

Supplementary File S2

Experiment manifest and run verification

The final traceable run used the Qiskit statevector executor and was completed with 160 configurations and 6,400 completed games. The manifest also confirms that all trials were two-player, all configurations had 40 runs, and the paired layout packet was reused across rule levels.

Table S2.1. Final experiment manifest summary.

Field	Value
Run command	python -m app.experiments.runner --trials 40 --seed 42 --output-dir backend/experiment outputs/final traceable manuscript_run 20260703 191713
Git commit	58491fe7046b647304cdfbfef29713778a66bb9
Timestamp	2026-07-03T13:52:38.434605+00:00
Board sizes	5×5, 10×10
Policies	Random, Checkerboard + follow-up, Belief-state, Greedy max-probability
Rule levels	L0, L1, L2, L3
Start orders	P1, P2
Main measurement probability	0.5
Sensitivity measurement probabilities	0.25, 0.75
Layout seed	42
Game seed	42
Qiskit available	True
Executor used	qiskit statevector
Number of configurations	160
Number of completed games	6400
Number of failed games	0
All trials two-player	True
Runs per configuration	40
Paired layout packet reused across levels	True
Duplicate branch pairs after fix	5×5: 0; 10×10: 0

Table S2.2. Manifest verification checks.

The verification block records the completed and expected configuration counts, completed and expected game counts, any bad configuration counts, and two-player confirmation.

Verification field	Value
completed configurations	160
expected configurations	160
completed games	6400
expected games	6400
bad configuration counts	{}
all configurations have 40 runs	True
10 probability values	['baseline']
all two player	True

Belief-policy validation checks

The belief-state policy validation file reports that all validation checks passed. These checks confirm normalization of the initial belief distribution, elimination of impossible states after misses, correct branch-resolution matching, and selection of the highest-scoring legal coordinate with deterministic lexicographic tie handling in the reported tied case.

Table S2.3. Belief-state policy validation checks.

Check	Passed	Additional details
initial_normalizes_to_1	True	
miss_removes_impossible_states	True	
branch_resolution_matches_selected_branch	True	
belief selects highest scoring legal coordinate	True	target=B2; score=0.06; tie count=2; tie break=lexicographic

Raw-data and generated-file inventory

The following inventory lists the machine-readable files provided with the trace run. Row and column counts are from the uploaded CSV files. SHA256 hashes are computed from the uploaded files and match the tracefile values for the generated CSV files.

Table S2.4. Raw-data and generated-file inventory.

File	Rows	Columns	SHA256	Role in trace run
trial_logs_main_p050.csv	2560	49	0ac763766185a64219768620e621f1e08277cf57667a880ca2a76f6b304d005e	Main p = 0.50 raw trial log for all board-size, policy, level, and start-order conditions.
trial_logs_sensitivity_p025_p075.csv	3840	49	d64ffb269b2e0900257713d5259101501de32d4d0c3ee5979054c1fd6d8f38e	Sensitivity raw trial log for p = 0.25 and p = 0.75 modified-level conditions.
summary_main_p050.csv	32	21	bccd9b86ec4881e55c541fce0f7ba81911a06a06fb6023d8cb329bae6f35a4e3	Main p = 0.50 summary table with terminal-turn, paired comparison, outcome, and L3 rate fields.
summary_sensitivity.csv	80	21	914f4353fd0f1b76449358b080767ebbadcf23c769117b41372b192b7d830106	Sensitivity summary table across baseline, p = 0.25, p = 0.50, and p = 0.75 conditions.
fairness_summary.csv	80	13	4f84c1b482f87d1a2ab43357dbc7800c30ba78bb39436bfc6134103ee0bf1e7	Terminal outcome and fairness-related summary table.
l3_decomposition_summary.csv	24	11	6fcbd263aa964ca217388ba7839cc0c5c838499effbc78ac11d235498dd2b5cc	L3 direct-hit, correlated-damage, sunk-ship, and terminal-correlated summary table.
manuscript_table_values_from_trace_run.csv	752	6	008898f5158de4211ad82769c7a3c84ce2cf564ba76e8d82f420fb2f06e27317	Long-format export of manuscript table values from the trace run.
table1_trace_values.csv	8	15	b2b59af761bb185b9a647353d7d73742ba762163f32b43171952695be982448a	Manuscript Table 1 trace export.
table2_trace_values.csv	32	12	541fb305e349150c504cf1c38c77efea9ec28ab1aedfd46b297f3dbf612457af	Manuscript Table 2 trace export.
table3_trace_values.csv	24	7	8119c53e63c658d697b67f74d43f2239a45e16e90947a93ee6d28dc1bb9d1c88	Manuscript Table 3 trace export.
table4_l3_trace_values.csv	8	10	40cad4de9f8222bb3e9f23bea6b93d91259c08b3f4b29898780d1f167278fe7	Manuscript Table 4 trace export.
table4_difference_report.csv	8	17	92113963d7367e034dc8ee3bec61d2b284bba23226b381cb9ead4189cfd8b662	Audit file comparing older Table 4 values to final trace values.

Raw trial-log field inventory

The main and sensitivity raw trial logs share the same 49 exported fields. The field names below are listed exactly as they appear in the CSV files.

Table S2.5. Raw trial-log column fields.

No.	Column field
1	board_size
2	trial_id
3	layout_id
4	p1_layout_id
5	p2_layout_id
6	policy
7	level
8	start_order
9	measurement_probability
10	effective_probability
11	executor_used
12	quantum_backend_source
13	terminal_result
14	terminal_winner
15	winner
16	draw
17	draw_reason
18	terminal_turns
19	p1_shots
20	p2_shots
21	p1_moves
22	p2_moves
23	p1_shot_sequence
24	p2_shot_sequence
25	p1_first_hit_turn

No.	Column field
26	p2 first hit turn
27	first hit turn
28	direct target hits
29	direct target destructions
30	correlated damage events
31	correlated damage coordinates
32	sunk ship events
33	total damage events
34	terminal caused or accelerated by correlated damage
35	quantum events count
36	selected branch indices
37	correlated link candidates used
38	quantum event log
39	branch selection log
40	link resolution log
41	shot log
42	l3 direct hit rate
43	l3 correlated damage rate
44	l3 direct hits per move
45	l3 direct destructions per move
46	l3 correlated damage events per move
47	l3 total damage events per move
48	all trials two player
49	errors warnings

Code-state and package traceability

The tracefile records the git commit, branch, run command, local output directory, Python version, and core package versions used for the trace run.

Table S2.6. Code and environment traceability.

Trace field	Value
Git commit hash	58491fe7046b647304cdfbfeff29713778a66bb9
Branch	main
Run command	PYTHONPATH=backend python -m app.experiments.runner --trials 40 --seed 42 --output-dir backend/experiment outputs/final traceable manuscript run 20260703 191713
Run timestamp	2026-07-03T19:17:13+05:30
Output directory	/Users/samarthlamba/Documents/Playground/quantum-battleship-mvp-work/backend/experiment outputs/final traceable manuscript run 20260703 191713
Seed	42
Trials per configuration	40
Python version	Python 3.13.5
Captured package versions	scipy==1.16.3; qiskit==2.3.0; numpy==2.3.5; matplotlib==3.10.8

Full trace-run summary tables

The following tables reproduce the final trace-run values used by the manuscript. Board labels use 5×5 and 10×10. L0 is treated as the fixed classical baseline because it has no circuit-measured transition. In the $p = 0.50$ tables, each board-policy-level condition combines both start orders and has $n = 80$ completed games.

Table S2.8. Main $p = 0.50$ terminal move counts and L0 comparisons.

Values are mean terminal turns $\pm$ 95% CI. The Δ values are trial id- and start order-matched $L_k - L_0$ differences. Positive Δ indicates more terminal turns than L0.

Board	Policy	L0 mean $\pm$ CI	L1 mean $\pm$ CI	L1 comparison	L2 mean $\pm$ CI	L2 comparison	L3 mean $\pm$ CI	L3 comparison
5×5	Random	22.43 $\pm$ 0.41	31.74 $\pm$ 0.78	$\Delta=+9.31$; $d=2.62$; $p<0.001$	31.74 $\pm$ 0.72	$\Delta=+9.31$; $d=2.55$; $p<0.001$	28.36 $\pm$ 0.77	$\Delta=+5.94$; $d=1.45$; $p<0.001$
5×5	Checkerboard + follow-up	17.65 $\pm$ 0.62	25.12 $\pm$ 0.86	$\Delta=+7.47$; $d=2.33$; $p<0.001$	25.45 $\pm$ 0.84	$\Delta=+7.80$; $d=1.62$; $p<0.001$	22.23 $\pm$ 0.88	$\Delta=+4.58$; $d=0.99$; $p<0.001$
5×5	Belief-state	16.55 $\pm$ 0.36	23.73 $\pm$ 0.77	$\Delta=+7.17$; $d=2.31$; $p<0.001$	23.40 $\pm$ 0.72	$\Delta=+6.85$; $d=1.98$; $p<0.001$	20.99 $\pm$ 0.76	$\Delta=+4.44$; $d=1.03$; $p<0.001$
5×5	Greedy max-probability	16.55 $\pm$ 0.36	24.40 $\pm$ 0.79	$\Delta=+7.85$; $d=2.26$; $p<0.001$	24.18 $\pm$ 0.88	$\Delta=+7.62$; $d=1.95$; $p<0.001$	20.49 $\pm$ 0.81	$\Delta=+3.94$; $d=0.91$; $p<0.001$
10×10	Random	88.10 $\pm$ 1.88	104.42 $\pm$ 1.08	$\Delta=+16.32$; $d=1.79$; $p<0.001$	104.41 $\pm$ 0.87	$\Delta=+16.31$; $d=1.67$; $p<0.001$	99.95 $\pm$ 1.39	$\Delta=+11.85$; $d=1.09$; $p<0.001$
10×10	Checkerboard + follow-up	41.48 $\pm$ 2.13	50.11 $\pm$ 2.26	$\Delta=+8.64$; $d=2.50$; $p<0.001$	52.98 $\pm$ 2.25	$\Delta=+11.50$; $d=0.89$; $p<0.001$	51.74 $\pm$ 1.74	$\Delta=+10.26$; $d=0.88$; $p<0.001$
10×10	Belief-state	40.80 $\pm$ 1.75	49.48 $\pm$ 1.92	$\Delta=+8.68$; $d=1.94$; $p<0.001$	49.90 $\pm$ 2.05	$\Delta=+9.10$; $d=0.77$; $p<0.001$	44.50 $\pm$ 2.09	$\Delta=+3.70$; $d=0.32$; $p=0.005$
10×10	Greedy max-probability	40.80 $\pm$ 1.75	49.76 $\pm$ 1.89	$\Delta=+8.96$; $d=2.29$; $p<0.001$	48.51 $\pm$ 2.01	$\Delta=+7.71$; $d=0.65$; $p<0.001$	45.31 $\pm$ 2.35	$\Delta=+4.51$; $d=0.36$; $p=0.002$

Table S2.9. Main $p = 0.50$ terminal outcomes and fairness-related metrics.

Starting-player win rate combines both start orders. Win-rate disparity is $|P1 \text{ wins} - P2 \text{ wins}| / n$, where n includes draws. Mirrored-start parity measures whether outcomes matched after reversing the first mover.

Board	Policy	Level	Prob.	n	P1 wins	P2 wins	Draws	Starting-player win rate	Win-rate disparity	Draw rate	Mirrored-start parity
5×5	Random	L0	baseline	80	42	38	0	60.0%	5.0%	0.0%	95.0%
5×5	Random	L1	0.50	80	44	36	0	60.0%	10.0%	0.0%	90.0%
5×5	Random	L2	0.50	80	46	34	0	55.0%	15.0%	0.0%	85.0%
5×5	Random	L3	0.50	80	34	42	4	41.2%	10.0%	5.0%	92.5%
5×5	Checkerboard + follow-up	L0	baseline	80	45	35	0	56.2%	12.5%	0.0%	87.5%
5×5	Checkerboard + follow-up	L1	0.50	80	53	27	0	48.7%	32.5%	0.0%	67.5%
5×5	Checkerboard + follow-up	L2	0.50	80	36	44	0	42.5%	10.0%	0.0%	90.0%
5×5	Checkerboard + follow-up	L3	0.50	80	44	36	0	52.5%	10.0%	0.0%	90.0%
5×5	Belief-state	L0	baseline	80	42	38	0	57.5%	5.0%	0.0%	95.0%
5×5	Belief-state	L1	0.50	80	40	40	0	47.5%	0.0%	0.0%	100.0%
5×5	Belief-state	L2	0.50	80	43	37	0	61.3%	7.5%	0.0%	92.5%
5×5	Belief-state	L3	0.50	80	43	34	3	40.0%	11.3%	3.7%	90.0%
5×5	Greedy max-probability	L0	baseline	80	42	38	0	57.5%	5.0%	0.0%	95.0%
5×5	Greedy max-probability	L1	0.50	80	41	39	0	48.7%	2.5%	0.0%	97.5%
5×5	Greedy max-probability	L2	0.50	80	42	38	0	57.5%	5.0%	0.0%	95.0%
5×5	Greedy max-probability	L3	0.50	80	35	39	6	42.5%	5.0%	7.5%	95.0%
10×10	Random	L0	baseline	80	44	36	0	47.5%	10.0%	0.0%	90.0%
10×10	Random	L1	0.50	80	44	36	0	57.5%	10.0%	0.0%	90.0%
10×10	Random	L2	0.50	80	32	48	0	50.0%	20.0%	0.0%	80.0%

Board	Policy	Level	Prob.	n	P1 wins	P2 wins	Draws	Starting-player win rate	Win-rate disparity	Draw rate	Mirrored-start parity
10×10	Random	L3	0.50	80	30	48	2	58.8%	22.5%	2.5%	77.5%
10×10	Checkerboard + follow-up	L0	baseline	80	52	28	0	50.0%	30.0%	0.0%	70.0%
10×10	Checkerboard + follow-up	L1	0.50	80	47	33	0	51.2%	17.5%	0.0%	82.5%
10×10	Checkerboard + follow-up	L2	0.50	80	40	40	0	60.0%	0.0%	0.0%	100.0%
10×10	Checkerboard + follow-up	L3	0.50	80	48	31	1	48.7%	21.2%	1.3%	77.5%
10×10	Belief-state	L0	baseline	80	53	27	0	51.2%	32.5%	0.0%	67.5%
10×10	Belief-state	L1	0.50	80	53	27	0	53.7%	32.5%	0.0%	67.5%
10×10	Belief-state	L2	0.50	80	43	37	0	58.8%	7.5%	0.0%	92.5%
10×10	Belief-state	L3	0.50	80	45	33	2	46.3%	15.0%	2.5%	87.5%
10×10	Greedy max-probability	L0	baseline	80	53	27	0	51.2%	32.5%	0.0%	67.5%
10×10	Greedy max-probability	L1	0.50	80	54	26	0	50.0%	35.0%	0.0%	65.0%
10×10	Greedy max-probability	L2	0.50	80	49	31	0	46.3%	22.5%	0.0%	77.5%
10×10	Greedy max-probability	L3	0.50	80	38	36	6	48.7%	2.5%	7.5%	92.5%

Table S2.10. Terminal-move sensitivity across measurement-probability conditions.

Values are mean terminal turns. L0 is shown as a fixed baseline because it has no circuit-measured transition and is not measurement-probability-dependent.

Board	Policy	Level	L0 baseline	p = 0.25	p = 0.50	p = 0.75
5×5	Random	L1	22.43	46.05	31.74	25.29
5×5	Random	L2	22.43	46.51	31.74	25.66
5×5	Random	L3	22.43	42.67	28.36	23.36
5×5	Checkerboard + follow-up	L1	17.65	39.39	25.12	20.16
5×5	Checkerboard + follow-up	L2	17.65	39.55	25.45	20.68
5×5	Checkerboard + follow-up	L3	17.65	36.46	22.23	18.89
5×5	Belief-state	L1	16.55	37.59	23.73	19.31
5×5	Belief-state	L2	16.55	39.58	23.40	18.93
5×5	Belief-state	L3	16.55	35.21	20.99	16.35
5×5	Greedy max-probability	L1	16.55	38.65	24.40	19.49
5×5	Greedy max-probability	L2	16.55	37.74	24.18	18.57
5×5	Greedy max-probability	L3	16.55	34.69	20.49	16.31
10×10	Random	L1	88.10	120.96	104.42	93.25
10×10	Random	L2	88.10	121.51	104.41	96.62
10×10	Random	L3	88.10	116.04	99.95	88.62
10×10	Checkerboard + follow-up	L1	41.48	67.45	50.11	44.27
10×10	Checkerboard + follow-up	L2	41.48	66.99	52.98	47.86
10×10	Checkerboard + follow-up	L3	41.48	67.55	51.74	48.10
10×10	Belief-state	L1	40.80	63.51	49.48	43.58
10×10	Belief-state	L2	40.80	64.47	49.90	44.11
10×10	Belief-state	L3	40.80	60.54	44.50	39.39
10×10	Greedy max-probability	L1	40.80	66.64	49.76	43.71
10×10	Greedy max-probability	L2	40.80	65.36	48.51	43.14
10×10	Greedy max-probability	L3	40.80	59.39	45.31	40.33

Table S2.11. L3 direct damage, correlated damage, sunk-ship events, and terminal effects under p = 0.50.

Counts and rates are calculated from L3 trace-run trial logs under p = 0.50, combining both start orders, with n = 80 games per board-policy condition. Correlated-damage events count successful linked-damage applications.

Board	Policy	n	Direct target hits	Direct-hit rate	Correlated-damage events	Correlated-damage rate	Sunk-ship events	Terminal correlated count	L3 Δ vs L0 turns
5×5	Random	80	1062	23.6%	240	5.3%	366	62	5.94
5×5	Checkerboard + follow-up	80	1049	29.9%	240	6.8%	369	23	4.58
5×5	Belief-state	80	1096	33.0%	240	7.2%	389	69	4.44
5×5	Greedy max-probability	80	1086	33.5%	240	7.4%	385	67	3.94
10×10	Random	80	1061	6.7%	240	1.5%	364	59	11.85
10×10	Checkerboard + follow-up	80	1034	12.6%	240	2.9%	376	32	10.26
10×10	Belief-state	80	1069	15.1%	240	3.4%	379	67	3.70
10×10	Greedy max-probability	80	1086	15.1%	240	3.3%	387	66	4.51

Table S2.12. Complete L3 decomposition across measurement-probability conditions.

This table reports all L3 decomposition rows from l3_decomposition_summary.csv, including p = 0.25, p = 0.50, and p = 0.75 conditions.

Board	Policy	Prob.	n	Direct target hits	Direct target destructions	Correlated-damage events	Sunk-ship events	Direct-hit rate	Correlated-damage rate	Terminal correlated count
5×5	Random	0.25	80	1024	1024	240	342	15.1%	3.5%	50
5×5	Random	0.50	80	1062	1062	240	366	23.6%	5.3%	62
5×5	Random	0.75	80	1079	1079	240	369	29.2%	6.5%	61
5×5	Checkerboard + follow-up	0.25	80	1023	1023	240	356	17.7%	4.1%	27
5×5	Checkerboard + follow-up	0.50	80	1049	1049	240	369	29.9%	6.8%	23
5×5	Checkerboard + follow-up	0.75	80	1067	1067	240	387	35.8%	8.1%	21
5×5	Belief-state	0.25	80	1071	1071	240	371	19.1%	4.3%	71
5×5	Belief-state	0.50	80	1096	1096	240	389	33.0%	7.2%	69
5×5	Belief-state	0.75	80	1093	1093	240	386	42.5%	9.3%	70
5×5	Greedy max-probability	0.25	80	1068	1068	240	372	19.4%	4.4%	72
5×5	Greedy max-probability	0.50	80	1086	1086	240	385	33.5%	7.4%	67
5×5	Greedy max-probability	0.75	80	1099	1099	240	391	42.8%	9.3%	71
10×10	Random	0.25	80	1041	1041	240	358	5.6%	1.3%	59
10×10	Random	0.50	80	1061	1061	240	364	6.7%	1.5%	59
10×10	Random	0.75	80	1067	1067	240	370	7.6%	1.7%	57
10×10	Checkerboard + follow-up	0.25	80	1045	1045	240	382	9.7%	2.2%	47
10×10	Checkerboard + follow-up	0.50	80	1034	1034	240	376	12.6%	2.9%	32
10×10	Checkerboard + follow-up	0.75	80	1046	1046	240	381	13.7%	3.1%	37
10×10	Belief-state	0.25	80	1087	1087	240	390	11.3%	2.5%	60
10×10	Belief-state	0.50	80	1069	1069	240	379	15.1%	3.4%	67
10×10	Belief-state	0.75	80	1074	1074	240	384	17.2%	3.8%	59
10×10	Greedy max-probability	0.25	80	1087	1087	240	390	11.5%	2.5%	64
10×10	Greedy max-probability	0.50	80	1086	1086	240	387	15.1%	3.3%	66
10×10	Greedy max-probability	0.75	80	1069	1069	240	378	16.7%	3.7%	60

Manifest output-file paths and visualization files

Table S2.13. Manifest output files.

Manifest key	Output path
--------------	-------------

Manifest key	Output path
trial logs main	backend/experiment outputs/final traceable manuscript run 20260703 191713/trial logs main p050.csv
trial logs sensitivity	backend/experiment outputs/final traceable manuscript run 20260703 191713/trial logs sensitivity p025 p075.csv
summary main	backend/experiment outputs/final traceable manuscript run 20260703 191713/summary main p050.csv
summary sensitivity	backend/experiment outputs/final traceable manuscript run 20260703 191713/summary sensitivity.csv
fairness summary	backend/experiment outputs/final traceable manuscript run 20260703 191713/fairness summary.csv
l3 decomposition summary	backend/experiment outputs/final traceable manuscript run 20260703 191713/l3 decomposition summary.csv
belief policy validation	backend/experiment outputs/final traceable manuscript run 20260703 191713/belief policy validation.json
manifest	backend/experiment outputs/final traceable manuscript run 20260703 191713/experiment run manifest.json

Table S2.14. Layout visualization files listed in the manifest.

The visualization file paths are listed for traceability. The image files themselves are not embedded in this supplementary document.

No.	Visualization file path
1	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/paired layout visual 5x5 trial 0001.png
2	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/paired layout comparison 5x5.png
3	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/l2 branch resolution example 5x5 trial 0001.png
4	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/l3 link example 5x5 trial 0001.png
5	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/paired layout visual 10x10 trial 0001.png
6	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/paired layout comparison 10x10.png
7	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/l2 branch resolution example 10x10 trial 0001.png
8	backend/experiment outputs/final traceable manuscript run 20260703 191713/layout visuals/l3 link example 10x10 trial 0001.png