Original Research Article

# Identifying Quantum Teleportation Error Mechanisms Through Bitstring-Level Noise Signatures

Nathan B. Niklewski

*Wyoming High School, 106 Pendery Ave, Cincinnati, Ohio, 45215, United States*

## ABSTRACT

Identifying dominant noise mechanisms is a central challenge in today's noisy quantum devices, where every operation is affected by hardware imperfections. To investigate how noise manifests in quantum operations, this study employed a three-qubit teleportation circuit as an experimental probe. Ideal simulations, noisy simulations in Qiskit Aer (readout misassignment, depolarizing, biased-Pauli, thermal relaxation), and runs on IBM Brisbane were compared after applying standard X and Z corrections in post-processing. Distinct patterns in the bitstring distributions showed that IBM hardware is dominated by readout-driven effects with mild bias in Pauli error probabilities. A short teleportation "probe" thus provides a simple, circuit-specific diagnostic test to guide noise mitigation. This study aims to determine whether bitstring distributions can serve as fingerprints for identifying dominant noise sources in quantum teleportation circuits.

**Keywords:** Quantum Information; Quantum Computing; Qiskit; Quantum Teleportation

## INTRODUCTION

Researchers need a fast way to determine which noise mechanism dominates their circuit before committing to lengthy, expensive identification processes. Calibration tables list error rates, but they do not indicate which errors leave the most significant impact on the results of a given circuit on a particular day.

Quantum computers are advancing quickly, but there are significant roadblocks. Current quantum hardware operates in the Noisy Intermediate-Scale Quantum (NISQ) era, where every operation in a circuit is affected by errors (1). NISQ refers to devices with tens to a few thousand qubits and no full error correction. These errors do not form a uniform, single blur. Different mechanisms (e.g., loss of energy, random phase drift, or crosstalk) leave distinct signatures (2). This results in the hardware leaving fingerprints on experimental outcomes. In this study, 'outcomes' refers to the measured bitstrings whose frequency patterns could reveal which errors are dominant.

Providers of quantum devices publish calibration data, but these snapshots are only part of the picture. Error rates change over time, vary across qubits, and can be correlated across operations (3). For example, two qubits may calibrate well individually yet form a poorer entangling pair. Dephasing during idle periods (randomization of phase), amplitude damping ($|1\rangle$ shifting into $|0\rangle$), crosstalk during two-qubit operations (one qubit affecting a nearby one), and mistakes in measuring all contribute to the cumulative error (2, 4,

**Corresponding author:** Nathan B. Niklewski, E-mail: nbniklewski@gmail.com.



**Accepted** November 4, 2025
https://doi.org/10.70251/HYJR2348.36353363

5). What is missing is a low-overhead, circuit-specific diagnostic that converts observed outcome patterns into a hypothesis about the dominant noise mechanism, without extensive techniques. This work proposes using bitstring distributions, rather than fidelity alone, as a circuit-specific diagnostic for identifying dominant noise mechanisms on NISQ hardware. Bitstrings are the results; the bitstring distribution is the histogram of measured outcomes across many shots. Two runs can share the same fidelity (e.g., ~93%) yet demand opposite mitigations: one readout-dominated (fix measurement), the other relaxation-leaning (shorten idles) (1). Fidelity collapses these cases into a single number, while the bitstring structure separates them.

Prior studies on noise characterization typically rely on gate-level diagnosis, randomized benchmarking, or heavy-weight channel estimation. These techniques yield valuable data but often require extensive resources and do not isolate which mechanisms most affect a given circuit. By contrast, this paper treats the measured bitstring distribution itself as data, turning the "wrong" outcomes into a compact, circuit-specific fingerprint. This reframes characterization from hardware-wide calibration to circuit-specific inference, using information already produced by the experiment rather than additional calibration runs.

This study investigates whether simple noise models can accurately reproduce the patterns observed on real IBM hardware, enabling the identification of the dominant sources of noise (6). Rather than attempting to define parameters meticulously, this study analyzes distinct patterns within bitstrings and identifies the error rate that yields similar fidelity. Specifically, this study examines the asymmetries within bitstring distributions (e.g., a bias towards 111 or correlated errors). This study tunes a small set of simple noise models and matches their bitstring signature to hardware, prioritizing pattern agreement over exhaustive parameter identification.

Quantum teleportation was used to test this idea. It involves a short circuit (so time-dependent decoherence still matters), utilizes entangling gates (to capture two-qubit errors), and concludes with a measurement (to capture readout errors). In its standard form, teleportation utilizes a shared entangled pair (a Bell state), a joint measurement on the sender's side, and two classical bits that instruct the receiver on which simple correction (X and/or Z) to apply (7). In a noise-free world, the receiver's qubit ends up in precisely the state the sender started with. On real hardware, however, imperfections in entanglement, decoherence during idle periods, and gate infidelities all introduce noise. Fidelity, in this context, refers to the percentage of correctly teleported states. Teleportation is a means, not the end. It is a compact, well-studied circuit that exposes common errors (8, 9). Intuitively, the sender and receiver share an entangled pair, the sender performs a joint measurement, and two classical bits tell the receiver which simple Pauli correction (X and/or Z) to apply.

This study compares three types of data: ideal simulations (with no noise), noisy simulations in Qiskit Aer under various noise models, and executions on IBM Quantum hardware (10). To isolate hardware errors, this study applied standard teleportation corrections in post-processing, rather than as physical gates. This keeps the circuit depth minimal and avoids current limitations around mid-circuit measurement and feedback. The noisy simulations use Qiskit Aer with small sets of simple noise models.

This study contributes three things. First, a pattern-matching framework that links teleportation bitstring signatures to dominant noise mechanisms using simple, interpretable models. Second, evidence from IBM hardware suggests that a readout-driven signature with slight bias toward specific Pauli-type errors explains most, but not all, of the observed structure. Third, a practical hypothesis tool uses a one-shot teleportation probe to prioritize mitigation before attempting larger, more expansive experiments.

## THEORETICAL BACKGROUND

To ground the analysis, this section reviews the basic elements of quantum teleportation and the noise models used, establishing the framework used throughout the study.

### Qubits, measurement, and notation

$$|\psi\rangle = a|0\rangle + b|1\rangle, \;\; |a|^2 + |b|^2 = 1 \tag{1}$$

Here, a and b are complex numbers called probability amplitudes, and $|a|^2$ is the probability of measuring the state $|0\rangle$ and $|b|^2$ is the probability of measuring the state $|1\rangle$.

Single-qubit Pauli operations represent basic transformations, including bit flips and phase flips. The X gate swaps $|0\rangle$ with $|1\rangle$ (bit-flip), and the Z gate changes the sign of $|1\rangle$ (a phase flip). Two-qubit operations, such as a controlled-NOT (CNOT) gate, create correlations and are essential for entanglement.

A controlled gate applies a single-qubit operation on the target qubit only when the control qubit is |1⟩; otherwise, the identity gate is applied. When a qubit is measured, its state collapses to a definite outcome of either 0 or 1. A single shot refers to one complete run of preparing and measuring the system. For a system of n qubits, a single shot produces an n-bit string such as 010…1, with each qubit's state collapsing to a classical 0 or 1 (2). Repeating many shots builds up a distribution over all possible bitstrings.

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \tag{2}$$

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \tag{3}$$

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{4}$$

X (bit flip) and Z (phase flip) gates are Pauli gates, which, in the context of quantum teleportation, are used to correct the circuit. Pauli gates include X, Z, and Y (a combination of X and Z), and together they form a complete set of single-qubit gates. CNOT gates are crucial to the teleportation protocol and facilitate the ability to relay information.

**Entanglement and Bell states**

Two qubits can be prepared in a maximally entangled state known as Bell states (7). One example, used in quantum teleportation, is:

$$|\Phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}} \tag{5}$$

In this state, each individual qubit, when considered on its own, does not correspond to a pure state but rather to a maximally mixed state. Only the two qubits together form a definite joint state with strong correlations, known as entanglement. If the first qubit is measured as 0, the second will also be 0, if the first is 1, the second will also be 1 (2).

**Teleportation as a measurement-based protocol**

The teleportation protocol used in this study is the standard three-qubit version, where Alice serves as the sender and Bob serves as the receiver (Figure 1).

The bottom two qubits are placed in a Bell state by the Hadamard and CNOT gates. The Hadamard gate creates an equal superposition of |0⟩ and |1⟩, turning a definite input into an equal probability of reading 0 or 1. Alice then performs a CNOT gate on the first two qubits, with the first qubit serving as the control and the second qubit serving as the target, followed by a Hadamard gate on the first qubit. Alice then measures both of the first two qubits to collapse their states into either 0 or 1. Alice sends the two classical results to Bob.

Bob receives the information about the first two qubits. If the second reads a one, an X gate is applied to the final state on the third qubit. If the first reads a 1, a Z gate is applied after the X Gate (7). This step, performed by Bob, is known as correction. The quantum state has already been transferred to Bob through entanglement and measurement, but without applying the appropriate Pauli gates his qubit may not yet match Alice's original state. In an ideal, noise-free world, Bob's qubit exactly matches Alice's original state after correction. In the experiments in this study, corrections are applied in post-processing using Python, which is equivalent to quantum gates without noise.

The outcome of the first two qubits signifies to Bob which corrections to apply in the experiments.

**Bitstrings**

In the teleportation protocol, the results, which are three classical bits, serve distinct roles. The first two bits specify which correction Bob must apply to recover the original state. The third bit indicated the outcome of the teleportation itself.

When teleporting |0⟩, correct shots produce Bob's corrected bit equals 0. The fraction of shots with corrected Bob's bit = 0 is the fidelity of the teleportation. For |0⟩ inputs, ideal outcomes occupy the subset {000,001,010,011} after correction. For clarity, the bitstrings are written in the order of $\{q_2, q_1, q_0\}$ where the leading qubit indicates Bob's teleported result.

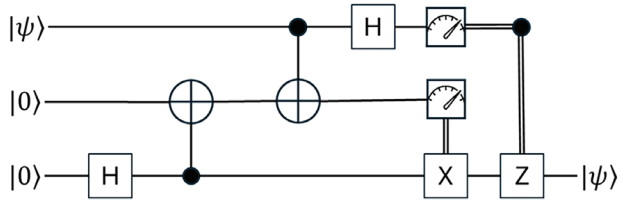


**Figure 1.** Simple Three-Qubit Quantum Teleportation Circuit: is the initial and teleported state; H represents a Hadamard gate; CNOT is represented by a plus sign in the circle; the analog dial denotes measurement of the qubit.

Noise-only shots produce Bob's corrected bit = 1. The distribution among the four erroneous bitstrings {100, 101, 110, 111} can act like a fingerprint of the underlying noise process. Readout error often flips one of Alice's classical bits, which leads Bob to apply the wrong correction. Thermal relaxation drives qubits toward |0⟩, creating error strings that are biased toward zeros. Depolarizing noise, in contrast, spreads probability more evenly across all erroneous outcomes. Their characteristic patterns are central to the analysis in this paper because by comparing the shapes of the error-only distributions, it becomes possible to infer which noise mechanisms are most influential on the hardware.

Reviewing the entire result bitstring allows simple patterns to be reviewed, and how "wrong" results are spread, and which type of noise likely caused the failures.

**Gate correction vs. Physical correction**

Quantum teleportation requires a correction step to transfer the initial state properly. On real quantum hardware, performing conditional correction requires mid-circuit measurement and feedback, which several quantum devices (such as the simpler IBM backend system used in this study) do not support. An alternative is to encode the corrections into larger three-qubit gates (e.g., CCY, CCZ, controlled-phase).

A controlled gate applies a single-qubit operation to the target only when the control is |1⟩. A doubly controlled gate (CCY, CCZ, CCNOT, and control-controlled-phase (CCP)) applies the operation only when both controls are |1⟩. In Figure 2, CCY is a doubly controlled gate on the target qubit. If the two controls (q[0] and q[1]) both read |1⟩, then the target (q[2]) undergoes Y (both X and Z). CCZ is the doubly controlled Pauli-Z, when the controls are both |1⟩, the joint basis state |111⟩ undergoes a phase flip. The CCNOT gate applies an X gate if the two controls read |1⟩. CCP denotes a controlled-controlled-phase gate, which applies a phase rotation to the target state only when both control qubits are |1⟩.

Mathematically, using custom 3-qubit gates such as control-control-Y (ccy), control-control-Z (ccz), and control-control-phase (ccp), properly teleports q[0]. However, these are hardware-intensive and increase noise.

In a noiseless system, only four outcomes are possible: 000, 001, 010, and 011. Any of the other four bitstrings (100, 101, 110, and 111) is a result of noise.

In practice, this study instead applied corrections in post-processing. Measurement results were adjusted classically rather than by adding gates in the quantum circuit. Since for this study, only teleportation of the basis state |0⟩ was analyzed, Z corrections would not alter measurement outcomes, leaving only X corrections relevant (11).

This study applies the correction in post-processing instead of adding gates, which keeps the circuit short (reducing noise) and facilitates sending the initial state. This produces the same classical outcomes as performing the corrections directly on hardware, but without adding gates.

**Noise as quantum channels**

Real devices are noisy. Every idle period, gate, and measurement can transform an ideal state into a different one. Different noise mechanisms leave distinct fingerprints in the distribution of measurement bitstrings. Readout error flips measured bits, often misassigning Alice's results. Thermal relaxation biases qubits toward |0⟩, resulting in predominantly zero-heavy outcomes. Depolarizing noise spreads probability more evenly across erroneous outcomes (2). Teleportation functions as a great probe because it exercises several common noise

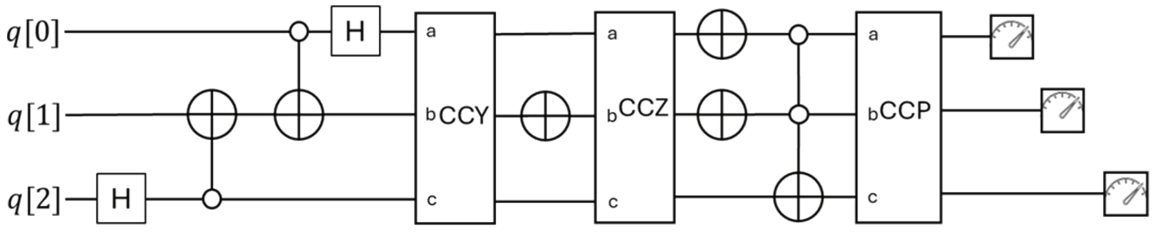


**Figure 2.** Quantum Teleportation Circuit with In-Circuit Correction: This circuit applies corrections as complex quantum gates (CCY, CCZ, CCNOT, and CCP) to work around hardware limitations. This study instead applied correction in post-processing to keep the circuit less complicated.

channels. The protocol requires multi-qubit gates (which are often the noisiest operation), includes idle periods that expose qubits to decoherence, and concludes with measurements that can introduce readout errors (8, 9).

In this study, the distributions of erroneous bitstrings are treated as a kind of fingerprint, meaning that different noise mechanisms consistently distort the histogram in characteristic ways. By comparing these shapes, it becomes possible to infer which noise mechanism is most dominant.

## METHODS AND MATERIALS

This study implements a standard three-qubit teleportation circuit and compares ideal simulations, noisy simulations, and IBM hardware runs. The circuit used is short, corrections are applied in post-processing, and bitstring patterns are analyzed.

### Circuit and Corrections

A standard three-qubit teleportation circuit was implemented in Qiskit (10). The circuit created a Bell pair between two qubits, entangled Alice's input qubit with her half of the pair using a CNOT and Hadamard, and then measured those two qubits and recorded the results, which indicate whether Bob's qubit required corrections (X and/or Z gates). Because IBM hardware does not support mid-circuit measurement and conditional feedback, the corrections were not applied directly within the circuit. Instead, equivalent corrections were applied in post-processing using Python. For the chosen input state $|0\rangle$, Z corrections affect only phase and do not alter measurement outcomes; therefore, only the X correction was applied by flipping Bob's reported bit as needed. Instead of physically applying X or Z on hardware, this study computes the effects on the classical results. In the absence of extra gate noise, this is equivalent to computational basis readout, thereby avoiding hardware limitations. Because the state $|0\rangle$ was teleported, this circuit is highly sensitive to X-like errors and much more resistant to Z-only effects.

### Bitstrings

Each measurement shot was stored as a three-bit string ordered as {q[2], q[1], q[0]}. After post-processing, bitstrings beginning with 1 (e.g., 1xx) represent error-only outcomes. The distribution among these four erroneous bitstrings (100, 101, 110, 111) served as the signature for comparing noise models against hardware data.

### Noise Scenarios and Parameters

The teleportation circuit was executed under eight scenarios (Table 1). An ideal noiseless simulation established the baseline for maximum fidelity, while executions on the IBM Brisbane backend without error mitigation provided the experimental reference. The remaining cases were modeled in Qiskit Aer. To make meaningful comparisons with hardware, each simulated noise model was assigned representative parameters that produced a similar fidelity (~93%) to the IBM runs. A readout error channel was applied with independent flip probabilities of 5% for 0→1 and 3% for 1→0. Depolarizing noise was simulated with 3.5% probability on single-qubit gates and 7% on two-qubit gates. A biased-Pauli channel introduced non-uniform errors with probabilities of 7% for X, 2% for Y, and 1% for Z. Thermal relaxation was modeled using Qiskit's

**Table 1.** Summary of Experimental and Simulated Noise Scenarios

| Scenario | Execution/Model | Noise Channel | Parameters |
|---|---|---|---|
| 1 | Ideal noiseless simulation | None | Baseline for maximum fidelity |
| 2 | IBM Brisbane backend | Native hardware noise | Experimental reference |
| 3 | Readout error | Independent readout flips | Flipping 0 to 1: 5%; Flipping 1 to 0: 3% |
| 4 | Depolarizing noise | Depolarizing channel | Single-qubit gates: 3.5%; Two-qubit gates: 7% |
| 5 | Biased Pauli noise | Pauli X/Y/Z with non-uniform odds | X: 7%; Y: 2%; Z: 1% |
| 6 | Thermal relaxation | Thermal_relaxation_error | with $T_1$ = 40 μs; $T_2$ = 60 μs; Gate durations: 0.25 μs (single-qubit); 1.875 μs (two-qubit) |
| 7 | Bit-flip channel | Pauli-X flips | Probability = 5% |
| 8 | Phase-flip channel | Pauli-Z flips | Probability = 5% |

thermal_relaxation_error with $T_1 = 40$ μs and $T_2 = 60$ μs, with gate durations of 0.25 μs (single-qubit) and 1.875 μs (two-qubit), biasing qubits toward $|0\rangle$ and degrading coherence (IBM Build Noise Models). Two additional channels were tested: a bit-flip noise model (Pauli-X with 5% probability) and a phase-flip noise model (Pauli-Z with 5% probability). For $|0\rangle$ inputs, the latter has only minor effects on the outcome distribution. These values were not meant to reproduce exact calibration data from IBM Brisbane; rather, they were chosen to yield an overall fidelity of about 93% to match the hardware runs, making comparison and analysis simpler.

In the composite model, this study combines asymmetric readout misassignment on Alice's measured qubits with light gate noise. For readout, q[1] reports 1 instead of the true 0 with a probability of 2.8% and reports 0 instead of the true 1 with a probability of 1.4%; q[0] uses 3.2% for 0 to 1 and the same 1.4% for 1 to 0. Bob's qubit (q[2]) has a mild, symmetric readout error of 0.9% in each direction. For gates, all single-qubit operations carry 0.2% depolarizing noise. Each two-qubit CX gate carries 2.0% depolarizing noise and is composed of a small X-biased Pauli component that applies IX (meaning identity is applied on the first qubit and X is applied on the second) or XI (meaning X is applied on the first qubit and I is applied on the second) with probabilities of 0.006 each.

### Execution Details and Repetitions

For each simulated scenario, ten independent runs were performed, each consisting of 4,096 shots. Hardware data were collected in the same manner, distributed across multiple days to account for calibration drift. Repeated runs improved consistency and reduced the influence of outliers. All experiments were conducted in Python 3.13.5 using Qiskit 2.1.1 (Aer 0.17.1) and IBM Brisbane (IBM Backend) (6, 10). This study assesses patterns in the error-only sector by recording the standard distribution for each bitstring.

### Statistics

Data are presented using descriptive statistics (mean, standard deviation). $\chi^2$ test was used to test for homogeneity. The test was conducted on a 2×8 contingency table (real vs. simulation by bitstring outcome), with 7 degrees of freedom. A significance threshold of $\alpha = 0.05$ was applied. 95% confidence intervals are presented graphically to present differences in populations. Excel (version 2508) was used for all statistical analyses.

## RESULTS

### Hardware observations

Across 10x4096 shots per model, overall fidelities were intentionally similar (~93%) because they were all designed to match the hardware noise; however, the pattern of erroneous outcomes significantly differed by noise model. This study, therefore, focuses on the error-only bitstrings (1xx) after the correction has been applied.

Each simulation shot creates a three-bit outcome ordered as $\{q_2, q_1, q_0\}$. After the classical processing step is completed, shots with Bob's corrected bit equal to 1 are entirely erroneous. These are the four "1xx" outcomes (100, 101, 110, 111). Table 2 reports, for each model, the mean rank of those four distributions across ten runs (rank 1 = most frequent, rank 4 = least). An asterisk denotes the smallest mean in each row, the dominant bitstring in that noise model.

The ideal simulator produced only correct outputs (000–011), as expected. On the IBM Brisbane Device, error-only counts per run averaged 264.9/4096 (range 215-372), corresponding to an average fidelity of 93.5%.

### Simulation results

Simulated noise models produced similar error totals: readout-only model averaged 273/ 4096 (fidelity 93.3%), depolarizing averaged 270.4/4096 (fidelity 93.4%), biased-Pauli 284.2/4096 (fidelity 93.1%), and thermal relaxation 267.6/4096 (fidelity 93.5%). The comparative simulation had a fidelity of 93.5% so it will best match the hardware's results.

Although overall fidelities were comparable, the distributions of erroneous outcomes revealed distinct

**Table 2.** Bitstring Distribution Mean Rank

| Noise Model | Bitstring | | | |
|---|---|---|---|---|
| | **100** | **101** | **110** | **111** |
| IBM Brisbane Device | 3.3 | 2.7 | 2.6 | 1.2* |
| Readout | 2.3 | 2.7 | 2.5 | 1.7* |
| Depolarizing | 2.6 | 2.4 | 2.3* | 2.5 |
| Biased Pauli | 2.8 | 2.3 | 2.7 | 2.2* |
| Thermal Relaxation | 1.3* | 1.7 | 3.3 | 3.6 |
| Composite Model | 3.2 | 3.2 | 2.1 | 1.3* |

Mean rank (1=most frequent) of error-only bitstrings (1xx) of 10x4096 shots. *Denotes most frequent bitstring.

signatures. Hardware runs consistently showed a strong bias toward 111, with a mean rank of 1.2; while 100 was consistently suppressed (3.3). The readout error model partially reproduced this structure, with 111 leading (rank 1.7) and 100 read higher than the hardware (rank 2.3). Depolarizing noise produced a more balanced distribution, with 110 slightly favored (rank 2.3), while the biased-Pauli model shifted weight toward 111 (rank 2.2) but did not fully reproduce the hardware's extreme bias. Thermal relaxation generated the opposite pattern: 100 dominated (rank 1.3) and 111 was strongly suppressed (rank 3.6), reflecting its zero-leaning tendency. Thus, hardware readout-dominated, depolarizing, uniform, thermal, zero-biased.

**Composite model performance**

The composite of readout and biased Pauli reduces the residual mismatch relative to readout-only, but the agreement was not exact. Figure 3 shows the bitstring distribution for the composite simulation. A chi-squared test for homogeneity comparing all eight bitstring outcomes between the hardware and the composite simulation ($\chi^2 = 29.83$, df = 7, $p = 1.0 \times 10^{-4}$) confirmed that the simulated distribution differed significantly from the hardware. A simpler 2×2 comparison of correct versus incorrect teleportation outcomes was not significant (p = 0.777), showing that the total error rate was similar even though the pattern of outcomes differed.

**Correlation structure of bitstrings**

Correlation structure within the erroneous outcomes reinforced these distinctions. On hardware, 100 and 110 move in opposite directions across runs, as do 101 and 111 (negative correlations), while 110 and 111 showed a mild positive correlation. Readout noise produced a similar pattern, including co-movement of 110 and 111, consistent with both of Alice's measurement bits being affected. Depolarizing noise displayed the most symmetric trades, particularly between 101 and 100 and between 111 and 110, whereas thermal relaxation shifted weight primarily between 100 and 101, reflecting its bias toward $|0\rangle$. Run-level correlations for the composite model followed the readout pattern (e.g., co-movement of 110 with 111), consistent with measurement-driven effects.

**Additional checks and validation**

For completeness, isolated X-flip and Z-flip models on q[0] and q[1] with the input $|0\rangle$ produced no 1xx outcomes and were excluded from rank comparisons.

Although the analysis emphasizes the 1xx sector, the 0xx counts provide a proper cross-check. The ideal simulator distributed weight almost evenly among 000, 001, 010, and 011, while hardware data showed slight elevation of 001 and 011 (26.6% and 25.8% of the 0xx outcomes, respectively) and suppression of 000 and 010 (24.0% and 23.6%). Readout errors 0xx outcomes distribute as 000 = 23.5%, 001 = 24.6%, 010 = 25.0%,

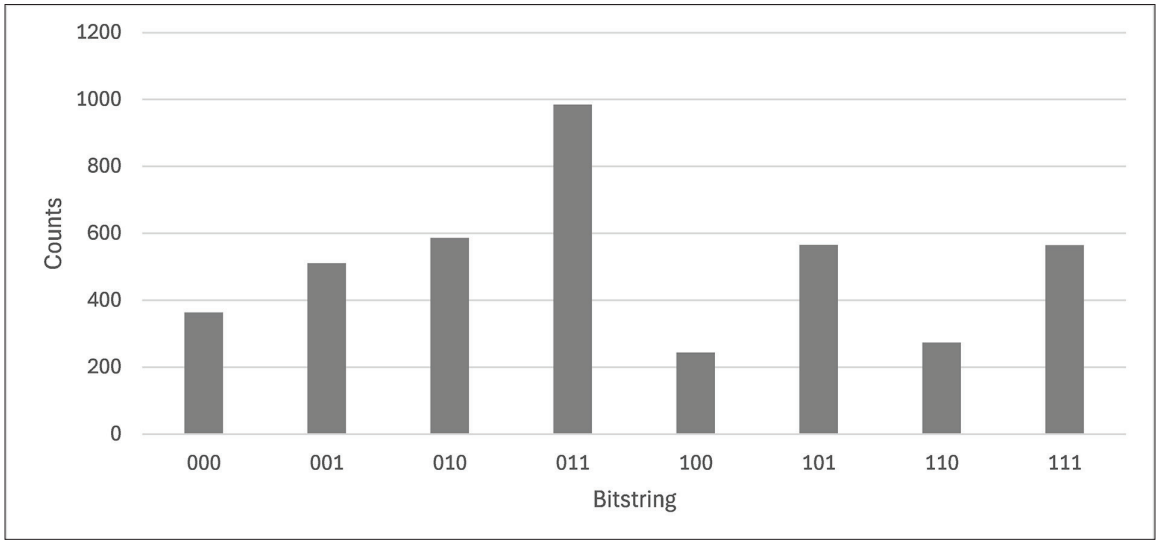


**Figure 3.** Bitstring Distribution of the Custom Circuit: Histogram of the eight computational basis outcomes from 4096 shots using the in-circuit correction. Demonstration only; the analysis uses post-processed corrections.

and 011 = 26.9%, showing a relative elevation of 011 and suppression of 000 compared to an even split. Depolarizing and biased-Pauli models preserved near symmetry, whereas thermal relaxation shifted weight toward 000 (27.1%) and away from 011 (22.8%), mirroring the zeros-leaning signature observed in its 1xx distribution. These shifts mirror the 1xx signatures: depolarizing remains near symmetric, thermal relaxation raises 000 and lowers 011, and hardware shows a mild preference for 001/011, which is consistent with readout skew.

**Comparison of simulation and hardware distributions**

A direct comparison of the confidence intervals between the recreated noise and the real hardware noise (Figure 4) reveals that the composite model accurately captures the shape of the hardware distribution, albeit with reduced contrast. In the success sector (0xx), hardware places more weight on 001 and 011 and less on 000/010. In the error-only sector (1xx), hardware concentrates failures on 111 and suppresses 100, whereas the simulation distributes these outcomes more evenly. The net effect is that the hardware histogram has higher peaks and deeper lows than the simulation, mirroring the readout-dominated model.

**Summary of findings**

Taken together, the results indicate that IBM hardware is characterized by a readout error-dominant distribution, a pattern most closely reproduced by readout error alone, while depolarizing noise yields balanced errors, and thermal relaxation produces the opposite of the hardware pattern, driving outcomes toward bitstrings with leading zeros (such as 100 and 101) and strongly suppressed 111. A mildly biased Pauli component nudges simulated results toward the hardware distribution but does not fully capture its structure. These findings support the use of erroneous bitstring distributions as a compact fingerprint of noise mechanisms, offering interpretability beyond fidelity alone.

## DISCUSSION

Using the full bitstring distribution turns the experiment into a diagnostic. The pattern of erroneous outcomes reveals which noise mechanism dominates, while the total error rate indicates the overall level of noise present. In the IBM runs, a 111-dominant signature with suppressed 100 aligned with readout-driven effects and a mild biased-Pauli component, whereas depolarizing and relaxation models produced opposite shapes. This supports the use of bitstring-level noise fingerprints as fast, circuit-specific diagnostics. For future hardware calibration, such short teleportation fingerprint runs could identify the system's dominant noise type and guide whether to prioritize readout recalibration, gate retuning, or circuit-depth reduction.

Most existing ways of describing noise focus on averages across hardware or calibration numbers listed by the provider. Those are useful, but they do not show which mechanism affects a specific circuit. The idea here is different. In this study, the bitstring pattern itself becomes the diagnostic, turning the experiment into its own test for identifying noise models. Instead of

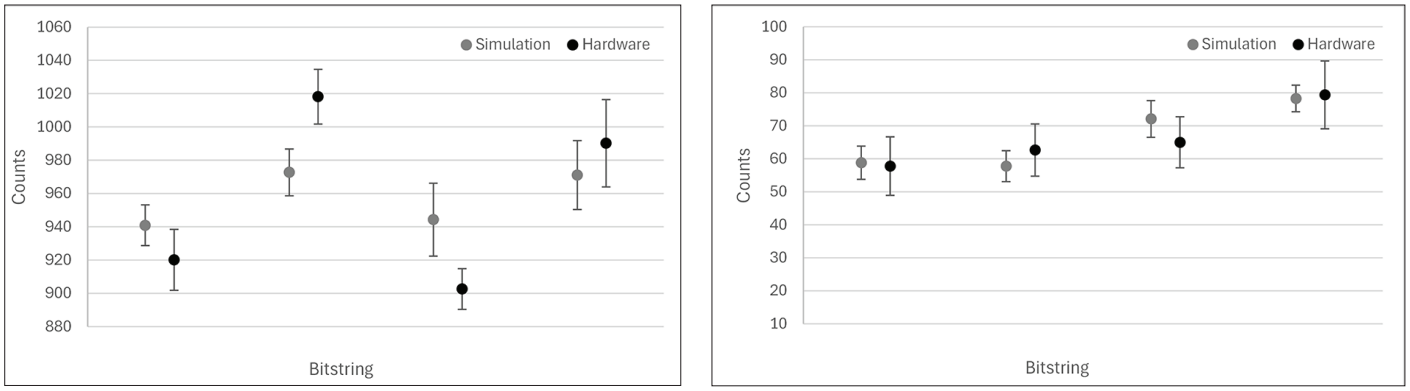


**Figure 4.** Comparison of Simulated Noise and Hardware Noise: The means with 95% confidence intervals (n=10) are presented to compare the composite simulated and hardware outcomes. The left shows the first 4 bitstring outcomes (000, 001, 010, 011), and the right shows the last 4 (100, 101, 110, 111).

adding separate calibration or lengthy experiments, the same results that give fidelity also reveal which noise mechanism is likely the strongest. The difference is not in collecting more data but in looking differently at the data already recorded.

A composite readout and biased-Pauli model reproduced the ordering observed on hardware in the majority of runs (111 as the modal, i.e., most frequent, with suppressed 100), supporting the interpretation that measurement misassignment, together with bias, could account for most, but not all, of the observed structure. The remaining discrepancies (e.g., occasional over-weighting of 110 in simulation) suggest additional contributions beyond two independent channels.

A simple readout error, specifically, 111 frequently leads with 100 being second, but not far ahead of 110 or 101. This skew is expected when misassignment occurs slightly more often when 0 is misinterpreted as 1 than when 1 is misinterpreted as 0. When a classical measurement error flips one of Alice's bits, the correction Bob thinks he needs to apply is wrong in a way that increases the 1xx mass, and it tends to do so in a way that favors Alice equal to 11 rather than Alice equal to 00. Applying Pauli corrections in post-processing mirrors how misassignment would misdirect Bob's correction on hardware, without adding extra gates. The biased-Pauli model nudges the simulated signature even closer to hardware by introducing a slight preference for specific Pauli errors during gate operations. Empirically, the rank of 111 improves (toward 1) and the gap to 100 widens, matching the hardware more closely. Together, these two mechanisms—measurement error and biased gate errors—account for the majority of observed hardware behavior. Together, readout misassignment plus mild gate bias account for the ordering and overall shape of the error-only bitstring distribution observed on hardware.

Equally important is what the hardware distribution does not resemble. Purely, depolarizing noise spreads probability evenly, with 101 and 110 frequently trading places. Although 101 did not have the highest mean rank, it often had the most number of shots in a specific run. Thermal relaxation produces the opposite of the hardware shape: it leans toward zeros on Alice's bits, so 100 and 101 dominate while 111 is suppressed. The hardware, therefore, does not support the view that relaxation dominates this circuit, although relaxation effects cannot be ruled out, and it could present itself more subtly. The composite model's success, therefore, narrows the plausible mechanism to measurement-dominant noise with minor gate bias, rather than relaxation or depolarizing-dominated behavior.

Rank analysis and within-sector comparisons strengthened these conclusions. Because ranks do not depend on total error rate, they remain stable even when absolute error counts fluctuate. Across repeated runs, hardware consistently reproduced the exact ordering of erroneous outcomes, as did readout-only simulations. Depolarizing remained balanced, and thermal relaxation consistently leaned toward zero. These robust patterns indicate that the observed structures reflect underlying mechanisms.

The way the bits trade places is also notable information. Within the erroneous outcomes, the hardware data shows a clear negative correlation between 100 and 110, as well as a similar negative correlation between 101 and 111. The readout shows similar behavior, and in addition, 110 often rises and falls with 111, which is consistent with both of Alice's bits being affected during measurement. Depolarization produces very symmetrical changes, as expected. Thermal relaxation trades mainly between 000 and 101. These correlations argue against a single, unadaptable "noise strength" and instead point to the noise model's specific effects.

This study investigates $|0\rangle$ inputs with post-processed corrections, isolates common noise models, and hardware runs across days in order to reduce the effects of drift, changing the results over time. The goal is to attribute the noise to models rather than complete identification. These simplifications attempt to make the patterns interpretable. Future work can layer complexity once the basic signature is established. There are limits to what this minimal setup can show. This study tested only the $|0\rangle$ input state, which captures bit-flip–type effects but not phase-related errors. As a result, conclusions about bias or readout skew apply mainly to noise acting in the computational basis. To form a more complete view, future work could include $|1\rangle$, $|+\rangle$, and other superposition inputs that reveal Z and phase errors that remain invisible here. Testing across these states would show whether the same fingerprint structure could hold when both bit and phase noise types are active. In this study, only isolated models were tested. Real devices can include correlations across qubits and time (for example, readout cross-talk), which could change the relative counts of bitstrings without necessarily changing the totals. The analysis also combines real hardware runs across days, which helps balance out the effects of confounding variables. Finally, ranks summarize central tendencies.

## CONCLUSION

This study investigated the hypothesis that the shape of the wrong-outcome distribution in a three-qubit teleportation circuit can convey information. Focusing on the error-only sector (the 1xx bitstrings after correction), the distribution in the data was stable across runs and different across models. Under the conditions used here, IBM hardware showed a 111-dominant pattern with 100 being less common. Readout error reproduced much of that pattern, and a mild biased-Pauli component improved the match further, while depolarizing noise remained balanced and thermal relaxation produced an almost opposite distribution (zeros leaning). These results suggest that composition can provide practical guidance about what is likely the most influential noise mechanism on quantum hardware.

The practical takeaway is guidance, not a verdict. A strong 111 mode is consistent with measurement-driven effects, making readout-error mitigation a sensible first option to assess. A balanced distribution, similar to depolarizing noise, suggests attention to gate quality, while a zero-leaning distribution implicates idle-time relaxation and could be improved by reducing circuit depth. These interpretations are not absolute, as different error combinations may produce similar histograms; however, the approach remains useful because it relies on error patterns rather than absolute rates.

This idea can serve as a straightforward method for device fingerprinting. A small reference library of canonical error patterns could be maintained. Before a larger experiment, a brief teleportation circuit could be run, its 1xx composition compared to the library, and the day's dominant error mechanism estimated. Composite models (e.g., readout plus a mildly biased-Pauli model) could then be tuned to match the observed shape. Such a model should be viewed as a working approximation and ideally cross-checked with more advanced systems.

The broader message is that bitstring structure carries more information than fidelity alone. Whereas fidelity reduces outcomes to a single number, distributions reveal mechanism-specific signatures. The approach may be extended to other compact circuits, although validation would be required on a case-by-case basis. Fidelity collapses outcomes into a single number; the bitstring structure preserves mechanism-specific signatures. Under these conditions, a readout-driven pattern with a slight bias toward specific Pauli error types explains most of the observed shape, whereas pure relaxation or pure depolarization does not. Targeted experiments (with additional inputs and bases) would help separate the remaining alternatives.

It is important to note that different error combinations or schedules may mimic similar histograms. Even with this confounding potential, a short teleportation circuit offers a low-resource approach to estimate the dominant influence and choose the mitigations to attempt before more resource-intensive executions. Because the method depends on the distribution rather than the total error, it stays useful as calibrations drift. Taken together, these results support the view that bitstring-level patterns are an accessible, mechanism-aware summary that can help guide mitigation and model tuning before more advanced techniques are employed.

## ACKNOWLEDGEMENT

The author gratefully acknowledges Dr. Aleksandra Ziolkowska, Postdoctoral Research Fellow, University of St. Andrew's, for her guidance and support as the supervisor of this paper.

## CONFLICT OF INTEREST

The author declares that there are no conflicts of interest regarding the publication of this article.

## REFERENCES

1. Preskill J. Quantum Computing in the NISQ era and beyond. *Quantum*. 2018; 2: 79. https://doi.org/10.22331/q-2018-08-06-79
2. Nielsen MA & Chuang IL. Quantum Computation and Quantum Information. (Cambridge University Press, 2012). doi:10.1017/CBO9780511976667.
3. IBM Quantum. Get backend information with Qiskit. https://quantum.cloud.ibm.com/docs/en/guides/get-qpu-information (2025).
4. IBM Quantum. Build noise models. https://quantum.cloud.ibm.com/docs/en/guides/build-noise-models (2025).
5. Sung Y, Ding L, Braumüller J, Vepsäläinen A, *et al*. Realization of High-Fidelity CZ and ZZ-Free iSWAP Gates with a Tunable Coupler. *Phys Rev X*. 2021; 11: 021058. https://doi.org/10.1103/PhysRevX.11.021058
6. IBM Quantum. IBM Quantum Platform. Online: https://quantum.cloud.ibm.com (2025).
7. Bennett CH, Brassard G, Crépeau C, Jozsa R, *et*

*al*. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys Rev Lett.* 1993; 70: 1895-1899. https://doi.org/10.1103/PhysRevLett.70.1895

8. Bowen G & Bose S. Teleportation as a Depolarizing Quantum Channel, Relative Entropy, and Classical Capacity. *Phys Rev Lett*. 2001; 87: 267901. https://doi.org/10.1103/PhysRevLett.87.267901
9. Fortes R & Rigolin G. Fighting noise with noise in realistic quantum teleportation. *Phys Rev A (Coll Park).* 2015; 92: 012338. https://doi.org/10.1103/PhysRevA.92.012338
10. Javadi-Abhari A, Treinish M, Krsulich K, Wood CJ, *et al*. *Quantum computing with Qiskit*. arXiv preprint arXiv:2405.08810 (2024).
11. Gottesman D. The Heisenberg representation of quantum computers. arXiv preprint arXiv:quant-ph/9807006 (1998).