Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Experimental Implementation of the Quantum Volunteer's Dilemma on NISQ Hardware: Noise Analysis and Digital-Twin Validation

Germán Darío Díaz Agreda

germandiaz@unicauca.edu.co

University of Cauca

Jhon Alejandro Andrade Hoyos

University of Cauca

Carlos Andrés Durán Paredes

University of Cauca

Sebastián Andrés Cajas Ordoñez

Massachusetts Institute of Technology

Noah Dane Hebdon

A*STAR Graduate Academy

Siong Thye Goh

A*STAR Graduate Academy

Dax Enshan Koh

A*STAR Graduate Academy

Research Article

Keywords: Quantum game theory, Volunteer’s dilemma, NISQ hardware, digital twin, quantum Nash equilibrium, readout error mitigation

Posted Date: June 22nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9918591/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Experimental Implementation of the Quantum Volunteer's Dilemma on NISQ Hardware: Noise Analysis and Digital-Twin Validation

Germán D. Díaz Agreda [1*], Jhon Alejandro Andrade Hoyos [1], Carlos Andrés Durán Paredes [1], Sebastián Andrés Cajas Ordoñez [2], Noah Dane Hebdon [3], Siong Thye Goh [4,5], Dax Enshan Koh [3,4,6]

[1]Department of Physics, Universidad del Cauca, Popayán, Colombia.
[2]MIT Critical Data, Massachusetts Institute of Technology, Cambridge, MA, USA.
[3]Quantum Innovation Centre (Q.InC), Agency for Science, Technology and Research (A*STAR), 2 Fusionopolis Way, Innovis #08-03, Singapore 138634, Republic of Singapore.
[4]Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR), 1 Fusionopolis Way, #16-16 Connexis, Singapore 138632, Republic of Singapore.
[5]Lee Kong Chian School of Business, Singapore Management University, 50 Stamford Rd, Singapore 178899, Republic of Singapore.
[6]Science, Mathematics and Technology Cluster, Singapore University of Technology and Design (SUTD), 8 Somapah Road, Singapore 487372, Republic of Singapore.

*Corresponding author(s). E-mail(s): germandiaz@unicauca.edu.co;
Contributing authors: jaandrade@unicauca.edu.co; caduran@unicauca.edu.co; sebasmos@mit.edu; nhebdon2@jh.edu; gohst2@a-star.edu.sg; dax_koh@a-star.edu.sg;

Abstract

We present an experimental implementation of the multiplayer Quantum Volunteer's Dilemma on noisy intermediate-scale quantum (NISQ) hardware, executed on the `ibm_kingston` backend via Qiskit Runtime. The game is evaluated for $N = 2$ to 9 players under four transpiler optimization levels, with 20 independent repetitions per configuration and 2048 shots per circuit, including post-processing readout error correction via `mthree`.

Target-state fidelity decays with system size but remains above 70% (corrected) through $N = 9$. With readout correction, the global average payoff reproduces

the quantum theoretical benchmark exactly for $N \leq 6$ and exceeds the classical Nash equilibrium across the full tested range. Optimization level 2 is selected as the reference configuration after gate count analysis reveals that levels 2 and 3 produce identical transpiled circuits, with level 2 achieving superior fidelity stability.
A Hamming distance analysis of raw measurement counts shows that single-qubit errors dominate at small N, with multi-qubit contributions growing beyond $N = 6$. A calibration-based digital twin captures global payoff trends but exhibits a linear fidelity decay profile that diverges from the hardware behavior at large N, exposing the limits of first-order independent per-qubit noise models.
These results demonstrate that aggregate quantum advantage in multiplayer games is robust to NISQ noise conditions across the full tested range, while the practical observability of state-level advantage is constrained to $N \leq 8$ under post-processed readout correction.

Keywords: Quantum game theory, Volunteer's dilemma, NISQ hardware, digital twin, quantum Nash equilibrium, readout error mitigation

1 Introduction

The Volunteer's Dilemma [1] is a classic problem in game theory in which each individual in a group of N players must simultaneously decide whether to "volunteer", incurring a personal cost to ensure a collective benefit, or "abstain", free-riding on others' contributions. It suffices for at least one player to volunteer for the group to obtain the benefit, which creates a fundamental strategic tension: each agent prefers that someone else bear the cost [2–4]. From a classical perspective, this dilemma leads to inefficient equilibria, particularly as the number of players grows [5].

Quantum game theory alters this dynamic by expanding the space of allowed strategies [6, 7]. By introducing quantum resources such as superposition and entanglement, quantum games allow players to employ non-classical strategies that can lead to equilibrium structures unattainable by their classical counterparts [8, 9]. Since the pioneering works of Meyer [6] and Eisert, Wilkens and Lewenstein (EWL) [7], this paradigm has been extended to a broad class of game-theoretic scenarios [10–25], generalized settings [26–29], and studies of how quantum resources [30–34] and noise [35–42] affect quantum strategic advantage. Experimental implementations on various platforms have further demonstrated quantum games as useful testbeds for near-term quantum devices [43–48].

Within this broader line of work, Koh, Kumar and Goh [18] recently introduced a quantum formulation of the dilemma, the Quantum Volunteer's Dilemma, within the EWL framework [7]. Their analysis shows that quantum strategies can reach Nash equilibria with higher expected payoffs than those attainable through classical mixed strategies, in particular under the cost-sharing variant studied by Weesie and Franzen [49]. However, these results are derived under idealized assumptions. A natural question then arises: does this quantum advantage persist when the game is implemented on real quantum hardware?

The present work is informed by our prior experimental implementation of the Battle of the Sexes on IBM Quantum hardware [48], which demonstrated that EWL-based quantum games are practically viable on NISQ devices and that hardware noise does not necessarily eliminate quantum advantage. Here we extend this line of empirical work to the multi-player Volunteer's Dilemma for $N = 2, \ldots, 9$, introducing a systematic transpiler optimization analysis and a calibration-based digital twin that were not part of that earlier study.

Current quantum devices operate in the noisy intermediate-scale quantum (NISQ) regime [50], characterized by significant noise including gate errors, decoherence, readout errors, and unwanted couplings between qubits [51]. These effects can degrade the quantum correlations responsible for the theoretical advantage, so their impact must be explicitly assessed.

In this work we address this question through an experimental approach complemented by data-driven noise modeling. We implement the Quantum Volunteer's Dilemma for $N = 2, \ldots, 9$ on IBM's `ibm_kingston` quantum processor, systematically exploring four transpiler optimization levels and analyzing the fidelity of the target state, noise redistribution, and payoff behavior. We additionally construct a *digital twin* of the device based on real calibration data, which we use to compare the model predictions against experimental results quantitatively.

Our results show that, although the fidelity of the target state decreases with system size, the global average payoff holds up well: it remains above the classical Nash equilibrium for all $N = 2, \ldots, 9$, with margins of approximately 0.16–0.17 at $N = 2$ and 0.24–0.28 at $N = 9$ (raw and readout-corrected results at optimization level L2). The quantum advantage at the aggregate payoff level persists even when the state preparation is imperfect. By comparing against the classical equilibrium, we identify the regimes in which this advantage holds and those in which it disappears, and establish practical limits for its observability in multi-agent systems.

2 Theoretical Framework

2.1 The Classical Volunteer's Dilemma

The Volunteer's Dilemma describes a situation in which a group of N agents must each decide whether to contribute to a public good. If at least one player volunteers, the collective benefit is obtained; otherwise, all players receive an unfavorable outcome. In the variant considered in this work, the provision cost is not borne by a single individual but is shared among all who decide to volunteer. This structure introduces a balance between individual incentives and collective cooperation, in which the marginal cost of contributing decreases as more players participate.

We consider the classic Volunteer's Dilemma under mixed strategies, following the formulation of [18]. Each player i adopts a probabilistic strategy $\pi_i \in [0, 1]$, where π_i denotes the probability of volunteering, while $1 - \pi_i$ corresponds to abstaining.

The game can be formalized as an n-tuple

$$G^{(n)}_{\mathrm{MVD}} = (T_1, T_2, \ldots, T_n; \$_1, \$_2, \ldots, \$_n),$$

where $T_i = [0,1]$ is the set of strategies of the player i and $\$_i : [0,1]^n \to \mathbb{R}$ is their expected payoff function, defined as the expected value of the underlying deterministic game:

$$\$_i^{\mathrm{MVD}}(\pi_1, \ldots, \pi_n) = \sum_{x \in \{0,1\}^n} \$_i^{\mathrm{VD}}(x)\, q_{\pi_1,\ldots,\pi_n}(x),$$

where $\$_i^{\mathrm{VD}}(x)$ is the payoff in the deterministic game. Let $w(x) = \sum_{j=1}^n x_j$ denote the number of volunteering players. The deterministic payoff is

$$\$_i^{\mathrm{VD}}(x) = \begin{cases} 0 & w(x) = 0, \\ b & x_i = 0,\ w(x) > 0, \\ b - c/w(x) & x_i = 1,\ w(x) > 0, \end{cases}$$

where $b > 0$ is the collective benefit and $c > 0$ is the provision cost. Throughout this work we use $b = 2$ and $c = 1$, following the parameterisation of [18]. The joint strategy distribution is

$$q_{\pi_1,\ldots,\pi_n}(x_1, \ldots, x_n) = \prod_{i=1}^n \pi_i^{x_i}(1 - \pi_i)^{1-x_i}.$$

Players are assumed to choose independently, so the distribution factorizes. This formulation naturally interpolates between pure strategies: $\pi_i = 1$ corresponds to volunteering with certainty, while $\pi_i = 0$ represents full abstention.

Unlike the pure-strategy case, this game admits a unique symmetric Nash equilibrium. Specifically, there exists a unique value $\alpha_n \in (0,1)$ such that the profile $(\alpha_n, \ldots, \alpha_n)$ constitutes an equilibrium. This value is obtained as the unique root in $(0,1)$ of the degree-n polynomial

$$g_n(\alpha) = (1 - \alpha)^{n-1}(2n\alpha + 1 - \alpha) - 1.$$

This result, due to Weesie and Franzen [49], implies that the probability of volunteering decreases with the number of players, with asymptotic behavior

$$\alpha_n \sim \frac{\omega^*}{n} + \mathcal{O}(n^{-2}),$$

where ω^* is the positive solution of $e^\omega = 1 + 2\omega$.

The expected payoff of each player at the symmetric equilibrium is:

$$\$_i^{\mathrm{MVD}}(\alpha_n, \ldots, \alpha_n) = \left(2 - \frac{1}{n}\right)\left[1 - (1 - \alpha_n)^n\right].$$

This is the classical benchmark against which quantum payoffs are compared throughout the paper.

2.2 The Quantum Volunteer's Dilemma

Following [18], we extend the Volunteer's Dilemma to the quantum domain using the EWL quantization protocol [7]. In this framework, each player controls a qubit and

may apply local unitary operations, expanding the strategy space beyond classical mixtures.

Game setup

The game begins with a maximally entangled initial state of n qubits:

$$|\psi_0\rangle = J|0\rangle^{\otimes n},$$

where the entangling operator J is defined as

$$J = e^{-i\frac{\pi}{4}Y^{\otimes n}} = \frac{1}{\sqrt{2}}(I - iY^{\otimes n}),$$

with Y the Pauli-Y matrix. Each player i selects a quantum strategy represented by a unitary $U(\theta_i, \phi_i)$ from the family

$$U(\theta, \phi) = \begin{pmatrix} e^{i\phi}\cos\left(\frac{\theta}{2}\right) & \sin\left(\frac{\theta}{2}\right) \\ -\sin\left(\frac{\theta}{2}\right) & e^{-i\phi}\cos\left(\frac{\theta}{2}\right) \end{pmatrix},$$

with $0 \leq \theta_i < 4\pi$ and $0 \leq \phi_i < 2\pi$. The final state prior to measurement is

$$|\psi_f\rangle = J^\dagger \left(\bigotimes_{i=1}^{n} U(\theta_i, \phi_i) \right) J|0\rangle^{\otimes n},$$

where $J^\dagger = J^{-1}$ is the Hermitian adjoint of the entangling operator J, applied after the players' individual strategy choices. In our implementation, an additional X gate is applied to each qubit before measurement to align the quantum encoding with the classical game convention, where $|1\rangle$ represents volunteering and $|0\rangle$ represents abstention.

Outcome distribution and payoff

Measurement in the computational basis induces a probability distribution over classical strategy profiles $x \in \{0, 1\}^n$:

$$p_{\theta,\phi}(x) = |\langle x|\psi_f\rangle|^2.$$

The expected payoff of each player is defined analogously to the classical case (Section 2.1):

$$\$_i(\theta_1, \phi_1, \ldots, \theta_n, \phi_n) = \sum_{x \in \{0,1\}^n} \$_i^{\mathrm{VD}}(x)\, p_{\theta,\phi}(x).$$

Symmetric quantum Nash equilibrium

For $n \leq 9$, it has been shown in [18] that a symmetric Nash equilibrium exists in which all players adopt the same strategy $Q = (0, \pi/n)$, i.e., the profile $(Q, Q, \ldots, Q)$. This corresponds to the unitary

$$U\left(0, \frac{\pi}{n}\right) = \begin{pmatrix} e^{\frac{i\pi}{n}} & 0 \\ 0 & e^{-\frac{i\pi}{n}} \end{pmatrix}.$$

Under this profile, quantum interference induced by the entangling operator enhances the probability of configurations with multiple volunteers, in contrast to the classical equilibrium where the individual probability of contributing decreases with n. As a result, the expected payoff at this quantum equilibrium can exceed that of the classical symmetric equilibrium. How well this advantage holds up on real hardware is what the following sections assess.

2.3 NISQ Hardware and Noise Sources

Current quantum processors operate in the *Noisy Intermediate-Scale Quantum* (NISQ) regime [50], where circuit execution is subject to multiple noise sources that degrade quantum correlations. These mechanisms directly shift the probability distribution over strategy profiles and thus determine whether game-theoretic predictions survive hardware execution.

Gate errors and depth accumulation

Each quantum gate introduces a finite probability of error that accumulates with circuit depth [52]. Two-qubit gates exhibit significantly higher error rates than single-qubit gates and are the primary source of degradation in circuits with entanglement [53]. In the EWL protocol, the CNOT chain implementing $\exp\left(-i\frac{\pi}{4}Z^{\otimes N}\right)$ grows linearly with N; this term is the dominant contributor to hardware-induced fidelity loss.

Decoherence: relaxation and dephasing

Decoherence describes the loss of quantum information due to interaction with the environment, characterized by the relaxation time T_1 and dephasing time T_2. For a gate of duration t_g, the system evolution is modeled by a thermal relaxation channel parameterized by (T_1, T_2, t_g), introducing both population loss and coherence decay [52].

Crosstalk and residual couplings

Beyond the noise sources described above, NISQ devices exhibit *crosstalk*: operations on one qubit can induce parasitic rotations on neighboring qubits through residual electromagnetic coupling [54, 55]. Unlike gate errors and decoherence, crosstalk is inherently non-local and cannot be captured by independent per-qubit noise models. Its magnitude and structure depend on the specific qubit layout and calibration state of the device at the time of execution.

Readout errors

The measurement process introduces classical misassignment errors, modeled by confusion probabilities $P(0|1)$ and $P(1|0)$, which are directly relevant here since the game payoff depends on the observed bitstring distribution.

The digital twin as a first-order approximation

To quantify the impact of these noise sources, we construct a *digital twin* of the device from real calibration data, incorporating thermal relaxation, gate errors, and readout confusion as independent per-qubit channels (Section 3.4). This model is an effective first-order approximation: it captures the dominant, local noise mechanisms but does not reproduce non-local effects such as crosstalk or temporal calibration drift. The comparison between digital twin predictions and experimental results therefore serves a dual purpose: validating global noise trends where the model is expected to hold, and revealing residual discrepancies that point toward the higher-order effects it omits. This distinction is central to the analysis in Section 4.

3 Methodology

3.1 Hardware and Execution Environment

All experiments were performed on the `ibm_kingston` quantum backend, accessed via Qiskit Runtime with IBM Quantum API authentication [56]. This device belongs to the family of superconducting processors and employs ECR gates as the native two-qubit operation.

To reduce variability associated with hardware recalibrations, all circuit executions were carried out in a single experimental session using the same backend and calibration conditions. In the NISQ regime, parameters such as T_1, T_2 times and gate errors can fluctuate on timescales of the order of hours; session consistency is therefore a prerequisite for controlled comparison.

To implement the EWL protocol consistently across all system sizes ($N = 2, \ldots, 9$), we selected a subset of physically connected qubits forming a linear chain (qubits 0 through 8). This choice allows the interactions required by the entangling operator to be directly mapped without introducing additional SWAP gates, thereby reducing effective circuit depth and error accumulation.

Circuit execution was performed using Qiskit Runtime's `SamplerV2` primitive. Four independent jobs were submitted, one per transpiler optimization level, each containing circuits for all player counts ($N = 2, \ldots, 9$) repeated 20 times, grouped into Primitive Unified Blocs (PUBs) to improve communication efficiency and reduce execution latency. Each circuit was executed with 2048 shots per repetition, providing sufficient statistical resolution for the comparative analysis of fidelities and payoffs across the 20-sample distribution.

Post-processing readout error correction was applied using the `mthree` library [57, 58], calibrated from the confusion probabilities $P(0|1)$ and $P(1|0)$ extracted from the backend properties at the time of execution. This correction yields a more conservative estimate of target-state fidelity and is reported alongside raw results throughout Section 4.

3.2 Circuit Construction

Using the Qiskit framework, we constructed the quantum circuit for the N-player game ($N = 2, \ldots, 9$) following the EWL quantization protocol of Koh, Kumar and Goh [18]. The circuit is structured into four main blocks:

1. **Entangled state preparation:** the global operator J is applied to $|0\rangle^{\otimes N}$. In this implementation it is decomposed as

$$J = S^{\otimes N} H^{\otimes N} \exp\Big(-i\frac{\pi}{4} Z^{\otimes N}\Big) H^{\otimes N} S^{\dagger \otimes N},$$

   where H is the Hadamard gate and S is the phase gate. The non-local term $\exp\big(-i\frac{\pi}{4} Z^{\otimes N}\big)$ is implemented via a CNOT chain that propagates the global parity to a single qubit, followed by an $R_z(\pi/2)$ rotation and the reverse chain:

$$\mathrm{CNOT}(0,1) \cdots \mathrm{CNOT}(N-2, N-1)\, R_z(\pi/2)\, \mathrm{CNOT}(N-2, N-1) \cdots \mathrm{CNOT}(0,1).$$

   This construction implements an effective $Z^{\otimes N}$ interaction using only single- and two-qubit gates.
2. **Individual strategies:** each player locally applies the unitary

$$U(\theta_i, \phi_i) = R_z(\pi - \phi_i)\, R_y(\theta_i)\, R_z(-\pi - \phi_i).$$

   All players adopt the symmetric Nash equilibrium strategy $(\theta_i, \phi_i) = (0, \pi/N)$, which reduces to a Z-axis phase rotation.
3. **Inverse Operator $J^\dagger$:** applied by sequentially inverting the gates composing J.
4. **Classical convention alignment:** an X gate is applied to each qubit before measurement, so that $|1\rangle$ corresponds to volunteering.

Under ideal conditions, this construction concentrates probability strongly on $|1\rangle^{\otimes N}$. However, the CNOT chain implementing $\exp\big(-i\frac{\pi}{4} Z^{\otimes N}\big)$ introduces a depth that grows with N; the resulting circuit is therefore sensitive to noise on NISQ devices.

Figure 1 shows the circuit representation for the multiplayer case, illustrating the entanglement structure, local strategy application, and final alignment block.

3.3 Transpiler Optimization Levels

Each circuit was compiled using optimization levels 0 through 3 provided by Qiskit's transpiler via `generate_preset_pass_manager` [59] (see Appendix B for a detailed description of each level). These levels correspond to progressively more aggressive transformation strategies:

- **Level 0:** minimal compilation, preserving the original circuit structure.
- **Level 1:** basic optimizations including redundant gate cancellation and initial qubit mapping.
- **Level 2:** more sophisticated routing heuristics and gate-count reduction.
- **Level 3:** aggressive transformations including subcircuit resynthesis and advanced two-qubit gate reduction.

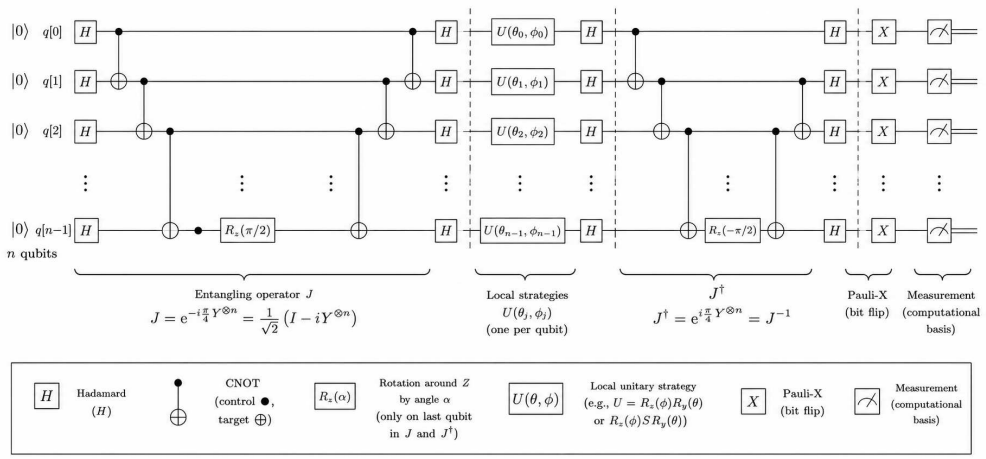


Fig. 1 Quantum circuit implementing the N-player Quantum Volunteer's Dilemma following the EWL protocol. The circuit consists of: (i) an entangling operator $J = e^{-i\frac{\pi}{4}Y^{\otimes N}}$, decomposed into Hadamard, controlled-NOT, and R_z rotations; (ii) local strategy unitaries $U(\theta_i, \phi_i)$ applied independently by each player; (iii) the inverse operator $J^\dagger$; and (iv) a final Pauli-X layer to align measurement outcomes with the classical convention (volunteer $= |1\rangle$). Measurements are performed in the computational basis. The entangling layer and its inverse each contribute $2(N-1)$ CNOT/ECR gates, giving $4(N-1)$ two-qubit gates in total (ranging from 4 at $N = 2$ to 32 at $N = 9$). Qubits 0 through $N-1$ are mapped to a linear chain on `ibm_kingston`

The comparison between optimization levels is relevant given the CNOT-chain structure of J, which makes depth, and consequently error accumulation, highly sensitive to compilation decisions. We do not assume higher levels yield uniformly better results; we evaluate their impact empirically on target-state fidelity and game payoff to identify trade-offs between circuit depth, structure, and error predictability. Results are discussed in Section 4.

3.4 Digital Twin Construction

To quantitatively assess noise impact, we constructed a *digital twin* of the quantum device from real calibration data extracted from the backend used in the experiments.

Calibration parameter extraction

For each physical qubit, we extracted the relaxation (T_1) and dephasing (T_2) times, single and two-qubit gate error rates, gate execution times (t_g), and readout confusion probabilities $P(0|1)$ and $P(1|0)$. All parameters were obtained from the backend properties of the same experimental job, so the noise model reflects the actual hardware conditions at the time of execution.

Noise model construction

A composite noise model was defined using Qiskit's `AerSimulator` [60], with three contributions:

- **Single-qubit gates:** composition of a depolarizing channel (stochastic gate errors) and a thermal relaxation channel (T_1, T_2, t_g).
- **Two-qubit gates (ECR):** for qubit pair (i, j), individual thermal relaxation channels per qubit (parameterised by $T_1^{(i)}, T_2^{(i)}, t_g^{(ij)}$) composed with a correlated ZZ Pauli error channel

$$\mathcal{E}_{ZZ}^{(ij)}(\rho) \;=\; (1 - \epsilon_{\mathrm{ECR},ij})\,\rho \;+\; \epsilon_{\mathrm{ECR},ij}\,(ZZ)\,\rho\,(ZZ),$$

  where $\epsilon_{\mathrm{ECR},ij}$ is the ECR gate error rate for pair (i, j) extracted directly from backend properties via the backend-reported ECR gate-error entry for pair (i, j).
- **Readout errors:** confusion matrices from backend-reported $P(0|1)$ and $P(1|0)$ values.

The same circuits, transpilation procedure, and qubit layouts used on hardware were applied to the simulation, so differences between experimental and simulated results reflect noise model limitations exclusively. The *digital twin* provides a first-order reference for trend identification and dominant error source estimation; it does not aim to reproduce all device-specific details such as spatial noise correlations or calibration drift.

3.5 Payoff Evaluation Method

Results are expressed as count distributions over the computational basis and processed by a dedicated evaluation engine. For an outcome $x \in \{0, 1\}^N$, the per-player payoff $\$_i^{\mathrm{VD}}(x)$ is the deterministic payoff defined in Section 2.1 with $b = 2$ and $c = 1$. Let $w(x)$ denote the number of volunteering players (bits equal to one); the rule reduces to: $w(x) = 0$: all receive 0; $w(x) > 0$: volunteers receive $2 - 1/w(x)$, abstainers receive 2.

The expected payoff for player i is:

$$\Pi_i = \sum_{x \in \{0,1\}^N} p(x)\,\$_i^{\mathrm{VD}}(x),$$

where $p(x)$ is the empirical probability estimated from measurement counts. Qiskit returns results in little-endian format, so bit order is reversed when processing each outcome to align with the player-qubit mapping.

From the probability distributions, the following metrics are computed for each (N, level) configuration:

- **Target-state fidelity:** $P(|1\rangle^{\otimes N})$.
- **Mean probability per non-target state:** average of $p(x)$ over the $2^N - 1$ non-target outcomes.
- **Target-state weighted payoff contribution:** $P(|1\rangle^{\otimes N}) \cdot (b - c/N)$, representing the contribution of the target state to the total expected payoff.
- **Global average payoff:** expected payoff Π_i across the full outcome distribution.

Statistical uncertainty is quantified using 95% confidence intervals based on the t-Student distribution across the 20 independent repetitions:

$$\mathrm{CI}_{95} = t_{0.975,\,19} \cdot \frac{\sigma}{\sqrt{20}},$$

where σ is the sample standard deviation across repetitions. For the target-state weighted payoff contribution, the confidence interval is derived analytically as $\mathrm{CI}_{\mathrm{target}} = (2 - 1/N) \cdot \mathrm{CI}_{P(|1\rangle^{\otimes N})}$, since this metric is an exact linear rescaling of the target-state probability.

4 Experimental Results

We present experimental results for the Quantum Volunteer's Dilemma on real hardware, alongside comparison with the digital twin, for $N = 2, \ldots, 9$ players. Section 4.1 characterizes fidelity decay and noise structure across all four optimization levels and presents the gate count analysis that motivates the selection of L2 as the reference configuration. Section 4.2 analyzes the Hamming distance structure of measurement outcomes. Section 4.3 validates the Digital Twin against hardware results. Section 4.4 presents the expected payoff analysis.

4.1 Target-State Fidelity and Scaling

Figure 2 presents the target-state fidelity $P(|1\rangle^{\otimes N})$ as a function of player count for all four optimization levels, both without (Figure 2 a) and with (Figure 2 b) readout error correction. The decay is approximately quasi-linear in all cases (Figure 2 a), with all levels producing comparable fidelities for $N \leq 5$.

Without readout correction, a divergence between levels becomes apparent from $N = 6$ onward: L0 and L1 exhibit a steeper and less regular decline, while L2 and L3 maintain a more stable and gradual decay through $N = 9$.

With readout correction applied, all four levels achieve near-complete target-state fidelity for $N \leq 5$. L3 begins to decay from $N = 6$, whereas L0, L1, and L2 remain stable until $N = 6$ before decreasing sharply. Among all levels, L2 exhibits the most stable trajectory and the least pronounced final decay, retaining the highest corrected fidelity at $N = 9$.

Gate count justification. Table 1 reports the ECR gate count and circuit depth after transpilation for each level and player count. L2 and L3 each reduce the ECR gate count by 2 relative to L0 and L1 across all N, and produce identical compiled circuits for all player counts—confirming that the aggressive resynthesis of L3 yields no further structural improvement for this topology. The two additional gate reductions at L2, combined with its superior fidelity stability, justify its selection as the reference configuration.

4.2 Hamming Distance Analysis of Undesired States

Figure 3 presents a heatmap of the occurrence counts for all non-target measurement outcomes, grouped by Hamming distance d_H from the target state $|1\rangle^{\otimes N}$. The analysis

Table 1 ECR gate count and circuit depth after transpilation for each optimization level (L0–L3) and player count ($N = 2, \ldots, 9$) on `ibm_kingston`. Values obtained via `generate_preset_pass_manager` with `initial_layout` $= [0, 1, \ldots, N-1]$.

N	L0		L1		L2		L3	
	ECR	Depth	ECR	Depth	ECR	Depth	ECR	Depth
2	4	55	4	28	2	14	2	14
3	8	71	8	48	6	36	6	36
4	12	87	12	62	10	46	10	46
5	16	103	16	80	14	60	14	60
6	20	119	20	96	18	72	18	72
7	24	135	24	112	22	84	22	84
8	28	151	28	128	26	96	26	96
9	32	167	32	144	30	108	30	108

is performed on the raw (uncorrected) hardware measurement counts, aggregated over all 20 repetitions at optimization level L2. Readout-corrected and simulated counts are not used here: the Hamming structure reflects the physical error process as recorded by the device, before any post-processing redistribution of probabilities. The Hamming distance between two bitstrings $x, x_{\text{target}} \in \{0,1\}^N$ is defined as the number of positions at which they differ:

$$d_H(x, x_{\text{target}}) = \sum_{i=1}^{N} \mathbb{1}\left[x_i \neq x_{\text{target},i}\right], \tag{1}$$

where $x_{\text{target}} = 11\ldots1$ denotes the theoretical equilibrium state.

In all player configurations, the distribution is dominated by $d_H = 0$ (the target state itself), with probability decreasing systematically for $d_H = 1, 2, \ldots$. This non-uniform structure, with states at distance 1 accumulating significantly more probability than those at larger distances, indicates that single-qubit errors dominate over correlated multi-qubit errors, providing empirical support for the independent per-qubit error channels used in the Digital Twin construction.

The heatmap further reveals that for $N \leq 6$, the target state retains a dominant probability and the redistribution toward $d_H \geq 1$ remains limited. From $N = 6$ onward, consistent with the fidelity behavior shown in Section 4.1, states at distances $d_H \geq 2$ begin to accumulate non-negligible counts, indicating a transition from single-qubit dominated errors to a regime where multi-qubit error contributions become relevant.

This redistribution explains why the target-state payoff contribution decays faster than the global average payoff: non-target states at $d_H = 1$ (a single abstainer among $N-1$ volunteers) still carry positive utility, partially compensating for the loss of $|1\rangle^{\otimes N}$ probability.

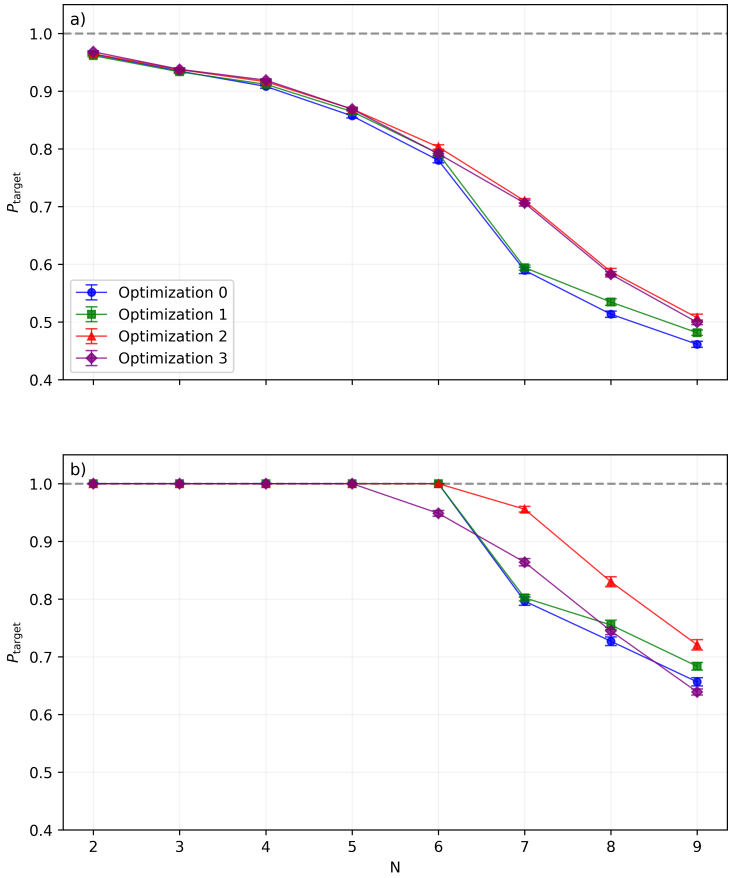


Fig. 2 Target-state fidelity $P(|1\rangle^{\otimes N})$ as a function of the number of players for all four optimization levels. Panel (a) shows raw QPU measurements; panel (b) shows readout-corrected measurements. Shaded bands represent 95% t-Student confidence intervals across 20 repetitions

4.3 Digital Twin Validation

Figure 4 compares the target-state fidelity $P(|1\rangle^{\otimes N})$ obtained from raw QPU measurements, readout-corrected QPU measurements, and the Digital Twin simulation, all at optimization level L2. The Digital Twin was executed with 20 independent random seeds to assess model stability under stochastic variation; the low variance across seeds confirms that the model's stochastic component is small relative to its systematic noise structure.

For $N \leq 6$, the Digital Twin overestimates the noise present in the raw QPU data, that is, it predicts a lower fidelity than what is experimentally observed. For $N > 6$, the model transitions to underestimating the noise in the raw counts: the hardware degrades faster than the model predicts.

The comparison between the Digital Twin and the readout-corrected QPU results deserves a methodological note: the Digital Twin incorporates readout confusion

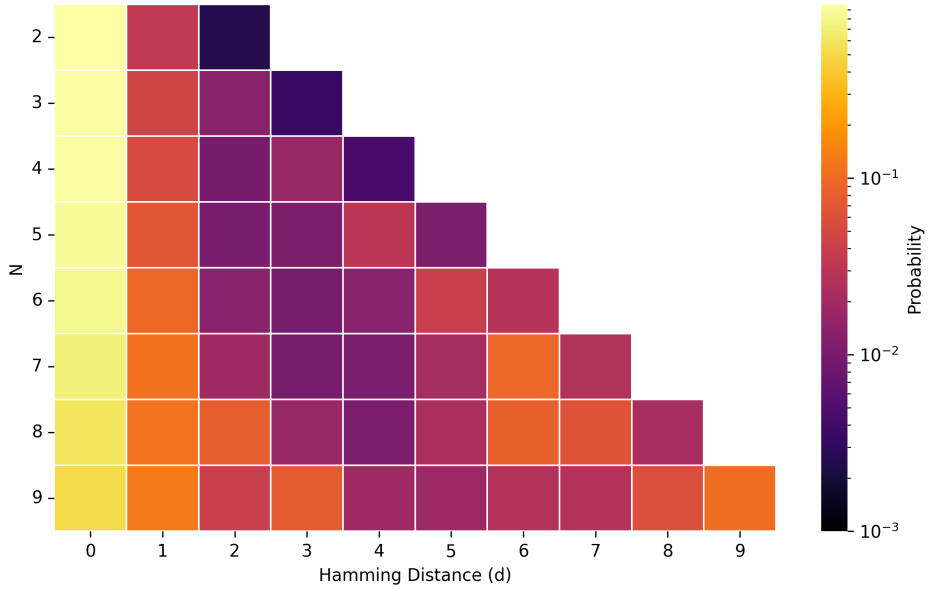


Fig. 3 Heatmap of measurement outcome counts grouped by Hamming distance from the target state $|1\rangle^{\otimes N}$, aggregated over 20 repetitions at optimization level 2. Each row corresponds to a player count N; columns represent Hamming distances $d_H = 0, 1, \ldots, N$. Color intensity encodes normalized count frequency

directly into the noise model via calibration-based confusion matrices, but no post-processing correction is applied to its output counts. The readout-corrected QPU results, by contrast, have undergone explicit probability redistribution via the `mthree` library. These two approaches address readout errors differently, so the comparison is not symmetric. Nevertheless, placing the Digital Twin alongside both raw and corrected QPU results is informative: it shows where the first-order noise model sits relative to the hardware before and after post-processing correction, and makes visible the fraction of the QPU-model gap that is attributable to readout noise specifically. With this framing, the observation that the Digital Twin consistently overestimates noise relative to the corrected data across the full range of N indicates that the model's intrinsic error channels predict more degradation than what the hardware exhibits once readout misassignment is removed.

4.4 Expected Payoffs

Figure 5 presents two complementary but distinct payoff metrics at optimization level L2. Panel (a) shows the target-state weighted payoff contribution $P(|1\rangle^{\otimes N}) \cdot (b - c/N)$: the share of the total expected payoff attributable solely to the $|1\rangle^{\otimes N}$ outcome, which decreases as that state loses probability to noise. Panel (b) shows the global average payoff Π_i summed over all 2^N outcomes: the full game-theoretic result that a player actually receives, which includes contributions from non-target states. These two metrics can behave very differently under noise: panel (a) tracks how well the

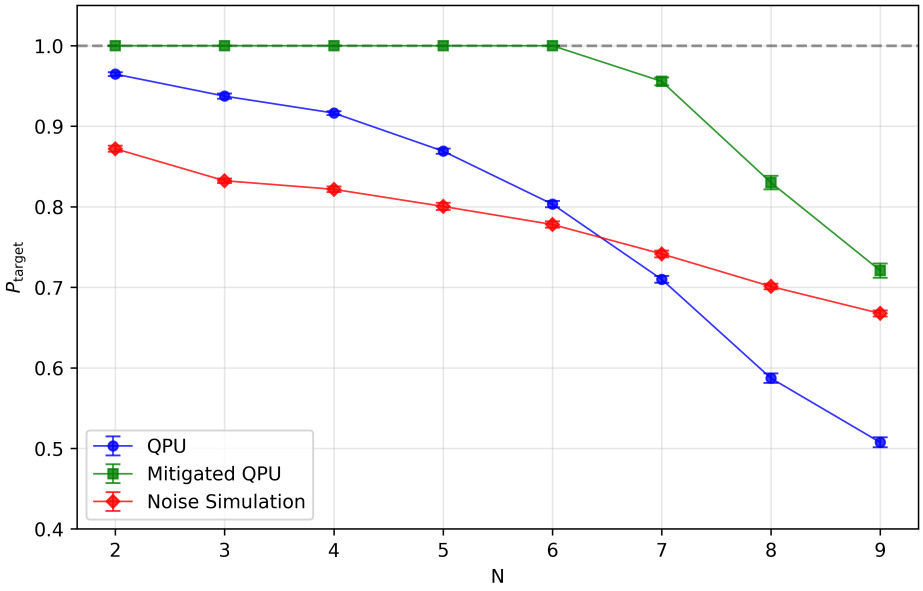


Fig. 4 Target-state fidelity $P(|1\rangle^{\otimes N})$ as a function of player count for three regimes: raw QPU (L2), readout-corrected QPU (L2), and Digital Twin (mean over 20 seeds). Error bars represent 95% t-Student confidence intervals

quantum strategy concentrates probability on the ideal outcome, while panel (b) measures whether the overall cooperative structure of the game survives.

Panel (a): Target-state contribution.
Because this metric equals $P(|1\rangle^{\otimes N})$ scaled by a fixed factor $(b - c/N)$, it degrades in direct proportion to target-state fidelity. Both the raw and readout-corrected results exceed the classical Nash equilibrium up to $N = 6$. The corrected values approximately match the theoretical quantum benchmark for small N. Beyond $N = 6$, the raw payoff falls below the classical baseline, while the corrected values remain above it through $N = 8$, dropping below at $N = 9$. The Digital Twin remains below the classical equilibrium for most of the range, with brief exceptions at $N = 3$ and $N = 4$.

Panel (b): Global average payoff.
All three regimes (raw, corrected, and Digital Twin) track the theoretical prediction closely up to $N = 5$. From $N = 5$ onward, the raw payoff decreases with its steepest drop at $N = 9$; the corrected values follow the theoretical curve through $N = 6$ before declining gradually. The Digital Twin decreases steadily from $N = 5$ without abrupt transitions. Across the full range, the global payoff of the raw and corrected hardware results remains above the classical Nash equilibrium, in contrast to the target-state contribution. This resilience arises from the positive utility of non-target states with at least one volunteer, as established by the Hamming distance structure in Section 4.2.

Tables 2 and A4 report the per-player readout correction impact (absolute and relative offset between raw and corrected global payoffs) and the absolute and relative error metrics (AE and RE) for raw, readout-corrected, and Digital Twin results against the theoretical benchmark, respectively.

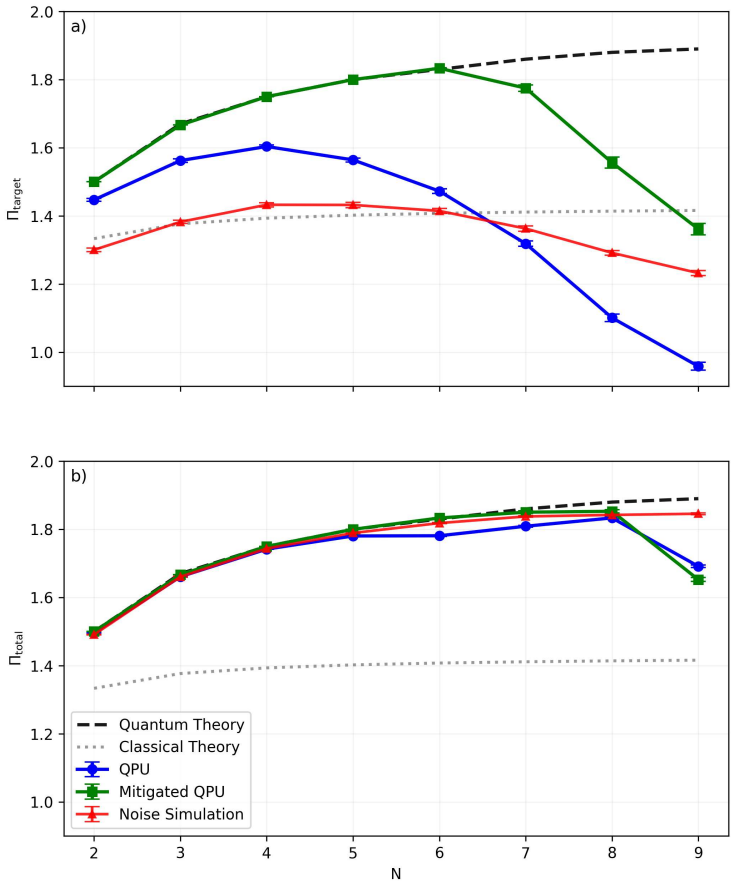


Fig. 5 Expected payoffs at optimization level L2 as a function of player count. Panel (a) shows the target-state weighted payoff contribution $P(|1\rangle^{\otimes N}) \cdot (b - c/N)$ for raw QPU, readout-corrected QPU, and Digital Twin, with theoretical benchmark and classical Nash equilibrium. Panel (b) shows the global average payoff for the same four regimes. Error bars represent 95% t-Student confidence intervals across 20 repetitions

5 Discussion

The central finding of this work is a systematic decoupling between state-level and aggregate-level behavior. With readout error correction, the global average payoff reproduces the quantum theoretical benchmark exactly for $N \leq 6$ and remains above

Table 2 Per-player readout correction impact: absolute and relative offset between raw and readout-corrected global average payoff at optimization level L2.

N	Absolute offset	Relative offset (%)
2	0.0038	0.25
3	0.0056	0.34
4	0.0081	0.47
5	0.0195	1.09
6	0.0522	2.93
7	0.0412	2.27
8	0.0195	1.06
9	-0.0390	-2.30

the classical Nash equilibrium across the full range $N = 2, \ldots, 9$. This aggregate advantage persists even as target-state fidelity degrades substantially with system size, reaching approximately 51% (raw) and 72% (corrected) at $N = 9$. The payoff structure of the Volunteer's Dilemma, in which $2^N - 1$ of the 2^N possible outcomes contribute positively to expected utility, is what sustains this advantage under noise.

5.1 Persistence of Quantum Advantage Under Noise

Two distinct thresholds emerge from the data, depending on the metric considered. For the *global average payoff*, readout-corrected results remain above the classical Nash equilibrium across the entire tested range ($N = 2, \ldots, 9$), with confidence intervals that do not overlap the classical baseline at any player count. Raw results also exceed the classical equilibrium for $N \leq 8$; at $N = 9$ the raw global payoff drops to 1.691 ± 0.004, which still clears the classical value of 1.416 by a statistically significant margin.

For the *target-state weighted payoff contribution*, the picture is more restrictive. Raw results fall below the classical baseline at $N = 7$; readout-corrected results maintain the advantage through $N = 8$ but drop below the classical value at $N = 9$. This metric therefore sets a more conservative practical threshold: quantum advantage is unambiguously resolvable up to $N = 6$ by either metric, with the global payoff extending the observable range further.

The Hamming distance analysis (Section 4.2) provides a structural explanation for this difference: probability redistributed to $d_H = 1$ states still contributes positive utility, acting as a natural damping mechanism for the global metric. The practical limit of quantum advantage therefore depends both on the total probability lost from the target state and on the specific metric used to assess it.

5.2 Digital Twin as a Noise Modeling Tool

The Digital Twin, as a first-order approximation built from independent per-qubit calibration channels, provides a useful but limited representation of the hardware behavior. Its primary strength lies in reproducing the qualitative structure of aggregate metrics: the global average payoff profile closely tracks the experimental trend,

and the model correctly identifies the direction and approximate magnitude of noise-induced degradation.

However, the Digital Twin's representation of target-state fidelity is less accurate. The model predicts a strictly linear decay, while the hardware exhibits a more gradual decline at lower player counts followed by a sharper degradation beyond $N = 6$. Quantitatively, the model overestimates noise for small N (predicting lower fidelity than observed) and underestimates it for large N in the raw data; for readout-corrected data, it consistently overestimates noise across the full range. This asymmetric behavior suggests that the non-local noise mechanisms omitted by the model, primarily crosstalk and temporal calibration drift, have a growing impact as circuit depth increases with N.

Despite these limitations, the low variance of the Digital Twin across 20 independent simulation seeds confirms that the model's stochastic component is well controlled, and that the discrepancies with hardware are systematic rather than random. The Digital Twin therefore remains a useful reference for trend identification, hypothesis testing, and estimating the contribution of individual noise channels, while its limitations motivate the development of more sophisticated models as discussed in Section 6.1.

5.3 Implications for Quantum Game Theory

The resilience of the global average payoff observed in Section 4.4 has a structural explanation rooted in the payoff geometry of the game. Any bitstring with at least one 1 contributes positively to the expected utility, so $2^N - 1$ of the 2^N possible outcomes pay positively. Under any noise model that spreads probability away from $|1\rangle^{\otimes N}$, the leaked probability mostly lands in states that still contribute. This payoff geometry, not a deep property of quantum mechanics, is what keeps the aggregate payoff well above the classical baseline.

The experiments do confirm that this payoff geometry is accessible on current hardware. The QPU global payoff exceeds the classical Nash equilibrium for all tested configurations, retaining its aggregate advantage over classical mixed strategies even at noise levels that reduce individual-outcome fidelity to approximately 51% at $N = 9$. Identifying which game structures will share this property on near-term devices is a useful open question for quantum game theory.

6 Conclusion

We have experimentally implemented the Quantum Volunteer's Dilemma for $N = 2, \ldots, 9$ players on the `ibm_kingston` quantum processor, providing a systematic characterization of this game under realistic NISQ noise conditions. Four transpiler optimization levels were evaluated across 20 independent repetitions with 2048 shots each, and post-processing readout error correction was applied throughout.

Two observability thresholds emerge depending on the metric. For the global average payoff, readout-corrected results exceed the classical Nash equilibrium across the entire tested range ($N = 2, \ldots, 9$), confirming that aggregate quantum advantage is robust to the noise levels of current NISQ hardware. For the target-state weighted

payoff contribution, the advantage holds through $N = 8$ with correction and through $N = 6$ without. The $N = 6$ boundary thus marks the point beyond which single-qubit dominated errors transition to a regime where multi-qubit contributions become relevant, rather than the disappearance of aggregate advantage.

The Hamming distance analysis of raw measurement counts reveals a non-uniform error redistribution dominated by single-qubit flips, consistent with the independent per-qubit noise channels of the Digital Twin and explaining why the global payoff is more resilient than target-state fidelity. Gate count analysis confirms that optimization level 2 minimizes ECR gates while providing the most stable fidelity across the full player range; level 3 produces identical transpiled circuits and offers no additional benefit for this topology.

The Digital Twin captures global aggregate trends but systematically misrepresents target-state fidelity, particularly at large N, revealing the limitations of first-order noise models in capturing non-local and depth-dependent hardware effects.

6.1 Future Work

The results obtained here open several directions for future research. The readout correction applied in this work represents only a minimal level of error mitigation; more advanced techniques [61], including zero-noise extrapolation (ZNE) [62–64] and probabilistic error cancellation (PEC) [62, 65], could improve the accuracy of absolute payoff estimates and extend the range of player numbers over which quantum advantage is statistically resolvable. Extending the experiment to larger games in a practically meaningful way would also require controlling the noise accumulated by the deeper EWL circuits, either through lower-depth implementations of the entangling operator or through hardware with lower gate error rates.

The Digital Twin's limitations at large N motivate the development of more sophisticated noise models that incorporate spatial qubit correlations, crosstalk between neighboring qubits in the CNOT chain, and temporal calibration drift. Such models would improve the predictive accuracy of the simulation and better characterize the transition from single-qubit dominated to multi-qubit correlated errors observed around $N = 6$.

Another direction is to extend the quantum framework to other variants of the Volunteer's Dilemma, including asymmetric models with player-dependent costs or benefits [66–70], threshold Volunteer's Dilemmas in which collective benefit requires sufficiently many volunteers [71–73], timing variants with delayed volunteering [67, 74], repeated variants involving sustained interaction among players [75–78], and cost-synergy models in which the volunteering cost decreases exponentially with the number of contributors [79]. Once such variants are formulated in the quantum setting, the experimental framework developed here could be used to test whether the payoff robustness observed in the Weesie and Franzen cost-sharing model [49] persists across different forms of the Volunteer's Dilemma. Adaptive versions of these games may also require mid-circuit measurements and feedforward [80], requiring mitigation methods tailored to such circuits [81].

Finally, applying this experimental framework to other multiplayer quantum games with different payoff structures [8, 82] would help identify which game-theoretic properties, beyond the specific cost-sharing mechanism studied here, favor noise robustness in NISQ implementations. This would help establish quantum games as a broader experimental setting for studying task-level robustness on NISQ hardware.

Acknowledgements. We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

Statements and Declarations

- **Funding:** The authors received no external funding for this work.
- **Competing Interests:** The authors declare no competing interests.
- **Data availability:** All experimental datasets, backend calibration files, and stored job results are publicly available at https://github.com/GermanDarDiaz/Quantum-Volunteer-s-Dilemma.
- **Code availability:** The complete Qiskit implementation, including circuit construction, noise modeling, and statistical analysis, is publicly available at https://github.com/GermanDarDiaz/Quantum-Volunteer-s-Dilemma.
- **Author contribution:** G.D.D.A. performed the experimental implementation and data analysis. J.A.A.H. performed formal analysis and manuscript review. C.A.D. coordinated the team and contributed to circuit design. S.A.C.O. contributed to the manuscript and research coordination. N.D.H. contributed to quantum game analysis. S.T.G. contributed to circuit design. D.E.K. provided the theoretical foundation.
- **Related work:** This article is partially based on the undergraduate thesis work of G.D.D.A. (in preparation, Universidad del Cauca). The present manuscript extends that work with additional experimental results, larger system sizes, and collaborative contributions.

Appendix A Complete Numerical Results

Table A1 Target-state probability results for optimization level L2. Reported values correspond to the mean over 20 executions with 95% confidence intervals.

N	P_{target}		
	Raw	Mitigated	Digital Twin
2	0.9647 ± 0.0025	$1.0000 \pm 1.78 \times 10^{-8}$	0.8721 ± 0.0037
3	0.9374 ± 0.0032	$1.0000 \pm 2.27 \times 10^{-8}$	0.8323 ± 0.0029
4	0.9164 ± 0.0024	$1.0000 \pm 3.04 \times 10^{-8}$	0.8217 ± 0.0036
5	0.8691 ± 0.0032	$1.0000 \pm 3.78 \times 10^{-8}$	0.8004 ± 0.0045
6	0.8033 ± 0.0039	$1.0000 \pm 4.23 \times 10^{-8}$	0.7780 ± 0.0040
7	0.7099 ± 0.0042	0.9558 ± 0.0049	0.7415 ± 0.0043
8	0.5871 ± 0.0059	0.8300 ± 0.0085	0.7010 ± 0.0034
9	0.5075 ± 0.0062	0.7208 ± 0.0089	0.6676 ± 0.0038

Table A2 Target-state weighted payoff contribution associated with the target state for optimization level L2. Reported values correspond to the mean over 20 executions with 95% confidence intervals.

N	Benchmark		Conditional payoff		
	Q_{th}	Q_{cl}	Raw	Mitigated	Digital Twin
2	1.50	1.333	1.4470 ± 0.0037	$1.5000 \pm 2.67 \times 10^{-8}$	1.3003 ± 0.0056
3	1.67	1.377	1.5623 ± 0.0053	$1.6667 \pm 3.78 \times 10^{-8}$	1.3824 ± 0.0048
4	1.75	1.393	1.6037 ± 0.0042	$1.7500 \pm 5.32 \times 10^{-8}$	1.4326 ± 0.0063
5	1.80	1.402	1.5643 ± 0.0057	$1.8000 \pm 6.81 \times 10^{-8}$	1.4322 ± 0.0080
6	1.83	1.408	1.4727 ± 0.0072	$1.8333 \pm 7.76 \times 10^{-8}$	1.4146 ± 0.0073
7	1.86	1.411	1.3185 ± 0.0078	1.7750 ± 0.0092	1.3628 ± 0.0079
8	1.88	1.414	1.1009 ± 0.0110	1.5563 ± 0.0160	1.2915 ± 0.0064
9	1.89	1.416	0.9586 ± 0.0117	1.3615 ± 0.0168	1.2322 ± 0.0072

Table A3 Global payoff results for optimization level L2 considering all measured basis states. Reported values correspond to the mean over 20 executions with 95% confidence intervals.

N	Benchmark		Global payoff		
	Q_{th}	Q_{cl}	Raw	Mitigated	Digital Twin
2	1.50	1.333	1.4962 ± 0.0008	$1.5000 \pm 2.67 \times 10^{-8}$	1.4910 ± 0.0011
3	1.67	1.377	1.6611 ± 0.0009	$1.6667 \pm 3.78 \times 10^{-8}$	1.6608 ± 0.0011
4	1.75	1.393	1.7419 ± 0.0009	$1.7500 \pm 5.32 \times 10^{-8}$	1.7435 ± 0.0011
5	1.80	1.402	1.7805 ± 0.0020	$1.8000 \pm 6.81 \times 10^{-8}$	1.7892 ± 0.0016
6	1.83	1.408	1.7811 ± 0.0024	$1.8333 \pm 7.76 \times 10^{-8}$	1.8183 ± 0.0014
7	1.86	1.411	1.8091 ± 0.0026	1.8503 ± 0.0040	1.8379 ± 0.0018
8	1.88	1.414	1.8334 ± 0.0033	1.8529 ± 0.0046	1.8423 ± 0.0020
9	1.89	1.416	1.6914 ± 0.0039	1.6525 ± 0.0058	1.8458 ± 0.0026

Table A4 Absolute error (AE) and relative error (RE) with respect to the theoretical quantum benchmark for the raw QPU results, readout-corrected results, and Digital Twin simulation. The final columns quantify the impact of readout correction relative to the raw experimental measurements.

N	**Raw QPU**		**Readout Corrected**		**Digital Twin**	
	AE	RE (%)	AE	RE (%)	AE	RE (%)
2	0.0038	0.2515	0.0000	0.0000	0.0090	0.6030
3	0.0089	0.5358	0.0033	0.1996	0.0092	0.5505
4	0.0081	0.4639	0.0000	0.0000	0.0065	0.3711
5	0.0195	1.0815	0.0000	0.0000	0.0108	0.5981
6	0.0489	2.6697	0.0033	0.1822	0.0117	0.6397
7	0.0509	2.7351	0.0097	0.5223	0.0221	1.1896
8	0.0466	2.4769	0.0271	1.4407	0.0377	2.0069
9	0.1986	10.5067	0.2375	12.5686	0.0442	2.3377

Appendix B Transpiler Optimization Levels

The Qiskit transpiler accepts an `optimization_level` parameter (integer, 0 to 3) that controls the degree of circuit transformation applied before hardware execution [83]. Level 0 applies no significant structural optimizations and preserves the original circuit structure. Level 1 introduces basic optimizations including redundant gate cancellation and initial mapping improvements. Level 2 incorporates additional heuristics for qubit assignment and gate reduction to better adapt the circuit to the backend topology. Level 3 employs more aggressive transformations including resynthesis and advanced two-qubit gate reduction. Higher optimization levels generally reduce effective circuit complexity, potentially decreasing error accumulation and improving experimental performance.

References

[1] Diekmann, A.: Volunteer's dilemma. Journal of Conflict Resolution **29**(4), 605–610 (1985) https://doi.org/10.1177/0022002785029004003

[2] Diekmann, A.: Volunteer's Dilemma. A Social Trap without a Dominant Strategy and some Empirical Results, pp. 187–197. Physica-Verlag HD, Heidelberg (1986). https://doi.org/10.1007/978-3-642-95874-8_13

[3] Campos-Mercade, P.: The volunteer's dilemma explains the bystander effect. Journal of Economic Behavior & Organization **186**, 646–661 (2021) https://doi.org/10.1016/j.jebo.2020.11.012

[4] Otsubo, H., Rapoport, A.: Dynamic volunteer's dilemmas over a finite horizon: An experimental study. Journal of Conflict Resolution **52**(6), 961–984 (2008) https://doi.org/10.1177/0022002708321401

[5] Goeree, J.K., Holt, C.A., Smith, A.M.: An experimental examination of the volunteer's dilemma. Games and Economic Behavior **102**, 303–315 (2017) https://doi.org/10.1016/j.geb.2017.01.002

[6] Meyer, D.A.: Quantum strategies. Physical Review Letters **82**(5), 1052–1055 (1999) https://doi.org/10.1103/PhysRevLett.82.1052

[7] Eisert, J., Wilkens, M., Lewenstein, M.: Quantum games and quantum strategies. Physical Review Letters **83**(15), 3077–3080 (1999) https://doi.org/10.1103/PhysRevLett.83.3077

[8] Khan, F.S., Solmeyer, N., Balu, R., Humble, T.S.: Quantum games: a review of the history, current state, and interpretation. Quantum Information Processing **17**(11), 309 (2018) https://doi.org/10.1007/s11128-018-2082-8

[9] Eisert, J., Wilkens, M.: Quantum games. Journal of Modern Optics **47**(14-15), 2543–2556 (2000) https://doi.org/10.1080/09500340008232180

[10] Du, J., Xu, X., Li, H., Zhou, X., Han, R.: Nash equilibrium in the quantum battle of sexes game. arXiv preprint quant-ph/0010050 (2000) https://doi.org/10.48550/arXiv.quant-ph/0010050

[11] Du, J., Xu, X., Li, H., Zhou, X., Han, R.: Playing prisoner's dilemma with quantum rules. Fluctuation and Noise Letters **2**(04), 189–203 (2002) https://doi.org/10.1142/S0219477502000993

[12] Piotrowski, E.W., Sładkowski, J.: Quantum market games. Physica A: Statistical Mechanics and its Applications **312**(1), 208–216 (2002) https://doi.org/10.1016/S0378-4371(02)00842-7

[13] Maioli, A.C., Passos, M.H.M., Balthazar, W.F., Souza, C.E.R., Huguenin, J.A.O., Schmidt, A.G.M.: Quantization and experimental realization of the Colonel Blotto game. Quantum Information Processing **18**(1), 10 (2018) https://doi.org/10.1007/s11128-018-2113-5

[14] Szopa, M.: Efficiency of classical and quantum games equilibria. Entropy **23**(5), 506 (2021) https://doi.org/10.3390/e23050506

[15] Kastampolidou, K., Andronikos, T.: Quantum tapsilou—a quantum game inspired by the traditional greek coin tossing game tapsilou. Games **14**(6) (2023) https://doi.org/10.3390/g14060072

[16] Frąckiewicz, P., Szopa, M.: Permissible extensions of classical to quantum games combining three strategies. Quantum Information Processing **23**(3), 75 (2024) https://doi.org/10.1007/s11128-024-04283-3

[17] Andronikos, T., Bitsakos, C., Nikas, K., Goumas, G.I., Koziris, N.: A GHZ-based protocol for the dining information brokers problem. Future Internet **17**(9) (2025) https://doi.org/10.3390/fi17090408

[18] Koh, D.E., Kumar, K., Goh, S.T.: Quantum volunteer's dilemma. Physical Review Research **7**(1), 013104 (2025) https://doi.org/10.1103/PhysRevResearch.7.013104

[19] Khan, F.S., Linke, N.M., Than, A.T., Baron, D.: Quantum advantage in trading: A game-theoretic approach. Quantum Economics and Finance **2**(1), 40–51 (2025) https://doi.org/10.1177/29767032251333418

[20] Frąckiewicz, P., Gorczyca-Goraj, A., Szopa, M.: Permissible four-strategy quantum extensions of classical games. Quantum Information Processing **24**(7), 201 (2025) https://doi.org/10.1007/s11128-025-04808-4

[21] Tiago, G., Naskar, J., Maioli, A.C., Balthazar, W.F., Schmidt, A.G.M., Huguenin, J.A.O.: Classical and quantum multiplayer Colonel Blotto game in all-optical setup. Quantum Information Processing **24**(5), 152 (2025) https://doi.org/10.

1007/s11128-025-04760-3

[22] Weeks, C., Strange, P., Drmota, P., Quintanilla, J.: Quantum-assisted graph domination games. arXiv preprint arXiv:2511.15802 (2025) https://doi.org/10.48550/arXiv.2511.15802

[23] Makram-Allah, A.T.M., Abd-Rabbou, M.Y., Metwally, N.: The possibility of fairly mitigating penalties in the quantum prisoner's dilemma via classical noise channels. The European Physical Journal Plus **141**(4), 456 (2026) https://doi.org/10.1140/epjp/s13360-026-07650-4

[24] Ahmad, R.: Native quantum games from interacting discrete-time quantum walks. arXiv preprint arXiv:2604.20455 (2026) https://doi.org/10.48550/arXiv.2604.20455

[25] Essalmi, K., Garrido, F., Nashashibi, F.: Multi-player, multi-strategy quantum game model for interaction-aware decision-making in autonomous driving. arXiv preprint arXiv:2602.03571 (2026) https://doi.org/10.48550/arXiv.2602.03571

[26] Benjamin, S.C., Hayden, P.M.: Multiplayer quantum games. Phys. Rev. A **64**, 030301 (2001) https://doi.org/10.1103/PhysRevA.64.030301

[27] Jiang-Feng, D., Hui, L., Xiao-Dong, X., Xian-Yi, Z., Rong-Dian, H.: Multi-player and multi-choice quantum game. Chinese Physics Letters **19**(9), 1221 (2002) https://doi.org/10.1088/0256-307X/19/9/301

[28] Li, H., Du, J., Massar, S.: Continuous-variable quantum games. Physics Letters A **306**(2-3), 73–78 (2002) https://doi.org/10.1016/S0375-9601(02)01628-6

[29] Ikeda, K., Aoki, S.: Infinitely repeated quantum games and strategic efficiency. Quantum Information Processing **20**(12), 387 (2021) https://doi.org/10.1007/s11128-021-03295-7

[30] Du, J., Xu, X., Li, H., Zhou, X., Han, R.: Entanglement playing a dominating role in quantum games. Physics Letters A **289**(1-2), 9–15 (2001) https://doi.org/10.1016/S0375-9601(01)00575-8

[31] Du, J., Li, H., Xu, X., Zhou, X., Han, R.: Entanglement enhanced multiplayer quantum games. Physics Letters A **302**(5), 229–233 (2002) https://doi.org/10.1016/S0375-9601(02)01144-1

[32] Li, A., Yong, X.: Entanglement guarantees emergence of cooperation in quantum prisoner's dilemma games on networks. Scientific reports **4**(1), 6286 (2014) https://doi.org/10.1038/srep06286

[33] Wei, Z., Zhang, S.: Quantum game players can have advantage without discord. Information and Computation **256**, 174–184 (2017) https://doi.org/10.1016/j.ic.2017.07.004

[34] Mohamed, N.M.A.A., Taisheng, H., Jinhui, P.: Quantum game theory on entangled players. In: International Conference on Intelligent Information Technologies for Industry, pp. 291–301 (2023). https://doi.org/10.1007/978-3-031-43789-2_27 . Springer

[35] Johnson, N.F.: Playing a quantum game with a corrupted source. Physical Review A **63**(2), 020302 (2001) https://doi.org/10.1103/PhysRevA.63.020302

[36] Chen, L.K., Ang, H., Kiang, D., Kwek, L.C., Lo, C.F.: Quantum prisoner dilemma under decoherence. Physics Letters A **316**(5), 317–323 (2003) https://doi.org/10.1016/S0375-9601(03)01175-7

[37] Flitney, A.P., Abbott, D.: Quantum games with decoherence. Journal of Physics A: Mathematical and General **38**(2), 449 (2004) https://doi.org/10.1088/0305-4470/38/2/011

[38] Shuai, C., Mao-Fa, F., Xiao-Juan, Z.: The effect of quantum noise on multiplayer quantum game. Chinese Physics **16**(4), 915–918 (2007) https://doi.org/10.1088/1009-1963/16/4/008

[39] Huang, Z., Qiu, D.: Quantum games under decoherence. International Journal of Theoretical Physics **55**, 965–992 (2016) https://doi.org/10.1007/s10773-015-2741-9

[40] Khan, S., Alam, S.: The dynamics of Nash equilibrium under non-Markovian classical noise in quantum prisoners' dilemma. Reports on Mathematical Physics **81**(3), 399–413 (2018) https://doi.org/10.1016/S0034-4877(18)30056-9

[41] Kairon, P., Thapliyal, K., Srikanth, R., Pathak, A.: Noisy three-player dilemma game: robustness of the quantum advantage. Quantum information processing **19**, 1–18 (2020) https://doi.org/10.1007/s11128-020-02830-2

[42] Legón, A.R., Medina, E.: Joint probabilities approach to quantum games with noise. Entropy **25**(8), 1222 (2023) https://doi.org/10.3390/e25081222

[43] Lu, J., Zhou, L., Kuang, L.-M.: Linear optics implementation for quantum game with two players. Physics Letters A **330**(1-2), 48–53 (2004) https://doi.org/10.1016/j.physleta.2004.07.063

[44] Buluta, I.M., Fujiwara, S., Hasegawa, S.: Quantum games in ion traps. Physics Letters A **358**(2), 100–104 (2006) https://doi.org/10.1016/j.physleta.2006.04.114

[45] Mitra, A., Sivapriya, K., Kumar, A.: Experimental implementation of a three qubit quantum game with corrupt source using nuclear magnetic resonance quantum information processor. Journal of Magnetic Resonance **187**(2), 306–313 (2007) https://doi.org/10.1016/j.jmr.2007.05.013

[46] Schmid, C., Flitney, A.P., Wieczorek, W., Kiesel, N., Weinfurter, H., Hollenberg, L.C.: Experimental implementation of a four-player quantum game. New Journal of Physics **12**(6), 063031 (2010) https://doi.org/10.1088/1367-2630/12/6/063031

[47] Xu, K., Wu, Z.: Experimental implementation of quantum prisoner dilemma on IBM quantum computers. In: 2022 15th International Conference on Advanced Computer Theory and Engineering (ICACTE), pp. 13–18 (2022). https://doi.org/10.1109/ICACTE55855.2022.9943732 . IEEE

[48] Díaz Agreda, G., Durán Paredes, C.A., Buenaventura Samboni, M., Andrade, J.A., Cajas Ordoñez, S.: Bridging theory and practice in quantum game theory: Optimized implementation of the battle of the sexes with error mitigation on NISQ hardware. In: 2025 IEEE CHILEAN Conference on Electrical, Electronics Engineering, Information and Communication Technologies (CHILECON), pp. 1–10 (2025). https://doi.org/10.1109/CHILECON66915.2025.11476221

[49] Weesie, J., Franzen, A.: Cost sharing in a volunteer's dilemma. Journal of Conflict Resolution **42**(5), 600–618 (1998) https://doi.org/10.1177/0022002798042005004

[50] Preskill, J.: Quantum computing in the NISQ era and beyond. Quantum **2**, 79 (2018) https://doi.org/10.22331/q-2018-08-06-79

[51] Cheng, B., Deng, X.-H., Gu, X., He, Y., Hu, G., Huang, P., Li, J., Lin, B.-C., Lu, D., Lu, Y., Qiu, C., Wang, H., Xin, T., Yu, S., Yung, M.-H., Zeng, J., Zhang, S., Zhong, Y., Peng, X., Nori, F., Yu, D.: Noisy intermediate-scale quantum computers. Frontiers of Physics **18**(2), 21308 (2023) https://doi.org/10.1007/s11467-022-1249-z

[52] Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press, Cambridge (2010). https://doi.org/10.1017/CBO9780511976667

[53] Gambetta, J.M., Chow, J.M., Steffen, M.: Building logical qubits in a superconducting quantum computing system. npj Quantum Information **3**(1), 2 (2017) https://doi.org/10.1038/s41534-016-0004-0

[54] Murali, P., McKay, D.C., Martonosi, M., Javadi-Abhari, A.: Software mitigation of crosstalk on noisy intermediate-scale quantum computers. In: Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '20), pp. 1001–1016. ACM, New York (2020). https://doi.org/10.1145/3373376.3378477

[55] Sarovar, M., Proctor, T., Rudinger, K., Young, K., Nielsen, E., Blume-Kohout, R.: Detecting crosstalk errors in quantum information processors. Quantum **4**, 321 (2020) https://doi.org/10.22331/q-2020-09-11-321

[56] IBM Quantum: QiskitRuntimeService API reference. https://quantum.cloud.ibm.com/docs/en/api/qiskit-ibm-runtime/qiskit_ibm_runtime.QiskitRuntimeService. Accessed: 2026-04-27 (n.d.)

[57] Nation, P.D., Kang, H., Sundaresan, N., Gambetta, J.M.: Scalable mitigation of measurement errors on quantum computers. PRX Quantum **2**(4), 040326 (2021) https://doi.org/10.1103/PRXQuantum.2.040326

[58] Qiskit Contributors: mthree: Matrix-free Measurement Mitigation. Accessed: Apr. 2026 (2024). https://qiskit.github.io/qiskit-addon-mthree/

[59] IBM Quantum: qiskit.transpiler.generate_preset_pass_manager. https://quantum.cloud.ibm.com/docs/en/api/qiskit/qiskit.transpiler.generate_preset_pass_manager. Accessed: 2026-04-27 (n.d.)

[60] IBM Quantum: Exact and noisy simulation with Qiskit Aer primitives. https://quantum.cloud.ibm.com/docs/en/guides/simulate-with-qiskit-aer. Accessed: 2026-04-27 (n.d.)

[61] 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. **95**, 045005 (2023) https://doi.org/10.1103/RevModPhys.95.045005

[62] Temme, K., Bravyi, S., Gambetta, J.M.: Error mitigation for short-depth quantum circuits. Phys. Rev. Lett. **119**, 180509 (2017) https://doi.org/10.1103/PhysRevLett.119.180509

[63] Li, Y., Benjamin, S.C.: Efficient variational quantum simulator incorporating active error minimization. Phys. Rev. X **7**, 021050 (2017) https://doi.org/10.1103/PhysRevX.7.021050

[64] Giurgica-Tiron, T., Hindy, Y., LaRose, R., Mari, A., Zeng, W.J.: Digital zero noise extrapolation for quantum error mitigation. In: 2020 IEEE International Conference on Quantum Computing and Engineering (QCE), pp. 306–316 (2020). https://doi.org/10.1109/QCE49297.2020.00045

[65] Endo, S., Benjamin, S.C., Li, Y.: Practical quantum error mitigation for near-future applications. Phys. Rev. X **8**, 031027 (2018) https://doi.org/10.1103/PhysRevX.8.031027

[66] Diekmann, A.: Cooperation in an asymmetric volunteer's dilemma game theory and experimental evidence. International Journal of Game Theory **22**(1), 75–85 (1993) https://doi.org/10.1007/BF01245571

[67] Weesie, J.: Asymmetry and timing in the volunteer's dilemma. Journal of conflict resolution **37**(3), 569–590 (1993) https://doi.org/10.1177/0022002793037003008

[68] He, J.-Z., Wang, R.-W., Li, Y.-T.: Evolutionary stability in the asymmetric volunteer's dilemma. PloS one **9**(8), 103931 (2014) https://doi.org/10.1371/journal.pone.0103931

[69] Healy, A.J., Pate, J.G.: Cost asymmetry and incomplete information in a volunteer's dilemma experiment. Social Choice and Welfare **51**(3), 465–491 (2018) https://doi.org/10.1007/s00355-018-1124-6

[70] Guo, Z.-X., He, J.-Z., Li, Q.-M., Shi, L., Wang, R.-W.: Asymmetric interaction and diverse forms in public goods production in volunteer dilemma game. Chaos, Solitons & Fractals **166**, 112928 (2023) https://doi.org/10.1016/j.chaos.2022.112928

[71] Chen, X., Gross, T., Dieckmann, U.: Shared rewarding overcomes defection traps in generalized volunteer's dilemmas. Journal of theoretical biology **335**, 13–21 (2013) https://doi.org/10.1016/j.jtbi.2013.06.014

[72] Mago, S.D., Pate, J.: Greed and fear: Competitive and charitable priming in a threshold volunteer's dilemma. Economic Inquiry **61**(1), 138–161 (2023) https://doi.org/10.1111/ecin.13117

[73] O'Riordan, C.: The k-threshold volunteer's dilemma on networks. Available at SSRN 6765666 (2026) https://doi.org/10.2139/ssrn.6765666

[74] Weesie, J.: Incomplete information and timing in the volunteer's dilemma: A comparison of four models. Journal of Conflict Resolution **38**(3), 557–585 (1994) https://doi.org/10.1177/0022002794038003008

[75] Przepiorka, W., Bouman, L., de Kwaadsteniet, E.W.: The emergence of conventions in the repeated volunteer's dilemma: The role of social value orientation, payoff asymmetries and focal points. Social Science Research **93**, 102488 (2021) https://doi.org/10.1016/j.ssresearch.2020.102488

[76] Kloosterman, A., Mago, S.: The infinitely repeated volunteer's dilemma: An experimental study. Games and Economic Behavior **142**, 812–832 (2023) https://doi.org/10.1016/j.geb.2023.10.008

[77] Amir, R., Liu, H., Mayskaya, T., Tian, J.: On the repeated volunteer's dilemma with equal cost sharing. Available at SSRN 5634211 (2026) https://doi.org/10.2139/ssrn.5634211

[78] Mukhopadhyay, A., Sur, S., Saha, T., Sadhukhan, S., Chakraborty, S.: Repeated quantum game as a stochastic game: Effects of the shadow of the future and entanglement. Physica A: Statistical Mechanics and its Applications **637**, 129613 (2024) https://doi.org/10.1016/j.physa.2024.129613

[79] Amir, R., Machowska, D., Tian, J.: The volunteer's dilemma with cost synergies.

Journal of Conflict Resolution **70**(6), 1095–1117 (2026) https://doi.org/10.1177/00220027251388601

[80] Hothem, D., Hines, J., Baldwin, C., Gresh, D., Blume-Kohout, R., Proctor, T.: Measuring error rates of mid-circuit measurements. Nature Communications **16**(1), 5761 (2025) https://doi.org/10.1038/s41467-025-60923-x

[81] Koh, J.M., Koh, D.E., Thompson, J.: Readout error mitigation for mid-circuit measurements and feedforward. PRX Quantum **7**, 010317 (2026) https://doi.org/10.1103/cj89-4h5t

[82] Toni, B.: In: Pimentel, E.A., Toni, B. (eds.) Compendium of Advances in Game Theory: Classical, Differential, Algorithmic, Non-Archimedean and Quantum Games, pp. 1–47. Springer, Cham (2026). https://doi.org/10.1007/978-3-031-97733-6_1

[83] IBM Quantum: Set transpiler optimization level. IBM Quantum Documentation. Accessed: Apr. 2026 (2025). https://quantum.cloud.ibm.com/docs/en/guides/set-optimization