AN ENTROPIC ORBIT-STABILISER THEOREM FOR SYMMETRY-AWARE QUANTUM LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

The classical orbit–stabiliser theorem describes exact symmetry through membership in the stabiliser subgroup. We define an orbit stabilisation entropy $H_{\mathrm{stab}}(\psi)$ and an average stabilisation weight $\overline{W}(\psi)$, and establish the new bound $H_{\mathrm{stab}}(\psi) \geq \log|\mathrm{Stab}_G([\psi])|$, giving an entropic version of the orbit–stabiliser theorem. We further show that $\overline{W}(\psi)$ equals the purity of the group-twirled state. In a controlled quantum learning setting, we study whether state-level symmetry is sufficient to produce symmetry of the learned function. We first encourage state symmetry through soft regularisation using $H_{\mathrm{stab}}(\psi)$ and $\overline{W}(\psi)$, but do not observe a statistically robust improvement in generalisation to unseen symmetry-transformed inputs. We then derive sufficient conditions for exact function-level symmetry and implement them through hard-equivariant architectures, specifically through parameter tying. We test this across five finite groups $\mathbb{Z}_2$, $\mathbb{Z}_4$, V_4, $\mathbb{Z}_6$, and S_3. The additional orbit-transfer gain from symmetry-enforcing parameter tying is significant for $\mathbb{Z}_2$ and S_3, but not for the other three groups. Three controlled experiments test possible explanations for this group-dependent effect. Finally, we apply our entropic bound to prune candidate stabiliser subgroups. In subgroup-rich group families, this can substantially reduce the number of candidates that require direct verification, giving a measured $4.12\times$ end-to-end speedup at group order 256. We also extend our framework to mixed states and study its behaviour under controlled noise and symmetry breaking. Our work opens the door to further study of entropic symmetry measures and their applications in symmetry-aware quantum learning.

1 INTRODUCTION

Let a finite group G act on a set X. For $x \in X$, recall that its orbit and stabiliser are

$$\mathrm{Orb}_G(x) = \{g \cdot x : g \in G\}, \qquad \mathrm{Stab}_G(x) = \{g \in G : g \cdot x = x\}.$$

The classical orbit–stabiliser theorem states

$$|\,\mathrm{Orb}_G(x)|\,|\,\mathrm{Stab}_G(x)| = |G|.$$

The classical orbit-stabiliser theorem connects the symmetries that leave a state unchanged with the size of its orbit under the group action. However, stabiliser membership is binary: a group element either fixes the state exactly or it does not. For quantum states, this can hide useful information. A state may not be exactly symmetric, but it can still be approximately invariant under some group transformations. Stabiliser membership does not capture this behaviour.

This motivates a continuous description of stabilisation. We define the orbit stabilisation entropy $H_{\mathrm{stab}}(\psi)$, which measures how evenly a quantum state is stabilised across the group. Rather than recording only exact stabiliser membership, it uses the overlap between a state and its transformed versions to quantify approximate stabilisation. We then propose and prove the bound

$$H_{\mathrm{stab}}(\psi) \geq \log|\,\mathrm{Stab}_G([\psi])|,$$

which gives an entropic analogue of the classical orbit–stabiliser theorem.

Orbit stabilisation entropy describes how stabilisation is distributed across the group, but not how much stabilisation is present overall. We therefore also consider the average stabilisation weight

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

$\overline{W}(\psi)$, which reveals a direct connection between approximate symmetry and the group-twirled state. $H_{\mathrm{stab}}(\psi)$ and $\overline{W}(\psi)$ can contain different information once $|G| > 2$, which we demonstrate using states produced by a verified quantum circuit.

These state-level quantities naturally raise a learning question: does encouraging symmetry of the quantum state also produce symmetry of the learned function, so that symmetry-related inputs x and gx satisfy $f_\theta(x) = f_\theta(gx)$?

Symmetry has increasingly been incorporated into quantum learning architectures through group-invariant and equivariant models (Larocca et al., 2022; Meyer et al., 2023; Nguyen et al., 2024; Schatzki et al., 2024). Quantum-state symmetry has also been studied through group twirling and resource-theoretic notions of asymmetry (Bartlett et al., 2007; Marvian & Spekkens, 2014b;a;c; Yamaguchi et al., 2026), including approximate symmetry in quantum error correction and open quantum systems (Liu & Zhou, 2023; Bandyopadhyay et al., 2023). Stabilizer Rényi entropies instead quantify nonstabilizerness or "magic" (Leone et al., 2022; Haug et al., 2024; Bittel & Leone, 2026; Odavić et al., 2025). Our work is complementary: H_{stab} measures how stabilisation is distributed across an arbitrary finite group action and connects this distribution directly to the size of the exact stabiliser through an entropic orbit–stabiliser bound.

In our paper, we study the learning question across five finite groups: $\mathbb{Z}_2$, $\mathbb{Z}_4$, V_4, $\mathbb{Z}_6$, and S_3. The pairs $\mathbb{Z}_4$ and V_4, and $\mathbb{Z}_6$ and S_3, allow us to compare different group structures at the same order. They represent all abstract group isomorphism classes of orders 4 and 6. We further investigate what might explain the different empirical behaviour across these groups through three controlled experiments.

Our bound also raises a direct computational question. Since

$$|\operatorname{Stab}_G([\psi])| \le e^{H_{\mathrm{stab}}(\psi)},$$

candidate stabiliser subgroups whose orders exceed $e^{H_{\mathrm{stab}}(\psi)}$ can be excluded without subgroup-specific verification. We therefore investigate whether this observation can be used to efficiently prune candidate stabiliser searches.

Finally, we extend the framework to noisy mixed states and ask how function-level symmetry behaves when noise is introduced. Additional experiments examine controlled symmetry breaking and coherent errors, with full details reported in the Appendix.

We begin with the entropic orbit–stabiliser theorem and its proof.

2 Entropic Orbit–Stabiliser Theorem

Let $U : G \to \mathcal{U}(\mathcal{H})$ be a unitary representation of a finite group G on a finite-dimensional complex Hilbert space $\mathcal{H}$, and let $|\psi\rangle \in \mathcal{H}$ be normalised. Since physical pure states are defined up to global phase, write

$$[\psi] = \{e^{i\theta}|\psi\rangle : \theta \in \mathbb{R}\}.$$

Its orbit and stabiliser are

$$\operatorname{Orb}_G([\psi]) = \{[U_g\psi] : g \in G\},$$

$$\operatorname{Stab}_G([\psi]) = \{g \in G : U_g|\psi\rangle = e^{i\theta_g}|\psi\rangle \text{ for some } \theta_g \in \mathbb{R}\}.$$

Thus g stabilises $[\psi]$ exactly when it leaves the physical state unchanged. The classical orbit–stabiliser theorem gives

$$|\operatorname{Orb}_G([\psi])|\,|\operatorname{Stab}_G([\psi])| = |G|.$$

For each $g \in G$, define the stabilisation strength

$$s_g(\psi) = |\langle\psi|U_g|\psi\rangle|^2.$$

Then $0 \le s_g \le 1$, with $s_g = 1$ exactly when $g \in \operatorname{Stab}_G([\psi])$. Define

$$W(\psi) = \sum_{g\in G} s_g(\psi), \qquad p_g(\psi) = \frac{s_g(\psi)}{W(\psi)},$$

so that $\{p_g\}_{g\in G}$ is a probability distribution. The orbit stabilisation entropy is

$$H_{\mathrm{stab}}(\psi) = -\sum_{g\in G} p_g(\psi)\log p_g(\psi).$$

It measures how evenly stabilisation weight is distributed across the group, while $e^{H_{\mathrm{stab}}(\psi)}$ can be interpreted as an effective number of transformations carrying significant stabilisation weight.

Theorem 1 (Entropic Orbit–Stabiliser)**.** *For any normalised* $|\psi\rangle$*,*

$$H_{\mathrm{stab}}(\psi) \geq \log|\operatorname{Stab}_G([\psi])|.$$

Consequently,

$$e^{H_{\mathrm{stab}}(\psi)}|\operatorname{Orb}_G([\psi])| \geq |G|.$$

Proof. Let $S = \operatorname{Stab}_G([\psi])$. Since $s_g = 1$ for every $g \in S$,

$$W = \sum_{g\in G} s_g = |S| + \sum_{g\notin S} s_g \geq |S|.$$

Since $p_g = s_g/W$ and $s_g \leq 1$, with equality for every $g \in S$,

$$p_{\max} = \frac{1}{W}.$$

Shannon entropy is bounded below by min-entropy, hence

$$H_{\mathrm{stab}}(\psi) \geq -\log p_{\max} = \log W \geq \log|S|.$$

This proves the first claim. Using $|\operatorname{Orb}_G([\psi])||S| = |G|$ and exponentiating gives

$$e^{H_{\mathrm{stab}}(\psi)}|\operatorname{Orb}_G([\psi])| \geq |G|.$$

□

H_{stab} varies continuously with the state and can therefore retain information about approximate symmetry that the exact stabiliser does not capture.

To illustrate this, let $G = \mathbb{Z}_4 = \langle r\rangle$ act on a four-dimensional Hilbert space by the cyclic shift

$$U_r|j\rangle = |j+1 \pmod 4\rangle,$$

and consider the fully symmetric state

$$|u\rangle = \frac{1}{2}\sum_{j=0}^{3}|j\rangle, \qquad |\operatorname{Stab}_G([u])| = 4, \qquad H_{\mathrm{stab}}(u) = \log 4.$$

For a normalised symmetry-breaking direction $|v_\perp\rangle \perp |u\rangle$, define

$$|\psi(\epsilon)\rangle = \frac{|u\rangle + \epsilon|v_\perp\rangle}{\sqrt{1+\epsilon^2}}.$$

Since each overlap $s_g(\epsilon) = |\langle\psi(\epsilon)|U_g|\psi(\epsilon)\rangle|^2$ varies continuously with ϵ, so does $H_{\mathrm{stab}}(\psi(\epsilon))$.

We test 25 independently sampled perturbation directions. Already at $\epsilon = 0.01$, the exact stabiliser collapses from size 4 to 1 for all 25 directions, while the mean $H_{\mathrm{stab}} = 1.38629 \approx \log 4$ and $e^{H_{\mathrm{stab}}} \approx 4$. Thus, exact stabiliser membership loses the degree of approximate symmetry under a tiny perturbation, whereas H_{stab} retains it. The entropic orbit–stabiliser bound holds at all 25×101 tested perturbation points. Further details are given in Appendix F.

3 Average Stabilisation and Group Twirling

While H_{stab} measures how stabilisation weight is distributed, we define its average amount by

$$\overline{W}(\psi) = \frac{1}{|G|} \sum_{g \in G} |\langle \psi | U_g | \psi \rangle|^2 = \frac{W(\psi)}{|G|}.$$

For the group-twirled state

$$\rho_G = \frac{1}{|G|} \sum_{g \in G} U_g |\psi\rangle\langle\psi| U_g^\dagger,$$

we obtain

$$\begin{aligned} \mathrm{Tr}(\rho_G^2) &= \frac{1}{|G|^2} \sum_{g,h \in G} |\langle \psi | U_g^\dagger U_h | \psi \rangle|^2 \\ &= \frac{1}{|G|} \sum_{k \in G} |\langle \psi | U_k | \psi \rangle|^2 = \overline{W}(\psi), \end{aligned}$$

since each $k = g^{-1}h$ occurs for exactly $|G|$ ordered pairs (g, h). Hence

$$\overline{W}(\psi) = \mathrm{Tr}(\rho_G^2).$$

Thus, $\overline{W}$ measures the average amount of stabilisation, while H_{stab} measures how that stabilisation is distributed.

4 From State Symmetry to Function Symmetry

The quantities above describe quantum states, whereas a quantum learning model computes a function. Let

$$|\phi(x)\rangle = E(x)|0\rangle, \qquad |\psi_\theta(x)\rangle = V_\theta |\phi(x)\rangle, \qquad f_\theta(x) = \langle \psi_\theta(x) | O | \psi_\theta(x) \rangle.$$

State-level symmetry alone does not generally imply $f_\theta(gx) = f_\theta(x)$, because function symmetry also depends on the encoder, trainable circuit, and readout.

Proposition 1 (Sufficient condition for exact function symmetry)**.** *Suppose that, for every $g \in G$,*

$$|\phi(gx)\rangle = U_g |\phi(x)\rangle, \qquad [U_g, V_\theta] = 0, \qquad U_g^\dagger O U_g = O.$$

Then

$$f_\theta(gx) = f_\theta(x)$$

for every x, g, and θ.

Proof. Encoder equivariance and $[U_g, V_\theta] = 0$ give

$$|\psi_\theta(gx)\rangle = V_\theta U_g |\phi(x)\rangle = U_g |\psi_\theta(x)\rangle.$$

Therefore,

$$f_\theta(gx) = \langle \psi_\theta(x) | U_g^\dagger O U_g | \psi_\theta(x) \rangle = f_\theta(x).$$

□

5 Methodology

5.1 Symmetry-Aware Classification Task

We use 6-dimensional inputs $x \in \mathbb{R}^6$, encoded into six qubits by

$$E(x) = \bigotimes_{j=0}^{5} R_X(x_j).$$

The binary label is

$$y(x) = \mathbf{1}\left\{\sum_{j=0}^{5} x_j > 0\right\}.$$

Since the label depends only on the sum of the coordinates, it is unchanged by permutations of them. For each group G, the training set contains one representative from each input orbit. Other symmetry-related members of the same orbit are withheld from training and used for evaluation. We call accuracy on these transformed inputs orbit-transfer accuracy. We separately evaluate fresh inputs to measure ordinary IID generalisation.

IID-test accuracy measures performance on new examples, while orbit-transfer accuracy measures whether knowledge learned from one member of an orbit transfers to unseen symmetry-related members with the same label.

5.2 Soft and Hard Symmetry

We first test whether encouraging symmetry of the quantum state is enough to improve orbit transfer. The baseline circuit uses three layers of trainable R_Y rotations and entangling gates. We compare it with three softly regularised variants using H_{stab}, $\overline{W}$, or both:

$$\mathcal{L}_H = \mathcal{L}_{\text{task}} - \lambda_H H_{\text{stab}}, \qquad \mathcal{L}_W = \mathcal{L}_{\text{task}} - \lambda_W \overline{W},$$

and

$$\mathcal{L}_{H+W} = \mathcal{L}_{\text{task}} - \lambda_H H_{\text{stab}} - \lambda_W \overline{W},$$

with $\lambda_H = \lambda_W = 0.08$. These losses encourage state-level symmetry, but do not impose function-level symmetry.

We then enforce function symmetry directly through the architecture. For each group, qubits related by the group action share trainable parameters, the entangling pattern respects the same action, and the readout is invariant. By the conditions derived in the previous section, the resulting hard-equivariant model satisfies $f_\theta(gx) = f_\theta(x)$ for every $g \in G$ and every θ.

To separate parameter tying from the rest of the architecture, we also use a matched-untied control. It has the same entangling structure and invariant readout as the hard-equivariant model, but symmetry-related qubits have independent trainable parameters. The hard-equivariant versus matched-untied comparison therefore directly tests whether symmetry-enforcing parameter tying adds an effect beyond the shared architecture and readout.

5.3 Five Group Experiment

We test the construction on $\mathbb{Z}_2$, $\mathbb{Z}_4$, V_4, $\mathbb{Z}_6$, and S_3. This includes both group structures of order 4, $\mathbb{Z}_4$ and V_4, and both group structures of order 6, $\mathbb{Z}_6$ and S_3. The same training and evaluation protocol is used for all five groups, with the corresponding group action defining the input orbits and equivariant models.

We numerically verify the equivariance conditions for every hard-equivariant construction, including the group representation, circuit commutation, and end-to-end orbit consistency. These checks verify that exact orbit consistency follows from the implemented architecture rather than being an empirical effect of training. Further details are given in Appendix A.

5.4 Training and Evaluation

All models are simulated using PennyLane's `default.qubit` backend with PyTorch automatic differentiation. Unless stated otherwise, models are trained with Adam for 15 epochs using a learning rate of 0.03. Each main condition is evaluated over 25 random seeds, with matched datasets across conditions within each seed.

We report training accuracy, IID-test accuracy, and orbit-transfer accuracy. To measure function symmetry directly, we also compute

$$\Delta_{\text{orbit}} = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \frac{1}{|G| - 1} \sum_{g \neq e} |f_\theta(x) - f_\theta(gx)|,$$

where $\mathcal{D}$ denotes the evaluation set. This measures the average change in model output across non-identity transformations of the same input orbit. Smaller values indicate greater function-level consistency. For the hard-equivariant models, $\Delta_{\text{orbit}} = 0$ by construction.

Statistical comparisons are paired across seeds. We report paired t-tests, Wilcoxon signed-rank tests, effect sizes, and 95% confidence intervals. Comparisons with the baseline are corrected for multiple testing within each group. The hard-equivariant versus matched-untied comparison directly asks whether parameter tying adds an effect beyond the shared architecture and readout. Full statistical results are given in Appendix D.

5.5 Comparing the Two Symmetry Measures

We test whether H_{stab} and $\overline{W}$ contain different information. For $|G| = 2$, we show analytically that H_{stab} and $\overline{W}$ are monotonically related. We then construct two $\mathbb{Z}_4$ circuit states with similar $\overline{W}$ but different H_{stab} to directly test whether the two measures can distinguish states once the group contains more than one non-identity stabilisation weight. Further details are given in Appendices B and C.

5.6 Controlled Tests of the Tying Effect

We test three possible explanations for results of the five-group experiment. First, we vary the amount of symmetry already built into the untied architecture. Second, we replace the matched-untied predictor at inference time with an orbit-averaged predictor. Third, we train the matched-untied model with stochastic orbit augmentation, while keeping the training budget fixed. These three experiments therefore test architectural symmetry, inference-time consistency, and training-time symmetry exposure as possible explanations for the tying effect. Further details are given in Appendix I.

5.7 Entropic Stabiliser Pruning

It directly follows from our theorem that any candidate stabiliser subgroup $H \leq G$ satisfying $|H| > e^{H_{\text{stab}}(\psi)}$ can be ruled out without verifying its elements individually. We evaluate this on three finite-group families: $\mathbb{Z}_{2^n}$, $(\mathbb{Z}_2)^n$, and $D_{2^{n-1}}$ using their regular representations and matched group orders 64, 128, 256, and 512. The reference state is chosen to have an order-8 stabiliser, so that the unperturbed state satisfies the equality case $e^{H_{\text{stab}}} = 8$.

After computing H_{stab}, candidate subgroups whose orders exceed the resulting entropic upper bound are removed before direct stabiliser verification. We measure both the fraction of candidate subgroups removed and the end-to-end runtime of the remaining verification procedure. Computing H_{stab} itself still requires the group overlaps, so this experiment tests candidate-subgroup screening rather than faster generic stabiliser identification. Its computational benefit is realised in subgroup-rich settings where many candidate subgroups would otherwise require separate verification. Further details are given in Appendix H.

5.8 Mixed States, Noise, and Coherent Errors

We extend the stabilisation overlaps to a mixed state ρ using

$$s_g(\rho) = \frac{\text{Tr}\big(\rho U_g \rho U_g^\dagger\big)}{\text{Tr}(\rho^2)}.$$

These overlaps define mixed-state versions of H_{stab} and $\overline{W}$. We then test how these quantities behave under controlled noise and whether exact function symmetry survives noise channels that commute with the group action. Further details are given in Appendix E.

We separately test whether H_{stab} contains information about the structure of coherent errors that is not contained in $\overline{W}$ alone. Across three stabiliser-state families, we generate localized and distributed coherent rotations while matching $\overline{W}$ between them, and compare the resulting H_{stab}. Matching $\overline{W}$ holds the average stabilisation level fixed, allowing us to test whether H_{stab} distinguishes how the symmetry breaking is distributed. Further details are given in Appendix G. The code used for the experiments in this work is available on our **GitHub repository**.

6 RESULTS

Table 1: Soft state-symmetry regularisation compared with the unregularised baseline across all five groups over 25 seeds. Differences are paired changes in orbit-transfer accuracy.

	H_{stab}		$\overline{W}$		Combined	
Group	Difference	p	Difference	p	Difference	p
$\mathbb{Z}_2$	$+0.042$	0.146	-0.002	0.946	-0.002	0.940
$\mathbb{Z}_4$	$+0.050$	0.118	$+0.036$	0.372	$+0.064$	0.105
V_4	$+0.026$	0.180	-0.009	0.720	-0.004	0.867
$\mathbb{Z}_6$	$+0.024$	0.205	$+0.018$	0.374	$+0.009$	0.679
S_3	$+0.027$	0.150	$+0.023$	0.298	$+0.020$	0.357

From Table 1, we observe that soft state-symmetry regularisation does not produce a statistically significant improvement in orbit-transfer accuracy in any of the five tested group constructions. Although several mean differences are positive, none of the three soft objectives is significant for any group. The largest observed increase is $+0.064$ for the combined objective at $\mathbb{Z}_4$, but this is not statistically significant ($p = 0.105$).

Thus, encouraging symmetry of the quantum state alone does not reliably improve symmetry transfer of the learned function in these experiments. We next enforce symmetry directly at the level of the learned function through hard-equivariant circuits, specifically parameter tying.

Table 2: Orbit inconsistency Δ_{orbit} across all five groups over 25 seeds.

Group	Hard-equivariant	Matched-untied (mean $\pm$ SD)
$\mathbb{Z}_2$	0 (exact)	0.159 ± 0.098
$\mathbb{Z}_4$	0 (exact)	0.136 ± 0.061
V_4	0 (exact)	0.139 ± 0.051
$\mathbb{Z}_6$	0 (exact)	0.129 ± 0.038
S_3	0 (exact)	0.146 ± 0.068

From Table 2, we observe that hard equivariance eliminates orbit inconsistency in all five group constructions. Exact orbit consistency does not, however, guarantee higher predictive accuracy, as we examine next.

Table 3: Five-group orbit-transfer results over 25 seeds (mean $\pm$ sample SD). The difference is hard-equivariant minus matched-untied accuracy. The matched-untied control has the same architecture and readout as the hard-equivariant model but independent parameters.

Group	Baseline	Matched-untied	Hard-equivariant	Difference [95% CI]
$\mathbb{Z}_2$	0.530 ± 0.119	0.514 ± 0.131	0.602 ± 0.079	$+0.088$ $[+0.033, +0.143]$
$\mathbb{Z}_4$	0.440 ± 0.159	0.554 ± 0.094	0.560 ± 0.087	$+0.006$ $[-0.043, +0.055]$
V_4	0.516 ± 0.088	0.543 ± 0.084	0.570 ± 0.092	$+0.027$ $[-0.019, +0.072]$
$\mathbb{Z}_6$	0.484 ± 0.079	0.541 ± 0.067	0.544 ± 0.094	$+0.003$ $[-0.036, +0.041]$
S_3	0.508 ± 0.092	0.522 ± 0.101	0.582 ± 0.071	$+0.060$ $[+0.015, +0.105]$

From Table 3, we observe that parameter tying significantly improves orbit-transfer accuracy over the matched-untied architecture for $\mathbb{Z}_2$ and S_3, but not for $\mathbb{Z}_4$, V_4, or $\mathbb{Z}_6$. This group-dependent effect is not explained by group order or by a simple abelian-versus-nonabelian distinction. We test three possible explanations for this variation next.

First, increasing architectural symmetry does not systematically reduce the tying benefit: the architecture-by-tying interaction is not significant ($\beta = -0.008$, 95% CI $[-0.032, 0.017]$, $p = 0.54$). Second, post-hoc orbit averaging enforces orbit consistency at inference time but decreases

accuracy in all five groups (-0.018 to -0.085). Third, stochastic orbit augmentation improves accuracy in all five groups ($+0.021$ to $+0.065$), but these gains do not track the tying-specific gains across groups (Spearman $\rho = 0.30$, $p = 0.62$). Architectural symmetry, inference-time orbit consistency, and training-time symmetry exposure therefore do not individually explain the observed variation. Equivariant parameter sharing can reduce model complexity and sample complexity (Cohen & Welling, 2016; Li et al., 2026), but our results show that this does not translate into a uniform accuracy gain across groups. What determines when parameter tying improves predictive accuracy remains an open question. Further details are given in Appendix I.

We next compare the information captured by H_{stab} and $\overline{W}$. For $|G| = 2$, the two quantities are monotonically related (Appendix B). For larger groups, however, states can have similar $\overline{W}$ but different H_{stab}.

Table 4: Two circuit-generated states under the $\mathbb{Z}_4$ action with nearly matched $\overline{W}$ but different H_{stab}.

State	$(s_e, s_g, s_{g^2}, s_{g^3})$	$\overline{W}$	H_{stab}
A	(1.000, 0.492, 0.011, 0.492)	0.4985	1.065
B	(1.000, 0.335, 0.332, 0.335)	0.5004	1.243

From Table 4, we observe that two circuit-generated states can have nearly the same average stabilisation but different stabilisation entropies. Therefore, $\overline{W}$ alone does not determine how stabilisation weight is distributed across the group.

The coherent-error experiment supports the same distinction. At matched $\overline{W}$, distributed coherent errors have larger H_{stab} than localised errors in all 15 tested cases, with an aggregate difference of 0.056 (95% CI $[0.019, 0.093]$). Thus, H_{stab} can distinguish different distributions of coherent symmetry breaking even when average stabilisation is matched. Further details are given in Appendix G.

Identifying stabilizer structure in quantum states is itself an active algorithmic problem (Hinsche et al., 2026; Arunachalam & Dutt, 2026). Our entropic orbit–stabiliser bound provides a complementary computational tool: it can exclude impossible stabiliser candidates before direct verification. This substantially reduces the candidate set in subgroup-rich group families. The strongest pruning occurs for the elementary abelian family $(\mathbb{Z}_2)^n$, where the number of candidate subgroups grows rapidly with group order.

Table 5: Entropic candidate-subgroup pruning for $G = (\mathbb{Z}_2)^n$. Pruning fractions are evaluated up to group order 1024. End-to-end runtime is benchmarked up to order 256; larger orders exceed the memory available on our experimental hardware.

$\lvert G\rvert$	Total subgroups	Pruned	Pruning	Speedup
64	2,825	715	25.3%	1.36×
128	29,212	14,606	50.0%	2.10×
256	417,199	308,993	74.1%	4.12×
512	8,283,458	7,451,476	90.0%	not benchmarked
1024	229,755,605	223,232,615	97.2%	not benchmarked

From Table 5, we observe that at group order 256, the entropic bound removes 74.1% of the 417,199 candidate subgroups before direct verification. This produces a measured $4.12\times$ end-to-end speedup, with a 95% bootstrap interval of $[4.09, 4.21]$. At orders 512 and 1024, the pruning fraction increases further to 90.0% and 97.2%, respectively.

The computational benefit depends strongly on subgroup-lattice richness. At matched group orders, pruning in the cyclic family $\mathbb{Z}_{2^n}$ grows only from 42.9% to 60.0%, while the dihedral family $D_{2^{n-1}}$ remains near 13%. In contrast, the elementary abelian family $(\mathbb{Z}_2)^n$ reaches 97.2% pruning at group order 1024. Note that computing H_{stab} still requires the group overlaps. The speedup applies when

verifying a large family of candidate subgroups, where the entropic bound can remove impossible candidates before direct verification. Further details are given in Appendix H.

Finally, we test whether the framework extends beyond pure, noiseless states. The entropic orbit–stabiliser bound extends to mixed states, and our noise experiments reveal the same distinction at both the state and function levels. Under symmetry-preserving noise, state purity can decrease while state symmetry remains preserved. Similarly, the hard-equivariant model maintains zero orbit inconsistency even as purity decreases. Thus, state degradation does not necessarily imply symmetry degradation, either for the quantum state or for the learned function. Further details are given in Appendix E.

7 Limitations

Our experiments cover five finite-group constructions, so the effects of parameter tying may not generalise to other groups, representations, or architectures. Soft regularisation is tested at a fixed strength, while the mixed-state and noise experiments use controlled simulations. Future work should test larger physical systems and hardware noise, and develop theory for when parameter tying improves predictive accuracy.

8 Further Scope

We believe the next step is to extend our entropic orbit–stabiliser framework to continuous groups such as $\mathrm{SO}(2)$ and $\mathrm{SO}(3)$. Continuous quantum symmetry has been studied through resource-theoretic approaches (Yamaguchi et al., 2026) and continuous symmetry measures in quantum chemistry (Alemany et al., 2017). Such an extension could allow H_{stab} to quantify approximate rotational symmetry, with potential applications in molecular modelling, materials science, and geometric quantum learning, where rotational equivariance is important (Batzner et al., 2022).

9 Broader Impact

Symmetry is central to many physical systems and increasingly to quantum learning models (Smidt, 2021; Tüysüz et al., 2024; West et al., 2024; Li et al., 2024; Bradshaw et al., 2025). Our framework provides a way to quantify approximate state symmetry while distinguishing it from function-level equivariance. This may help analyse symmetry in quantum learning models and, through the entropic bound, reduce the computational cost of candidate stabiliser searches. At the same time, symmetry should not be treated as a proxy for model quality. Our results show that exact equivariance guarantees orbit consistency but does not necessarily improve predictive accuracy.

10 Conclusion

In this work, we asked how approximate quantum-state symmetry can be quantified and how it relates to symmetry of the learned function. We introduced H_{stab} and the average stabilisation weight $\overline{W}$ as complementary measures of approximate symmetry, and proved an entropic orbit–stabiliser bound.

Our experiments reveal a clear distinction between state-level and function-level symmetry. Soft state-symmetry regularisation does not produce statistically robust improvements in orbit-transfer accuracy, whereas hard equivariance guarantees exact orbit consistency across all five tested groups. Parameter tying improves predictive accuracy for $\mathbb{Z}_2$ and S_3, but not for the other three groups. We also demonstrate a computational application of the entropic bound, obtaining a measured $4.12\times$ end-to-end speedup through candidate-subgroup pruning at group order 256.

Overall, approximate state symmetry, exact function-level symmetry, and predictive benefit are distinct properties in symmetry-aware quantum learning. Our framework also extends to mixed states and remains informative under controlled symmetry breaking and noise. Understanding when parameter tying improves predictive accuracy, and extending the framework to continuous symmetry groups, are natural directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools to assist with programming tasks, the generation of synthetic datasets, and grammatical and sentence-level editing of the manuscript. Generative AI tools were not used to develop the novel contribution, mathematical proofs, methodology, interpretation of results, or conclusions. All AI-assisted work was reviewed by the author, and LLM-generated code was verified and tested for correctness. The author takes full responsibility for the final content of this work, including all text, claims, results, code, and other artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work is theoretical and computational in nature and does not involve human subjects, personal data, or other ethically sensitive data. There are no ethical concerns arising from this work.

REPRODUCIBILITY STATEMENT

We provide the complete proofs of our theoretical results in the main text and appendices. The experimental setup, model constructions, training protocol, statistical tests, and additional numerical experiments are described in the main text through the Methodology section and Appendices. The code for our project is available on our **GitHub repository**.

REFERENCES

Pere Alemany, David Casanova, Santiago Alvarez, Chaim Dryzun, and David Avnir. Continuous symmetry measures: A new tool in quantum chemistry. In Abby L. Parrill and Kenny B. Lipkowitz (eds.), *Reviews in Computational Chemistry*, volume 30, pp. 289–352. Wiley, 2017. doi: 10.1002/9781119356059.ch7.

Srinivasan Arunachalam and Arkopal Dutt. Learning stabilizer structure of quantum states. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pp. 1949–1959. Association for Computing Machinery, 2026. doi: 10.1145/3798129.3800900.

Rahul Bandyopadhyay, Alex H. Rubin, Marina Radulaski, and Mark M. Wilde. Efficient quantum algorithms for testing symmetries of open quantum systems. *Open Systems & Information Dynamics*, 30(3):2350017, 2023. doi: 10.1142/S1230161223500178.

Stephen D. Bartlett, Terry Rudolph, and Robert W. Spekkens. Reference frames, superselection rules, and quantum information. *Reviews of Modern Physics*, 79(2):555–609, 2007. doi: 10.1103/RevModPhys.79.555.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022. doi: 10.1038/s41467-022-29939-5.

Lennart Bittel and Lorenzo Leone. Operational interpretation of the stabilizer entropy. *Quantum*, 10:2069, 2026. doi: 10.22331/q-2026-04-15-2069.

Zachary P. Bradshaw, Ethan N. Evans, Matthew Cook, and Margarite L. LaBorde. Learning equivariant maps with variational quantum circuits. *Physical Review Applied*, 23(4):044007, 2025. doi: 10.1103/PhysRevApplied.23.044007.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999. PMLR, 2016.

Tobias Haug, Soovin Lee, and M. S. Kim. Efficient quantum algorithms for stabilizer entropies. *Physical Review Letters*, 132(24):240602, 2024. doi: 10.1103/PhysRevLett.132.240602.

Marcel Hinsche, Jens Eisert, and Jose Carrasco. Abelian state hidden subgroup problem: Learning stabilizer groups and beyond. *PRX Quantum*, 7:020337, 2026. doi: 10.1103/6frk-891j.

Martín Larocca, Frédéric Sauvage, Faris M. Sbahi, Guillaume Verdon, Patrick J. Coles, and M. Cerezo. Group-invariant quantum machine learning. *PRX Quantum*, 3(3):030341, 2022. doi: 10.1103/PRXQuantum.3.030341.

Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. Stabilizer rényi entropy. *Physical Review Letters*, 128(5):050402, 2022. doi: 10.1103/PhysRevLett.128.050402.

Hui Li, Zhiguo Wang, Li Sheng, and Bohui Chen. Generalization bounds for equivariant networks on markov data. *SIAM Journal on Mathematics of Data Science*, 8(2):382–410, 2026. doi: 10.1137/25M1739777.

Zhelun Li et al. Enforcing exact permutation and rotational symmetries in the application of quantum neural networks on point cloud datasets. *Physical Review Research*, 6(4):043028, 2024. doi: 10.1103/PhysRevResearch.6.043028.

Zi-Wen Liu and Sisi Zhou. Approximate symmetries and quantum error correction. *npj Quantum Information*, 9:119, 2023. doi: 10.1038/s41534-023-00788-4.

Iman Marvian and Robert W. Spekkens. Asymmetry properties of pure quantum states. *Physical Review A*, 90(1):014102, 2014a. doi: 10.1103/PhysRevA.90.014102.

Iman Marvian and Robert W. Spekkens. Modes of asymmetry: The application of harmonic analysis to symmetric quantum dynamics and quantum reference frames. *Physical Review A*, 90(6):062110, 2014b. doi: 10.1103/PhysRevA.90.062110.

Iman Marvian and Robert W. Spekkens. Extending noether's theorem by quantifying the asymmetry of quantum states. *Nature Communications*, 5:3821, 2014c. doi: 10.1038/ncomms4821.

Johannes Jakob Meyer, Marian Mularski, Elies Gil-Fuster, Antonio Anna Mele, Francesco Arzani, Alissa Wilms, and Jens Eisert. Exploiting symmetry in variational quantum machine learning. *PRX Quantum*, 4(1):010328, 2023. doi: 10.1103/PRXQuantum.4.010328.

Quynh T. Nguyen, Louis Schatzki, Paolo Braccia, Michael Ragone, Patrick J. Coles, Frédéric Sauvage, Martín Larocca, and M. Cerezo. Theory for equivariant quantum neural networks. *PRX Quantum*, 5(2):020328, 2024. doi: 10.1103/PRXQuantum.5.020328.

Jovan Odavić, Michele Viscardi, and Alioscia Hamma. Stabilizer entropy in nonintegrable quantum evolutions. *Physical Review B*, 112(10):104301, 2025. doi: 10.1103/y9r6-dx7p.

Louis Schatzki, Martín Larocca, Quynh T. Nguyen, Frédéric Sauvage, and M. Cerezo. Theoretical guarantees for permutation-equivariant quantum neural networks. *npj Quantum Information*, 10: 12, 2024. doi: 10.1038/s41534-024-00804-1.

Tess E. Smidt. Euclidean symmetry and equivariance in machine learning. *Trends in Chemistry*, 3 (2), 2021.

Cenk Tüysüz, Su Yeon Chang, Maria Demidik, Karl Jansen, Sofia Vallecorsa, and Michele Grossi. Symmetry breaking in geometric quantum machine learning in the presence of noise. *PRX Quantum*, 5(3):030314, 2024. doi: 10.1103/PRXQuantum.5.030314.

Maxwell T. West, Jamie Heredge, Martin Sevior, and Muhammad Usman. Provably trainable rotationally equivariant quantum machine learning. *PRX Quantum*, 5(3):030320, 2024. doi: 10.1103/PRXQuantum.5.030320.

Koji Yamaguchi, Yosuke Mitsuhashi, Tomohiro Shitara, and Hiroyasu Tajima. Quantum geometric tensor determines the pure-state i.i.d. conversion rate in the resource theory of asymmetry for any compact lie group. *Physical Review X*, 16(3):031028, 2026. doi: 10.1103/qqf6-x85b.

A Encoder and Equivariant-Circuit Verification

We numerically verify the equivariance conditions used in the main text for all five group constructions. The checks are performed independently of model training. For each group, we verify the implemented group action, commutation of the hard-equivariant circuit with the group representation, and end-to-end orbit consistency. The results are summarized in Table 6.

Table 6: Numerical verification of the hard-equivariant constructions. Errors at the scale of 10^{-16} are at machine precision. Control discrepancies are reported under the corresponding end-to-end symmetry test.

Group	Group-action check	Max. commutation error	Hard-equivariant discrepancy	Control discrepancy
$\mathbb{Z}_2$	$\mathrm{SWAP}(0,1)^2 = I$	2.22×10^{-16}	$\leq 2.22 \times 10^{-16}$	$\approx 3 \times 10^{-4}$ (unconstrained)
$\mathbb{Z}_4$	$r^4 = e$	4.44×10^{-16}	$\leq 1.67 \times 10^{-16}$	≈ 0.54 (unconstr.); ≈ 0.17 (untied)
V_4	$a^2 = b^2 = c^2 = e,\ ab = ba = c$	1.95×10^{-16}	$\leq 10^{-10}$	≈ 0.037 (untied)
$\mathbb{Z}_6$	$r^6 = e$	1.62×10^{-16}	$\leq 4.44 \times 10^{-16}$	≈ 0.052 (untied)
S_3	all 6 elements checked	$\leq 2.22 \times 10^{-16}$	$\leq 4.44 \times 10^{-16}$	≈ 0.15 (unconstr.); ≈ 0.10 (untied)

For $\mathbb{Z}_2$, encoder equivariance is additionally verified directly: encoding a transformed input agrees with transforming the encoded state. End-to-end discrepancies are evaluated over random inputs and parameter draws; for S_3 and $\mathbb{Z}_6$, all five non-identity transformations are included.

The hard-equivariant constructions therefore satisfy the required symmetry relations up to numerical precision for all five groups, whereas the unconstrained and matched-untied controls give nonzero discrepancies. This confirms that the zero orbit inconsistency reported in the main text follows from the equivariant construction rather than from training.

B Relationship Between H_{stab} and $\overline{W}$ for $|G| = 2$

For $|G| = 2$, we have $s_e = 1$, so the stabilisation profile has only one free quantity, $s = s_g \in [0, 1]$. For $s \in (0, 1]$,

$$W = 1 + s$$

and

$$H_{\mathrm{stab}}(s) = -\frac{1}{1+s} \log \frac{1}{1+s} - \frac{s}{1+s} \log \frac{s}{1+s}.$$

Differentiating gives

$$H'_{\mathrm{stab}}(s) = -\frac{\log s}{(1+s)^2} \geq 0,$$

since $\log s \leq 0$ for $s \in (0, 1]$. At the same time,

$$\overline{W} = \frac{1+s}{2}$$

is strictly increasing in s.

Thus, for a group of order two, H_{stab} and $\overline{W}$ are monotonically related through the single nontrivial stabilisation weight. This is no longer true in the same way for larger groups, where the stabilisation weight can be distributed across several non-identity elements.

C $\mathbb{Z}_4$ Construction with Matched Average Stabilisation

To test whether H_{stab} can distinguish states with nearly the same average stabilisation, we construct two circuit-generated states with $\overline{W} \approx 0.5$ but different stabilisation profiles. We use the $\mathbb{Z}_4$ representation, the unconstrained six-qubit ansatz, and a fixed input x. Circuit parameters are optimized with Adam for 400 steps using

$$200 \left(\overline{W}(\theta) - 0.5\right)^2 + \eta \left(\pm H_{\mathrm{stab}}(\theta)\right), \qquad \eta = 0.15.$$

The entropy term is introduced after step 100, once $\overline{W}$ is close to its target. We then minimize H_{stab} for one state and maximize it for the other while maintaining $\overline{W} \approx 0.5$.

Table 7: Two circuit-generated $\mathbb{Z}_4$ states with nearly equal average stabilisation but different stabilisation entropies.

	Stabilisation profile	$\overline{W}$	H_{stab}
State A	(1.000, 0.492, 0.011, 0.492)	0.4985	1.065
State B	(1.000, 0.335, 0.332, 0.335)	0.5004	1.243

The two states have nearly identical $\overline{W}$, but State B distributes its stabilisation weight more evenly across the non-identity group elements and therefore has larger H_{stab}. Since both profiles are direct outputs of converged quantum circuits, this demonstrates that for $|G| > 2$, nearly equal average stabilisation can correspond to different distributions of stabilisation weight that are distinguished by H_{stab}.

D Five-Group Statistical Details

Table 3 in the main text reports results over the same 25 paired random seeds for each group. For each condition, we report the mean orbit-transfer accuracy and sample standard deviation across seeds. The reported difference is the paired hard-equivariant minus matched-untied accuracy, and the corresponding 95% confidence interval is computed from the paired seed-wise differences.

As an additional robustness check, we compare the hard-equivariant and matched-untied models using both paired t-tests and Wilcoxon signed-rank tests. The resulting p-values are shown in Table 8.

Table 8: Paired statistical tests ($N = 25$ paired seeds) for the hard-equivariant versus matched-untied orbit-transfer comparison.

Group	Paired t-test p	Wilcoxon p
$\mathbb{Z}_2$	0.0030	0.0037
$\mathbb{Z}_4$	0.8000	0.9700
V_4	0.2396	0.0740
$\mathbb{Z}_6$	0.8819	0.5426
S_3	0.0111	0.0149

These targeted within-group comparisons are evaluated at $\alpha = 0.05$. Both tests support a tying-specific orbit-transfer improvement for $\mathbb{Z}_2$ and S_3, but not for $\mathbb{Z}_4$, V_4, or $\mathbb{Z}_6$. Code for reproducing these results is provided on our **GitHub repository**.

E Mixed States and Controlled Noise

We extend the state-level quantities to mixed states and then study their behaviour under controlled noise.

E.1 Mixed-State Extension

For a density matrix ρ and unitary representation $U : G \to U(\mathcal{H})$, define

$$s_g(\rho) = \frac{\mathrm{Tr}(\rho U_g \rho U_g^\dagger)}{\mathrm{Tr}(\rho^2)}.$$

For a pure state $\rho = |\psi\rangle\langle\psi|$, this reduces to

$$s_g(\psi) = |\langle\psi|U_g|\psi\rangle|^2.$$

Proposition 2 (Mixed-state stabilisation strength)**.** *For every $g \in G$,*

$$0 \le s_g(\rho) \le 1, \qquad s_e(\rho) = 1.$$

Moreover,

$$s_g(\rho) = 1 \quad \Longleftrightarrow \quad U_g \rho U_g^\dagger = \rho.$$

Proof. The identity $s_e(\rho) = 1$ follows immediately from $U_e = I$. Using the Hilbert–Schmidt inner product and Cauchy–Schwarz,

$$\mathrm{Tr}(\rho U_g \rho U_g^\dagger) \le \sqrt{\mathrm{Tr}(\rho^2)} \sqrt{\mathrm{Tr}\Big((U_g \rho U_g^\dagger)^2\Big)} = \mathrm{Tr}(\rho^2),$$

because unitary conjugation preserves purity. Hence $s_g(\rho) \le 1$. Equality in Cauchy–Schwarz requires ρ and $U_g \rho U_g^\dagger$ to be proportional. Both have unit trace, so they must be equal. □

Define

$$W(\rho) = \sum_{g \in G} s_g(\rho), \qquad p_g(\rho) = \frac{s_g(\rho)}{W(\rho)},$$

and

$$H_{\mathrm{stab}}(\rho) = -\sum_{g \in G} p_g(\rho) \log p_g(\rho), \qquad \overline{W}(\rho) = \frac{W(\rho)}{|G|}.$$

The mixed-state stabiliser is

$$\mathrm{Stab}_G(\rho) = \{g \in G : U_g \rho U_g^\dagger = \rho\}.$$

Theorem 2 (Mixed-state entropic orbit–stabiliser bound)**.**

$$H_{\mathrm{stab}}(\rho) \ge \log |\mathrm{Stab}_G(\rho)|.$$

Proof. Let $S = \mathrm{Stab}_G(\rho)$. For every $g \in S$, $s_g(\rho) = 1$, so

$$W(\rho) \ge |S|.$$

Since $s_g(\rho) \le 1$, the largest probability is $p_{\max} = 1/W(\rho)$. Therefore,

$$H_{\mathrm{stab}}(\rho) \ge -\log p_{\max} = \log W(\rho) \ge \log |S|.$$

□

The group-twirling identity also extends directly.

Proposition 3 (Mixed-state twirling identity)**.** *Let*

$$\rho_G = \frac{1}{|G|} \sum_{g \in G} U_g \rho U_g^\dagger.$$

Then

$$\overline{W}(\rho) = \frac{\mathrm{Tr}(\rho_G^2)}{\mathrm{Tr}(\rho^2)}.$$

Proof. Expanding the purity gives

$$\mathrm{Tr}(\rho_G^2) = \frac{1}{|G|^2} \sum_{g,h \in G} \mathrm{Tr}\left(U_g \rho U_g^\dagger U_h \rho U_h^\dagger\right).$$

Using cyclicity of the trace and $k = g^{-1}h$,

$$\mathrm{Tr}(\rho_G^2) = \frac{1}{|G|} \sum_{k \in G} \mathrm{Tr}(\rho U_k \rho U_k^\dagger) = \overline{W}(\rho)\, \mathrm{Tr}(\rho^2).$$

The result follows. □

We additionally verify these identities numerically on random mixed states of ranks 1–4 in dimension 5 under an order-four unitary representation, including a state with an exact stabiliser of size 2. No violations are observed.

E.2 Symmetry-Preserving Noise

Exact function symmetry is preserved by noise that respects the same group action.

Proposition 4 (Noise-covariant channels preserve orbit consistency)**.** *Suppose the encoder, trainable circuit, and observable satisfy the equivariance conditions of the main text. Let $\mathcal{N}$ satisfy*

$$\mathcal{N}(U_g \rho U_g^\dagger) = U_g \mathcal{N}(\rho) U_g^\dagger$$

for every $g \in G$. Then

$$F_\theta(x) = \mathrm{Tr}\left(\mathcal{N}(\rho_\theta(x))O\right)$$

satisfies

$$F_\theta(gx) = F_\theta(x).$$

Proof. Equivariance gives

$$\rho_\theta(gx) = U_g \rho_\theta(x) U_g^\dagger.$$

Noise covariance therefore implies

$$\mathcal{N}(\rho_\theta(gx)) = U_g \mathcal{N}(\rho_\theta(x)) U_g^\dagger.$$

Using $U_g^\dagger O U_g = O$,

$$F_\theta(gx) = \mathrm{Tr}\left(\mathcal{N}(\rho_\theta(x)) U_g^\dagger O U_g\right) = F_\theta(x).$$

□

E.3 Controlled Noise Experiment

We test the $\mathbb{Z}_2$ hard-equivariant and matched-untied models under noise applied only at inference time. Training is unchanged. This isolates the effect of noise from changes in optimization.

We consider single-qubit depolarizing noise at

$$p \in \{0, 0.025, 0.05, 0.10, 0.15, 0.20\}.$$

In the symmetric condition, identical channels are applied to the two symmetry-related qubits. In the asymmetric condition, noise is applied only to one of them. Results are averaged over 10 seeds.

Table 9: Controlled noise experiment averaged over 10 seeds. Noise is applied only at inference time.

Model	Noise	0	0.025	0.05	0.10	0.15	0.20
Purity $\mathrm{Tr}(\rho^2)$							
—	—	1.000	0.910	0.827	0.680	0.558	0.457
Orbit inconsistency Δ_{orbit}							
Equiv.	symmetric	0.000	0.000	0.000	0.000	0.000	0.000
Equiv.	asymmetric	0.000	0.004	0.009	0.017	0.026	0.034
Untied	symmetric	0.180	0.174	0.168	0.156	0.144	0.132
Untied	asymmetric	0.180	0.178	0.176	0.172	0.168	0.165
$H_{\mathrm{stab}}(\rho)$*, equivariant model under symmetric noise*							
—	—	0.562	0.566	0.570	0.580	0.592	0.606

Symmetric noise reduces purity from 1.000 to 0.457, while the hard-equivariant model maintains zero orbit inconsistency at every tested noise strength. This agrees with the noise-covariance proposition above. At the same time, H_{stab} increases from 0.562 to 0.606 as the depolarizing channel moves the state toward the maximally mixed state, which is itself symmetric under the group action.

Asymmetric noise behaves differently. For the hard-equivariant model, Δ_{orbit} increases smoothly from 0 to 0.034 as the noise strength increases. The matched-untied model is already inconsistent at $p = 0$ and remains more inconsistent than the hard-equivariant model throughout the tested range.

These results separate state degradation from symmetry degradation. A state can lose substantial purity while preserving exact function-level symmetry when the noise respects the group action. Conversely, symmetry-breaking noise produces increasing orbit inconsistency. This experiment characterizes the behaviour of the proposed quantities under controlled noise.

F Controlled Symmetry Breaking

We study how H_{stab} behaves when an exact symmetry is broken by a small perturbation.

Let $G = \mathbb{Z}_4$ act on a four-dimensional Hilbert space through

$$U_r|j\rangle = |j + 1 \pmod 4\rangle.$$

The reference state

$$|u\rangle = \frac{1}{2}\sum_{j=0}^{3} |j\rangle$$

is invariant under the full group. We perturb it along a normalized direction $|v_\perp\rangle \perp |u\rangle$:

$$|\psi(\epsilon)\rangle = \frac{|u\rangle + \epsilon|v_\perp\rangle}{\sqrt{1+\epsilon^2}}, \qquad \epsilon \in [0, 1].$$

We evaluate 101 perturbation strengths for each of 25 independently sampled directions. Exact stabilisation is evaluated using the numerical criterion $s_g > 1 - 10^{-10}$.

Table 10: Controlled symmetry breaking across 25 perturbation directions. $N_{\mathrm{eff}} = e^{H_{\mathrm{stab}}}$.

ϵ	N_{eff} mean (SD)	H_{stab} mean (SD)	$\overline{W}$ mean (SD)	Fidelity	Prop. $\lvert\mathrm{Stab}_G([\psi])\rvert = 1$
0.01	$4.0000\ (8.0 \times 10^{-9})$	$1.38629\ (2.0 \times 10^{-9})$	$0.99980\ (8.7 \times 10^{-10})$	0.99990	1.00
0.05	$3.9999\ (5.0 \times 10^{-6})$	$1.38629\ (1.3 \times 10^{-6})$	$0.99502\ (5.4 \times 10^{-7})$	0.99751	1.00
0.25	$3.9872\ (3.1 \times 10^{-3})$	$1.38309\ (7.7 \times 10^{-4})$	$0.88744\ (3.0 \times 10^{-4})$	0.94118	1.00
0.50	$3.8047\ (4.3 \times 10^{-2})$	$1.33619\ (1.1 \times 10^{-2})$	$0.65886\ (3.5 \times 10^{-3})$	0.80000	1.00

At the smallest tested nonzero perturbation, $\epsilon = 0.01$, the exact stabiliser has collapsed from size 4 to 1 for all 25 directions. However,

$$H_{\mathrm{stab}} = 1.38629 \approx \log 4, \qquad e^{H_{\mathrm{stab}}} \approx 4.$$

Thus, exact stabiliser membership records the loss of exact symmetry but does not quantify how small the perturbation is. The entropic quantity retains this graded information.

Across all $25 \times 101 = 2525$ tested points, the entropic orbit–stabiliser bound holds with no violations. This result is a numerical illustration of the theorem and of the continuity of H_{stab}.

G Coherent Error Structure

We test whether H_{stab} distinguishes different coherent error structures when average stabilisation is matched.

We use three stabilizer-state families: the three-qubit GHZ state ($|G| = 8$), a four-qubit linear cluster state ($|G| = 16$), and a six-qubit ring graph state ($|G| = 64$). For each state we apply coherent rotations $R_\alpha(\theta)$, $\alpha \in \{X, Y, Z\}$, to localized, intermediate, or distributed subsets of qubits.

For each state, rotation axis, location, and target

$$\overline{W}^\star \in \{0.95, 0.90, 0.80, 0.70, 0.60\},$$

we solve for θ numerically so that $\overline{W}$ matches the target. All 600 configurations are reachable, and the states remain pure throughout.

At matched $\overline{W}$, distributed coherent errors produce larger H_{stab} than localized errors in all 15 (state, $\overline{W}^\star$) combinations. The mean differences are 0.037 for GHZ_3, 0.063 for the four-qubit cluster state, and 0.067 for the six-qubit ring state. Pooling the 15 combinations gives

$$\Delta H = 0.056, \qquad 95\%\ \mathrm{CI} = [0.019, 0.093].$$

The sign of the difference remains unchanged under finite-shot estimation, weak depolarizing noise, and an independent fidelity-matched control.

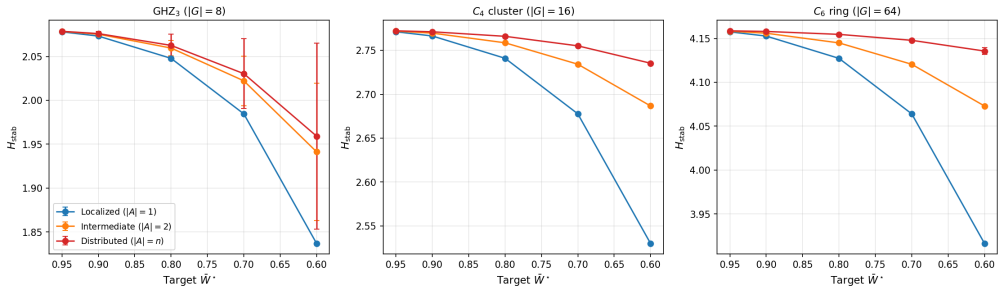


Figure 1: H_{stab} at matched target $\overline{W}$ for localized, intermediate, and distributed coherent errors. Error bars show one standard deviation across locations and rotation axes.

This shows that H_{stab} can distinguish different distributions of coherent symmetry breaking even when $\overline{W}$ is matched. However, this additional state-level information does not automatically improve downstream classification. In grouped five-fold cross-validation, a classifier using $(\overline{W}, F, H_{\mathrm{stab}})$ obtains AUC 0.72, compared with 0.81 using only $(\overline{W}, F)$. In a transfer test to the unseen six-qubit ring state, both obtain AUC 0.84.

We therefore interpret H_{stab} as providing complementary information about coherent error structure.

H Entropic Candidate-Subgroup Pruning

We test how the entropic orbit–stabiliser bound can reduce the number of candidate subgroups that require direct verification.

We study three group families,

$$\mathbb{Z}_{2^n}, \qquad (\mathbb{Z}_2)^n, \qquad D_{2^{n-1}},$$

using their regular representations. The group constructions are verified by checking closure over all $|G|^2$ pairs at each tested order.

For each group, an order-eight subgroup $S \leq G$ defines the reference state

$$|\psi_0\rangle = \frac{1}{\sqrt{|S|}} \sum_{s \in S} |s\rangle.$$

Across all tested family/order combinations,

$$|\operatorname{Stab}_G([\psi_0])| = 8, \qquad e^{H_{\mathrm{stab}}(\psi_0)} = 8,$$

confirming equality in the entropic bound.

We then perturb the reference state along random directions. Across 1080 perturbed states, the exact stabiliser reduces to the identity in every trial, while the entropic bound is never violated. Across all tested families, orders, and perturbation strengths,

$$\max e^{H_{\mathrm{stab}}} = 9.81 < 16.$$

The bound therefore certifies that candidate stabiliser subgroups of order 16 or larger can be excluded before direct verification.

For the elementary abelian family, subgroup counts are obtained from the Gaussian binomial coefficients

$$\binom{n}{k}_2.$$

The resulting total numbers of subgroups for $n = 6, \ldots, 10$ are

$$2825, \quad 29212, \quad 417199, \quad 8283458, \quad 229755605.$$

Table 11: Fraction of candidate subgroups removed by the entropic bound across the three tested group families.

n	$\lvert G\rvert$	Cyclic	Dihedral	Elementary abelian
6	64	42.9%	13.04%	25.3%
7	128	50.0%	13.43%	50.0%
8	256	55.6%	13.31%	74.1%
9	512	60.0%	13.08%	90.0%
10	1024	–	–	97.2%

For $n = 6, 7, 8$, these counts are independently reproduced by explicit enumeration using reduced-row-echelon generating matrices over $\mathbb{F}_2^n$.

The pruning fraction therefore depends strongly on subgroup-lattice structure, rather than group order alone. The elementary abelian family shows the strongest scaling, reaching 97.2% pruning at order 1024.

We benchmark end-to-end verification at group orders 64, 128, and 256, where complete enumeration is tractable on our hardware. Each candidate subgroup is directly checked using

$$V(H) = \max_{h \in H} |1 - s_h|.$$

Timing uses warm-up runs, randomized full/pruned execution order, and bootstrap confidence intervals for the speedup.

Table 12: End-to-end candidate verification with and without entropic pruning.

$\lvert G\rvert$	T_{full}	T_{pruned}	Speedup	95% CI
64	12.53 ms	9.16 ms	1.36×	$[1.35, 1.36]$
128	135.12 ms	64.40 ms	2.10×	$[2.07, 2.11]$
256	1986.36 ms	481.97 ms	**4.12**×	$[4.09, 4.21]$

At group order 256, pruning removes 74.1% of the 417199 candidate subgroups and reduces measured end-to-end verification time from 1986.36 ms to 481.97 ms, giving a $4.12\times$ speedup. We do not report end-to-end runtime at orders 512 and 1024 because complete unpruned enumeration was not tractable on our hardware.

Moreover, the computational benefit of our candidate-subgroup screening procedure arises when a large collection of candidate subgroups would otherwise require direct verification.

I Testing Explanations for the Group-Dependent Tying Gain

Parameter tying improves orbit-transfer accuracy for $\mathbb{Z}_2$ and S_3, with 95% confidence intervals excluding zero, but no corresponding improvement is established for $\mathbb{Z}_4$, V_4, or $\mathbb{Z}_6$. We test three possible explanations for this group-dependent pattern.

Architectural symmetry. We first test whether stronger architectural symmetry reduces the additional benefit of parameter tying. Within S_3, we construct three levels of entangling density and compare tied and untied variants over 30 seeds. The architecture-by-tying interaction is not significant,

$$\beta = -0.008, \qquad 95\%\ \text{CI} = [-0.032, 0.017], \qquad p = 0.54.$$

The tying gains are also non-monotonic across the three levels.A separate $\mathbb{Z}_2$ replication shows the same direction of effect, but its confidence interval includes zero.

Post-hoc orbit averaging. We next test whether the accuracy benefit follows simply from making predictions consistent across an orbit. For each matched-untied model, we replace its prediction at inference time by the average prediction across the corresponding group orbit.

Table 13: Controlled tests of three candidate explanations for the group-dependent tying gain.

Explanation	Intervention	Result
Architectural symmetry substitutes for tying	S_3 entangling-density ladder with tied and untied variants	Interaction $\beta = -0.008$, 95% CI $[-0.032, 0.017]$, $p = 0.54$. No systematic reduction in tying benefit.
Inference-time orbit consistency produces the gain	Post-hoc orbit averaging of matched-untied predictions	Accuracy decreases in every group, from -0.018 to -0.085.
Training-time symmetry exposure produces the gain	Stochastic orbit augmentation with matched training compute	Accuracy increases by $+0.021$ to $+0.065$, but the gains do not track the tying pattern ($\rho = 0.30$, $p = 0.62$).

Although this enforces output-level orbit consistency, accuracy decreases in every group, with changes between -0.018 and -0.085. Thus, inference-time orbit consistency does not reproduce the predictive effect of parameter tying.

Training-time orbit augmentation. Finally, we expose the matched-untied model to symmetry-transformed examples during training by sampling $g \sim \mathrm{Uniform}(G)$ for each training example. The architecture, number of epochs, optimizer, and number of updates remain unchanged.

Orbit augmentation improves orbit-transfer accuracy in all five groups, with gains from $+0.021$ to $+0.065$. However, these gains do not follow the group-dependent tying pattern. Their cross-group correlation with the tying-specific gain is

$$\rho = 0.30, \qquad p = 0.62.$$

For example, $\mathbb{Z}_4$ has the smallest tying-specific gain but one of the largest augmentation gains. Residual orbit variance also increases after augmentation in four of the five groups, showing that correct classification of transformed examples does not require numerically identical predictions across the orbit.