | x, y)^2 - \frac{1}{2} \sum_{b'} P_0(0, b' | x, y)^2 \right. \\ &\quad \left. - \frac{1}{2} \sum_{a'} P_0(a', 0 | x, y)^2 + P_0(0, 0 | x, y)^2 \right] \\ &= \sigma_y^A \sigma_x^B \left[\frac{1+E_{xy}^2}{16} - \frac{1+E_{xy}^2}{16} - \frac{1+E_{xy}^2}{16} + \frac{(1+E_{xy})^2}{16} \right] \\ &= \sigma_y^A \sigma_x^B \cdot \frac{(1+E_{xy})^2 - (1+E_{xy}^2)}{16} = \sigma_y^A \sigma_x^B \cdot \frac{E_{xy}}{8}, \end{aligned} \quad (39)$$

where the last equality uses $(1+E)^2 - (1+E^2) = 2E$. Substituting $E_{xy} \in \{1/\sqrt{2}, -1/\sqrt{2}\}$ and the appropriate signs gives the six off-diagonal entries: $(N_{A,0}, N_{B,0})$ on common block $(0,0)$ yields $E_{00}/8 = \sqrt{2}/16 = 2\sqrt{2}/32$; $(N_{A,0}, N_{B,1})$ on common block $(0,1)$ with $\sigma_1^A \sigma_0^B = -1$ yields $-E_{01}/8 = -\sqrt{2}/16 = -2\sqrt{2}/32$; and analogously for the remaining pairs. The pattern shown in (33) follows.

It remains to compute the eigenvalues of (33). The 4×4 matrix $32G$ is real symmetric with trace $\text{tr}(32G) = 24$. By inspection,

$$v_{\pm}^{(1)} = (1, 1, \pm 1, \mp 1)/2, \quad v_{\pm}^{(2)} = (1, -1, \pm 1, \pm 1)/2, \quad (40)$$

form an orthonormal basis of $\mathbb{R}^4$. Direct substitution gives

$$32G v_+^{(1)} = 10 v_+^{(1)}, \quad 32G v_-^{(1)} = 2 v_-^{(1)}, \quad 32G v_+^{(2)} = 10 v_+^{(2)}, \quad 32G v_-^{(2)} = 2 v_-^{(2)}, \quad (41)$$

so G has eigenvalues $\{10/32, 2/32, 10/32, 2/32\} = \{5/16, 1/16, 5/16, 1/16\}$, each with multiplicity 2. Their square roots give the singular values of DN_{P_0} ,

$$\sigma_1 = \sigma_2 = \frac{\sqrt{5}}{4}, \quad \sigma_3 = \sigma_4 = \frac{1}{4},$$

and all four are non-zero, so $\text{rank}(DN_{P_0}) = 4$. $\square$

The Gram matrix (33) and singular values were also verified numerically by independent central differences; the analytic and numerical Jacobians agreed to $\|J_{\text{num}} - J_{\text{an}}\|_F = 1.86 \times 10^{-10}$. The full 4×16 matrix DN_{P_0} in closed form, together with its block-by-block structure, is recorded in Appendix A.

7. Dimension Count and Parametrization

The residual array F_{abxy} has $2^4 = 16$ real components, and the redundancy subspace $\mathcal{G}$ has dimension 4, so the residual quotient $\mathbb{R}^{16}/\mathcal{G}$ has dimension 12. This matches the dimension of the normalized CHSH box space.

Proposition 2 (Residual coordinates parameterize positive boxes). *Let P_0 be a strictly positive box and let Q be any strictly positive normalized box on the same setting/outcome alphabet. Define*

$$F_{abxy} = -\log \frac{Q(a, b | x, y)}{P_0(a, b | x, y)}. \quad (42)$$

Then $P_F = Q$, and the equivalence class $[F] \in \mathbb{R}^{16}/\mathcal{G}$ is uniquely determined by Q .

Proof. Setting-wise,

$$\sum_{a,b} P_0(a, b | x, y) \exp(-F_{abxy}) = \sum_{a,b} Q(a, b | x, y) = 1,$$

so $Z_{xy} = 1$ and $P_F(a, b|x, y) = Q(a, b|x, y)$. If F' is any other residual with $P_{F'} = Q$, then $\exp(-(F'_{abxy} - F_{abxy})) = Z'_{xy}/Z_{xy}$ for each (x, y) , so $F' - F$ depends only on (x, y) and lies in $\mathcal{G}$. $\square$

In the binary case, the no-signalling polytope is parameterized by

$$Q(a, b|x, y) = \frac{1}{4} [1 + A_x s_a + B_y t_b + C_{xy} s_a t_b], \quad (43)$$

with eight real parameters $A_0, A_1, B_0, B_1, C_{00}, C_{01}, C_{10}, C_{11}$ and the positivity constraints

$$1 + A_x s_a + B_y t_b + C_{xy} s_a t_b \geq 0. \quad (44)$$

The interior is therefore eight-dimensional [5,8]. This dimension count is known; the point of the present analysis is that it is recovered from residual exponential-tilt coordinates after quotienting the setting-wise redundancy. Concretely, combining Propositions 1 and 2, the local dimension count at the Tsirelson reference is

quantity	dimension
residual space $\mathbb{R}^{16}$	16
$\ker D NS_{P_0}$	12
setting-only redundancy $\mathcal{G}$	4
$\ker D NS_{P_0} / \mathcal{G}$	8

(45)

The bottom line agrees dimensionally with the eight-parameter description of binary no-signalling boxes. No-signalling does not eliminate all nontrivial residual tilts; it removes the four first-order signalling directions and leaves an eight-dimensional gauge-free tangent space of locally admissible residual deformations.

This final eight-dimensional space should therefore be interpreted carefully. It is not proposed as a new no-signalling polytope dimension; that dimension is standard. The new point is that the same dimension is recovered from exponential residual coordinates after removing the setting-only redundancy and imposing the first-order no-signalling constraint.

8. Scaling and Residual Classes

The Taylor expansion of the no-signalling map gives the local diagnostic used in our verification.

Proposition 3 (Scaling of no-signalling leakage). *Let P_0 be no-signalling and $f \in \mathbb{R}^{16}$ a residual direction. Then, for any norm $\|\cdot\|$ on $\mathbb{R}^4$,*

$$NS(P_{\epsilon f}) = \epsilon D NS_{P_0}[f] + O(\epsilon^2), \quad (46)$$

and consequently

$$\|NS(P_{\epsilon f})\| = \Theta(\epsilon) \quad \text{when } D NS_{P_0}[f] \neq 0, \quad \|NS(P_{\epsilon f})\| = O(\epsilon^2) \quad \text{when } f \in \ker D NS_{P_0}. \quad (47)$$

This follows immediately from $NS(P_0) = 0$ and Theorem 1; the map $F \mapsto P_F$ is smooth in a neighbourhood of $F = 0$ since $P_0 > 0$. We test the prediction on four residual classes.

The distinction in (47) is local. A residual in $\ker D NS_{P_0}$ suppresses signalling leakage to the second order, but this does not mean that the finite- ϵ tilted box is exactly no-signalling. Exact no-signalling at finite tilt strength is still governed by the nonlinear compatibility Equations (21) and (22).

(i) Generic random residuals

A random $f \in \mathbb{R}^{16}$ has $DN_{P_0}[f] \neq 0$ almost surely, so Proposition 3 predicts $\|NS(P_{\epsilon f})\| = \Theta(\epsilon)$. Throughout, “generic” refers to a property holding outside a proper Lebesgue-null linear subspace of the relevant residual space; for the random class this is the kernel $\ker DN_{P_0}$, which is twelve-dimensional in $\mathbb{R}^{16}$.

(ii) Admissible tangent residuals

We sample $f \in \ker DN_{P_0}$ by drawing a random vector and projecting onto the kernel. The linear term vanishes and Proposition 3 predicts $\|NS(P_{\epsilon f})\| = O(\epsilon^2)$. Such residuals are not exactly no-signalling at finite ϵ ; they are locally admissible.

(iii) Setting-only redundancy residuals

$F_{abxy} = G_{xy}$ produces $P_F = P_0$ exactly. These are not physical deformations of the box but redundancies of the residual representation.

This class is qualitatively different from the admissible tangent class: setting-only residuals vanish exactly after setting-wise normalization, whereas tangent-admissible residuals generally produce non-zero higher-order effects.

(iv) Local-additive residuals

A local-additive residual is

$$F_{abxy} = A_{ax} + B_{by}. \quad (48)$$

The form is local: an Alice-only piece plus a Bob-only piece. Setting-wise exponential reweighting, however, acts on the joint distribution $P_0(a, b|x, y)$, so algebraic locality of the residual is not, by itself, sufficient for no-signalling compatibility of the resulting probability table.

This distinction is central for the interpretation of the examples below. The residual expression may separate into Alice and Bob terms, but the normalization is performed on the correlated joint distribution in each setting block. As a result, a local-additive residual can shift a remote marginal even though it contains no explicit nonlocal term.

Proposition 4 (Local-additive residuals are not no-signalling safe). *Let P_0 be the unbiased Tsirelson CHSH box. The subset of local-additive residuals $f_{abxy} = A_{ax} + B_{by}$ satisfying the linearized no-signalling condition $DN_{P_0}[f] = 0$ is a proper linear subspace of the local-additive subspace; in particular there exist non-zero local-additive residuals with $DN_{P_0}[f] \neq 0$.*

Proof. We exhibit an explicit local-additive element outside $\ker DN_{P_0}$; this implies that the admissible subset is a proper linear subspace of the local-additive directions.

Take $A \equiv 0$ and $B_{0,0} = +1, B_{1,0} = -1, B_{0,1} = B_{1,1} = 0$, so

$$f_{abxy} = B_{by} = \begin{cases} +1 & \text{if } y = 0, b = 0, \\ -1 & \text{if } y = 0, b = 1, \\ 0 & \text{if } y = 1. \end{cases} \quad (49)$$

Using (30), the setting-wise expectation

$$\langle f \rangle_{xy} = \sum_{a,b} P_0(a, b|x, y) B_{by} = \frac{1}{2}(B_{0y} + B_{1y}) = 0 \quad (50)$$

vanishes for both y , so the centred residual at (x, y) is $f_{abxy} - \langle f \rangle_{xy} = B_{by}$. Then, according to (26),

$$(D NS_{P_0}[f])_{N_{A,x}} = - \sum_b P_0(0, b|x, 0) B_{b0} = - [P_0(0, 0|x, 0) - P_0(0, 1|x, 0)] = - \frac{E_{x0}}{2}, \quad (51)$$

where the last equality uses $P_0(0, b|x, y) = \frac{1}{4}[1 + E_{xy}t_b]$ at the unbiased Tsirelson box. With $E_{00} = E_{10} = 1/\sqrt{2}$ this gives

$$(D NS_{P_0}[f])_{N_{A,0}} = (D NS_{P_0}[f])_{N_{A,1}} = -1/(2\sqrt{2}), \quad (D NS_{P_0}[f])_{N_{B,0}} = (D NS_{P_0}[f])_{N_{B,1}} = 0. \quad (52)$$

Hence

$$D NS_{P_0}[f] = \left(-\frac{1}{2\sqrt{2}}, -\frac{1}{2\sqrt{2}}, 0, 0\right) \neq 0, \quad \|D NS_{P_0}[f]\| = \frac{1}{2}. \quad (53)$$

To see that this remote-marginal shift is observable at finite ϵ , direct evaluation of (14) at $\epsilon = 0.01$ gives

$$P_{\epsilon f}(a = 0|x = 0, y = 0) = 0.4964645839, \quad P_{\epsilon f}(a = 0|x = 0, y = 1) = 0.5. \quad (54)$$

The displayed marginal difference is 3.535×10^{-3} , in agreement with the first-order component prediction $\epsilon/(2\sqrt{2}) \approx 3.535 \times 10^{-3}$. The Euclidean norm of the full no-signalling vector is predicted to be $\epsilon\|D NS_{P_0}[f]\| = \epsilon/2 = 5 \times 10^{-3}$.

The CHSH correlator is unchanged at first order, $\text{CHSH}(P_{\epsilon f}) = 2\sqrt{2} + O(\epsilon^2)$: the first-order correction to each individual correlator $E_{xy}^F = \sum_{a,b} s_a t_b P_F(a, b|x, y)$ equals

$$- \sum_{a,b} s_a t_b P_0(a, b|x, y) (B_{by} - \langle f \rangle_{xy}) = - \sum_a s_a P_0(a|x, y) \sum_b t_b B_{by}, \quad (55)$$

using that B_{by} is independent of a and $\langle f \rangle_{xy} = 0$. This vanishes because the unbiased marginal gives

$$\sum_a s_a P_0(a|x, y) = 0. \quad (56)$$

□

The example separates two notions that are easily confused. A residual may be local as an algebraic expression in outcomes and settings while still failing to preserve no-signalling when applied to a non-factorising reference box. The numerical verification below shows that local-additive residuals behave like generic residuals on the no-signalling axis while leaving the CHSH value unchanged at first order.

Operationally, this means that CHSH stability alone is not enough to certify no-signalling stability under residual tilts. A deformation can leave the correlator unchanged at first order while still producing a detectable remote-marginal shift. This is why the no-signalling vector, rather than only the CHSH value, is the central diagnostic in Section 9.

9. Numerical Verification

The following computations are consistency checks of the analytic scaling law; the proofs are in Sections 6 and 8.

The computations are exact-probability calculations, not simulations of loophole-free Bell experiments. Their purpose is to verify the local scaling predicted by Proposition 3 and to illustrate how different residual classes behave under the no-signalling diagnostic.

The Jacobian $D NS_{P_0}$ at the unbiased Tsirelson reference was evaluated both numerically (central differences) and from the closed form derived in Section 5; the two agreed to

$$\|J_{\text{num}} - J_{\text{an}}\|_F = 1.86 \times 10^{-10}. \quad (57)$$

The singular values of the analytic Jacobian were $\{\sqrt{5}/4, \sqrt{5}/4, 1/4, 1/4\}$, matching Proposition 1.

For the scaling test we drew 200 unit-norm residual directions from each of the four classes and evaluated $\|NS(P_{\epsilon f})\|$ on the geometric grid

$$\epsilon \in \{10^{-4}, 3 \times 10^{-4}, 10^{-3}, \dots, 10^{-1}\}. \quad (58)$$

For each realisation we fitted

$$\log_{10} \|NS(P_{\epsilon f})\| = m \log_{10} \epsilon + c \quad (59)$$

on the small- ϵ portion $\epsilon \leq 10^{-2}$ to avoid the non-asymptotic regime; slopes were averaged across realisations. The results are summarized in Table 1.

Table 1. Fitted log–log slopes of $\|NS(P_{\epsilon f})\|$, averaged over 200 realisations per class. The fit uses the small- ϵ window $\epsilon \leq 10^{-2}$. Redundancy slopes are not informative because $\|NS\|$ is at machine precision throughout.

Class	Predicted Slope	Fitted Slope (Mean $\pm$ Std)
generic random	1	1.0000 ± 0.0001
admissible tangent	2	2.0000 ± 0.0003
local-additive	1	1.0000 ± 0.0000
setting-only redundancy	—	$\ NS\ \sim 3 \times 10^{-17}$ at all ϵ

A mean sweep over the same 200 realisations is shown in Table 2.

Table 2. Mean $\|NS(P_{\epsilon f})\|$ across 200 realisations per class at the unbiased Tsirelson reference. The CHSH column reports the mean CHSH value of $P_{\epsilon f}$ ($\text{CHSH}(P_0) = 2\sqrt{2} \approx 2.828427$).

ϵ	$\ NS\ _{\text{rand}}$	$\text{CHSH}_{\text{rand}}$	$\ NS\ _{\text{adm}}$	CHSH_{adm}	$\ NS\ _{\text{loc}}$	CHSH_{loc}	$\ NS\ _{\text{red}}$
10^{-4}	2.09×10^{-5}	2.828428	8.59×10^{-11}	2.828427	1.22×10^{-5}	2.828427	$\sim 10^{-17}$
10^{-3}	2.09×10^{-4}	2.828435	8.59×10^{-9}	2.828428	1.22×10^{-4}	2.828427	$\sim 10^{-17}$
10^{-2}	2.09×10^{-3}	2.828503	8.59×10^{-7}	2.828431	1.22×10^{-3}	2.828427	$\sim 10^{-17}$
10^{-1}	2.09×10^{-2}	2.829003	8.60×10^{-5}	2.828044	1.22×10^{-2}	2.828437	$\sim 10^{-17}$

The fitted slopes match the predictions of Proposition 3 to four decimal places. Two features deserve comment.

First, redundancy residuals produce no-signalling residuals at the level of machine precision ($\sim 3 \times 10^{-17}$), confirming that they leave the box invariant exactly, not only to first order.

Second, the CHSH columns expose the structural difference between random and local-additive residuals. Random residuals shift CHSH at first order in ϵ , while local-additive residuals leave CHSH invariant up to $O(\epsilon^2)$, even though their no-signalling leakage is itself $\Theta(\epsilon)$ and indistinguishable in slope from the random class. This confirms Proposition 4: local-additive residuals preserve correlators but break no-signalling at first order through the remote marginals.

This is the numerical counterpart of the local-additive counterexample: a deformation can be nearly invisible in the CHSH value while still being visible in the marginal no-signalling diagnostic. The no-signalling norm therefore captures information that the CHSH scalar alone does not.

The corresponding log–log plot is shown in Figure 1.




Figure 1. No-signalling leakage $\|NS(P_{\epsilon f})\|$ for the four residual classes at the unbiased Tsirelson reference, averaged over 200 realisations. The dashed and dotted reference lines have slope 1 and 2. Generic random and local-additive residuals follow slope 1 asymptotics; admissible tangent residuals follow slope 2; setting-only redundancy residuals are at machine precision. The fitted slopes reported in Table 1 are computed on the small- ϵ window $\epsilon \leq 10^{-2}$.

10. Sampling Noise

The scaling laws above are statements about exact-probability boxes. This paper does not analyse loophole-free Bell data; it studies the analytic structure of exponential tilts. Nevertheless, the finite-shot detectability of the two regimes is worth noting.

This section should therefore be read only as a detectability estimate for marginal no-signalling leakage, not as an experimental analysis of Bell-test data. The quantities considered here are empirical deviations from exact-probability tables under finite sampling.

Suppose N shots are taken for each setting pair (x, y) . Even if the exact box is perfectly no-signalling, binomial fluctuations in the empirical marginals give

$$\|NS(\hat{P})\| = O(N^{-1/2}). \quad (60)$$

A generic residual produces exact leakage of order ϵ , which becomes visible above the sampling floor when

$$\epsilon \gg N^{-1/2}, \quad \text{that is,} \quad N \gg \epsilon^{-2}. \quad (61)$$

An admissible tangent residual produces only $O(\epsilon^2)$ exact leakage; resolving it against the same floor requires

$$\epsilon^2 \gg N^{-1/2}, \quad \text{that is,} \quad N \gg \epsilon^{-4}. \quad (62)$$

Admissible tangent residuals are therefore much harder to detect through finite-shot no-signalling leakage. Finite-shot tests should be read primarily as recoverability tests for sufficiently large residual signals, not as the cleanest evidence for the differential scaling of Proposition 3.

The practical message is that first-order signalling leakage is statistically much easier to detect than second-order leakage. This supports the use of the linearized no-signalling Jacobian as a diagnostic for identifying residual directions that are dangerous at leading order.

11. Tsirelson and Quantum-Admissibility Caution

No-signalling is necessary for causal admissibility but not sufficient for quantum realizability. The set of quantum correlations is a proper subset of the no-signalling set [5,7,10,11]. A tilted box may preserve no-signalling, or preserve it to first order, while moving outside the set of boxes obtainable from measurements on a single quantum state.

This distinction is essential for the interpretation of the results. The present paper studies probability-level transformations of CHSH tables. It does not assert that every admissible residual tilt can be implemented by a quantum state and fixed measurement operators. Nor does it claim that a no-signalling tilt is automatically allowed by quantum mechanics.

This is visible in the CHSH columns of Table 2. Random residual directions can move the table above $2\sqrt{2}$. Such a value is not a physical violation of Tsirelson's theorem; the tilted boxes are setting-wise reweightings of probability tables and need not correspond to a quantum state with a fixed family of measurement operators. They are probability-level deformations, not automatic quantum realizations. The implication chain is

$$\text{quantum realizable} \implies \text{no-signalling}, \quad (63)$$

but not conversely.

The present paper studies the no-signalling part of this implication chain. A further quantum-admissibility analysis would require imposing additional constraints, for example through the geometry of the quantum set or through hierarchy-based tests such as the NPA hierarchy. That problem is outside the scope of the present first-order no-signalling analysis.

12. Conclusions

We have given a first-order admissibility analysis of exponential residual tilts of CHSH boxes. The four main results are as follows:

- (1) The residual representation has a four-dimensional setting-only redundancy $\mathcal{G}$, in which different residual arrays describe the same probability box;
- (2) The linearized no-signalling Jacobian at the unbiased Tsirelson box has the closed-form Gram matrix (33) with singular values $\{\sqrt{5}/4, \sqrt{5}/4, 1/4, 1/4\}$, hence rank 4 and kernel of dimension 12; the gauge-free admissible space has dimension $\dim(\ker D NS_{P_0}/\mathcal{G}) = 8$, recovering the eight-parameter no-signalling description from a residual-coordinate perspective;
- (3) The scaling dichotomy $\|NS(P_{ef})\| = \Theta(\epsilon)$ for generic f and $\|NS(P_{ef})\| = O(\epsilon^2)$ for $f \in \ker D NS_{P_0}$ holds exactly and is verified to four decimals on 200 realisations per class;
- (4) Local-additive residuals $f_{abxy} = A_{ax} + B_{by}$ are not automatically no-signalling safe; an explicit non-zero counterexample shifts a remote marginal at first order while leaving the CHSH correlator invariant.

These results should be read as a local no-signalling admissibility criterion for exponential tilts, not as a full classification of all exactly no-signalling tilts. The exact nonlinear admissibility problem is given by the compatibility Equations (21) and (22); the Jacobian analysis identifies the tangent directions that do not create first-order signalling leakage.

We do not claim that Born statistics are experimentally false, that every no-signalling residual tilt is quantum-realizable, or that the CHSH value can physically exceed $2\sqrt{2}$. The probability boxes studied here are mathematical exponential tilts of a reference box; they are not, by themselves, loophole-free Bell experiments or guaranteed quantum realizations. Exponential tilts of Bell boxes require additional admissibility constraints if they are to be interpreted physically.

The main conceptual lesson is that setting-wise normalization is not enough in a Bell scenario. A transformation can be locally well defined inside each measurement context and still fail to preserve the cross-setting marginal constraints. Conversely, a deformation can leave the CHSH value nearly unchanged while producing first-order leakage in the no-signalling vector. For this reason, no-signalling preservation is a separate admissibility condition that must be checked directly.

The natural next steps are nonlinear classification of exactly no-signalling residuals; the identification of directions simultaneously tangent to the no-signalling and quantum sets at the Tsirelson box; and studying the structured residuals $F(W)$ generated by a comparison geometry rather than arbitrary arrays F_{abxy} .

A further direction is to relate the residual-coordinate criterion to device-independent and resource-theoretic settings in which Bell probability tables are reweighted, post-processed, or used as diagnostic objects. Such applications would require combining the present no-signalling analysis with additional quantum-admissibility or protocol constraints.

Funding: The author received no specific funding for this work.

Data Availability Statement: No empirical data were generated or analysed in this work. The numerical illustration in Section 9 is fully specified by the deterministic construction in Appendix A and reproduces the reported values when implemented with any standard numerical linear-algebra library. The numerical checks reported in this manuscript are fully specified by the finite-dimensional formulas stated in Appendix A, together with the Tsirelson reference box and the fixed random seed 20260428. No external data are required.

Acknowledgments: The author is the sole author of this manuscript. The text occasionally uses “we” in the conventional mathematical sense, to include the reader in the development of the argument. AI-assisted writing and editing tools were used to help organize draft material, improve clarity, and check the consistency of notation. The author reviewed and takes full responsibility for all mathematical claims, interpretations, and final manuscript content.

Conflicts of Interest: The author declares no competing interests, financial or non-financial, related to the content of this manuscript.

Abbreviations

Symbol	Type/Domain	Meaning/Assumptions
Basic indices and scalars		
$\mathbb{R}$	field	Real numbers.
ϵ	$\mathbb{R}_{>0}$	Small perturbation parameter in $F = \epsilon f$.
x, y	$\{0, 1\}$	Measurement settings for Alice and Bob.
a, b	$\{0, 1\}$	Measurement outcomes for Alice and Bob.
s_a, t_b	$\{\pm 1\}$	Sign variables, with $s_0 = t_0 = +1$ and $s_1 = t_1 = -1$.
N	integer	Number of finite-shot samples per setting pair, used only in the sampling-noise discussion.
CHSH boxes		
$P(a, b x, y)$	probability box	Conditional probability of outcomes (a, b) given settings (x, y) .
P_0	probability box	Strictly positive reference CHSH box. Usually assumed normalized and no-signalling.
P_F	probability box	Exponentially tilted box obtained from P_0 using residual array F .
$P_{\epsilon f}$	probability box	Tilted box with residual $F = \epsilon f$, used for linearized analysis.

Symbol	Type/Domain	Meaning/Assumptions
Q	probability box	Generic strictly positive normalized box used in residual-coordinate parametrization. Alice marginal:
$P_A(a x, y)$	marginal	$P_A(a x, y) = \sum_b P(a, b x, y).$
		Bob marginal:
$P_B(b x, y)$	marginal	$P_B(b x, y) = \sum_a P(a, b x, y).$
Exponential tilts and residuals		
F	$\mathbb{R}^{16}$	Residual array $F = (F_{abxy})$, one real value for each outcome-setting tuple.
f	$\mathbb{R}^{16}$	Infinitesimal residual direction, with $F = \epsilon f$.
F_{abxy}	scalar	Residual cost assigned to outcome (a, b) under settings (x, y) . Setting-wise normalization factor:
$Z_{xy}(F)$	$\mathbb{R}_{>0}$	$Z_{xy}(F) = \sum_{a,b} P_0(a, b x, y) e^{-F_{abxy}}.$
G_{xy}	scalar	Setting-only residual. If $F_{abxy} = G_{xy}$, the tilt leaves the box unchanged. Local-additive residual components in
A_{ax}, B_{by}	scalars	$F_{abxy} = A_{ax} + B_{by}.$
		Setting-wise P_0 expectation:
$\langle f \rangle_{xy}$	scalar	$\langle f \rangle_{xy} = \sum_{a,b} P_0(a, b x, y) f_{abxy}.$
$f_{abxy} - \langle f \rangle_{xy}$	centered residual	The only part of f that affects the first-order tilt at setting (x, y) . Setting-only redundancy:
$\mathcal{G}$	subspace of $\mathbb{R}^{16}$	$\mathcal{G} = \{F \in \mathbb{R}^{16} : F_{abxy} = G_{xy} \text{ for some } G \in \mathbb{R}^4\}.$
$[F]$	element of $\mathbb{R}^{16}/\mathcal{G}$	Residual equivalence class modulo setting-only redundancy.
No-signalling constraints		
		No-signalling vector:
$NS(P)$	$\mathbb{R}^4$	$NS(P) = (N_{A,0}, N_{A,1}, N_{B,0}, N_{B,1})^\top.$
		Alice-side signalling component:
$N_{A,x}(P)$	scalar	$N_{A,x}(P) = \sum_b P(0, b x, 0) - \sum_b P(0, b x, 1).$

Symbol	Type/Domain	Meaning/Assumptions
		Bob-side signalling component:
$N_{B,y}(P)$	scalar	$N_{B,y}(P) = \sum_a P(a, 0 0, y) - \sum_a P(a, 0 1, y).$
$\ NS(P)\ $	scalar	Size of no-signalling violation; Euclidean norm unless stated otherwise.
$D NS_{P_0}$	linear map $\mathbb{R}^{16} \rightarrow \mathbb{R}^4$	Linearized no-signalling Jacobian at P_0 in residual coordinates.
$\ker D NS_{P_0}$	subspace of $\mathbb{R}^{16}$	Residual directions with no first-order no-signalling leakage.
		Local admissible tangent space:
$T_{\text{adm}}(P_0)$	subspace	$T_{\text{adm}}(P_0) = \ker D NS_{P_0}.$
$\ker D NS_{P_0} / \mathcal{G}$	quotient space	Gauge-free locally admissible residual space; dimension 8 at the unbiased Tsirelson box.
Tsirelson reference box and CHSH quantities		
		Correlator of a box P :
$E_{xy}(P)$	correlator	$E_{xy}(P) = \sum_{a,b} s_a t_b P(a, b x, y).$
		Unbiased Tsirelson box, used as the main special choice of P_0 :
P_{Ts}	probability box	$P_{\text{Ts}}(a, b x, y) = \frac{1}{4} (1 + E_{xy} s_a t_b).$
		Correlators of the chosen Tsirelson reference:
$E_{00}, E_{01}, E_{10}, E_{11}$	scalars	$E_{00} = E_{01} = E_{10} = \frac{1}{\sqrt{2}}, \quad E_{11} = -\frac{1}{\sqrt{2}}.$
		CHSH value:
$\text{CHSH}(P)$	scalar	$\text{CHSH}(P) = E_{00}(P) + E_{01}(P) + E_{10}(P) - E_{11}(P).$
2	scalar	Bell-local CHSH bound for local hidden-variable boxes.
$2\sqrt{2}$	scalar	Tsirelson bound for quantum-realizable CHSH boxes.
		Singular values of $D NS_{P_0}$. At the unbiased Tsirelson box:
σ_i	$\mathbb{R}_{\geq 0}$	$\left\{ \frac{\sqrt{5}}{4}, \frac{\sqrt{5}}{4}, \frac{1}{4}, \frac{1}{4} \right\}.$
$\text{rank}(D NS_{P_0})$	integer	Rank of the no-signalling Jacobian; equals 4 at the unbiased Tsirelson box.
Residual classes and diagnostics		
generic random residual	class	Random unit vector $f \in \mathbb{R}^{16}$; generically gives first-order no-signalling leakage.

Symbol	Type/Domain	Meaning/Assumptions
admissible tangent residual	class	Unit vector projected onto $\ker D NS_{P_0}$; first-order leakage vanishes.
setting-only redundancy residual	class	Residual of the form $F_{abxy} = G_{xy}$; gives exactly $P_F = P_0$.
local-additive residual	class	Residual of the form $F_{abxy} = A_{ax} + B_{by}$; local in form but not generically no-signalling safe.
		Fitted log-log slope in
m	scalar	$\log_{10} \ NS(P_{ef})\ = m \log_{10} \epsilon + c.$
c	scalar	Intercept in the log-log fit.
$\ \cdot\ $	norm	Euclidean norm on $\mathbb{R}^4$ or $\mathbb{R}^{16}$ unless stated otherwise.
$\ \cdot\ _F$	norm	Frobenius norm, used for comparing analytic and numerical Jacobians.
$O(\epsilon)$	asymptotic order	Generic first-order no-signalling leakage.
$O(\epsilon^2)$	asymptotic order	Quadratic no-signalling leakage for tangent-admissible residuals.
$O(N^{-1/2})$	asymptotic order	Finite-shot empirical marginal noise floor.

Appendix A. Closed Form of the Jacobian $D NS_{P_0}$

For reproducibility, we record the full 4×16 matrix $D NS_{P_0}$ at the unbiased Tsirelson reference.

We order the residual columns block-wise by

$$\text{idx}(a, b, x, y) = 4(2x + y) + 2a + b. \quad (\text{A1})$$

Thus the four columns belonging to a fixed setting pair (x, y) appear consecutively, in the outcome order

$$(a, b) = (0, 0), (0, 1), (1, 0), (1, 1).$$

This is the ordering used in the block matrix below.

For the setting-local block structure derived in Section 5, each row of $D NS_{P_0}$ has support on exactly two of the four (x, y) -blocks. Within block (x, y) , the entries take the closed form

$$J_{a'b'}^{(x,y)} = \sigma_{xy} \cdot \begin{cases} -\frac{1 + E_{xy}t_{b'}}{8}, & a' = 0, \\ +\frac{1 - E_{xy}t_{b'}}{8}, & a' = 1, \end{cases} \quad (\text{A2})$$

where $t_0 = +1$, $t_1 = -1$, and the sign σ_{xy} is set by the no-signalling row: for the $N_{A,x}$ row, $\sigma_{xy} = +1$ if $y = 0$ and -1 if $y = 1$ (and the row is supported only on blocks with the matching x); for the $N_{B,y}$ row, the analogous formula with the roles of the parties exchanged gives $\sigma_{xy} = +1$ if $x = 0$ and -1 if $x = 1$ (with the row supported only on blocks with the matching y).

At $E_{xy} = \pm 1/\sqrt{2}$, the four numerical entry magnitudes per block are $\{(1 + 1/\sqrt{2})/8, (1 - 1/\sqrt{2})/8\}$ each appearing twice, i.e., approximately $\{0.21339, 0.03661\}$. Substituting into (A2) and assembling block by block produces the full Jacobian

$$D NS_{P_0} = \frac{1}{8\sqrt{2}} M, \quad (\text{A3})$$

where M is the integer-coefficient matrix

$$M = \begin{matrix} & \begin{matrix} \text{block (0,0)} & \text{block (0,1)} & \text{block (1,0)} & \text{block (1,1)} \end{matrix} \\ \begin{matrix} N_{A,0} \\ N_{A,1} \\ N_{B,0} \\ N_{B,1} \end{matrix} & \begin{pmatrix} (-\alpha, -\beta, +\beta, +\alpha) & (+\alpha, +\beta, -\beta, -\alpha) & 0 & 0 \\ 0 & 0 & (-\alpha, -\beta, +\beta, +\alpha) & (+\beta, +\alpha, -\alpha, -\beta) \\ (-\alpha, +\beta, -\beta, +\alpha) & 0 & (+\alpha, -\beta, +\beta, -\alpha) & 0 \\ 0 & (-\alpha, +\beta, -\beta, +\alpha) & 0 & (+\beta, -\alpha, +\alpha, -\beta) \end{pmatrix} \end{matrix}$$

with $\alpha = \sqrt{2} + 1$ and $\beta = \sqrt{2} - 1$, and where each block of four numbers lists the entries in column order

$$(a', b') = (0, 0), (0, 1), (1, 0), (1, 1).$$

Direct substitution gives the entries of Section 6 and reproduces the Gram matrix (33).

The block-wise ordering in (A1) is used only for displaying the matrix. Any other column ordering gives the same linear map after the corresponding permutation of columns, and therefore the same Gram matrix $D NS_{P_0} (D NS_{P_0})^\top$, rank, and singular values.

Reproducible verification.

The matrix above is also produced by the analytic Jacobian routine in the supplementary `numerics.py`, agreeing with the numerical central-difference Jacobian to Frobenius distance 1.86×10^{-10} . A self-contained check is, in pseudocode ,

```
# Block-wise residual column order:
# idx(a,b,x,y) = 4*(2*x + y) + 2*a + b

P0[a,b,x,y] = (1 + E[x,y] * s[a] * t[b]) / 4          # Tsirelson box
J_an[k, idx(a',b',x',y')] =                          # analytic Jacobian
    sum over (a,b,x,y) appearing in NS_k of
    delta_{x,x'} delta_{y,y'} * P0[a,b,x,y] *
    (P0[a',b',x,y] - delta_{a,a'} delta_{b,b'})
J_num[k, c] = (NS(P_{+eps e_c}) - NS(P_{-eps e_c})) / (2*eps)
assert ||J_an - J_num||_F < 1e-9
G = J_an @ J_an.T
assert 32 * G == [[6, 0, 2*sqrt(2), -2*sqrt(2)],
                  [0, 6, -2*sqrt(2), -2*sqrt(2)],
                  [2*sqrt(2), -2*sqrt(2), 6, 0],
                  [-2*sqrt(2), -2*sqrt(2), 0, 6]]      # up to floating point
assert sorted(eigvals(G)) == [1/16, 1/16, 5/16, 5/16]
```

The full implementation, including the scaling sweep of Section 9, is in `numerics.py` and `gram_analysis.py`; both run from a fixed seed (20260428) and reproduce all numerical values reported in this paper.

References

1. Jaynes, E.T. Information theory and statistical mechanics. *Phys. Rev.* **1957**, *106*, 620–630. [CrossRef]
2. Bell, J.S. On the Einstein–Podolsky–Rosen paradox. *Phys. Phys. Fiz.* **1964**, *1*, 195–200. [CrossRef]

3. Clauser, J.F.; Horne, M.A.; Shimony, A.; Holt, R.A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **1969**, *23*, 880–884. [[CrossRef](#)]
4. Popescu, S.; Rohrlich, D. Quantum nonlocality as an axiom. *Found. Phys.* **1994**, *24*, 379–385. [[CrossRef](#)]
5. Brunner, N.; Cavalcanti, D.; Pironio, S.; Scarani, V.; Wehner, S. Bell nonlocality. *Rev. Mod. Phys.* **2014**, *86*, 419–478. [[CrossRef](#)]
6. Arbel E.; Israel, N.; Belgorodsky, M.; Shafrir, Y.; Maslennikov, A.; Gandelman, S.P.; Rozenman, G.G. Optical emulation of quantum state tomography and bell test—A novel undergraduate experiment. *Results Opt.* **2025**, *21*, 100847. [[CrossRef](#)]
7. Tsirelson, B.S. Quantum generalizations of bell’s inequality. *Lett. Math. Phys.* **1980**, *4*, 93–100.
8. Barrett, J. Nonlocal correlations as an information-theoretic resource. *Phys. Rev. A* **2005**, *71*, 022101. [[CrossRef](#)]
9. Piani, M.; Horodecki, M.; Horodecki, P.; Horodecki, R. Properties of quantum nonsignaling boxes. *Phys. Rev. A* **2006**, *74*, 012305. [[CrossRef](#)]
10. Navascués, M.; Pironio, S.; Acín, A. Bounding the set of quantum correlations. *Phys. Rev. Lett.* **2007**, *98*, 010401. [[CrossRef](#)] [[PubMed](#)]
11. Navascués, M.; Pironio, S.; Acín, A. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.* **2008**, *10*, 073013. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.