

Article

Not peer-reviewed version

Reconstructing Quantum Noise from Symmetry-Violating Generators

[Sebastian Raubitze](#), Krzysztof Werner, [Sebastian Schrittwieser](#)^{*}, [Kamil Wereszczyński](#), Krzysztof A. Cyran, [Kevin Mallinger](#)

Posted Date: 27 February 2026

doi: 10.20944/preprints202602.1809.v1

Keywords: quantum computing; process tomography; symmetry violation; lie groups; group theory; IBM quantum; noise modeling; quantum simulation



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Reconstructing Quantum Noise from Symmetry-Violating Generators

Sebastian Raubitsek¹, Krzysztof Werner², Sebastian Schrittwieser^{3,*}, Kamil Wereszczyński², Krzysztof A. Cyran² and Kevin Mallinger¹

¹ SBA Research gGmbH, Complexity and Resilience Research Group, Floragasse 7/5.0G, 1040 Vienna, Austria

² merQlab, Department of Computer Graphics, Vision and Digital Systems, Silesian University of Technology, Gliwice, Poland

³ Christian Doppler Laboratory for Assurance and Transparency in Software Protection, Faculty of Computer Science, University of Vienna, Kolingasse 14–16, 1090 Vienna, Austria

* Correspondence: sebastian.schrittwieser@univie.ac.at

Abstract

By systematically introducing small symmetry-violating terms into quantum gates, we examine whether such perturbations can reproduce the tomographic noise patterns observed on actual quantum hardware. This study fits into the broader context of quantum computing and quantum information processing, where realistic noise modeling remains essential for understanding device behavior. Our analytical framework mirrors Qiskit's tomography procedures and shot model, using transparent noise control. Our results show that weak generator perturbations can reproduce average deviations in Pauli Transfer Matrices with high fidelity (given the recorded hardware noise characteristics). These findings suggest that symmetry-violating generators provide a phenomenological model for relevant classes of hardware noise. This approach offers a reproducible analytic baseline for testing noise hypotheses and calibrating simplified models against experimental data. Beyond this descriptive role, the appearance of stable symmetry-violating perturbations in the fitted models supports their physical relevance: they indicate that real devices operate slightly outside the symmetry-preserving regime assumed in standard noise models, and that these deviations leave identifiable signatures in tomographic reconstructions. Because the analytical framework isolates these violations at the generator level, it enables principled diagnosis of non-unitarity and related effects. As a result, the model not only reproduces experimental process tomography data but also identifies which symmetry assumptions fail and by what magnitude, providing actionable information for calibration, simulation, and the development of symmetry-sensitive mitigation and error-correction strategies.

Keywords: quantum computing; process tomography; symmetry violation; lie groups; group theory; IBM quantum; noise modeling; quantum simulation

1. Introduction

Real quantum hardware deviates from idealized unitary dynamics due to relaxation, dephasing, leakage, and control errors that collectively break the mathematical idealization of quantum information processing. We aim to model these imperfections as deviations at the generator level, where small symmetry-violating terms modify the effective dynamics. In this work, we propose an analytic description of such non-ideal behavior by introducing *symmetry-violating generator noise* directly into quantum gate's Lie algebra representation.

Quantum process tomography provides the quantitative framework to study these deviations. Process tomography reconstructs the full dynamical map that represents a quantum operation, enabling direct comparison between the expected and realized transformations [1–3]. In this context, it answers a concrete question: how does an implemented gate act on arbitrary input states compared to its theoretical target? Applying process tomography to our circuits therefore allows us to identify and classify aspects of hardware imperfections.

In this study, we rebuild single-qubit process tomography analytically using a finite-dimensional matrix pipeline. Gates are defined as exponentials of $\mathfrak{su}(2)$ generators, evaluated through matrix exponentials. Each gate $U(\theta) = \exp[-i(\theta/2)G]$ is replaced by

$$U_{\Delta}(\theta) = \exp\left[-i\frac{\theta}{2}(G + \Delta)\right],$$

where Δ represents a tunable additive matrix perturbation. A Hermitian Δ corresponds to coherent miscalibration such as axis tilt or rotation error, while a general complex Δ produces effective non-unitary behavior that emulates dissipation, leakage, or other open-system effects at the circuit level [1,2]. Here, θ denotes the physical rotation angle of the gate, $G \in \mathfrak{su}(2)$ is the ideal generator associated with the target operation, $\Delta \in \mathbb{C}^{2 \times 2}$ is the additive perturbation matrix modeling symmetry-violating noise, and $U_{\Delta}(\theta)$ is the resulting noisy gate obtained by exponentiating the perturbed generator.

Preparation and measurement follow the Pauli schedules implemented in Qiskit Experiments, and expectation values are sampled through a binomial shot model to reproduce hardware statistics [3,4]. We then compare these analytic reconstructions against runs on an IBM Quantum system (IBM Brisbane), averaging over repeated executions to suppress statistical noise and temporal drift.

The analysis focuses on two central questions:

1. Can small symmetry-violating perturbations introduced at the generator level reproduce features observed in experimental process tomographies?
2. Can such analytic perturbations be tuned to represent realistic hardware noise profiles, at least in an average and statistically consistent sense?

Our results demonstrate that generator-level perturbations can be tuned to reproduce the average tomographic signatures of noise observed on IBM Quantum hardware for single-qubit circuits. Using the same Pauli preparation and measurement schedules as in Qiskit's process tomography, the analytic reconstructions consistently reproduce deviations found in experimental Pauli-transfer matrices. The study shows that symmetry violation at the generator level is one possible mechanism contributing to these deviations, but it is not the only one; additional physical modeling using completely positive trace-preserving constraints, Kraus or Lindbladian formulations, and explicit treatment of crosstalk and drift are other potential sources for these deviations [5–7]. The framework presented here provides an explicit and controllable analytic counterpart to the Qiskit tomography workflow, allowing direct comparison between simulated and experimental process reconstructions under both ideal and noisy conditions. Using a genetic algorithm, we can reproduce features of non-trace-preserving behavior and the component-wise variance patterns observed in hardware data while retaining full analytical transparency of the underlying reconstruction process.

These findings are important because they reveal that symmetry-violating perturbations might not be merely a mathematical inconvenience but measurable (within the shown uncertainty) structured features of real hardware noise. By isolating these effects at the generator level, the framework shows which assumed symmetries in standard noise models break down and quantifies the scale of this symmetry violation directly from experimental data. This provides a principled route for developing related noise models, calibration tools, and error-mitigation strategies that explicitly account for symmetry-violating behavior of real device imperfections.

2. Related Work

Foundational formulations of quantum state and process tomography were established by Chuang and Nielsen [1] and by Poyatos, Cirac, and Zoller [2], with standard textbook treatments collected in [3]. The Pauli-transfer-matrix (PTM) and related affine Bloch-vector form are widely used to represent single-qubit channels and compare operations; see, e.g., Gilchrist, Langford, and Nielsen [8].

Beyond direct linear inversion, self-consistent and overcomplete approaches seek robustness against state-prep and measurement (SPAM) errors. Merkel et al. introduced self-consistent process tomography [9], while gate set tomography (GST) provides a comprehensive, SPAM-robust charac-

terization across a gate family [10]. These methods complement our analytic reconstructions, which explicitly mirror the minimal Pauli schedules but can also be evaluated on overcomplete schedules.

Noise benchmarking and diagnostics provide orthogonal evidence about coherent and stochastic error structure. Randomized benchmarking (RB) offers SPAM-robust error-rate estimates [11], and comprehensive introductions and tutorials now standardize the methodology and its practical nuances [12]. Clarifying the interpretation of RB data, Proctor et al. show precisely what RB actually measures and under which assumptions [13]. RB-style protocols have also been combined with tomographic techniques to obtain robust gate reconstructions [14]. Related notions such as *unitarity* help separate coherent miscalibration from stochastic noise [15]. These tools contextualize the symmetry-violating generator model explored here: Hermitian perturbations emulate coherent miscalibration, while general (non-Hermitian) perturbations emulate effective non-unitary behavior at the circuit level.

In parallel, error-mitigation techniques—notably zero-noise extrapolation (ZNE)—adjust circuit-level parameters to infer noise-free estimates [16]. Although our goal is not error mitigation, the idea of *controlled* circuit-level modifications aligns with our analytic perturbation model. The Qiskit-Experiments framework operationalizes these tomography and calibration workflows in practice [4,17], which we reproduce analytically via explicit matrix products.

Finally, symmetry-related generator perturbations have been explored as a modeling tool beyond characterization, including for privacy-preserving data obfuscation in quantum-inspired pipelines [18]. Our study adapts this generator-level perspective to investigate whether small symmetry violations can mimic averaged tomographic signatures observed on IBM hardware, thereby providing a compact phenomenological model to test hypotheses about realistic device noise.

Genetic algorithms (GAs) are population-based stochastic search methods inspired by biological evolution and have been developed and analyzed since the foundational work of Holland and the engineering codification by Goldberg [19,20]. Core GA operators (selection, recombination, and mutation) enable global exploration with problem-agnostic objective functions, and their behavior has been surveyed extensively in the optimization literature [21–23]. Multiobjective extensions such as NSGA-II introduced elitist selection and fast nondominated sorting, improving convergence and solution diversity for vector-valued objectives [24]. In our context, we use a GA to fit per-gate perturbation parameters by minimizing a tomography-derived discrepancy; the same class of metaheuristics has been shown effective when the search space is nonconvex, noisy, or only implicitly defined. Relatedly, GA-driven optimization has been coupled with stochastic generative models in prior work by Raubitzeck *et al.*, where a GA selects among fractional Brownian-bridge candidates using a smoothness-based fitness to interpolate sparse time series [25]. This demonstrates a practical pattern, combine a stochastic proposal mechanism with evolutionary selection under a task-specific fitness, that we also follow here for calibrating symmetry-violating generator parameters against experimental process-tomography targets.

3. Theoretical Foundations

In this work we consider quantum information processing on single qubits. A qubit is a two-level quantum system with Hilbert space $\mathcal{H} \simeq \mathbb{C}^2$ and computational basis $\{|0\rangle, |1\rangle\}$ [3]. A pure single-qubit state is a normalized vector

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \quad (1)$$

To include classical uncertainty and decoherence, we describe the state by a density operator ρ , which is a positive semidefinite, unit-trace 2×2 matrix. Any single-qubit density operator can be written in Bloch form

$$\rho = \frac{1}{2} \left(I_2 + r_x \sigma_x + r_y \sigma_y + r_z \sigma_z \right), \quad \mathbf{r} = (r_x, r_y, r_z)^\top \in \mathbb{R}^3, \quad \|\mathbf{r}\|_2 \leq 1, \quad (2)$$

where $\sigma_x, \sigma_y, \sigma_z$ are the Pauli operators and \mathbf{r} denotes the Bloch vector of the state [3].

Quantum gates, i.e., transformations acting on the qubit, are represented by unitary operators acting on \mathcal{H} . In the Schrödinger picture, the action of a single-qubit gate $U \in U(2)$ on ρ is

$$\rho \mapsto U\rho U^\dagger, \quad (3)$$

and for pure states $|\psi\rangle$ this corresponds to $|\psi\rangle \mapsto U|\psi\rangle$. A convenient basis for single-qubit operations is given by rotations generated by Pauli operators. For σ_i with $i \in \{x, y, z\}$, the corresponding rotation is

$$R_i(\theta) = e^{-i(\theta/2)\sigma_i}. \quad (4)$$

On superconducting-qubit platforms such as IBM Quantum, native single-qubit gates are expressed through Z-axis rotations and $\pi/2$ -rotations around the X-axis [26,27]. A common basis is $\{R_z(\theta), SX, X\}$, defined as

$$R_z(\theta) = e^{-i(\theta/2)\sigma_z}, \quad (5)$$

$$SX = e^{-i(\pi/4)\sigma_x}, \quad (6)$$

$$X = e^{-i(\pi/2)\sigma_x}. \quad (7)$$

Standard gates such as the Hadamard and the phase gate can be decomposed in this basis, for example

$$H = R_z\left(\frac{\pi}{2}\right) SX R_z\left(\frac{\pi}{2}\right), \quad S = R_z\left(\frac{\pi}{2}\right), \quad (8)$$

consistent with typical backend calibrations [27].

A single-qubit circuit is a time-ordered sequence of gates applied to an initial state. Let ρ_{in} be the input state and $\{U_k\}_{k=1}^L$ the unitaries corresponding to the gates from first applied (U_1) to last applied (U_L). The total unitary is

$$U_{\text{tot}} = U_L U_{L-1} \cdots U_1, \quad (9)$$

and the ideal output state is

$$\rho_{\text{out}}^{(\text{ideal})} = U_{\text{tot}} \rho_{\text{in}} U_{\text{tot}}^\dagger. \quad (10)$$

Logical circuits are decomposed into this native gate sequence to the actual physical sequence executed on the backend [26,27].

A typical single-qubit experiment consists of state preparation, application of the target gate sequence, and measurement. The device initializes the qubit in $|0\rangle$, after which a preparation unitary U_{prep} creates the desired input state,

$$\rho_{\text{in}} = U_{\text{prep}} |0\rangle\langle 0| U_{\text{prep}}^\dagger. \quad (11)$$

The target gate sequence U_{target} then acts on this state,

$$\rho_{\text{mid}} = U_{\text{target}} \rho_{\text{in}} U_{\text{target}}^\dagger. \quad (12)$$

Before measurement, an optional basis-rotation unitary U_{meas} may be inserted to access observables in different measurement bases, resulting in

$$\rho_{\text{out}} = U_{\text{meas}} \rho_{\text{mid}} U_{\text{meas}}^\dagger. \quad (13)$$

The device then performs a projective measurement in the computational basis $\{|0\rangle, |1\rangle\}$. The full evolution factorizes as

$$\rho_{\text{out}} = U_{\text{meas}} U_{\text{target}} U_{\text{prep}} |0\rangle\langle 0| U_{\text{prep}}^\dagger U_{\text{target}}^\dagger U_{\text{meas}}^\dagger. \quad (14)$$

Measurements in other Pauli bases are implemented through suitable pre-measurement rotations. Measuring in the X basis uses $U_{\text{meas}} = H$, since $H\sigma_z H = \sigma_x$. Measuring in the Y basis uses $U_{\text{meas}} = HS^\dagger$.

Let $\Pi_0 = |0\rangle\langle 0|$ and $\Pi_1 = |1\rangle\langle 1|$ denote the computational basis projectors. For a given output state ρ_{out} , the Born rule assigns the probabilities

$$p(b) = \text{Tr}(\rho_{\text{out}}\Pi_b), \quad b \in \{0, 1\}. \quad (15)$$

In an experiment, the circuit is executed N times, yielding empirical frequencies $\hat{p}(b) = n_b/N$. These estimates exhibit binomial fluctuations with variance

$$\text{Var}[\hat{p}(b)] = \frac{p(b)(1-p(b))}{N}. \quad (16)$$

Expectation values of observables follow from

$$\langle O \rangle = \text{Tr}(\rho_{\text{out}}O). \quad (17)$$

For Pauli operators,

$$\langle \sigma_i \rangle = \text{Tr}(\rho_{\text{out}}\sigma_i) = r_i^{(\text{out})}, \quad i \in \{x, y, z\}, \quad (18)$$

providing direct access to the components of the output Bloch vector when the appropriate measurement basis is used.

Real devices deviate from the ideal unitary description due to control errors, relaxation, dephasing, and other noise sources. The effective action of an implemented single-qubit gate is thus modeled by a quantum channel \mathcal{E} , a linear, completely positive, trace-preserving map on density operators [3]. A standard representation is the Kraus form

$$\mathcal{E}(\rho) = \sum_k E_k \rho E_k^\dagger, \quad \sum_k E_k^\dagger E_k = I_2, \quad (19)$$

where $\{E_k\}$ are Kraus operators. An ideal unitary corresponds to the special case $\mathcal{E}_U(\rho) = U\rho U^\dagger$. When a noisy channel acts on an input state, the output becomes

$$\rho_{\text{out}}^{(\text{noisy})} = \mathcal{E}(\rho_{\text{in}}), \quad (20)$$

and measurement probabilities follow as

$$p(b) = \text{Tr}(\mathcal{E}(\rho_{\text{in}})\Pi_b). \quad (21)$$

The combination of state preparation, channel action, and projective measurement fully determines the observed statistics. This formalism—single-qubit states, unitary evolution, quantum channels, and measurement statistics—forms the basis for the characterization and analysis of implemented single-qubit operations.

3.1. Quantum Process Tomography

Quantum process tomography (QPT) aims to reconstruct the dynamical map \mathcal{E} that describes how input states are transformed by a physical operation [1,3,9]. In the single-qubit setting considered here, \mathcal{E} is modeled as a completely positive, trace-preserving (CPTP) map acting on density operators, such that for any input state ρ_{in} the corresponding output state is

$$\rho_{\text{out}} = \mathcal{E}(\rho_{\text{in}}). \quad (22)$$

The goal of QPT is to determine \mathcal{E} from experimental data obtained by preparing a set of input states, applying the process, and measuring the resulting output states using the state-preparation and measurement framework introduced in Sec. 3.

Formally, QPT exploits the linearity of \mathcal{E} . A tomographically complete set of input states $\{\rho_k\}$ spans the space of single-qubit density operators. In practice, one chooses a small set of linearly independent states, such as the eigenstates of the Pauli operators: the computational basis states $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$, together with superposition states such as $|+\rangle\langle +|$ and $|+i\rangle\langle +i|$, where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|+i\rangle = (|0\rangle + i|1\rangle)/\sqrt{2}$. The device prepares each ρ_k by applying an appropriate U_{prep} to $|0\rangle\langle 0|$, then applies the physical gate or gate sequence described by \mathcal{E} , and finally performs measurements after a basis-rotation unitary U_{meas} , as in Sec. 3. For each input state and each measurement basis (for example, Z , X , and Y), the experiment yields outcome frequencies which estimate the corresponding probabilities and expectation values.

To represent \mathcal{E} , several equivalent formalisms are available, including the Kraus representation, the Choi matrix, and the Pauli transfer matrix (PTM) [3,9]. In this work we use the PTM, which is convenient for single-qubit processes and matches the conventions of the Qiskit tomography tools [26,27]. We expand states and channels in the operator basis $\{I_2, \sigma_x, \sigma_y, \sigma_z\}$. For any single-qubit state ρ , we define the Bloch vector $\mathbf{r} = (r_x, r_y, r_z)^\top$ through

$$\rho = \frac{1}{2} \left(I_2 + r_x \sigma_x + r_y \sigma_y + r_z \sigma_z \right), \quad (23)$$

and introduce the augmented vector

$$\mathbf{s} = \begin{bmatrix} 1 \\ \langle X \rangle \\ \langle Y \rangle \\ \langle Z \rangle \end{bmatrix} = \begin{bmatrix} 1 \\ r_x \\ r_y \\ r_z \end{bmatrix}, \quad (24)$$

where $\langle \sigma_i \rangle = \text{Tr}(\rho \sigma_i)$ are the expectation values defined in Sec. 3. The PTM $R \in \mathbb{R}^{4 \times 4}$ of a channel \mathcal{E} is then defined by its action on these vectors,

$$\mathbf{s}_{\text{out}} = R \mathbf{s}_{\text{in}}, \quad \mathbf{s}_{\text{in}} = \begin{bmatrix} 1 \\ \langle X \rangle_{\text{in}} \\ \langle Y \rangle_{\text{in}} \\ \langle Z \rangle_{\text{in}} \end{bmatrix}, \quad \mathbf{s}_{\text{out}} = \begin{bmatrix} 1 \\ \langle X \rangle_{\text{out}} \\ \langle Y \rangle_{\text{out}} \\ \langle Z \rangle_{\text{out}} \end{bmatrix}. \quad (25)$$

Trace preservation implies that the first component of \mathbf{s}_{out} is always equal to one, so the first row of R is fixed to $[1, 0, 0, 0]^\top$. The remaining 3×3 submatrix encodes how the Bloch vector is transformed and captures both coherent rotations and incoherent effects such as dephasing and relaxation [9].

In an experiment, the PTM elements are inferred from measured expectation values. For each prepared input state ρ_k and each measurement basis corresponding to a Pauli operator σ_i , the protocol estimates $\langle \sigma_i \rangle_{\text{in},k}$ before the process and $\langle \sigma_i \rangle_{\text{out},k}$ after the process. These expectation values are obtained from the observed outcome frequencies using the Born rule and the finite-shot statistics from Sec. 3. Collecting all input and output Bloch vectors into a data set yields a linear system whose solution determines the matrix R . For single-qubit processes, a minimal informationally complete scheme uses four input states (for example $|0\rangle\langle 0|$, $|1\rangle\langle 1|$, $|+\rangle\langle +|$, $|+i\rangle\langle +i|$) and measurements in the three Pauli bases X , Y , and Z . This results in twelve distinct circuits on hardware: four state preparations combined with three measurement bases. In the IBM Quantum and Qiskit framework, these circuits are built from the native gate set $\{R_z(\theta), SX, X\}$ and the standard pre-measurement rotations implementing X - and Y -basis measurements [26,27]. The Qiskit tomography routines then perform a linear inversion or a constrained reconstruction to obtain an estimate \hat{R} that is consistent with the CPTP constraints to the extent allowed by statistical noise and systematic errors [4,9].

Once reconstructed, the PTM \hat{R} provides a compact and operational description of the implemented channel. It can be compared directly to theoretical PTMs derived from ideal unitary gates or from explicit noise models, and it allows one to quantify coherent miscalibrations and incoherent error mechanisms at the level of single-qubit processes. In the rest of this work, these reconstructed PTMs serve as the main interface between analytic models of noisy single-qubit evolution and the experimental data obtained from IBM Quantum hardware.

3.2. Qiskit Basis and the 12-Circuit Process Tomography

Single-qubit process tomography in Qiskit is implemented using a minimal native gate basis $\{R_z(\theta), SX, X\}$, which matches the hardware-level operations used on IBM Quantum devices [26,27]. This basis is universal for single-qubit unitaries and is compatible with the state-preparation and measurement operations introduced in Sec. 3. In the QPT experiment, the process \mathcal{E} is characterized by preparing four linearly independent input states and measuring each of the corresponding output states in three Pauli bases. These twelve configurations provide complete information to reconstruct the Pauli transfer matrix (PTM) of the implemented channel [1,9].

The four input states are chosen as eigenstates of the Pauli operators. Using the notation from Sec. 3, they are

$$|Z_+\rangle = |0\rangle, \quad |Z_-\rangle = |1\rangle, \quad |X_+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \quad |Y_+\rangle = \frac{|0\rangle + i|1\rangle}{\sqrt{2}}.$$

Each is prepared by an appropriate unitary U_{prep} acting on $|0\rangle$. After the process under study is applied, the output state is measured in one of the three Pauli bases Z , X , or Y . Measurements in the X and Y bases are implemented using the pre-measurement rotations H and HS^\dagger defined in Sec. 3. Combining the four preparations with the three measurement bases yields the twelve circuits

$$U_{p,m} = M_m U_T P_p, \quad p \in \{Z_+, Z_-, X_+, Y_+\}, \quad m \in \{Z, X, Y\},$$

where P_p prepares the input state, U_T denotes the target operation, and M_m implements the measurement basis.

Explicitly, the twelve circuits are

$$\begin{aligned} (1) & M_Z U_T |Z_+\rangle, & (2) & M_X U_T |Z_+\rangle, & (3) & M_Y U_T |Z_+\rangle, \\ (4) & M_Z U_T |Z_-\rangle, & (5) & M_X U_T |Z_-\rangle, & (6) & M_Y U_T |Z_-\rangle, \\ (7) & M_Z U_T |X_+\rangle, & (8) & M_X U_T |X_+\rangle, & (9) & M_Y U_T |X_+\rangle, \\ (10) & M_Z U_T |Y_+\rangle, & (11) & M_X U_T |Y_+\rangle, & (12) & M_Y U_T |Y_+\rangle. \end{aligned}$$

For each preparation p and measurement basis m , the experiment collects measurement outcomes that estimate the expectation value $\langle \sigma_m \rangle_p$. These twelve expectation values form a linear system relating input Bloch vectors \mathbf{s}_{in} to output Bloch vectors \mathbf{s}_{out} . Solving this system yields the PTM

$$R = \begin{pmatrix} 1 & 0 & 0 & 0 \\ t_x & T_{xx} & T_{xy} & T_{xz} \\ t_y & T_{yx} & T_{yy} & T_{yz} \\ t_z & T_{zx} & T_{zy} & T_{zz} \end{pmatrix},$$

where the first row is fixed by trace preservation and the remaining elements describe the affine map of the Bloch vector. The translation vector (t_x, t_y, t_z) and the linear block T are inferred from the measured expectation values, using the Born rule and the finite-shot statistics discussed in Sec. 3. These twelve circuits thus provide a complete and minimal dataset for reconstructing the implemented single-qubit channel, both in Qiskit's tomography routines and in the simulator used in this work [9,28].

The entries of the reconstructed PTM have a direct operational interpretation. The 3×3 block T describes how the input Bloch vector is rotated and contracted by the process. For an ideal unitary

operation, T is an orthogonal matrix with determinant $+1$, corresponding to a rigid rotation of the Bloch sphere. Deviations from orthogonality quantify incoherent effects such as dephasing, amplitude damping, or general relaxation processes. The translation vector (t_x, t_y, t_z) describes an affine shift of the Bloch vector and therefore captures nonunital effects. A process is unital if and only if it maps the maximally mixed state to itself; in PTM form this corresponds to $(t_x, t_y, t_z) = (0, 0, 0)$. Nonunital behavior thus reflects population relaxation or asymmetries in the underlying noise mechanisms. These components, together with the Pauli block T , specify the complete linear transformation of the Bloch vector under the implemented quantum channel.

In the trace-preserving case, the first row of the PTM must take the form $(1, 0, 0, 0)$, reflecting the fact that total probability is conserved. In practice, experimental reconstructions often produce a first row with nonzero off-diagonal elements, indicating apparent deviations from strict trace preservation. These deviations may arise from finite sampling, readout errors, leakage out of the computational subspace, or effective non-trace-preserving dynamics caused by state preparation and measurement imperfections. In the PTM representation, nonzero components in the first row quantify how the total probability mass associated with each Pauli basis element fails to be conserved. Allowing the first row to be reconstructed without enforcing trace preservation provides additional diagnostic information: it reveals whether an implemented operation exhibits loss, leakage, or preparation-dependent detection biases. This perspective is central to our analysis, as it enables us to distinguish coherent miscalibration effects from nonunitary and potentially non-trace-preserving contributions observed on hardware.

3.3. Shot-Based Simulator for Single-Qubit Process Tomography

To compare analytic gate models with experimentally reconstructed channels, we develop a simulator that reproduces the twelve-circuit process tomography experiment introduced in Sec. 3.1. The simulator operates directly in the Pauli–Liouville representation and uses the same augmented Bloch-vector formalism $\mathbf{s} = [1, \langle X \rangle, \langle Y \rangle, \langle Z \rangle]^\top$. Its output is a reconstructed Pauli transfer matrix R obtained from synthetic measurement data generated through finite-sample statistics. The simulator does not impose trace preservation or complete positivity during reconstruction; instead, it reproduces the linear inversion structure used in Qiskit and therefore allows nonunital and non-trace-preserving behaviour to appear explicitly in the PTM, mirroring what is observed on hardware.

The simulator evaluates exactly the twelve circuits that make up single-qubit process tomography: four state preparations $\{|Z_+\rangle, |Z_-\rangle, |X_+\rangle, |Y_+\rangle\}$ combined with measurements in the three Pauli bases Z, X , and Y . For each preparation p , the simulator computes or samples the expectation values $\langle X \rangle_p, \langle Y \rangle_p$, and $\langle Z \rangle_p$, as well as a basis-independent survival statistic s_p , which corresponds to the total detection probability for that input state. These four survival probabilities determine the first row of the PTM. The remaining twelve expectation values determine the affine translation vector (t_x, t_y, t_z) and the 3×3 linear block T , following the same relations used during experimental reconstruction:

$$t = \frac{1}{2}(r_{Z_+} + r_{Z_-}), \quad T = [r_{X_+} - t \quad r_{Y_+} - t \quad r_{Z_+} - t],$$

where $r_p = [\langle X \rangle_p, \langle Y \rangle_p, \langle Z \rangle_p]^\top$. The first row $(R_{00}, R_{01}, R_{02}, R_{03})$ is obtained from the four survival values through

$$s_{Z_+} = R_{00} + R_{03}, \quad s_{Z_-} = R_{00} - R_{03}, \quad s_{X_+} = R_{00} + R_{01}, \quad s_{Y_+} = R_{00} + R_{02},$$

which uniquely determines the coefficients even if the process is not trace preserving. In the trace-preserving limit, all $s_p = 1$ and the row reduces to $(1, 0, 0, 0)$. In general, allowing $s_p \neq 1$ provides access to loss and leakage mechanisms that affect probability mass, enabling a reconstruction that includes non-trace-preserving behaviour directly in the PTM.

All measurement results are produced through a two-stage binomial sampling procedure that mirrors hardware behaviour. For a given shot budget N and preparation p , a single effective detection count N_{eff} is drawn from a binomial distribution with success probability equal to the true survival

s_p times a readout-efficiency factor. This detection count determines the number of detected events shared across all three measurement axes for the same input state. Conditional on detection, each axis measurement draws from a binomial model whose success probability equals the analytic conditional probability for that axis. This construction ensures that survival fluctuations affect only the first row of the PTM, while axis-dependent fluctuations affect only the affine translation and the Pauli block. In the limit $N \rightarrow \infty$, both the survival estimates and the expectation values converge to their analytic values, and the reconstructed PTM converges to the true PTM of the underlying channel.

The target operation is generated by exponentiating native single-qubit generators of the IBM gate set. For an ideal gate $U(\theta) = \exp[-i(\theta/2)G]$, the simulator introduces controlled deviations by replacing the generator with $G + \Delta$, yielding

$$U_{\Delta}(\theta) = \exp\left[-i\frac{\theta}{2}(G + \Delta)\right].$$

Here, Δ may be Hermitian, capturing coherent miscalibrations such as axis tilts or over- and underrotations, or fully general (non-Hermitian), producing effective nonunitary behaviour. These modifications alter both the Bloch-sphere rotation and the overall state norm, allowing the model to express nonunital and non-trace-preserving contributions. Because the simulator does not constrain the PTM to satisfy trace-preserving or positivity conditions, any such generator-level perturbation can propagate into all entries of the reconstructed PTM. This makes the simulator sensitive to symmetry-violating effects, including those that alter the first row or induce translations in the Bloch vector.

Repeating the twelve-circuit simulation with independent random seeds produces an ensemble of PTMs that reflects both sampling fluctuations and the statistical variability caused by generator perturbations. The resulting per-entry means and standard deviations are compared to hardware PTMs obtained through Qiskit's tomography routines. This enables a direct, elementwise comparison between model predictions and experiment, revealing which types of generator perturbations best reproduce the structure of noise observed on the device, including deviations from unitality, coherent miscalibrations, and non-trace-preserving effects.

3.4. Genetic Algorithm for Fitting Generator-Level Noise

To relate experimentally observed process matrices to effective microscopic noise parameters, we use a genetic algorithm (GA) to fit additive generator perturbations for the native single-qubit gates. The goal is to find a set of perturbations that, when inserted into the simulator described in Sec. 3.3 and reconstructed via the twelve-circuit QPT scheme of Sec. 3.1, reproduces the target Pauli transfer matrix R_{target} obtained from hardware. In particular, we do not impose trace preservation, unitality, or Hermiticity constraints on the fitted channels, so that symmetry-violating and non-trace-preserving effects can be expressed explicitly in the fitted PTMs.

We consider a fixed set of native gates $\mathcal{G} = \{R_z, SX, X\}$. Each gate $g \in \mathcal{G}$ is generated by an ideal Hermitian operator G_g and implemented as

$$U_g(\theta_g) = \exp\left[-i\frac{\theta_g}{2}G_g\right], \quad (26)$$

with a gate-dependent rotation angle θ_g . To model imperfect behaviour, we introduce an additive perturbation $\Delta_g \in \mathbb{C}^{2 \times 2}$ at the generator level and replace

$$U_g(\theta_g) \mapsto U_g^{\Delta}(\theta_g) = \exp\left[-i\frac{\theta_g}{2}(G_g + \Delta_g)\right]. \quad (27)$$

The matrices Δ_g are general complex 2×2 matrices and therefore allow for both Hermitian (coherent) and non-Hermitian (effectively nonunitary) contributions. For a given collection $\{\Delta_g\}_{g \in \mathcal{G}}$, the simulator constructs the noisy circuits, performs shot-based QPT-12 as in Sec. 3.3, and outputs a reconstructed PTM $R_{\text{fit}}(\{\Delta_g\})$.

Each GA individual encodes a full set of generator perturbations $\{\Delta_g\}_{g \in \mathcal{G}}$. To obtain a convenient, bounded parameterization, we represent each Δ_g by its eight real components (real and imaginary parts of its four entries) and stack them into a real vector $w_g \in \mathbb{R}^8$. The complete genome is

$$w = (w_{R_z}, w_{S_X}, w_X) \in \mathbb{R}^{8|\mathcal{G}|}. \quad (28)$$

These real parameters are mapped to bounded complex matrices through a componentwise saturating transform

$$\Delta_g(w_g) = \Delta_{\max} \tanh(w_g), \quad (29)$$

where $\Delta_{\max} > 0$ controls the maximum allowed magnitude of each matrix element. The tanh map ensures that the search space is continuous but bounded, preventing unphysical growth of generator entries while still allowing both small and moderate perturbations. The vector w thus parametrizes the entire noisy gate set, and through the simulator it induces a PTM $R_{\text{fit}}(w)$.

The fitting objective is formulated as a weighted least-squares discrepancy between the fitted PTM and the target PTM,

$$J(w) = \sum_{i,j} \omega_{ij} (R_{\text{fit}}(w)_{ij} - R_{\text{target}ij})^2, \quad (30)$$

where the indices i, j run over the PTM entries and $\omega_{ij} \geq 0$ are tunable weights. In the simplest case, $\omega_{ij} = 1$ for all entries, corresponding to an unweighted Frobenius norm. Optionally, the weights can incorporate information about the empirical standard deviation of each PTM entry, implementing an approximate χ^2 -type cost, or can be amplified for entries where $|R_{\text{target}ij}|$ is small. The latter increases sensitivity to weak but systematic features of the channel that might otherwise be dominated by larger entries. The GA uses the fitness function

$$F(w) = -J(w), \quad (31)$$

so that better agreement with the target corresponds to higher fitness.

The GA maintains a population of N_{pop} genomes $\{w^{(k)}\}_{k=1}^{N_{\text{pop}}}$. An initial population is created by sampling the components of each $w^{(k)}$ from a centered distribution and clipping them to a fixed range $[-W_{\max}, W_{\max}]$. For each individual, the simulator constructs the perturbed gates, runs the twelve-circuit QPT protocol with a fixed shot budget, reconstructs $R_{\text{fit}}(w^{(k)})$, and evaluates the cost $J(w^{(k)})$. Because the simulator operates in a shot-based regime, the objective $J(w)$ is stochastic: repeated evaluations of the same genome yield slightly different PTMs and therefore slightly different costs. This reflects the statistical uncertainty present in the hardware data and encourages solutions that are robust to sampling noise.

Each GA generation consists of selection, crossover, mutation, and elitism. First, the population is ranked by fitness, and a small elite set of the best-performing genomes is copied unchanged into the next generation. The remaining slots are filled with offspring generated from parent genomes drawn from the top-performing fraction of the current population. Offspring genomes are created by uniform crossover: for each component of the child vector, the entry is inherited from one of the two parents with equal probability. After crossover, mutation introduces diversity and enables exploration of the parameter space. Mutation acts on a random fraction of the non-elite individuals and perturbs a random subset of their components by adding Gaussian noise with fixed variance, followed by clipping back to $[-W_{\max}, W_{\max}]$. With a smaller probability, mutation can replace an entire genome by a new random vector, which helps escape local minima. Elitism ensures that the current best solutions are never degraded by mutation or crossover, while selection pressure and recombination guide the population towards regions of lower cost.

The GA is run for a fixed number of generations T . The best genome w^* observed over all generations defines the final set of generator perturbations $\{\Delta_g^*\}$ via the saturating map. These perturbations are then passed to the simulator to compute the corresponding PTM, both in the analytic

(shot-free) limit and in the shot-based QPT-12 setting. Repeating the shot-based reconstruction for $\{\Delta_g^*\}$ yields an ensemble $\{R_{\text{fit}}^{(r)}\}$ with mean \bar{R}_{fit} and standard deviation σ_{fit} for each matrix element. These quantities are compared element-wise to the hardware PTMs and their empirical uncertainties, for example through combined-uncertainty z-scores. In this way, the GA provides an effective inverse map from observed process-level asymmetries—including nonunital and non-trace-preserving behaviour visible in all entries of R —back to plausible generator-level noise contributions for the native gate set.

4. Experiments

The experiments in this work combine repeated process tomography on IBM Quantum hardware with simulations based on the shot-based QPT-12 framework introduced in Secs. 3.1–3.3. The hardware experiments provide empirical Pauli transfer matrices (PTMs) and their statistical variation, which we treat as reference data. The simulations implement generator-level noise models and use the same twelve-circuit process tomography to obtain synthetic PTMs that can be compared element-wise to the hardware results.

All hardware data were acquired on the `ibm_brisbane` backend, an Eagle r3 device with native single-qubit gate set $\{R_z(\theta), SX, X\}$ [27,29]. For a chosen physical qubit, we ran the full twelve-circuit single-qubit process tomography schedule for the X gate, using Qiskit’s tomography utilities [4,9,26]. Each tomography run prepared the four input states $\{|Z_+\rangle, |Z_-\rangle, |X_+\rangle, |Y_+\rangle\}$, measured in the three Pauli bases Z, X, Y , and used a fixed number of shots per circuit. This complete set of twelve circuits was executed many times as independent jobs under comparable calibration conditions. For each run we reconstructed a PTM $R^{(k)}$ from the twelve-circuit data. Aggregating over all runs yields a sample mean

$$R_{\text{target}} = \frac{1}{K} \sum_{k=1}^K R^{(k)} \quad (32)$$

and a component-wise sample standard deviation σ_{target} , which together serve as the empirical description of the implemented X -gate channel on hardware. All subsequent simulations are evaluated against R_{target} and σ_{target} using the PTM formalism of Sec. 3.1, with trace preservation not enforced so that possible loss and leakage effects remain visible in the first row of the PTM.

4.1. Genetic-Algorithm Reconstruction of Generator noise

The first simulation experiment uses the genetic algorithm (GA) of Sec. 3.4 to infer generator-level perturbations that reproduce the hardware PTM for the X -Gate. The optimization variables are the additive perturbations $\{\Delta_g\}_{g \in \mathcal{G}}$ to the native generators G_g of the gates $g \in \mathcal{G} = \{R_z, SX, X\}$. Each Δ_g is a complex 2×2 matrix parametrized by eight real numbers, which are mapped to a bounded range via a componentwise tanh-transform with scale $\Delta_{\text{max}} = 0.05$. The full genome $w \in \mathbb{R}^{8|\mathcal{G}|}$ thus encodes a complete noisy gate set, and through the shot-based QPT-12 simulator it induces a fitted PTM $R_{\text{fit}}(w)$.

For a given genome w , the simulator constructs the perturbed unitaries

$$U_g^\Delta(\theta_g) = \exp\left[-i \frac{\theta_g}{2} (G_g + \Delta_g(w))\right], \quad g \in \mathcal{G}, \quad (33)$$

inserts them into the twelve canonical tomography circuits, and runs the QPT-12 protocol with a shot budget of $N_{\text{shots}} = 10^4$ per circuit. The reconstruction routine of Sec. 3.3 then produces a PTM $R_{\text{fit}}(w)$ without imposing trace preservation (TP) or unitality. The cost function is a weighted least-squares discrepancy

$$J(w) = \sum_{i,j} \omega_{ij} (R_{\text{fit}}(w)_{ij} - R_{\text{target}ij})^2, \quad (34)$$

with default weights $\omega_{ij} = 1$ and amplification of entries with small $|R_{\text{target}ij}|$ (e.g., $|R_{\text{target}ij}| < 0.05$) by a factor of 3.5. This makes the optimizer more sensitive to weak but systematic features in the PTM. The GA maximizes the fitness $F(w) = -J(w)$.

The GA population consists of $N_{\text{pop}} = 10^4$ individuals. Initial genomes are drawn from a centered distribution and clipped to $[-W_{\text{max}}, W_{\text{max}}]$ with $W_{\text{max}} = 6$. In each generation, individuals are ranked by fitness. An elite set of $N_{\text{elite}} = 10^3$ genomes is copied unchanged into the next generation, ensuring that the current best solutions are preserved. The remaining population is filled with offspring created from parents selected from the top fraction of the population (here the best 50% by fitness). Offspring genomes are produced by uniform crossover: for each parameter, the child inherits the value from one of the two parents with equal probability. Mutation is applied to a random subset of non-elite individuals by adding Gaussian noise to a randomly selected subset of parameters and re-clipping to $[-W_{\text{max}}, W_{\text{max}}]$. With a smaller probability, mutation can replace an entire genome by a new random vector, which helps the population escape local minima. The GA is run for $T = 3 \times 10^4$ generations, and the best genome observed across all generations is retained as w^* .

For the final analysis, the perturbations $\{\Delta_g^*\}$ encoded by w^* are fed back into the simulator. We compute both the analytic PTM (shot-free limit) and shot-based PTMs using the QPT-12 scheme with $N_{\text{shots}} = 10^4$. Repeating the shot-based reconstruction many times yields an ensemble of PTMs $\{R_{\text{fit}}^{(r)}\}$ with mean \bar{R}_{fit} and standard deviation σ_{fit} . These are compared element-wise to the hardware PTM via the combined-uncertainty z-score

$$z_{ij} = \frac{|\bar{R}_{\text{fit},ij} - R_{\text{target},ij}|}{\sqrt{\sigma_{\text{target},ij}^2 + \sigma_{\text{fit},ij}^2}}, \quad (35)$$

which quantifies how well the fitted generator-level noise reproduces the structure of the experimentally observed channel, including deviations from unitality and trace preservation in the first row of R .

4.2. Random Hermitian Versus Non-Hermitian Noise Ensembles

The second simulation experiment studies the structural impact of Hermitian versus non-Hermitian generator perturbations by sampling large random ensembles. The objective is to determine whether purely coherent (Hermitian) generator noise can account for the symmetry-violating and non-trace-preserving features of the hardware PTM, or whether general non-Hermitian perturbations are required.

For both noise families, we consider the same set of native gates $\mathcal{G} = \{R_z, SX, X\}$. In the Hermitian ensemble, each Δ_g is drawn as a random Hermitian matrix with real diagonal entries and complex-conjugate off-diagonal entries, with all real parameters sampled uniformly from $[-\lambda, \lambda]$ for a fixed scale $\lambda = 0.05$. This produces small random axis tilts and over- or under-rotations while preserving Hermiticity at the generator level; the corresponding noisy gates $U_g^\Delta(\theta_g)$ are unitary, and any nonunitarity in the reconstructed PTM arises solely from the shot-based tomography and reconstruction procedure. In the non-Hermitian ensemble, each Δ_g is a fully complex 2×2 matrix with real and imaginary parts sampled independently and uniformly from $[-\lambda, \lambda]$. These perturbations generate effective nonunitary dynamics at the level of the gate, allowing for trace-decreasing and nonunital behaviour that can manifest directly in all entries of the PTM, including the first row.

For each family, we draw $N_{\text{samples}} = 10^3$ independent noise maps $\{\Delta_g^{(k)}\}$. For each sample k , we construct the perturbed gates U_g^Δ , run the QPT-12 simulator with $N_{\text{shots}} = 10^4$, and reconstruct a PTM $R^{(k)}$ using the same shot-based procedure as in Sec. 3.3, again without enforcing trace preservation. This yields two ensembles of PTMs,

$$\{R_{\text{herm}}^{(k)}\}_{k=1}^{N_{\text{samples}}} \quad \text{and} \quad \{R_{\text{nonherm}}^{(k)}\}_{k=1}^{N_{\text{samples}}}, \quad (36)$$

from which we compute the component-wise means and standard deviations

$$\bar{R}_{\text{herm}} = \frac{1}{N_{\text{samples}}} \sum_k R_{\text{herm}}^{(k)}, \quad \sigma_{\text{herm}}, \quad \bar{R}_{\text{nonherm}} = \frac{1}{N_{\text{samples}}} \sum_k R_{\text{nonherm}}^{(k)}, \quad \sigma_{\text{nonherm}}. \quad (37)$$

These ensemble means describe the typical process-level distortions produced by small Hermitian and non-Hermitian generator noise, respectively.

To compare the ensembles to the hardware data, we compute element-wise z-scores relative to the hardware mean and variance. For the Hermitian family, we define

$$z_{ij}^{(\text{herm})} = \frac{|\bar{R}_{\text{herm},ij} - R_{\text{target},ij}|}{\sqrt{\sigma_{\text{target},ij}^2 + \sigma_{\text{herm},ij}^2}}, \quad (38)$$

and analogously for the non-Hermitian family,

$$z_{ij}^{(\text{nonherm})} = \frac{|\bar{R}_{\text{nonherm},ij} - R_{\text{target},ij}|}{\sqrt{\sigma_{\text{target},ij}^2 + \sigma_{\text{nonherm},ij}^2}}. \quad (39)$$

Entries with $z_{ij} \lesssim 1$ are statistically compatible with the hardware within one combined standard deviation, while larger values indicate discrepancies that cannot be explained by the corresponding ensemble. Inspecting these z-score matrices reveals which aspects of the experimental PTM can be reproduced by purely Hermitian generator noise and which require general non-Hermitian contributions. In particular, systematic deviations in the first row of R and in the affine components associated with nonunitarity provide direct evidence for effective non-trace-preserving dynamics that cannot be captured by Hermitian perturbations alone.

5. Results

This section reports the experimentally reconstructed single-qubit process matrices obtained on `ibm_brisbane`, the fitted generator perturbations from the genetic optimization, and the corresponding model process matrices. All PTMs are defined in the Pauli transfer matrix formalism introduced in Sec. 3.1, where the channel acts as $\mathbf{s}_{\text{out}} = R \mathbf{s}_{\text{in}}$. Agreement between model and experiment is quantified using the element-wise z-scores defined in Secs. 4.1–4.2, based on the combined standard deviation of hardware and simulator statistics.

5.1. Hardware Target: Averaged PTM and Standard Deviation

We aggregated $n_{\text{runs}} = 88$ independent single-qubit process tomography runs for the physical implementation of the X gate on `ibm_brisbane`. Each run comprised the full twelve-circuit QPT-12 schedule described in Sec. 3.1. From the reconstructed PTMs $\{R^{(k)}\}_{k=1}^{n_{\text{runs}}}$ we computed the component-wise sample mean (Eq. 32) and the associated sample standard deviation $\text{STD}[R_{\text{target}}]$. These quantities summarize the effective single-qubit channel realized on hardware, including deviations from unitality and trace preservation that appear in the first row and in the Pauli block.

The aggregated hardware PTM and its standard deviation are

$$R_{\text{target}} = \begin{bmatrix} 1.00000 & -0.04382 & 0.04596 & -0.04950 \\ -0.00196 & 0.90412 & 0.01792 & -0.01845 \\ -0.00283 & 0.01961 & -0.90994 & 0.00360 \\ 0.00240 & 0.00184 & 0.00348 & -0.90794 \end{bmatrix} \quad (40)$$

$$\text{STD}[R_{\text{target}}] = \begin{bmatrix} 0.00000 & 0.02959 & 0.02583 & 0.02777 \\ 0.00278 & 0.02618 & 0.01462 & 0.01233 \\ 0.00349 & 0.01474 & 0.02588 & 0.01213 \\ 0.00299 & 0.00763 & 0.00841 & 0.02765 \end{bmatrix}. \quad (41)$$

For comparison, the ideal noiseless Pauli transfer matrix of the single-qubit X gate in the ordered basis (I, X, Y, Z) is

$$R_X = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \quad (42)$$

which represents a unitary, trace-preserving, and unital channel. The nontrivial entries in the first row of R_{target} indicate effective deviations from ideal trace preservation at the process level, while contractions along the Pauli-diagonal entries reveal relaxation- and dephasing-like behaviour relative to the ideal reference R_X . These features serve as the reference signatures that the simulator and generator-level noise models aim to reproduce.

5.2. Fitted Generator Perturbations

Using the genetic algorithm described in Sec. 4.1, we fitted additive generator perturbations Δ_g for the native single-qubit gate set $\{R_z, SX, X\}$. Each Δ_g is a complex 2×2 matrix that modifies the corresponding ideal generator G_g at the Hamiltonian level, $G_g \mapsto G_g + \Delta_g$, and thus defines a perturbed gate

$$U_g^\Delta(\theta_g) = \exp\left[-i \frac{\theta_g}{2} (G_g + \Delta_g)\right].$$

The GA operates on a bounded real parametrization of these matrices and optimizes the generator-level parameters to minimize the PTM discrepancy $J(w)$ in Eq. 30, using shot-based QPT-12 reconstruction at each evaluation.

The best-found perturbations are reported here as separate real and imaginary parts:

$$\begin{aligned} \Delta[R_z]_{\text{Re}} &= \begin{bmatrix} -0.04985 & 0.01624 \\ -0.04976 & 0.04994 \end{bmatrix}, & \Delta[R_z]_{\text{Im}} &= \begin{bmatrix} -0.01773 & 0.01577 \\ 0.03355 & -0.01346 \end{bmatrix}, \\ \Delta[SX]_{\text{Re}} &= \begin{bmatrix} 0.04971 & -0.00908 \\ 0.02513 & -0.04678 \end{bmatrix}, & \Delta[SX]_{\text{Im}} &= \begin{bmatrix} -0.01871 & 0.04548 \\ 0.02234 & -0.05000 \end{bmatrix}, \\ \Delta[X]_{\text{Re}} &= \begin{bmatrix} 0.01588 & 0.03059 \\ -0.01421 & -0.01726 \end{bmatrix}, & \Delta[X]_{\text{Im}} &= \begin{bmatrix} 0.02428 & -0.01665 \\ -0.00771 & 0.00382 \end{bmatrix}. \end{aligned}$$

These matrices encode both Hermitian and non-Hermitian contributions and therefore allow nonunital, non-trace-preserving, and symmetry-violating behaviour at the process level. The next subsection shows how these fitted perturbations translate into process matrices via QPT-12.

5.3. Model Reconstructions (QPT-12) with Shot Sampling

Given the fitted perturbations $\{\Delta_g\}$, we use the shot-based simulator of Sec. 3.3 together with the twelve-circuit scheme of Sec. 3.1 to reconstruct model PTMs. We consider a finite-shot regime with shots = 10^4 per circuit. In this regime, we perform $N = 1000$ independent QPT-12 reconstructions and compute the mean and standard deviation across these realizations, consistent with the statistical treatment in Sec. 4.1.

Sampled model (shots= 10^4).

The reconstructed PTM is random due to shot noise and the two-stage sampling of survival and conditional expectation values. Averaging over $N = 1000$ independent QPT-12 reconstructions yields

$$R_{\Delta, 10^4}^{\text{mean}} = \begin{bmatrix} 0.96576 & -0.03227 & 0.03424 & -0.03424 \\ -0.00314 & 0.91424 & 0.02030 & -0.01796 \\ -0.02546 & 0.02304 & -0.96596 & 0.00099 \\ 0.00001 & 0.00207 & 0.00169 & -0.99999 \end{bmatrix} \quad (43)$$

$$\text{STD}[R_{\Delta, 10^4}] = \begin{bmatrix} 0.00124 & 0.00295 & 0.00124 & 0.00124 \\ 0.00695 & 0.00809 & 0.01189 & 0.00708 \\ 0.00718 & 0.01276 & 0.00729 & 0.00733 \\ 0.00004 & 0.01004 & 0.01009 & 0.00004 \end{bmatrix}. \quad (44)$$

In this shot-based regime, the first row deviates from $(1, 0, 0, 0)$ and exhibits nontrivial off-diagonal entries consistent with the symmetry-violating signatures seen in R_{target} . The Pauli-diagonal entries in the lower 3×3 block show contraction relative to the ideal X-gate PTM, again in line with the hardware data.

5.4. Agreement with Hardware: Entry-Wise and Block-Wise

To quantify the agreement between model and hardware, we compare the sampled model PTM $R_{\Delta, 10^4}^{\text{mean}}$ and R_{target} entry-wise. Following the uncertainty framework in Sec. 4.1, we define the element-wise difference

$$\Delta R_{ij} = R_{\Delta, 10^4, ij}^{\text{mean}} - R_{\text{target}, ij}$$

and the combined standard deviation

$$\sigma_{\text{comb}, ij} = \sqrt{\text{STD}[R_{\text{target}}]_{ij}^2 + \text{STD}[R_{\Delta, 10^4}]_{ij}^2},$$

which treats the hardware and simulator uncertainties as independent. The z-score

$$z_{ij} = \frac{|\Delta R_{ij}|}{\sigma_{\text{comb}, ij}}$$

is identical in form to the definition used in Secs. 4.1–4.2 and measures how many combined standard deviations the model deviates from the hardware mean for each PTM entry. Entries with $z_{ij} \leq 1$ are statistically compatible at the one-sigma level; entries with $z_{ij} > 1$ indicate tension that may point to missing physical mechanisms in the model.

Table 1 summarizes the resulting z-scores for $R_{\Delta, 10^4}^{\text{mean}}$ relative to R_{target} ; entries with $z \leq 1$ are marked with \checkmark .

Table 1. Entry-wise comparison for the sampled model (shots= 10^4): $z = |\Delta R|/\sigma_{\text{comb}}$ and 1σ consistency. Outliers ($z > 1$) highlight where the model does not explain hardware within combined uncertainty. Rows and columns are ordered as (I, X, Y, Z) .

| | z (\checkmark if $z \leq 1$) | | | |
|--------------|------------------------------------|-------------------|-------------------|-------------------|
| $(0, \cdot)$ | 27.6 | \checkmark 0.39 | \checkmark 0.45 | \checkmark 0.55 |
| $(1, \cdot)$ | \checkmark 0.16 | \checkmark 0.37 | \checkmark 0.13 | \checkmark 0.03 |
| $(2, \cdot)$ | 2.84 | \checkmark 0.18 | 2.08 | \checkmark 0.19 |
| $(3, \cdot)$ | \checkmark 0.80 | \checkmark 0.02 | \checkmark 0.14 | 3.33 |

Most entries fall within one combined standard deviation. Four entries exceed 1σ and cluster in components strongly influenced by trace preservation and diagonal relaxation:

- *Affine row (0,1:3)*: The first-row off-diagonals are reproduced within uncertainty, with $z \in [0.39, 0.55]$. These entries encode the symmetry-violating affine shifts discussed in Sec. 3.1 and are well captured by the fitted generator perturbations.
- *Pauli-diagonal entries*: The largest deviations occur at (2,2) (Y, Y) and (3,3) (Z, Z), with $z \approx 2.08$ and $z \approx 3.33$, respectively. These entries summarize effective longitudinal and transverse damping and indicate that the generator-level model does not fully capture the dissipative structure present on hardware.
- *First column*: The entry (2,0) shows $z \approx 2.84$, while other entries in the first column remain within uncertainty. This is consistent with additional state-preparation, measurement, or leakage effects that are not fully represented by the generator perturbations alone.
- *Element (0,0)*: The hardware reconstruction enforces normalization such that $\text{STD}[R_{\text{target},00}] = 0$, while the sampled model yields $R_{\Delta, 10^4, 00}^{\text{mean}} \approx 0.966$. Any deviation at this entry is therefore counted as a significant outlier.

Viewed at the block level, the off-diagonal entries of the Pauli 3×3 block are largely consistent within combined uncertainty once shot noise is included. Disagreement is concentrated on the Pauli diagonals and on R_{00} . This pattern suggests that symmetry-violating generator perturbations are sufficient to reproduce the affine first-row structure and cross-axis mixing, while additional mechanisms—such as explicit dissipative channels, SPAM, crosstalk, or slow drift—are needed to match the observed diagonal damping in detail.

5.5. Ensemble Perturbations: Random Hermitian vs. Non-Hermitian Noise

The ensemble experiments of Sec. 4.2 complement the single-instance GA fit by testing whether random Hermitian or non-Hermitian generator perturbations can statistically reproduce the hardware PTM. The same z -score formalism

$$z_{ij} = \frac{|\bar{R}_{\text{ensemble},ij} - R_{\text{target},ij}|}{\sqrt{\text{STD}[R_{\text{target},ij}]^2 + \text{STD}[R_{\text{ensemble},ij}]^2}}$$

is used, with ensemble means $\bar{R}_{\text{ensemble}}$ and ensemble standard deviations as defined in Sec. 4.2.

For reference, the hardware mean and standard deviation for this comparison are

$$R_{\text{target}}^{\text{mean}} = \begin{bmatrix} 1.00000 & -0.04382 & 0.04596 & -0.04950 \\ -0.00196 & 0.90412 & 0.01792 & -0.01845 \\ -0.00283 & 0.01961 & -0.90994 & 0.00360 \\ 0.00240 & 0.00184 & 0.00348 & -0.90794 \end{bmatrix} \quad (45)$$

$$\text{STD}[R_{\text{target}}] = \begin{bmatrix} 0.00000 & 0.02959 & 0.02583 & 0.02777 \\ 0.00278 & 0.02618 & 0.01462 & 0.01233 \\ 0.00349 & 0.01474 & 0.02588 & 0.01213 \\ 0.00299 & 0.00763 & 0.00841 & 0.02765 \end{bmatrix}. \quad (46)$$

For the Hermitian ensemble, the mean PTM and its standard deviation over $N_{\text{samples}} = 10^3$ draws are

$$\bar{R}_{\text{Herm}} = \begin{bmatrix} 1.000000 & 0.000000 & 0.000000 & 0.000000 \\ -0.001507 & 0.984212 & -0.003987 & -0.000833 \\ -0.004144 & -0.001489 & -0.984492 & 0.014013 \\ -0.005702 & 0.003505 & -0.002475 & -0.989418 \end{bmatrix} \quad (47)$$

$$\text{STD}[R_{\text{Herm}}] = \begin{bmatrix} 0.000000 & 0.000000 & 0.000000 & 0.000000 \\ 0.021101 & 0.026124 & 0.109947 & 0.076930 \\ 0.044610 & 0.117438 & 0.046393 & 0.147799 \\ 0.005614 & 0.085740 & 0.108511 & 0.009184 \end{bmatrix}. \quad (48)$$

For the non-Hermitian ensemble, the corresponding quantities are

$$\bar{R}_{\text{NonHerm}} = \begin{bmatrix} 0.961351 & -0.012407 & -0.017592 & 0.006949 \\ 0.057300 & 0.906746 & 0.030684 & -0.008628 \\ 0.055049 & 0.012691 & -1.041182 & -0.012357 \\ -0.006668 & 0.078110 & 0.092247 & -0.987106 \end{bmatrix} \quad (49)$$

$$\text{STD}[R_{\text{NonHerm}}] = \begin{bmatrix} 0.052509 & 0.062923 & 0.068657 & 0.025109 \\ 0.145101 & 0.128538 & 0.146580 & 0.089535 \\ 0.093521 & 0.128767 & 0.092401 & 0.120268 \\ 0.020724 & 0.128284 & 0.139064 & 0.020815 \end{bmatrix}. \quad (50)$$

The element-wise z-scores for both ensembles are summarized in Tables 2 and 3. Entries with $z \leq 1$ are marked with \checkmark .

Table 2. Hermitian ensemble: z-scores vs. hardware (rows/cols ordered as I, X, Y, Z).

| | I | X | Y | Z |
|-----|-------------------|-------------------|-------------------|-------------------|
| I | \checkmark 0.08 | 1.48 | 1.78 | 1.78 |
| X | \checkmark 0.02 | 2.17 | \checkmark 0.20 | \checkmark 0.23 |
| Y | \checkmark 0.03 | \checkmark 0.18 | 1.40 | \checkmark 0.07 |
| Z | 1.27 | \checkmark 0.02 | \checkmark 0.05 | 2.80 |

Table 3. Non-Hermitian ensemble: z-scores vs. hardware (rows/cols ordered as I, X, Y, Z).

| | I | X | Y | Z |
|-----|-------------------|-------------------|-------------------|-------------------|
| I | \checkmark 0.74 | \checkmark 0.45 | \checkmark 0.87 | 1.51 |
| X | \checkmark 0.41 | \checkmark 0.02 | \checkmark 0.09 | \checkmark 0.11 |
| Y | \checkmark 0.62 | \checkmark 0.05 | 1.37 | \checkmark 0.13 |
| Z | \checkmark 0.43 | \checkmark 0.59 | \checkmark 0.64 | 2.29 |

The non-Hermitian ensemble explains the hardware PTM more comprehensively than the Hermitian ensemble: 13/16 entries satisfy $z \leq 1$ for Table 3, compared to 9/16 for Table 2. In both cases, the largest tensions occur on Pauli-diagonal elements, in particular (Z, Z) with $z = 2.80$ (Hermitian) and $z = 2.29$ (non-Hermitian), and (Y, Y) with $z = 1.40$ and $z = 1.37$, respectively. These entries encode effective longitudinal and transverse damping and show that diagonal relaxation is not fully captured by either ensemble. By contrast, affine first-row entries $(0, 1:3)$, which are sensitive to symmetry-violating and non-trace-preserving behaviour, are mostly within 1σ for the non-Hermitian ensemble and only partially for the Hermitian one.

The row and column structure provides additional interpretation. First-column elements $(\cdot, 0)$ are linked to survival and preparation biases; here the Hermitian ensemble exhibits a miss at (Z, I) with $z = 1.27$, while the non-Hermitian ensemble remains within 1σ . Within the Pauli 3×3 block, most off-diagonal entries are consistent once ensemble variability is accounted for, indicating that random gate-level perturbations—Hermitian or non-Hermitian—capture coherent cross-axis mixing. The remaining diagonal excess points to dissipative or drift mechanisms beyond static generator noise. Overall, the non-Hermitian ensemble provides a closer match to the hardware PTM, particularly for affine-row entries and off-diagonal Pauli couplings, supporting the conclusion that effective non-Hermitian contributions are needed to model the observed symmetry violations on the device.

5.6. Summary

- The fitted generator-level model reproduces the affine first-row signatures and most off-diagonal structure within combined uncertainties once finite-shot sampling is included (12/16 entries within 1σ).
- The main discrepancies are on (Y, Y) and (Z, Z) and, to a lesser extent, (Y, I) , indicating additional dissipative channels or SPAM/drift not captured by Δ alone.
- In magnitude, large differences (above error bars) are confined to the Pauli diagonals and $(0, 0)$; small differences (well within error bars) dominate the affine row and off-diagonal Pauli entries.
- Overall, symmetry-violating generator perturbations explain the components indicative of symmetry breaking; other sources (dissipation, SPAM, crosstalk, slow drift) are required to close the remaining gap on the diagonals.

6. Discussion

The comparison between the experimentally reconstructed and simulated process matrices reveals in quantitative terms how well the model captures the physical behavior of the hardware. Each element of the Pauli transfer matrix (PTM) describes how a given input Pauli operator is transformed by the gate, with diagonal entries corresponding to attenuation or amplification of specific basis components and off-diagonal terms representing cross-axis mixing. In the absence of noise, the PTM is ideally orthogonal and trace-preserving, corresponding to a completely positive and unital map [3,28,30]. Deviations from this structure therefore directly encode dissipative and coherent errors, and the element-wise z -score analysis $z = |\Delta|/\sigma_{\text{comb}}$ provides a rigorous statistical measure of consistency between model and experiment [9,26].

The hardware variability, characterized by the standard deviations of the reconstructed PTMs, sets the scale of meaningful comparison. Off-diagonal elements in the affine row exhibit larger statistical spread, indicating sensitivity to state-preparation and measurement fluctuations, while the Pauli-diagonal entries typically display smaller variance. Where the hardware variance is low, even small systematic deviations of the model lead to large z values. Conversely, entries with higher noise tolerate greater discrepancies before they become statistically significant. This behavior is consistent with the fundamental sensitivity of process tomography to trace-preserving and unitality violations, where normalization errors—associated with $\text{Tr}(\rho) \neq 1$ or non-unital channels—manifest most strongly in the first PTM row [10,31,32].

The fitted generator-based model, obtained via genetic optimization of gate-local perturbations, reproduces the main affine-row structure and off-diagonal mixing within one combined standard deviation for most matrix elements. The few remaining tensions appear primarily on the Pauli-diagonal components (Y, Y) and (Z, Z) , suggesting the presence of residual dissipative effects not captured by purely coherent perturbations. This observation aligns with the expected influence of relaxation and dephasing processes on IBM transmon hardware [33–35], where asymmetries in T_1 and T_2 times cause unequal contractions along the Y and Z axes of the Bloch sphere. Such effects are visible in the PTM as anisotropic diagonal damping and are often accompanied by small affine shifts from imperfect calibration or measurement crosstalk.

The ensemble analysis further clarifies the robustness and interpretability of different perturbation models. When random Hermitian perturbations ($\Delta = \Delta^\dagger$) are sampled across many realizations, the resulting mean PTM remains nearly orthogonal but underestimates dissipative asymmetries. In contrast, the non-Hermitian ensemble, which allows asymmetric and complex-valued perturbations, captures the hardware behavior more comprehensively, yielding a larger fraction of entries with $z \leq 1$. This pattern indicates that the real device noise cannot be described as purely unitary or symmetry-preserving, but rather exhibits mild non-Hermitian character—consistent with weak amplitude damping, residual leakage, and state-dependent relaxation [7,36,37]. The ensemble's standard deviation quantifies the model-side uncertainty, which broadens σ_{comb} and ensures statistically balanced comparison.

From a mathematical perspective, the non-Hermitian perturbations effectively emulate channels that are not trace-preserving or not unital, representing physically dissipative maps. These can be expressed in Lindblad form with non-Hermitian generators [5,38], which capture relaxation and dephasing processes beyond coherent over-rotations. The improved agreement of such models with the experimental PTMs suggests that the effective noise on the hardware includes contributions from these open-system dynamics. The statistical structure of the z -matrices supports this interpretation: off-diagonal PTM entries show consistent $z \leq 1$ behavior, implying that cross-axis couplings are well represented by the gate-local perturbation framework, while diagonal entries exceed $z > 2$ in some cases, revealing missing dissipative terms. This mirrors typical deviations observed in self-consistent tomography studies, where trace non-preservation or SPAM bias manifests as large diagonal residuals [11,39,40].

Finite-shot effects play a secondary yet quantifiable role. The inclusion of finite-sampling variance in σ_{comb} softens the statistical test by incorporating model-side uncertainty, improving the apparent agreement in high-precision entries. However, shot noise cannot remove systematic offsets; thus, persistent deviations on the Pauli diagonals reflect genuine model incompleteness rather than statistical limitations. This distinction underscores the utility of z -based metrics, which clearly separate stochastic fluctuations from structural mismatches—a key advantage over global error measures like fidelity or trace distance that conflate the two [9,28].

Overall, the results indicate that non-Hermitian, gate-local perturbations constitute a minimal yet powerful extension of purely coherent noise models. They capture affine shifts and cross-axis coupling while approximating dissipative contraction in a statistically meaningful way. The element-wise analysis provides a transparent map between experimental data and physical mechanisms: small z values along off-diagonal PTM components correspond to well-modeled coherent mixing, while high z values along the diagonal mark unmodeled relaxation or measurement-induced asymmetry. This quantitative structure connects naturally to the physical error channels of superconducting qubits—energy relaxation, dephasing, and residual measurement bias—and thus bridges microscopic noise sources with process-level observables [7,32,33].

These findings point toward a practical modeling strategy for calibration and simulation workflows: (i) introduce gate-local generator perturbations that include non-Hermitian terms to emulate open-system effects, (ii) use ensemble statistics to represent uncertainty in both model and experiment, and (iii) quantify agreement through entry-wise z -scores rather than aggregate norms. Such an approach aligns with recent proposals for noise spectroscopy and model-based error mitigation [16,41,42]. Extending this framework to multi-qubit PTMs would allow correlated and crosstalk-induced noise to be analyzed under the same statistical principle, providing a more granular picture of how symmetry breaking propagates through realistic quantum processors.

In summary, the detailed analysis of element-wise deviations, combined with ensemble perturbation modeling, demonstrates that non-Hermitian noise components play a measurable role in the effective dynamics of current superconducting qubits. This supports a broader view of quantum hardware as an open, weakly non-unitary system whose deviations from ideal symmetry can be meaningfully parameterized, quantified, and statistically validated using process tomography.

6.1. Significance of Symmetry-Violating Generator Perturbations

The presence of symmetry-violating generator perturbations in the fitted models provides a direct and experimentally backed indication that real superconducting qubit hardware operates outside the idealized symmetry-preserving regime usually assumed in analytical noise models. In standard descriptions of gate noise, one often assumes that small imperfections manifest as coherent over-rotations, Pauli-stochastic mixtures, or weak amplitude/phase damping. These channels preserve the algebraic structure of the ideal generator set and therefore respect the native symmetries of the intended gate operations. Our results show that this assumption does not generally hold: the reconstructed PTMs display first-row deviations, affine shifts, and non-unital terms that cannot be captured by symmetric generator perturbations. The fact that our fitted non-Hermitian Δ models reproduce these

deviations with statistically meaningful precision demonstrates that symmetry-violating errors are not negligible artifacts but systematic features of the device-level dynamics.

This matters because symmetry-violating noise reshapes the operational regime in which calibration, verification, and fault tolerance techniques must function. Many analytical guarantees in error mitigation and quantum control implicitly rely on noise channels being approximately trace-preserving, unital, or symmetry-respecting in their action on the Bloch sphere. When the physical device violates these properties, then common simplifications—such as depolarizing approximations, Pauli error models, or uniform amplitude-damping assumptions—no longer capture the relevant error signatures. The fitted generator perturbations therefore reveal a more complete and physically motivated error landscape. In this landscape, certain deviations (e.g. affine drifts or survival loss) originate directly from symmetry-breaking microscopic mechanisms rather than calibration inaccuracy alone. Consequently, the noise model becomes richer, but also more actionable: it quantifies which structural assumptions fail and by what magnitude.

What follows from this is a clearer route toward actionable diagnostic and mitigation strategies. Because symmetry-violating components produce distinct, measurable patterns in the PTM (especially in the first row and Pauli-diagonal elements), they can serve as stable indicators of leakage, population loss, or state-dependent relaxation. These signatures make it possible to design targeted mitigation routines that address precisely those deviations the symmetric model cannot absorb. For example, calibration strategies could include explicit compensation of affine Bloch-vector shifts; tomography pipelines could flag symmetry-violating structure as an additional hardware metric; and simulation workflows can incorporate non-Hermitian perturbation sets as a more realistic baseline noise model. Moreover, these insights motivate the following subsection: if symmetry-violating perturbations are intrinsic to the physical gate generators, then quantum error-correction schemes that rely on gauge structures or encoded symmetries must be reconsidered in light of these violations. Understanding how symmetry breaking manifests at the generator level is therefore essential for evaluating how well symmetry-based QEC strategies will perform on actual hardware.

Beyond error modelling, these symmetry-violating generator perturbations connect directly to recent work on data obfuscation based on noisy Lie-group feature maps. Raubitzeck et al. [18] show that injecting controlled perturbations into Lie-algebra generators inside the exponential map yields transformations that are deliberately non-invertible while preserving downstream classification performance. In their construction, noise is added at the generator level such that the resulting Lie-group elements no longer implement perfectly symmetric, one-to-one maps on the data manifold, thereby preventing reliable reconstruction of the original features from the obfuscated representation while keeping the geometric structure relevant for learning. Our fitted non-Hermitian Δ play an analogous structural role: they break exact unitarity, unitality, and trace preservation at the generator level, so that distinct classical inputs can be mapped to process signatures that are no longer uniquely invertible. If such symmetry-violating patterns are stable device characteristics rather than transient calibration artifacts, then quantum information processing can in principle implement hardware-level obfuscation channels in the spirit of Ref. [18], where native generator noise is leveraged as an intrinsic mechanism for making encoded data practically unrecoverable while maintaining its utility for inference.

6.2. Implications for Error Correction: Exploiting and Mitigating Symmetry-Violating Noise

In our framework we assume that the (or at least some) symmetry-violating generator perturbations are indeed present in the actual hardware implementation of the gate set. Given this assumption, the observation of affine shifts, cross-axis mixing, and non-trace-preserving components in the reconstructed Pauli transfer matrices (PTMs) suggests potential opportunities—and constraints—for error-correction strategies.

On the one hand, the presence of identifiable symmetry violations means that the effective noise channel deviates from the ideal symmetric form and carries additional structure (for example, generator-level perturbations of the form $\Delta \neq 0$). If this structure is stable and measurable, one could

imagine tailoring error-correction or mitigation protocols that specifically address those symmetry-violating contributions. For example, by measuring syndrome outcomes that are sensitive to deviations from the symmetry subspace, one may detect leakage or loss channels that break the expected symmetry and then apply corrective operations or post-selection to restore the logical space. The literature on codes that exploit physical symmetries (or their violation) indicates that detecting violations of gauge or conservation laws can serve as error flags [43,44]. In other words, the symmetry-violating components could be turned into a resource for error detection: when a measured process shows non-ideal affine translation or survival loss (first row deviation), this signals physical error mechanisms (leakage, state-dependent loss) that standard stabilizer codes may not explicitly detect.

On the other hand, the existence of symmetry violation also poses constraints on error-correction. Recent work shows that continuous symmetries and quantum error correction (QEC) exhibit a trade-off: codes that enforce or exploit symmetries may have limited ability to correct arbitrary errors, and symmetries can impose constraints on transversal gate sets and error-correcting performance (e.g., approximate versions of the Eastin–Knill theorem) [45,46]. In our context, the mismatch between the hardware channel and the ideal symmetry-preserving generator model means that standard error-correction assumptions (unitary noise, trace-preserving CPTP maps, unital blocks) may fail. Consequently, tailoring codes to correct these symmetry-violating deviations may require extensions beyond conventional stabilizer frameworks.

In practical terms, a reverse mechanism to reduce errors might proceed as follows: first **characterize** the symmetry-violating generator perturbations (as we have done via tomography and fitting); second **augment** the encoding or recovery step with operations that specifically project away or mitigate the identified violations (for example by applying a conditional rotation or channel that cancels the affine translation, or by post-selecting on survival events); third **monitor** the first row and affine translation of the PTM as diagnostic metrics for residual symmetry violation and adjust the recovery accordingly.

However, critical caveats apply:

- The symmetry-violating generator model may only capture part of the full noise budget. Effects such as time-varying drift, crosstalk, SPAM (state preparation and measurement) errors, or correlated multi-qubit noise may remain unaddressed.
- The reverse mechanism presumes that the symmetry violation is *stable*, *repeatable*, and *measurable* on the timescale of encoding and correction. If the symmetry-violating perturbations fluctuate more rapidly than the correction cycle, then mitigation becomes more challenging.
- Introducing recovery operations targeted at symmetry violations must itself avoid introducing further asymmetries or coupling errors back into the system. Recovery operations may need to themselves commute with the native gate symmetries or be fault-tolerant under the same noise model.

In summary, the fact that symmetry-violating generator perturbations manifest in the reconstructed PTM offers both a diagnostic handle and a potential lever for error-correction strategies. While we cannot claim that the identified deviations are *directly* correctable with standard QEC codes, they point to an enriched error model where affine shifts, survival loss, and cross-axis mixing are explicitly addressed. Future work should investigate how to embed these symmetry-violation diagnostics into fault-tolerant encoding schemes and whether tailored recovery operations can systematically reduce the impact of such deviations.

7. Conclusions and Outlook

In this work, we introduced and experimentally validated a framework for modeling symmetry-violating noise in single-qubit quantum operations. The central idea was to treat noisy gates as perturbations of their ideal Hermitian generators and to examine how controlled symmetry-breaking effects manifest in reconstructed process matrices. By explicitly distinguishing Hermitian and non-Hermitian

perturbations, we demonstrated that different classes of noise—coherent and dissipative—leave characteristic signatures in the experimentally observed Pauli transfer matrices (PTMs).

Using quantum process tomography on IBM Quantum hardware, we compared the hardware-averaged PTM with sampled model reconstructions derived from a fitted perturbation map. The genetic optimization procedure produced generator perturbations that, once combined with finite-shot variability, reproduced most affine-row and off-diagonal PTM components within one combined standard deviation. The remaining discrepancies were concentrated on the Pauli-diagonal entries, consistent with missing dissipative structure or SPAM-related drift. This showed that coherent, gate-local symmetry violations explain a substantial portion of the experimentally observed asymmetries without overfitting the sampling noise.

We need to emphasize that performing these experiments on IBM quantum hardware represents a strict practical limit. Acquiring the required data is costly, and collecting the 88 hardware runs used in this study required more than half a year until IBM Brisbane was no longer available. Any extension of the work to larger-scale experiments—for example, runs on the order of 100,000 executions, would incur costs in the range of several tens of thousands of euros. This constitutes a substantial barrier for further large-scale hardware-based evaluation.

In addition, ensemble experiments with random Hermitian and non-Hermitian perturbations confirmed that non-Hermitian, symmetry-violating noise reproduces the experimental PTM more comprehensively than Hermitian perturbations alone. The non-Hermitian ensemble captured affine shifts and cross-axis couplings within uncertainty for the majority of matrix elements, while Hermitian ensembles left residual deviations on the diagonal entries associated with relaxation and dephasing. This result supports the interpretation that effective hardware noise contains both coherent and dissipative components, with the latter breaking generator-level Hermiticity and trace preservation.

These findings have several implications. First, controlled inclusion of symmetry-violating perturbations provides a transparent and quantitative way to account for affine and cross-axis effects that standard unitary error models miss. Second, the element-wise z-score analysis links deviations directly to the underlying uncertainty scales of both model and experiment, providing a diagnostic tool that complements global distance measures. Third, the systematic improvement under non-Hermitian perturbations indicates that the effective noise processes in current superconducting qubits cannot be fully described by symmetry-preserving models.

At the same time, introducing non-Hermitian components raises conceptual challenges. Once Hermiticity and unitality are relaxed, familiar assumptions about trace preservation and orthogonality in the PTM representation no longer apply. This limits direct interpretability in conventional error metrics but enhances descriptive accuracy when modeling asymmetric or dissipative channels. The observed behavior highlights the necessity of combining generator-based modeling with uncertainty-aware comparison metrics to avoid overinterpreting statistical fluctuations as physical effects.

Overall, our study demonstrates that symmetry-violating generator perturbations, especially non-Hermitian ones, can systematically reproduce the experimentally observed deviations in single-qubit process tomography. The approach clarifies which PTM components—primarily affine and off-diagonal entries—stem from coherent symmetry breaking and which residual discrepancies point to dissipative or SPAM-driven processes. Extending this framework to multi-qubit systems will allow characterization of correlated and crosstalk-induced symmetry violations, offering a route toward more physically grounded noise models for quantum simulation, calibration, and error mitigation.

The full code is available in a corresponding GitHub repository at <https://github.com/Raubkatz/QuantumSymmetryViolations>.

Author Contributions: S.R.: mathematics, main manuscript, numerical simulation, validation, review; Kr.W.: main manuscript, numerical simulation, experiments, review; S.S.: main manuscript, experiments, review, funding; Ka.W.: main manuscript, experiments, review, funding; K.C.: main manuscript, experiments, review, funding; K.M.: main manuscript, experiments, review, funding

Funding: The authors would like to acknowledge that this paper has been written based on the results achieved within the OptiQ project. This Project has received funding from the European Union's Horizon Europe programme under the grant agreement No 101080374-OptiQ. Supplementarily, the Project is co-financed from the resources of the Polish Ministry of Science and Higher Education in a frame of programme International Co-financed Projects. Disclaimer. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA – granting authority). Neither the European Union nor the granting authority can be held responsible for them. The financial support by the Austrian Federal Ministry of Economy, Energy and Tourism, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged. SBA Research (SBA-K1 NGC) is a COMET Center within the COMET – Competence Centers for Excellent Technologies Programme and funded by BMIMI, BMWET, and the federal state of Vienna. The COMET Programme is managed by FFG. Open access funding provided by University of Vienna.

Use of Artificial Intelligence: We used Grammarly and ChatGPT to review the manuscript for grammar issues, typographical errors, and clarity. The tools were also used to improve the overall LaTeX presentation. All scientific content, interpretations, and conclusions reflect the authors' work and judgment.

Conflicts of Interest: The authors have no competing interests as defined by Nature Portfolio, or other interests that might be perceived to influence the results and/or discussion reported in this paper.

References

1. Chuang, I.L.; Nielsen, M.A. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics* **1997**, *44*, 2455–2467, <https://doi.org/10.1080/09500349708231894>.
2. Poyatos, J.F.; Cirac, J.I.; Zoller, P. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. *Phys. Rev. Lett.* **1997**, *78*, 390–393. <https://doi.org/10.1103/PhysRevLett.78.390>.
3. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*; Cambridge University Press, 2010.
4. Qiskit-Experiments Team. Qiskit Experiments: State and Process Tomography (documentation). <https://qiskit.org/>, 2025. Accessed 2025-10-22.
5. Breuer, H.P.; Petruccione, F. *The Theory of Open Quantum Systems*; Oxford University Press, 2007. <https://doi.org/10.1093/acprof:oso/9780199213900.001.0001>.
6. Mundada, P.; Zhang, G.; Hazard, T.; Houck, A. Suppression of Qubit Crosstalk in a Tunable Coupling Superconducting Circuit. *Phys. Rev. Appl.* **2019**, *12*, 054023. <https://doi.org/10.1103/PhysRevApplied.12.054023>.
7. McKay, D.C.; Sheldon, S.; Smolin, J.A.; Chow, J.M.; Gambetta, J.M. Three-Qubit Randomized Benchmarking. *Phys. Rev. Lett.* **2019**, *122*, 200502. <https://doi.org/10.1103/PhysRevLett.122.200502>.
8. Gilchrist, A.; Langford, N.K.; Nielsen, M.A. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A* **2005**, *71*, 062310. <https://doi.org/10.1103/PhysRevA.71.062310>.
9. Merkel, S.T.; Gambetta, J.M.; Smolin, J.A.; Poletto, S.; Córcoles, A.D.; Johnson, B.R.; Ryan, C.A.; Steffen, M. Self-consistent quantum process tomography. *Phys. Rev. A* **2013**, *87*, 062119. <https://doi.org/10.1103/PhysRevA.87.062119>.
10. Blume-Kohout, R.; Gamble, J.K.; Nielsen, E.; Rudinger, K.; Mizrahi, J.; Fortier, K.; Maunz, P. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nature Communications* **2017**, *8*, 14485. <https://doi.org/10.1038/ncomms14485>.
11. Magesan, E.; Gambetta, J.M.; Emerson, J. Scalable and Robust Randomized Benchmarking of Quantum Processes. *Phys. Rev. Lett.* **2011**, *106*, 180504. <https://doi.org/10.1103/PhysRevLett.106.180504>.
12. Silva, A.; Greplova, E. Hands-on introduction to randomized benchmarking. *SciPost Physics Lecture Notes* **2025**, *97*, [2410.08683]. Received 2024-11-18; accepted 2025-07-07; published 2025-07-21, <https://doi.org/10.21468/SciPostPhysLectNotes.97>.
13. Proctor, T.; Rudinger, K.; Young, K.; Sarovar, M.; Blume-Kohout, R. What Randomized Benchmarking Actually Measures. *Phys. Rev. Lett.* **2017**, *119*, 130502. <https://doi.org/10.1103/PhysRevLett.119.130502>.
14. Johnson, B.R.; da Silva, M.P.; Ryan, C.A.; Kimmel, S.; Chow, J.M.; Ohki, T.A. Demonstration of Robust Quantum Gate Tomography via Randomized Benchmarking. *New Journal of Physics* **2015**, *17*, 113019, [1505.06686]. Journal ref: New J. Phys. 17, 113019 (2015), <https://doi.org/10.1088/1367-2630/17/11/113019>.
15. Wallman, J.; Granade, C.; Harper, R.; Flammia, S.T. Estimating the coherence of noise. *New Journal of Physics* **2015**, *17*, 113020. <https://doi.org/10.1088/1367-2630/17/11/113020>.

16. Temme, K.; Bravyi, S.; Gambetta, J.M. Error Mitigation for Short-Depth Quantum Circuits. *Phys. Rev. Lett.* **2017**, *119*, 180509. <https://doi.org/10.1103/PhysRevLett.119.180509>.
17. Qiskit-Experiments Team. API Reference. <https://qiskit-community.github.io/qiskit-experiments/apidocs/index.html>, 2025. Accessed 2025-10-22.
18. Raubitzek, S.; Schrittwieser, S.; Schatten, A.; Mallinger, K. Data Obfuscation for Privacy-Preserving Machine Learning Using Quantum Symmetry Properties. *Big Data and Cognitive Computing* **2025**, *9*. <https://doi.org/10.3390/bdcc9090223>.
19. Holland, J.H. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*; Complex Adaptive Systems, The MIT Press: Cambridge, MA, 1992. Bradford Books imprint, <https://doi.org/10.7551/mitpress/1090.001.0001>.
20. Goldberg, D.E.; Holland, J.H. Genetic Algorithms and Machine Learning. *Machine Learning* **1988**, *3*, 95–99. <https://doi.org/10.1023/A:1022602019183>.
21. Whitley, D. A Genetic Algorithm Tutorial. *Statistics and Computing* **1994**, *4*, 65–85. <https://doi.org/10.1007/BF00175354>.
22. Mitchell, M. *An Introduction to Genetic Algorithms*; Complex Adaptive Systems, The MIT Press: Cambridge, MA, 1996. Bradford Books imprint, <https://doi.org/10.7551/mitpress/3927.001.0001>.
23. Eiben, A.E.; Smith, J.E. *Introduction to Evolutionary Computing*, 2 ed.; Natural Computing Series, Springer: Berlin, Heidelberg, 2015; pp. XII + 287. Includes 55 b/w and 12 color illustrations, <https://doi.org/10.1007/978-3-662-44874-8>.
24. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* **2002**, *6*, 182–197. <https://doi.org/10.1109/4235.996017>.
25. Raubitzek, S.; Neubauer, T.; Friedrich, J.; Rauber, A. Interpolating Strange Attractors via Fractional Brownian Bridges. *Entropy* **2022**, *24*. <https://doi.org/10.3390/e24050718>.
26. Qiskit Development Team. Qiskit: Open-source quantum software for large-scale execution. <https://qiskit.org/>, 2024. Accessed: 2025-11-03.
27. IBM Quantum. IBM Quantum Documentation. <https://quantum-computing.ibm.com/docs/>, 2024. Accessed: 2025-11-03.
28. Greenbaum, D. Introduction to Quantum Gate Set Tomography, 2015, [arXiv:quant-ph/1509.02921].
29. Thirunirai Selvam, J.; Saravana Veni, S. Quantum Circuit Benchmarking on IBM Brisbane: Performance Insights from Superconducting Qubit Models. *arXiv e-prints* **2025**, pp. arXiv–2508.
30. Wood, C.J.; Biamonte, J.D.; Cory, D.G. Tensor networks and graphical calculus for open quantum systems. *Quantum Information & Computation* **2015**, *15*, 759–811, [arXiv:quant-ph/1111.6950]. <https://doi.org/10.26421/QIC15.9-10-3>.
31. Magesan, E.; Gambetta, J.M.; Emerson, J. Characterizing quantum gates via randomized benchmarking. *Phys. Rev. A* **2012**, *85*, 042311. <https://doi.org/10.1103/PhysRevA.85.042311>.
32. Gambetta, J.M.; Chow, J.M.; Steffen, M. Building logical qubits in a superconducting quantum computing system. *npj Quantum Information* **2017**, *3*, 2. <https://doi.org/10.1038/s41534-016-0004-0>.
33. Krinner, S.; Lacroix, N.; Remm, A.; Paolo, A.D.; Genois, E.; Leroux, C.; Hellings, C.; Lazar, S.; Swiadek, F.; Herrmann, J.; et al. Realizing repeated quantum error correction in a distance-three surface code. *Nature* **2022**, *605*, 669–674. <https://doi.org/10.1038/s41586-022-04566-8>.
34. Chow, J.M.; Gambetta, J.M.; Córcoles, A.D.; Merkel, S.T.; Smolin, J.A.; Rigetti, C.; Poletto, S.; Keefe, G.A.; Rothwell, M.B.; Rozen, J.R.; et al. Universal Quantum Gate Set Approaching Fault-Tolerant Thresholds with Superconducting Qubits. *Phys. Rev. Lett.* **2012**, *109*, 060501. <https://doi.org/10.1103/PhysRevLett.109.060501>.
35. Martinis, J.M. Qubit metrology for building a fault-tolerant quantum computer. *npj Quantum Information* **2015**, *1*, 15005. <https://doi.org/10.1038/npjqi.2015.5>.
36. Kandala, A.; Temme, K.; Córcoles, A.D.; Mezzacapo, A.; Chow, J.M.; Gambetta, J.M. Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **2019**, *567*, 491–495. <https://doi.org/10.1038/s41586-019-1040-7>.
37. Sun, L.; Petrenko, A.; Leghtas, Z.; Vlastakis, B.; Kirchmair, G.; Sliwa, K.M.; Narla, A.; Hatridge, M.; Shankar, S.; Blumoff, J.; et al. Tracking photon jumps with repeated quantum non-demolition parity measurements. *Nature* **2014**, *511*, 444–448. <https://doi.org/10.1038/nature13436>.
38. Lindblad, G. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics* **1976**, *48*, 119–130. <https://doi.org/10.1007/BF01608499>.

39. Proctor, T.J.; Carignan-Dugas, A.; Rudinger, K.; Nielsen, E.; Blume-Kohout, R.; Young, K. Direct Randomized Benchmarking for Multiqubit Devices. *Phys. Rev. Lett.* **2019**, *123*, 030503. <https://doi.org/10.1103/PhysRevLett.123.030503>.
40. Blume-Kohout, R.; Gamble, J.K.; Nielsen, E.; Mizrahi, J.; Sterk, J.D.; Maunz, P. Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit, 2013, [[arXiv:quant-ph/1310.4492](https://arxiv.org/abs/1310.4492)].
41. Kim, Y.; Eddins, A.; Anand, S.; Wei, K.X.; van den Berg, E.; Rosenblatt, S.; Nayfeh, H.; Wu, Y.; Zaletel, M.; Temme, K.; et al. Evidence for the utility of quantum computing before fault tolerance. *Nature* **2023**, *618*, 500–505. <https://doi.org/10.1038/s41586-023-06096-3>.
42. Cai, Z.; Babbush, R.; Benjamin, S.C.; Endo, S.; Huggins, W.J.; Li, Y.; McClean, J.R.; O'Brien, T.E. Quantum error mitigation. *Rev. Mod. Phys.* **2023**, *95*, 045005. <https://doi.org/10.1103/RevModPhys.95.045005>.
43. Rajput, A.; Roggero, A.; Wiebe, N. Quantum error correction with gauge symmetries. *npj Quantum Information* **2023**, *9*, 41. <https://doi.org/10.1038/s41534-023-00706-8>.
44. Niu, M.Y.; Chuang, I.L.; Shapiro, J.H. Hardware-efficient bosonic quantum error-correcting codes based on symmetry operators. *Phys. Rev. A* **2018**, *97*, 032323. <https://doi.org/10.1103/PhysRevA.97.032323>.
45. Liu, Z.W.; Zhou, S. Approximate symmetries and quantum error correction. *npj Quantum Information* **2023**, *9*, 119. <https://doi.org/10.1038/s41534-023-00788-4>.
46. Faist, P.; Nezami, S.; Albert, V.V.; Salton, G.; Pastawski, F.; Hayden, P.; Preskill, J. Continuous Symmetries and Approximate Quantum Error Correction. *Phys. Rev. X* **2020**, *10*, 041018. <https://doi.org/10.1103/PhysRevX.10.041018>.

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.