Understanding the Bell Inequality

Quantum Foundations Series

PakCrypt Research Division
https://pakcrypt.org

Abstract. This tutorial is a self-contained guide to Bell's theorem, quantum entanglement, and device-independent randomness, written for students with a basic background in linear algebra and probability. We derive the CHSH bound algebraically, explain why quantum mechanics violates it, give a detailed treatment of entanglement, and discuss the landmark 2026 ETH Zurich demonstration of certified perfect randomness as a concrete application.

Keywords: Bell inequality · CHSH bound · quantum entanglement · singlet state · QRNG · randomness amplification · device-independent cryptography

1 What Bell's Inequality Is About

Bell's inequality is a mathematical test for whether nature can be explained by *local realism*—a classical picture of reality built on two interlocking assumptions.

1.1 The Two Pillars of Local Realism

Realism. Measurement outcomes are already determined by hidden properties of the particles, even before any measurement is made. Each particle carries a secret "instruction card" written at the moment the pair was created.

Locality. What happens at one location cannot instantly affect a distant location. No influence can travel faster than light. If Alice makes a measurement in her lab, it cannot change anything in Bob's lab across the room—let alone across the galaxy.

Bell showed that if both assumptions hold simultaneously, certain experimental correlations must obey a mathematical bound. In the most widely used form, the *CHSH inequality* [2] (named after Clauser, Horne, Shimony, and Holt), that bound is:

$$|S| \leq 2, \tag{1}$$

where

$$S = E(a, b) + E(a, b') + E(a', b) - E(a', b'). \tag{2}$$

Each $E(x, y)$ is a *correlation coefficient* between outcomes measured on the two particles when Alice uses detector setting x and Bob uses detector setting y. Outcomes are labelled $+1$ or -1 (e.g. spin-up or spin-down). A correlation of $+1$ means the two detectors always agree; -1 means they always disagree; 0 means no correlation at all.

1.2 A Primer: Measurements, Spin, and Bases

Key Concept: What Is a Measurement in Quantum Mechanics?

In classical physics, measuring a property simply reads out a pre-existing value. A ball's velocity is what it is, regardless of how you look at it.

In quantum mechanics, measurement is an *active, irreversible* process. Before measurement, a quantum particle such as a photon or electron can be in a *superposition* of several possible outcomes. The act of measurement forces the system into one definite result—chosen at random with probabilities prescribed by the quantum state.

Photon polarisation and electron spin are the most common properties used in Bell experiments. Both are two-valued: a measurement along any chosen axis yields either $+1$ or -1, never anything in between.

The direction along which you choose to measure is called the *measurement basis*. If you rotate your detector, you are choosing a different basis. Crucially, a particle that was spin-up along the vertical axis has a well-defined probability of being measured spin-up along a tilted axis—and that probability depends on the *angle* between the two axes. No pre-written value can account for all possible angles simultaneously, as Bell's theorem proves.

For photons, polarisation is the relevant property. A photon can be horizontally polarised, vertically polarised, or—in a quantum superposition—both at once. A polariser (the detector) is oriented at some angle θ; it transmits the photon with probability $\cos^2\theta$ relative to the photon's polarisation direction, and blocks it otherwise. The outcome is binary: pass $(+1)$ or block (-1).

If you choose the wrong basis—for example, trying to detect a diagonally polarised photon with a horizontal/vertical polariser—the outcome is completely random with equal probability. The information about the original diagonal polarisation is irretrievably lost. This is not a deficiency in your equipment; it is a fundamental feature of quantum mechanics.

2 Why the Classical Limit Is 2

The bound $|S| \leq 2$ follows from a simple algebraic argument. Assume each particle carries predetermined answers for all four possible measurement settings: a, a', b, b'. Call the hidden outcomes $A, A' \in \{-1, +1\}$ for Alice's settings and $B, B' \in \{-1, +1\}$ for Bob's settings.

2.1 The Algebraic Core

Consider the quantity:

$$S^* = AB + AB' + A'B - A'B'. \tag{3}$$

Factoring gives:

$$S^* = A(B + B') + A'(B - B'). \tag{4}$$

Since B and B' each take values in $\{-1, +1\}$, exactly one of two cases holds:

- **Case 1:** $B = B'$. Then $B + B' = \pm 2$ and $B - B' = 0$, giving $S^* = \pm 2A$ and $|S^*| = 2$.
- **Case 2:** $B = -B'$. Then $B + B' = 0$ and $B - B' = \pm 2$, giving $S^* = \pm 2A'$ and $|S^*| = 2$.

In every possible case $|S^*| = 2$. Since the measured correlation $E(a, b)$ is the average of AB over many trials, and the average of a quantity bounded by 2 in magnitude is also bounded by 2, we obtain $|S| \leq 2$ for any local hidden-variable theory.

Why Randomness Cannot Save Local Realism

One might think: what if the hidden variables are random? Does that create a loophole?

No. The argument above holds for *every single trial*, whatever values the hidden variables take. Averaging over any distribution of hidden-variable values preserves $|S| \leq 2$. The bound is a structural consequence of pre-assignment plus locality, not an artefact of any particular distribution.

The limit of 2 is therefore not about noise, imprecision, or incomplete knowledge. It is a ceiling written into any theory where results are both pre-assigned on the particle and locally independent of the distant detector's orientation.

3 Why Quantum Mechanics Can Exceed 2

Quantum mechanics does not treat measurement outcomes as pre-written answers read off a hidden card. For entangled states, the correlation between results at two distant detectors is determined by the *relative angle* between those detectors—a continuous, geometric quantity that no fixed hidden card can reproduce for all angles at once.

3.1 The Quantum Correlation Formula

For the spin singlet state (Sect. 5), the correlation between Alice's and Bob's outcomes is:

$$E(a, b) = -\cos\theta_{ab}, \tag{5}$$

where θ_{ab} is the angle between Alice's detector direction a and Bob's detector direction b. This smooth cosine dependence is the fingerprint of quantum entanglement.

3.2 The Optimal CHSH Violation

Choosing the four detector angles $a = 0°$, $b = 45°$, $a' = 90°$, $b' = 135°$ (equivalently $a = 0°$, $a' = 45°$, $b = 22.5°$, $b' = 67.5°$), quantum mechanics gives:

$$|S| = 2\sqrt{2} \approx 2.828. \tag{6}$$

This is the *Tsirelson bound*—the maximum Bell violation quantum mechanics allows. It exceeds the classical limit of 2 by a factor of $\sqrt{2}$.

Why Can't Local Realism Match the Cosine?

Imagine writing all four answers A, A', B, B' on a card before any measurement. The algebraic argument in Sect. 2 shows that any such card must give $|S^*| = 2$ exactly, not $2\sqrt{2}$.

The cosine curve $-\cos\theta$ is not piecewise constant—it bends smoothly. A hidden-variable card must commit to discrete ± 1 values for each angle. No matter how cleverly you arrange the ± 1 values on the card, you cannot reproduce the smooth cosine correlation for *all* possible angle choices simultaneously. There is a fundamental mismatch between the continuous geometry of quantum correlations and the discrete combinatorics of pre-assigned outcomes.

Experiments—most notably the loophole-free Bell tests of 2015 [3,4,5]—have confirmed Bell violations reaching close to $2\sqrt{2}$, ruling out local hidden-variable theories at high statistical confidence.

4 What Entanglement Is Doing

Entanglement is a property of the *joint* quantum state of two or more particles. It means the system cannot be fully described as two independent single-particle states.

4.1 Product States vs. Entangled States

A *product state* looks like:

$$|\Psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle. \tag{7}$$

Here particle 1 is in state $|\psi_1\rangle$ and particle 2 is in state $|\psi_2\rangle$, entirely independently. Measuring particle 1 tells you nothing new about particle 2. A product state is a joint state, but not an entangled one.

An entangled state *cannot* be written this way. The two particles share a single, inseparable quantum description. The Bell singlet state is the canonical example:

$$|\Psi^-\rangle = \frac{1}{\sqrt{2}}\big(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle\big). \tag{8}$$

There is no way to factor this into a state for particle 1 times a state for particle 2. The pair is genuinely one system, despite its parts potentially being light-years apart.

4.2 Why Entanglement Produces Stronger Correlations

When Alice measures her particle along direction a and gets $+1$, the joint state collapses to one where Bob's particle will give -1 if he measures along the same direction. But this does not allow faster-than-light signalling: Bob's outcomes still look completely random to him until he compares notes with Alice through a classical channel.

Why Can't We Simulate Entanglement Classically?

Suppose Alice and Bob meet beforehand, agree on a strategy (a hidden-variable model), and then travel to opposite ends of the lab. Can their pre-agreed strategy reproduce quantum correlations?

Bell's theorem proves: *no*. No pre-agreed strategy—however clever, however randomised—can produce correlations that violate $|S| \leq 2$. The quantum world generates correlations that are simply outside the reach of any classical coordination scheme.

Entanglement is not a hidden telephone line. It does not transmit signals. It creates a new kind of statistical dependency—one that is stronger than any classical correlation but still unable to carry information.

4.3 How Is Entanglement Created?

Entanglement arises whenever two quantum systems interact such that their joint state cannot be factored. Common laboratory methods include:

- **Spontaneous Parametric Down-Conversion (SPDC).** A laser photon passes through a nonlinear crystal and splits into two entangled photons of lower energy. This is the workhorse of photonic Bell experiments.
- **Beam-splitter interference.** Two photons entering a beam splitter from opposite sides can emerge in an entangled superposition.
- **Controlled quantum gates.** In superconducting or trapped-ion processors, a CNOT gate applied to two qubits produces entanglement from a product state.
- **Atomic emissions.** Some atomic transitions emit two photons in a cascade, entangled in polarisation.

4.4 Is Entanglement Only for Photons?

No. Entanglement is a general feature of quantum mechanics. It has been demonstrated in electrons (spin entanglement in solid-state devices), atoms and ions

(trapped-ion quantum computers), molecules, mechanical oscillators visible under a microscope, and superconducting circuits—the platform used in the 2026 ETH Zurich experiment discussed in Sect. 7.

There is no theoretical upper limit on the mass of an entangled system, though *decoherence*—the loss of quantum properties due to interaction with the environment—makes entanglement increasingly fragile as systems grow larger and warmer.

4.5 Why Are Spins Opposite in the Singlet State?

The singlet state has total spin zero. Angular momentum is conserved when the pair is created, so the two spins must add to zero: if one is up, the other must be down. However—and this is the quantum twist—neither spin is up or down *until measured*. Both are in superposition. The correlation is not the result of pre-existing opposite spins; it is established at the moment of measurement, as prescribed by the joint quantum state.

This is precisely what confounds classical intuition and what Bell's theorem quantifies.

5 Singlet State versus Joint State

These two terms are often confused. Let us be precise.

5.1 The Joint State

A *joint state* is any quantum state of a multi-particle system. It lives in the tensor-product Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2$. Every two-particle quantum state—entangled or not—is a joint state.

5.2 The Singlet State

The singlet state (8) is one specific joint state of two spin-$\frac{1}{2}$ particles. It is special for two reasons. First, the total spin is exactly zero, so the pair carries no net angular momentum. Second, it is *rotationally invariant*: no matter which direction you choose to measure, the statistics look the same. This symmetry yields the clean cosine correlation $E(a, b) = -\cos\theta$ and makes the singlet the ideal state for Bell experiments.

The hierarchy is:

- Every singlet state is a joint state.
- Not every joint state is a singlet.
- Not every joint state is entangled. (A product state $|\uparrow\rangle|\downarrow\rangle$ is a joint state but not entangled.)

Can More Than Two Particles Be Entangled?

Yes. Multi-particle entanglement is well established and is crucial for quantum computing.

GHZ states (Greenberger–Horne–Zeilinger) entangle three or more particles. *Cluster states* and *graph states* entangle many qubits in structured patterns that underpin measurement-based quantum computing. Quantum error-correcting codes use entanglement across dozens of physical qubits to protect one logical qubit.

As of 2026, quantum processors with hundreds of entangled qubits have been demonstrated, though maintaining coherence across all of them simultaneously remains a major engineering challenge.

6 Why Bell Violation Matters Physically

Observing a Bell violation—experimentally measuring $|S| > 2$—has a precise and profound implication: the world cannot be explained by a theory that combines pre-existing outcomes, locality, and classical hidden variables. At least one of these must be abandoned.

6.1 What Bell Violation Does Not Mean

- It does *not* mean faster-than-light signalling is possible. Alice and Bob's individual outcomes are still locally random; no information is transmitted by the correlations alone.
- It does *not* pick a unique alternative theory. Quantum mechanics is the current best description, but other non-local or non-realist interpretations (Bohmian mechanics, many-worlds, relational QM) also reproduce the violations—they simply give up locality or realism in different ways.

6.2 What Bell Violation Does Mean

It means the classical picture of reality—particles with pre-written properties, no faster-than-light influences—is provably incomplete. Quantum correlations are real, verified, and stronger than any classical model can achieve.

This has moved from philosophical speculation to engineering fact: Bell violations are now the foundation of device-independent cryptography and certified randomness generation (Sect. 7).

6.3 Quantum Computers: Superposition and Entanglement

Is Quantum Computing Just Classical Computing Without Entanglement?

Essentially, yes. A quantum circuit with no entanglement can be efficiently simulated on a classical computer. The exponential advantage of quantum computing derives from the combination of superposition *and* entanglement.

Superposition allows a qubit to represent $|0\rangle$ and $|1\rangle$ simultaneously. Entanglement correlates qubits in ways that cannot be decomposed qubit-by-qubit. Together, n entangled qubits can represent 2^n amplitudes simultaneously—a Hilbert space that grows *exponentially* with n.

Quantum algorithms such as Shor's factoring algorithm and Grover's search algorithm exploit this structure to solve certain problems exponentially or quadratically faster than the best known classical methods.

Without entanglement, a quantum computer degrades to a probabilistic classical computer. Entanglement is the resource that makes quantum speedup possible.

7 Bell Inequality and Quantum Random Number Generators

Not every quantum random number generator (QRNG) needs Bell inequality violation. But for the strongest possible certification of randomness, Bell tests become indispensable.

7.1 Types of QRNGs

Single-system QRNGs. A photon hitting a beam splitter takes one of two paths with equal probability. The randomness is quantum-mechanical but you must trust that the device works as specified. If the manufacturer has planted a bias, you cannot detect it from the outputs alone.

Semi-device-independent QRNGs. Some assumptions about the device are relaxed (e.g. the dimension of the Hilbert space), providing intermediate certification.

Device-independent QRNGs (DI-QRNGs). Bell violation is the certification tool. If $|S| > 2$ is observed, no local deterministic hidden-variable model can explain the outputs—randomness is certified without any trust in the internal device structure.

7.2 How Does Bell Violation Certify Randomness?

Here is the key logical step. A Bell violation certifies that the measurement outcomes cannot be predicted by any local hidden variable—not even one known

to the device manufacturer. If the outcomes were deterministic (pre-written on a hidden card), the CHSH value would satisfy $|S| \leq 2$. Observing $|S| > 2$ therefore implies the outcomes contain *irreducible, unpredictable randomness*. The device itself certifies its own unpredictability.

7.3 Breakthrough: ETH Zurich and Certified Perfect Randomness (2026)

Case Study: Perfect Randomness Realised for the First Time [6]

In May 2026, researchers at ETH Zurich led by Renato Renner and Andreas Wallraff published a landmark result in *Nature*: the first experimental demonstration of certified perfect randomness using a Bell test.

The setup. The experiment used two superconducting qubit chips cooled to temperatures near absolute zero, connected by a 30-metre cryogenically cooled tube through which microwave photons could travel, creating quantum entanglement between the two qubits. The 30-metre separation was not merely technical: it was large enough to ensure that no classical signal could travel between the chips during the measurement window, thereby closing the locality loophole.

Why was this hard? A Bell test requires three things simultaneously: high-quality entanglement (to produce CHSH values well above 2), high data rate (to accumulate enough statistics for rigorous certification), and space-like separation (to close the locality loophole). Achieving all three at once with superconducting qubits—which must be kept at millikelvin temperatures—was an exceptional engineering feat. Previous experiments had achieved two of the three conditions but not all simultaneously with sufficient statistical power.

Randomness amplification. The key conceptual insight is *randomness amplification*: you do not need perfect randomness to start with. Even if your initial random seed is slightly biased (up to 0.75% bias in this experiment), a Bell test can amplify it into certified-perfect randomness. The team applied a *two-source extractor*—a classical algorithm that takes two independent weak random inputs and produces one strongly random output. One input came from the Bell-test measurement outcomes; the other came from the original (imperfect) random seed used to choose the measurement settings.

The numbers. The experiment ran for about nine hours and performed 1,342,177,280 trials at a rate of 50 000 trials per second. The average CHSH value achieved was 2.271—comfortably above the classical bound of 2. The final output was a certified-random string of 45,025,658 bits, extracted from roughly 5.4×10^9 low-quality input bits. The failure probability was set at 10^{-12}.

Why it matters. Renner described the achievement as a threshold moment: the output is guaranteed to remain perfectly random for all time, regardless of what analytical methods are used to assess it. Possible applications include cryptographic key generation, public randomness services for lotteries and blockchains, and certified randomness beacons for digital security infrastructure—analogous to what atomic clocks provide for timekeeping.

8 The Main Picture in One Chain

Everything in this paper connects into a single logical chain:

1. **Local realism makes a prediction.** If particles carry pre-written answers and no influence travels faster than light, then the CHSH correlator satisfies $|S| \leq 2$.
2. **Entangled states defy the prediction.** The quantum singlet state produces cosine correlations $E(a, b) = -\cos\theta$. With the right choice of angles, these give $|S| = 2\sqrt{2} \approx 2.83$.
3. **Experiments confirm the violation.** Loophole-free Bell tests [3] and the ETH Zurich experiment [6] consistently measure $|S| > 2$, ruling out local hidden-variable theories.
4. **Conclusion about reality.** The world is not fully described by local pre-written answers. At least one of locality or realism must be abandoned.
5. **Practical payoff.** Quantum randomness can be certified without trusting any device. Bell violation proves unpredictability from the outside.

8.1 A Note on Interpretations

Giving up local realism does not uniquely fix which alternative is correct. Different interpretations of quantum mechanics make different choices:

- **Copenhagen.** Measurement outcomes are genuinely not determined until the act of measurement (gives up realism).
- **Bohmian mechanics.** Particles have definite positions at all times, guided by a non-local pilot wave (gives up locality).
- **Many-worlds.** All outcomes occur in branching universes; the apparent randomness is observer-relative (gives up a single outcome per experiment).

What Bell's theorem guarantees is that the simple classical picture—local, realist, no hidden message between particles—is *experimentally falsified*. The deeper story of what replaces it remains one of the great open questions in the foundations of physics.

Table 1. Key terms and definitions.

Term	Definition
Local Realism	Particles carry pre-written outcomes; no faster-than-light influence.
CHSH Bound	$\|S\| \leq 2$—the maximum correlation achievable by any local hidden-variable theory.
Tsirelson Bound	$\|S\| = 2\sqrt{2} \approx 2.83$—the quantum maximum, achieved by the singlet state at optimal angles.
Bell Violation	Experimental observation of $\|S\| > 2$, ruling out local hidden-variable theories.
Entanglement	Joint quantum state that cannot be factored into independent single-particle states.
Singlet State	$\|\Psi^-\rangle = \frac{1}{\sqrt{2}}(\|\uparrow\downarrow\rangle - \|\downarrow\uparrow\rangle)$ —total spin zero, rotationally invariant, maximally entangled.
DI-QRNG	Device-independent QRNG—randomness certified by Bell violation alone.
Randomness Amplification	Protocol converting weak (biased) randomness into certified-perfect randomness using a Bell test and a classical extractor.

Quick Reference Glossary

References

1. Bell, J.S.: On the Einstein-Podolsky-Rosen paradox. *Physics* **1**(3), 195–200 (1964)
2. Clauser, J.F., Horne, M.A., Shimony, A., Holt, R.A.: Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**(15), 880–884 (1969). doi:10.1103/PhysRevLett.23.880
3. Hensen, B., et al.: Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015). doi:10.1038/nature15759
4. Giustina, M., et al.: Significant-loophole-free test of Bell's theorem with entangled photons. *Phys. Rev. Lett.* **115**, 250401 (2015). doi:10.1103/PhysRevLett.115.250401
5. Shalm, L.K., et al.: Strong loophole-free test of local realism. *Phys. Rev. Lett.* **115**, 250402 (2015). doi:10.1103/PhysRevLett.115.250402
6. Kulikov, A., Storz, S., Schär, J.D., Sandfuchs, M., Wolf, R., Berterottiè, F., Hellings, C., Wallraff, A., Renner, R.: Experimental randomness amplification. *Nature* (27 May 2026). doi:10.1038/s41586-026-10521-8. See also: `https://ethz.ch/en/news-and-events/eth-news/news/2026/05/perfect-randomness-realised-for-the-first-time.html`