Quantum Computing, Information and Physics in qang Space

A treatise on the qg unit, the symmetry filter and their applications

Vicente Humberto Monteverde

2026-01-01

Table of contents

Preface **17**
How to read the results . . . 17
Code and data . . . 18
Status of this edition . . . 18

I I. Foundations: the angular–probabilistic metric (the qang unit) **19**

1 The qang unit: definition, milli-qang and the probability equation **20**
1.1 Definition of the qang . . . 20
1.2 Theory of the milli-qang . . . 22
1.2.1 Definition and scale . . . 22
1.2.2 Linear in probability, not linear in angle . . . 22
1.2.3 Resolution from shots . . . 23
1.2.4 Gate error budget . . . 23
1.3 The probability equation of the qang and the milli-qang . . . 26
1.3.1 More than two outcomes: the statistical distance . . . 27
1.4 The angular scale: the angular qang (qga) . . . 28
1.4.1 Unverified proposal: energy and antimatter in qga . . . 29
1.5 Why measure in qg: the geometry–probability gap . . . 31
1.6 Inversion and the branch sign . . . 31
1.7 The outcome entropy: qg_S . . . 31
1.8 Invertibility of qg_S and its error bars . . . 32
1.9 The qang angle as the arcsine transformation . . . 32
1.10 Mixed states and multi-qubit registers . . . 33
1.11 Derived quantities of the reading . . . 33
1.12 Scope and limits of the base . . . 36
1.13 Without qang and with qang . . . 38
1.14 In code . . . 40
1.15 Summary . . . 40
Bibliographic notes . . . 41

2 Native gates, closed algebra and the universal operator **42**
2.1 The elementary gate R_{qg} . . . 42
2.2 The universal operator in qg . . . 43

2.3 The gate as a probability equation . . . 43
2.4 Verification against standard libraries . . . 44
2.5 Beyond the qubit: the Cartan subalgebra . . . 45
2.6 The minimal pulse-area principle . . . 46
2.7 An exposure screen before running . . . 48
2.8 One qubit order for every kit . . . 49
2.9 Without qang and with qang . . . 49
2.10 In code . . . 50
2.11 Summary . . . 51
Bibliographic notes . . . 51

3 Reflections, Householder and Grover dynamics 52
3.1 Transition matrices without transcendental functions . . . 52
3.2 The Householder reflector in qg space . . . 53
3.3 The Grover diffuser as a qg reflection . . . 53
3.4 Transition probability between qg states . . . 54
3.5 The Grover trajectory as a Chebyshev polynomial . . . 55
3.5.1 The trajectory as a harmonic oscillation . . . 57
3.6 Analytic limits: coarse angles and overshoot . . . 57
3.6.1 Signal-to-noise and curvature along the trajectory . . . 58
3.6.2 Grover with noise . . . 59
3.7 Amplitude estimation . . . 59
3.8 Without qang and with qang . . . 60
3.9 In code . . . 61
3.10 Summary . . . 62
Bibliographic notes . . . 62

4 Sampling statistics, Bayesian inference and the Haar distribution 63
4.1 Propagating shot noise . . . 63
4.1.1 The azimuthal phase . . . 64
4.2 The Haar distribution in qg space . . . 64
4.3 Exact intervals near the poles . . . 65
4.4 Second-order unbiased estimators . . . 66
4.5 Planning the shots . . . 67
4.6 Few-shot tomography and classical shadows . . . 67
4.7 The foundation: Riesz–Fréchet . . . 68
4.8 The derived quantities in sampling . . . 69
4.9 Without qang and with qang . . . 70
4.10 In code . . . 72
4.11 Summary . . . 72
Bibliographic notes . . . 73

II II. Local geometry, quantum optics and geometric phases 74

5 The local Bloch radius as a measure of purity and entanglement 75
5.1 The probability equation with a radius . . . 75
5.2 The deficit as an entanglement measure . . . 76
5.3 The radius under gates and algorithms . . . 76
5.4 Definite weight: the radius from the Z basis . . . 77
5.5 Direction and radius apart . . . 78
5.6 Wave–particle duality and entropic uncertainty . . . 80
5.7 Estimating the radius from shots . . . 82
5.8 Without qang and with qang . . . 82
5.9 In code . . . 83
5.10 Summary . . . 83
Bibliographic notes . . . 84

6 Optical polarization and Stokes parameters in qg units 85
6.1 Malus's law as a probability equation . . . 85
6.2 The Stokes parameters are qg values . . . 87
6.3 Degree of polarization and purity . . . 88
6.4 Lossless elements: the Mueller matrix is the gate rule . . . 89
6.5 Photon-counting polarimetry . . . 89
6.6 Without qang and with qang . . . 91
6.7 In code . . . 92
6.8 Summary . . . 92
Bibliographic notes . . . 93

7 Geometric phases and Stokes' theorem on the sphere 94
7.1 The Berry connection and curvature . . . 94
7.2 Stokes on the sphere . . . 94
7.3 Loops at constant height: the phase is the probability . . . 95
7.4 Three independent computations . . . 95
7.5 Measuring the phase: the interferometer and the probability equation 96
7.6 Pancharatnam's phase and the octant . . . 97
7.7 Without qang and with qang . . . 98
7.8 In code . . . 98
7.9 Summary . . . 99
Bibliographic notes . . . 99

8 Fisher information, clocks, gravity and metrology 100
8.1 From the base: the information in one shot . . . 100
8.2 Information geometry and the natural gradient . . . 102
8.3 Two clocks and gravity . . . 102
8.4 Ramsey sensing: where to operate . . . 104

8.5 GHZ metrology: when entanglement wins . . . 105
8.6 Thermal states: $qg_Z = \tanh(\beta h)$ . . . 107
8.7 Without qang and with qang . . . 107
8.8 In code . . . 108
8.9 Summary . . . 109
Bibliographic notes . . . 109

III III. Neutrino physics, oscillations and quantum simulation (`qang.neutrino`) 110

9 Quantum kinematics of flavour oscillation in vacuum 111
9.1 From the base: the oscillation probability in m-qg . . . 111
9.2 Oscillation as a rotation on a cone . . . 112
9.3 The mixing angle is a quarter of the qang angle . . . 113
9.4 Error bars that do not depend on the angle . . . 113
9.5 Three flavours on three qubits . . . 114
9.6 Noise suppression with the qg filter . . . 114
9.7 The CP asymmetry in vacuum . . . 116
9.8 What the link between neutrinos and qang is, and is not . . . 118
9.9 Without qang and with qang . . . 118
9.10 In code . . . 119
9.11 Summary . . . 120
Bibliographic notes . . . 120

10 Matter effects, solar neutrinos and collective oscillations 121
10.1 From the base: the medium as a shift in m-qg . . . 121
10.2 The MSW effect in the Earth . . . 122
10.3 The mass ordering read through matter . . . 123
10.4 Solar neutrinos: the whole Sun is one rotation . . . 124
10.5 Collective oscillations: what is conserved is the mass number . . . 126
10.6 Without qang and with qang . . . 127
10.7 In code . . . 128
10.8 Summary . . . 128
Bibliographic notes . . . 129

11 Majorana phases, neutrinoless double beta decay and sterile neutrinos 130
11.1 From the base: what changes no probability changes no reading . . . 130
11.2 Majorana phases are invisible to oscillations . . . 130
11.3 Majorana phases in neutrinoless double beta decay . . . 131
11.4 The sterile neutrino: a fourth qubit in the weight-1 sector . . . 133
11.5 The four quantum numbers in qg terms . . . 134
11.6 Without qang and with qang . . . 135

11.7 In code . . . 137
11.8 Summary . . . 138
Bibliographic notes . . . 138

IV IV. Quantum teleportation in optical networks and entanglement distribution (`qang.qpon`) 139

12 Optical teleportation and the Bell-pair limit 140
12.1 From the base: the fidelity is a probability . . . 140
12.2 Teleportation with a Bell pair . . . 141
12.3 A channel that loses photons . . . 141
12.4 The design rule . . . 141
12.5 The rule on simulated devices . . . 143
12.6 Bell pairs in a network, read in qg . . . 145
12.7 Without qang and with qang . . . 146
12.8 In code . . . 147
12.9 Summary . . . 148
Bibliographic notes . . . 148

13 Dual rail, the sector filter and long-distance fidelity 149
13.1 From the base: three qg readings predict a fidelity . . . 149
13.2 The fidelity from three correlators . . . 149
13.3 Dual rail: every loss is heralded . . . 150
13.4 What dual rail measures on simulated devices . . . 151
13.5 Entanglement swapping along a chain . . . 153
13.6 What remains open . . . 153
13.7 Without qang and with qang . . . 155
13.8 In code . . . 155
13.9 Summary . . . 156
Bibliographic notes . . . 156

V V. Relativity, time loops and black-hole evaporation (`qang.postselection`) 157

14 Postselected closed timelike curves (P-CTC) and the resolution of paradoxes 158
14.1 From the base: the success probability has qang angle φ . . . 158
14.2 Two ways to close a time loop . . . 158
14.3 The success law and the grandfather paradox . . . 159
14.4 The rotation controlled by the system . . . 160
14.5 What a qg reading tells apart between D-CTC and P-CTC . . . 161
14.6 The cost of postselection . . . 163

14.7 What remains open . . . 163
14.8 Without qang and with qang . . . 163
14.9 In code . . . 164
14.10 Summary . . . 165
Bibliographic notes . . . 165

15 Black-hole thermodynamics, quantum chaos and the Yoshida–Kitaev decoder 166
15.1 From the base: success 1/4 and the equator of the fidelity . . . 166
15.2 The information paradox and complementarity . . . 166
15.3 Scrambling: a 3-qubit micro black hole . . . 167
15.4 The Hayden–Preskill protocol and the decoder . . . 167
15.5 The exact law $P \cdot F = 1/4$ . . . 168
15.6 On simulated noise models . . . 168
15.7 Larger black holes . . . 169
15.8 A toy Page curve . . . 171
15.9 Traversable wormholes . . . 172
15.10 What remains open . . . 173
15.11 Without qang and with qang . . . 173
15.12 In code . . . 174
15.13 Summary . . . 175
Bibliographic notes . . . 175

VI VI. Quantum cybersecurity, random-number generation and post-quantum frontiers 176

16 Demarcation: quantum cryptography against post-quantum cryptography (PQC) 177
16.1 From the base: which probabilities exist on each side . . . 177
16.2 The nature of PQC . . . 178
16.3 Decryption failure is counted, not measured on a qubit . . . 178
16.4 The quantum attack vector: Grover . . . 180
16.5 The quantum attack vector: Shor and phase estimation . . . 181
16.6 The demarcation . . . 182
16.7 What remains open . . . 183
16.8 Without qang and with qang . . . 183
16.9 In code . . . 184
16.10 Summary . . . 185
Bibliographic notes . . . 185

17 Quantum key distribution (QKD) and BB84 in qg space 186
17.1 From the base: four readings and the image of the sphere . . . 186

17.2 The monitor: drift or eavesdropper . . . 188
17.3 Finite blocks . . . 190
17.4 Beyond intercept-resend . . . 191
17.5 On IonQ's noisy simulator . . . 193
17.6 What remains open . . . 196
17.7 Without qang and with qang . . . 196
17.8 In code . . . 197
17.9 Summary . . . 198
Bibliographic notes . . . 198

18 Certified quantum random numbers (QRNG) and channel auditing 199
18.1 From the base: output and privacy are different readings . . . 199
18.2 The output-bias analysis fails; the qg reading does not . . . 200
18.3 Near the pole certification is shot-hungry . . . 202
18.4 Readout calibration: temperature or detector error . . . 203
18.5 On IonQ's noisy simulator, with a leak . . . 205
18.6 What remains open . . . 205
18.7 Without qang and with qang . . . 206
18.8 In code . . . 207
18.9 Summary . . . 208
Bibliographic notes . . . 208

19 Quantum batteries and networks 209
19.1 From the base: three distances between readings . . . 209
19.2 Storage with T1 and T2: inversion dies at the equator . . . 210
19.3 Certifying the charge from shots . . . 212
19.4 Registers: charge locked in the correlations . . . 212
19.5 Networks: a Bell pair stored in two memories . . . 214
19.6 What remains open . . . 215
19.7 Without qang and with qang . . . 216
19.8 In code . . . 218
19.9 Summary . . . 218
Bibliographic notes . . . 219

VII VII. Quantum machine learning (QML) and the symmetry filter 220

20 Weight-conserving quantum neural networks (WeightQNN) 221
20.1 From the base: a sector and its readings . . . 221
20.2 Architectures and encodings . . . 223
20.3 Readouts: qg_Z, qg_{ZZ} and two channels . . . 224
20.4 Classical simulability: no quantum advantage . . . 225
20.5 Scaling and multiclass . . . 226

20.6 Training under noise or with the filter . . . 226
20.7 The ecosystem: Qiskit, qiskit-machine-learning and PennyLane . . . 227
20.8 The qg encoding and quantum signal processing . . . 228
20.9 What remains open . . . 228
20.10 Without qang and with qang . . . 228
20.11 In code . . . 230
20.12 Summary . . . 230
Bibliographic notes . . . 230

21 The qg symmetry filter and exactness under T1 relaxation 232
21.1 From the base: decay pushes the readings to the north pole . . . 232
21.2 What it costs: the mean-squared-error rule . . . 233
21.3 Training through the filter . . . 234
21.4 F5: unequal T1 depends only on the ratios . . . 236
21.5 F6: dephasing, with the filter and extrapolation . . . 237
21.6 Gate and readout errors . . . 238
21.7 What remains open . . . 239
21.8 Without qang and with qang . . . 239
21.9 In code . . . 240
21.10 Summary . . . 240
Bibliographic notes . . . 241

22 Echo calibration and the noise bias (pull effect) 242
22.1 From the base: two errors the filter treats differently . . . 242
22.2 In-sector errors and the echo . . . 243
22.3 Does it reach the classifier? . . . 243
22.4 The pull: why the loss without qang varies between seeds . . . 244
22.5 The pull-free readout . . . 246
22.6 Calibrating with the device itself . . . 247
22.7 What remains open . . . 247
22.8 Without qang and with qang . . . 247
22.9 In code . . . 249
22.10 Summary . . . 249
Bibliographic notes . . . 249

23 Information geometry, the natural gradient (QNG) and kernels 250
23.1 From the base: the angle and the readout . . . 250
23.2 The singularity and its regularizations . . . 250
23.3 The angle is the Fisher-flat coordinate . . . 251
23.4 The natural gradient and the qg_YY metric . . . 252
23.5 Pole-damped descent . . . 254
23.6 Quantum kernels and the qg filter . . . 255
23.7 Barren plateaus: global cost against the qg local cost . . . 256

23.8 Three plateaus and what qang does with each . . . 257
23.9 What remains open . . . 261
23.10 Without qang and with qang . . . 261
23.11 In code . . . 263
23.12 Summary . . . 263
Bibliographic notes . . . 264

VIIIVIII. Many-body dynamics, open channels, matter and chemistry 265

24 Hamiltonian simulation, quantum walks and VarQTE 266
24.1 From the base: conserving N means having a filter . . . 266
24.2 Trotter and the quantum walk on the XY chain . . . 266
24.3 Several excitations: the XXZ chain . . . 267
24.4 Channels as affine maps of the readouts . . . 269
24.5 Variational evolution at fixed depth . . . 269
24.6 Differential equations . . . 270
24.7 Hubbard and a gauge theory: symmetry filters on every observable 271
24.8 Coherence in the transverse-field Ising model . . . 271
24.9 Boltzmann machines: the I–η plane . . . 272
24.10 What remains open . . . 272
24.11 Without qang and with qang . . . 273
24.12 In code . . . 274
24.13 Summary . . . 274
Bibliographic notes . . . 275

25 Quantum spectroscopy, PCA and molecular chemistry 276
25.1 From the base: counting electrons with the register readout . . . 276
25.2 PCA in the weight-1 sector . . . 277
25.3 H_2: the witness and the filter . . . 277
25.4 Why the filter helps H_2 and not LiH . . . 278
25.5 Checks that reach every group . . . 279
25.6 Overlaps: inversion test against swap test . . . 280
25.7 Constrained optimization: QAOA with “exactly K” . . . 280
25.8 What remains open . . . 281
25.9 Without qang and with qang . . . 281
25.10 In code . . . 283
25.11 Summary . . . 283
Bibliographic notes . . . 284

IX IX. Error correction, continuous syndromes and characterization 285

26 Error correction and continuous syndromes 286
26.1 From the base: the syndrome is a qg readout . . . 286
26.2 Which code to use under T1 and dephasing . . . 286
26.3 The syndrome as a continuous witness . . . 288
26.4 The surface code and the T1-aware decoder . . . 288
26.5 Qutrits and erasure qubits . . . 289
26.6 Cutting circuits: where and at what cost . . . 290
26.7 Randomized benchmarking, ZNE and qg_S as a signal . . . 290
26.8 What remains open . . . 291
26.9 Without qang and with qang . . . 291
26.10 In code . . . 293
26.11 Summary . . . 293
Bibliographic notes . . . 294

27 Hardware characterization in qg 295
27.1 From the base: an idle qubit is a readout . . . 295
27.2 T1, T2, temperature and readout in one sweep . . . 295
27.3 Which noise is it? The mean readout against XEB and HOP . . . 297
27.4 Coherent errors and the always-on ZZ . . . 298
27.5 Control quantization: an even grid in the angle or in qg . . . 298
27.6 What remains open . . . 299
27.7 Without qang and with qang . . . 299
27.8 In code . . . 301
27.9 Summary . . . 301
Bibliographic notes . . . 301

X X. Experimental evidence, real hardware and the falsification audit 302

28 Methodology: the pre-registered falsification protocol 303
28.1 From the base: a prediction is a readout that can be wrong . . . 303
28.2 Two ways of pinning a prediction . . . 303
28.3 The tally of hits and misses . . . 304
28.4 Audit of the “before” . . . 304
28.5 Testbeds . . . 305
28.6 Real hardware . . . 306
28.7 Without qang and with qang . . . 307
28.8 In code . . . 307
28.9 Summary . . . 308
Bibliographic notes . . . 308

29 The confirmed empirical advantages **309**
29.1 From the base: where the advantages come from . . . 309
29.2 Exact restoration under T1 . . . 309
29.3 Lower readout error on simulated devices . . . 310
29.4 In-sector errors: the echo . . . 311
29.5 Characterization and control . . . 311
29.6 Neutrinos: mitigation without extra circuits . . . 313
29.7 BB84: fewer false alarms under thermal drift . . . 313
29.8 What remains open . . . 314
29.9 Without qang and with qang . . . 314
29.10 In code . . . 316
29.11 Summary . . . 316
Bibliographic notes . . . 316

30 Structural limits, equivalences and negative results **317**
30.1 From the base: what the scale cannot see . . . 317
30.2 Graph states: entanglement invisible in Z . . . 317
30.3 qg_S is not monotonic under T1 . . . 318
30.4 The natural-gradient equivalence . . . 318
30.5 The filter does not reach deep circuits . . . 319
30.6 The weight-1 network is classical . . . 320
30.7 Where the filter does not apply: phases and classical encryption . . . 321
30.8 The failed predictions of §157–§164 . . . 321
30.9 Consolidated list of limitations . . . 322
30.10 Without qang and with qang . . . 323
30.11 In code . . . 324
30.12 Summary . . . 325
Bibliographic notes . . . 325

31 First real-hardware experiment: IQM Resonance **326**
31.1 From the base: what goes to the chip . . . 326
31.2 Why IQM, and what compiling costs . . . 326
31.3 The `qang.iqm` module . . . 327
31.4 A free rehearsal before spending . . . 328
31.5 The budget: planned and spent . . . 328
31.6 The main experiment: three-flavour neutrinos . . . 329
31.7 What came next, and what remains . . . 330
31.8 Result: hits, misses and comparison with simulation . . . 331
31.9 Without qang and with qang . . . 331
31.10 In code . . . 332
31.11 Summary . . . 332
Bibliographic notes . . . 333

32 Software architecture, the multi-SDK ecosystem and physical execution **334**
32.1 From the base: one count, many paths . . . 334
32.2 One qubit order . . . 334
32.3 V2 primitives and filtered observables . . . 335
32.4 Sector-aware compilation . . . 335
32.5 Static ZZ: simulated, not mitigated . . . 336
32.6 ZNE with the filter, and shot planning . . . 337
32.7 Guidelines by hardware type . . . 338
32.8 Without qang and with qang . . . 339
32.9 In code . . . 340
32.10 Summary . . . 340
Bibliographic notes . . . 340

XI XI. Nanoquantum **341**

33 Nanoquantum: transmission, partition noise and nanoscale qubits **342**
33.1 From the base: one probability per experiment . . . 342
33.2 Tunnelling transmission as a qang probability . . . 343
33.3 The Landauer formula: one channel, one reading . . . 343
33.4 Partition noise: the binomial variance, stabilized . . . 344
33.5 Coulomb blockade: the occupation of a level . . . 345
33.6 Charge and spin qubits in quantum dots . . . 347
33.7 NV centres: contrast as an m-qg budget . . . 347
33.8 Without qang and with qang . . . 348
33.9 In code . . . 349
33.10 Summary . . . 349
Bibliographic notes . . . 349

XII XII. Cross-domain applications **351**

34 Angular architecture: chemistry, condensed matter, relativistic optics, thermodynamics and geometry **352**
34.1 From the base: readings, not cycles . . . 352
34.2 Quantum chemistry . . . 353
34.2.1 Jordan–Wigner: the exchange sign is a reading . . . 353
34.2.2 UCC for H_2: the amplitude is half a qang angle . . . 353
34.2.3 Larmor and the Hahn echo . . . 354
34.3 Condensed matter . . . 355
34.3.1 Transverse-field Ising: the peak moves with size . . . 355

34.3.2 Kitaev chain: the protection is in the Hamiltonian, not in the angle . . . 355
34.3.3 Aharonov–Bohm: one fringe per flux quantum . . . 356
34.4 Relativity and optics . . . 357
34.4.1 Wigner rotation . . . 357
34.4.2 Fresnel: Brewster is a zero of probability, total reflection is a phase . . . 358
34.4.3 Bogoliubov and Unruh: the half angle is the Gudermannian . . . 358
34.5 Thermodynamics and metrology . . . 359
34.5.1 Jarzynski with two measurements . . . 359
34.5.2 Fisher information of a GHZ state . . . 359
34.6 Geometry and chaos . . . 360
34.6.1 Fubini–Study: pole to pole, half the qang angle . . . 360
34.6.2 OTOC: integrable against chaotic . . . 360
34.7 The test matrix, revised . . . 360
34.8 Without qang and with qang . . . 362
34.9 In code . . . 363
34.10 Summary . . . 364
Bibliographic notes . . . 364

XIIIXIII. Ecosystem: qang inside other libraries 365

35 qang as a module: changes to quantum and machine-learning libraries 366
35.1 From the base: what each library gets . . . 366
35.2 How it is built . . . 367
35.3 Mitigation: mitiq . . . 367
35.4 Calibration: qiskit-experiments . . . 368
35.5 Open systems: QuTiP . . . 369
35.6 Error correction: Stim . . . 369
35.7 Learning: scikit-learn, JAX and PyTorch . . . 370
35.8 Pulses: Qiskit Dynamics . . . 370
35.9 SDKs: Qibo, pytket, Braket and CUDA-Q . . . 370
35.10 What to propose to each project . . . 371
35.11 Environment conflicts . . . 372
35.12 Without qang and with qang . . . 373
35.13 In code . . . 374
35.14 Summary . . . 374
Bibliographic notes . . . 374

References 376

Appendices **396**

A Catalogue of the 15 fundamental gates in qg representation **396**
A.1 Fixed single-qubit gates 396
A.2 Parametric gates 397
A.3 Multi-qubit gates 398

B Transition formulas, Gell-Mann matrices and Chebyshev polynomials **400**
B.1 Single-qubit formulas 400
B.2 Diagonal Gell-Mann matrices 401
B.3 Chebyshev polynomials and the Grover trajectory 402
Bibliographic notes 403

C The echo-calibration protocol and inversion in Hamming sectors (§105–§108) **404**
C.1 The protocol 404
C.2 Cost 405
C.3 An exact example 405
C.4 The four pre-registered runs 406
Bibliographic notes 407

D Quick reference of the qang 0.14.0 API **408**
D.1 Public names by module 412
D.2 `qang.integrations` 415

E Falsification register: the predictions the repository allows one to check **416**
E.1 Falsification statement 416
E.2 Verdicts stored by the code 416
E.3 Pre-registrations by sha256 418
E.4 Pre-registrations by commit 419
E.5 What this register leaves out 421
Bibliographic notes 421

F Functional-analysis foundation: Riesz–Fréchet **422**
F.1 The setting 422
F.2 The theorem 422
F.3 Application: qg_Z has a representative 423
F.4 What has no representative: qg_S 423
F.5 On Riesz's lemma (1918) 424
Bibliographic notes 424

G Relativistic energy and antimatter as qang readings **425**
G.1 The relativistic angle 425

G.2 $E = mc^2$ as a reading: the expectation value of β 426
G.3 Antimatter: the negative branch and charge conjugation 427
G.4 The probabilities of the source document 428
G.5 Measuring the relativistic angle with shots 428
G.6 The reading of β on a real chip (§170) 428
G.7 What is claimed and what is not 430

H Catalogue of qang sub-measures **432**
H.1 Base readings 433
H.2 Derived quantities 434
H.3 States and registers 436
H.4 Gates and compilation 439
H.5 Noise and readout 441
H.6 Readings of known physical laws 444
H.7 Open and rejected proposals 447
H.8 Balance 449
H.9 In code 449

Preface

This book gathers, in reading order, the results of a research programme built on a simple idea: describing the qubit by what is measured. The qang unit (qg) expresses a qubit's state on the scale on which the experiment delivers it, the polar bias $qg_Z = \langle \sigma_z \rangle$, and that choice yields a practical tool, the **qg symmetry filter**: when a circuit conserves a number (excitations, photons, neutrino flavour, mass), the shots that violate it can only come from noise and are discarded.

The book goes from the foundations (Parts I and II) to the applications: neutrinos (Part III), teleportation in optical networks (Part IV), information-theoretic models of time loops and black holes (Part V), cryptography (Part VI), quantum machine learning (Part VII), dynamics and chemistry (Part VIII), error correction and characterization (Part IX). Part X gathers the evidence: the method, the confirmed advantages, the limits and the negative results, and the first run on real hardware.

How to read the results

Three conventions run through the book:

- **Predictions recorded before running.** Every experiment cited fixed its predictions before it was run (a commit or a sha256 hash of the script), and is reported as passed or failed. Failures are not left out: many of the book's lessons come from them.
- **Simulated noise unless stated.** Whatever is described "on a device" refers to calibrated noise models of IBM or IonQ processors, except for the real-hardware results, which are marked explicitly (Chapter 31).
- **References to the notes.** Citations such as (§151) point to the corresponding section of the project's research notes (`RESEARCH_NOTES.md`), where the code, the tests and the data are.

Code and data

All the material is open: the `qang` library (`pip install qang`), the script of every experiment, the runnable notebooks and the data are at github.com/Viny2030/qang. The software record has the DOI 10.5281/zenodo.22832150.

Status of this edition

The book is written and published chapter by chapter. Chapters marked "in preparation" show their planned contents and are completed in successive editions.

Vicente Humberto Monteverde — Universidad del Museo Social Argentino (UMSA), Argentina. ORCID 0000-0001-8884-4811.

Part I

I. Foundations: the angular–probabilistic metric (the qang unit)

1 The qang unit: definition, milli-qang and the probability equation

A qubit is designed with angles and observed with frequencies. Whoever writes a circuit chooses rotations $R_y(\theta)$; whoever runs it receives counts of zeros and ones. Between the two descriptions sits a trigonometric function that is rarely named, and that is the source of several practical difficulties: gradients that diverge, error bars that shrink or swell with the state, and parameters that cannot be read directly from the experiment. This chapter builds the foundation of the book in three steps: the **definition** of the qang unit (qg), the **theory of the milli-qang** (its submultiple) and the **probability equation** that ties both to what is measured. Every result is backed by a figure computed with the package itself (`qang.probability`, §167), so the figures and the text quote the same numbers.

1.1 Definition of the qang

Every pure state of a qubit can be written, up to a global phase, as

$$|\psi\rangle = \cos\frac{\theta}{2}\,|0\rangle + e^{i\varphi}\sin\frac{\theta}{2}\,|1\rangle, \qquad \theta \in [0,\pi],\ \varphi \in [0,2\pi).$$

The pair (θ,φ) is a point on the Bloch sphere and is the natural way to *prepare* the state: a rotation $R_y(\theta)$ followed by a phase $R_z(\varphi)$. A measurement in the computational basis does not return θ: it returns, on each shot, the outcome 0 with probability

$$p_0 = \cos^2\frac{\theta}{2} \qquad \text{(Born's rule).}$$

Definition. The qang unit takes as its primary coordinate the expectation value of σ_z, called the *polar bias*:

$$\boxed{qg_Z = \langle\sigma_z\rangle = p_0 - p_1 = 2p_0 - 1 = \cos\theta} \qquad qg_Z \in [-1,1].$$

The other two readings of the unit come from the same probability: the **qang angle** $\theta = \arccos(qg_Z) \in [0, \pi]$ and the **superposition spread** $qg_S = H(p_0) \in [0, 1]$, the Shannon entropy of the outcome (studied below). Figure 1.1 summarizes the definition: the point on the circle is the state, its projection on the vertical axis is qg_Z, and the scale on the right gives the same reading in milli-qang and in probability.

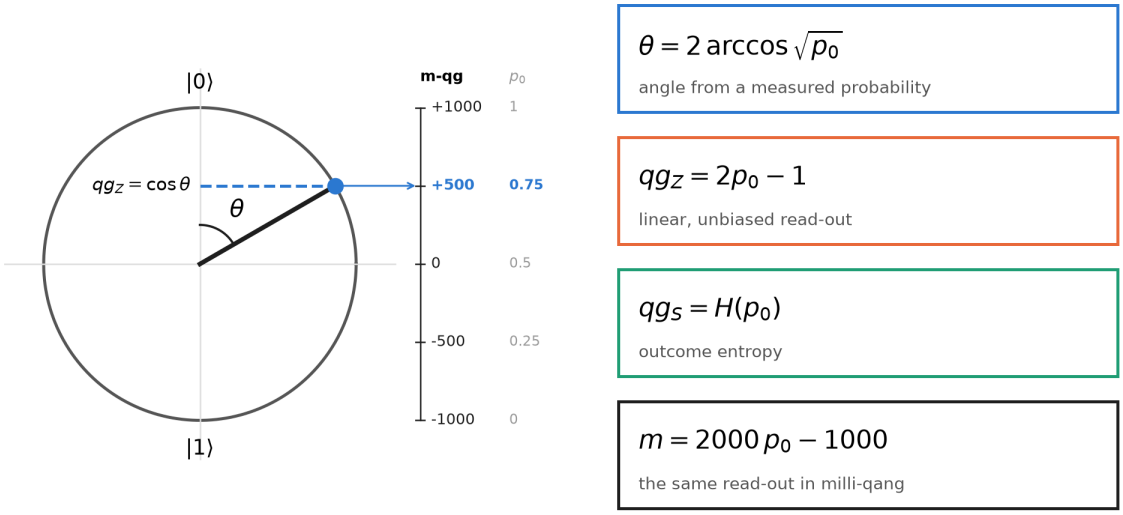


Figure 1.1: The qang circle. The state at angle $\theta = 60^\circ$ has $p_0 = 0.75$, $qg_Z = \cos\theta = 0.5$ and a reading of $+500$ m-qg. The four equations on the right are the same probability read in four ways.

Three properties make qg_Z a good working coordinate:

1. **It is a bijection.** On $\theta \in [0, \pi]$, $\cos\theta$ is strictly decreasing and covers $[-1, 1]$: each value of qg_Z corresponds to a single latitude of the sphere.
2. **It is directly measurable.** $\widehat{qg}_Z = (n_0 - n_1)/N$ is the average of the ± 1 outcomes of the shots; it needs no inversion and no fit.
3. **It holds for any state.** For a mixed state, $qg_Z = \mathrm{Tr}(\rho\, \sigma_z)$, which reduces to $\cos\theta$ when $\rho = |\psi\rangle\langle\psi|$ (§3).

The sign and range have an immediate physical reading: $qg_Z = +1$ is the pole $|0\rangle$, $qg_Z = -1$ the pole $|1\rangle$, and $qg_Z = 0$ any point of the equator, where the measurement outcome is maximally uncertain.

An engineering unit. A useful unit does not need to be a new physical quantity; it needs to make what is measured legible. The decibel adds no physics to power: it expresses it on a scale that makes the gains of a chain of amplifiers additive. The qang plays an analogous role for the qubit: it expresses its state

on the scale in which the experiment delivers it, so that preparation, noise and read-out are compared in the same units. It is an operational convention, and like every convention it is worth what it simplifies.

1.2 Theory of the milli-qang

1.2.1 Definition and scale

The deviations that matter in a real qubit are small: calibration errors, drift, read-out noise. A submultiple is convenient to express them, the **milli-qang** (m-qg):

$$1 \text{ qg} = 1000 \text{ m-qg}, \qquad m = 1000\, qg_Z \ \in [-1000,\, +1000] \text{ m-qg}.$$

The scale runs from -1000 m-qg at the pole $|1\rangle$ to $+1000$ m-qg at the pole $|0\rangle$, with zero on the equator. A state with $qg_Z = 0.5018$ sits at 501.8 m-qg; its difference from the ideal 0.5 is 1.8 m-qg. The same definition holds for the other axes, qg_X and qg_Y, so the whole Bloch vector can be expressed in m-qg.

1.2.2 Linear in probability, not linear in angle

Since $qg_Z = 2p_0 - 1$, the milli-qang is **linear in probability**: one m-qg is always $\Delta p_0 = 0.0005$, that is, 0.05 percentage points, anywhere on the scale. It is **not linear in angle**:

$$\frac{dm}{d\theta} = -1000 \, \sin\theta \qquad \Longrightarrow \qquad \Delta\theta \approx \frac{1}{1000 \, \sin\theta} \text{ rad per m-qg}.$$

On the equator ($\theta = 90^\circ$) one m-qg is exactly one milliradian; at $\theta = 30^\circ$ it is two. Near a pole the linear relation fails and the exact one must be used: a state δ m-qg below the pole $|0\rangle$ makes with it the angle

$$\theta_\delta = \arccos\!\left(1 - \frac{\delta}{1000}\right) \approx \sqrt{\frac{2\delta}{1000}}.$$

One m-qg from the pole is 0.0447 rad (2.56°); ten m-qg, 8.1°. Figure 1.2 shows the consequence: on the milli-qang ruler the probability marks are evenly spaced and the angle marks crowd at the poles. That is why the unit is defined on probability and not on angle: it is the scale in which the instrument delivers the data.

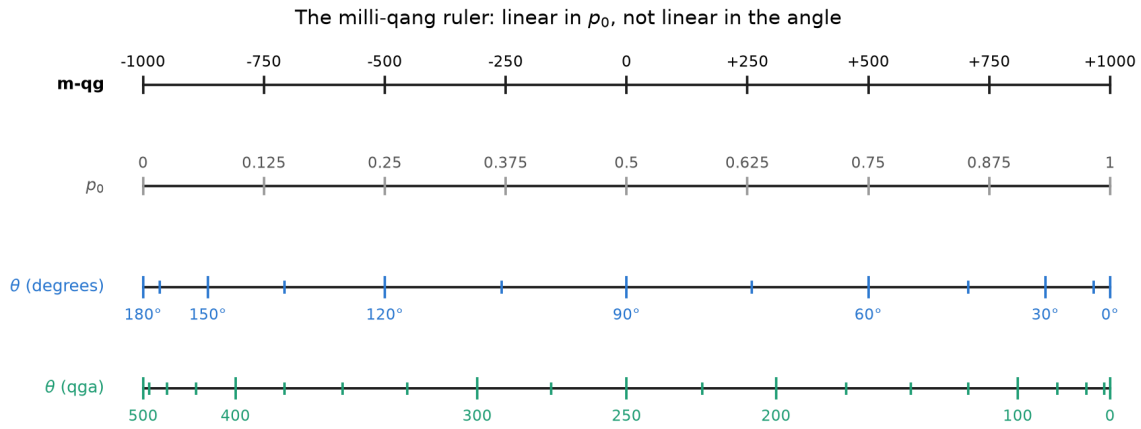


Figure 1.2: The milli-qang ruler. Top, the scale in m-qg; middle, the probability p_0, aligned and linear; bottom, the qang angle in degrees (marks every $15°$) and in qga, the angular scale of Section 1.4 (marks every 25 qga); both are compressed towards the poles.

1.2.3 Resolution from shots

With N shots, $\widehat{qg}_Z$ has variance $\sin^2\theta/N$ (derived in the arcsine section), so the standard deviation of the reading in milli-qang is

$$\sigma_m = \frac{1000\,\sin\theta}{\sqrt{N}}\ \text{m-qg}, \qquad N_\delta = \left(\frac{1000\,\sin\theta}{\delta}\right)^2$$

where N_δ is the number of shots needed to resolve δ m-qg at one sigma. On the equator, resolving 10 m-qg takes 10^4 shots and resolving 1 m-qg, 10^6. At $60°$ and 1000 shots, $\sigma_m = 27.4$ m-qg; to get down to 5 m-qg at that point takes 30,000 shots. Figure 1.3 compares σ_m with the uncertainty of the angle, $\sigma_\theta = 1/\sqrt{N}$: the m-qg reading is more precise near the poles, and the angle's is the same all around the circle. The two curves touch on the equator, where one m-qg is one milliradian.

1.2.4 Gate error budget

The scale lets a gate error be expressed in the same units as the read-out. If the amplitude of a pulse has a random relative error of standard deviation σ, the expected polar bias after $R_x(\theta)$ is

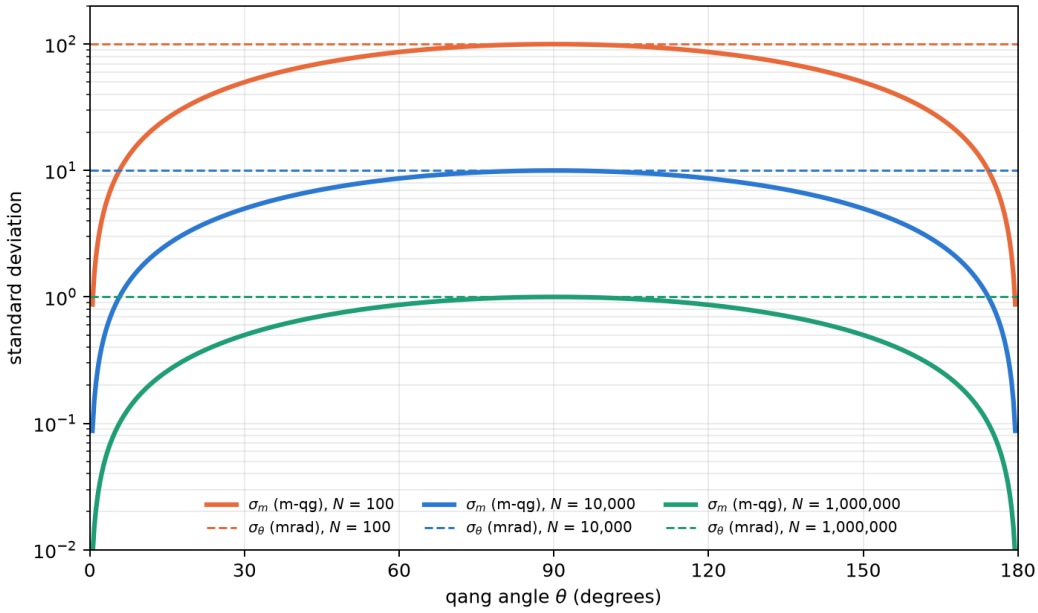


Figure 1.3: Resolution from shots. Solid lines: standard deviation σ_m of the m-qg reading; dashed lines: σ_θ of the qang angle in milliradians, which does not depend on the state. For each N, the two coincide at $\theta = 90°$.

$$\mathbb{E}[qg_Z] = \cos\theta\, e^{-\theta^2\sigma^2/2} \qquad \Longrightarrow \qquad \varepsilon_m = 1000\cos\theta\left(e^{-\theta^2\sigma^2/2} - 1\right) \text{ m-qg}$$

(§160). The error pulls the reading towards the equator: with $\sigma = 0.05$, an X gate ($\theta = \pi$) delivers $qg_Z = -0.988$ instead of -1, a shift of 12.3 m-qg. At multiples of π the shift grows with the square of the angle, always towards the equator; in general it goes as $\theta^2 \cos\theta$ (Figure 1.4). A $\pi/2$ rotation does not shift at all, because it already sits on the equator. Chapter 2 uses this formula to choose, among equivalent representations of the same gate, the one with the smallest pulse area.

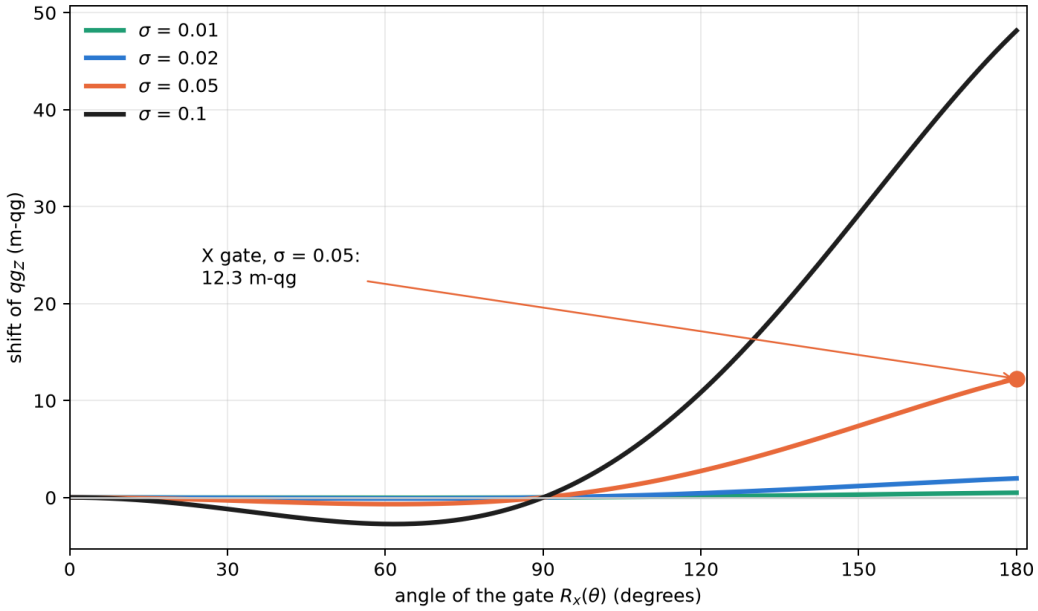


Figure 1.4: Error budget in milli-qang: shift of the polar bias after $R_x(\theta)$ for several relative amplitude errors σ.

Chapters 2 and 3 use two derivatives of the reading with respect to the angle, which are worth naming here. The **polar gradient** and the **curvature** (or *pulse impedance*) are

$$\nabla_{qg} = \frac{dm}{d\theta} = -1000\,\sin\theta \text{ m-qg/rad}, \qquad \kappa_{qg} = \left|\frac{d^2m}{d\theta^2}\right| = 1000\,|\cos\theta| \text{ m-qg/rad}^2.$$

They are not new readings: they are derivatives of $m = 1000\cos\theta$. The curvature explains the budget above: to second order in the error, $|\varepsilon_m| \approx \frac{1}{2}\theta^2\sigma^2\,\kappa_{qg}$, so a

random amplitude error shifts the reading in proportion to κ_{qg}, which is 1000 at the poles and 0 at the equator.

1.3 The probability equation of the qang and the milli-qang

The previous sections lead to a single statement, which is the foundation of the whole book: **every qang reading of a qubit is an exact function of a single probability**, p_0:

$$\theta = 2\arccos\sqrt{p_0}, \qquad qg_Z = 2p_0 - 1, \qquad qg_S = H(p_0), \qquad m = 2000\,p_0 - 1000 \text{ m-qg}$$

and, in the other direction, the probability equation of the milli-qang:

$$p_0 = \frac{1000 + m}{2000}, \qquad \theta = \arccos\frac{m}{1000}.$$

There are no approximations: they are four readings of one number. Figure 1.5 shows them together. qg_Z and the milli-qang are the same straight line on two scales; the angle is monotonic but steepens at the poles; the entropy rises up to the equator and falls again, which is why it needs a branch to be inverted. Table 1.1 gives the reference values.

Table 1.1: Anchor points of the probability equation.

θ	p_0	qg_Z	m-qg	state
0	1	$+1$	$+1000$	pole $\lvert 0\rangle$
$\pi/4$	0.854	$+0.707$	$+707$	
$\pi/3$	0.75	$+0.5$	$+500$	
$\pi/2$	0.5	0	0	equator ($\lvert +\rangle$, $\lvert +i\rangle$, ...)
$2\pi/3$	0.25	-0.5	-500	
π	0	-1	-1000	pole $\lvert 1\rangle$

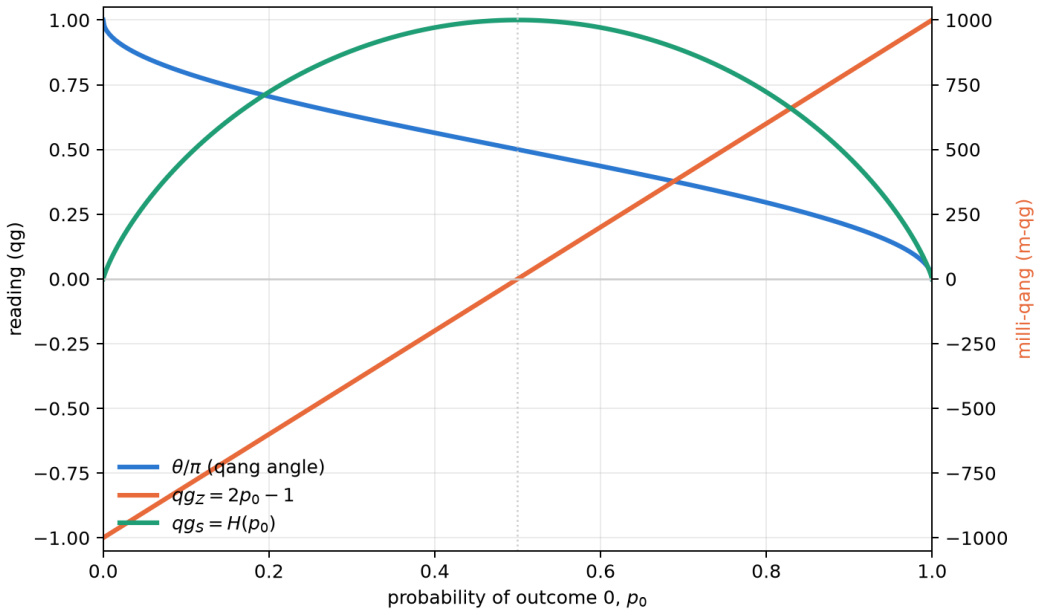


Figure 1.5: The readings of one probability. The right axis expresses qg_Z in milli-qang.

1.3.1 More than two outcomes: the statistical distance

The equation generalizes to a measurement with several outcomes. For two distributions $p = (p_k)$ and $q = (q_k)$, **Wootters' statistical distance** (1981), or Bhattacharyya angle (Bhattacharyya 1943), is

$$d(p, q) = \arccos \sum_k \sqrt{p_k\, q_k}.$$

For two outcomes, the distance between $(p_0, 1 - p_0)$ and the pole $(1, 0)$ is $\arccos \sqrt{p_0} = \theta/2$: up to the factor 2, the qang angle is the statistical distance between what is measured and the reference state. With three outcomes, for example, $p = (0.5, 0.3, 0.2)$ and $q = (0.4, 0.4, 0.2)$ are 0.113 rad apart. This is the door through which the unit leaves the qubit: Part II uses it for the Fisher information and the Fubini–Study metric, and the same geometry holds for any measured proportion, quantum or not.

1.4 The angular scale: the angular qang (qga)

The *Geometric and Dimensional Foundation of the Qang System* (Monteverde 2026a) proposes measuring angles on a thousand-based scale: one full turn is 1000 units. The book adopts that scale under another name, because throughout the book, the library and the chapters that follow, **m-qg names the reading** $m = 1000\,qg_Z$, which is linear in the probability and is not an angle (previous section). The angular scale is written **qga** (angular qang) and its thousandth **m-qga**:

$$1\,\text{turn} = 2\pi\,\text{rad} = 1000\,\text{qga} = 10^6\,\text{m-qga}, \qquad \theta[\text{rad}] = \frac{2\pi}{1000}\,\theta_{qga}, \qquad 1\,\text{m-qga} \approx 6.283{\cdot}10^{-6}\,\text{rad}.$$

With the Bloch polar angle in qga, the Born rule and the reading become

$$p_0 = \cos^2\frac{\pi\,\theta_{qga}}{1000} = \frac{1+\cos(2\pi\theta_{qga}/1000)}{2}, \qquad m = 1000\cos\frac{2\pi\,\theta_{qga}}{1000}\ \text{m-qg}.$$

The pole $|0\rangle$ is at 0 qga, the equator at 250 qga and the pole $|1\rangle$ at 500 qga; the phase φ runs over $[0, 1000)$ qga. The two scales complement each other: qga is uniform in the angle and m-qg is uniform in the probability. The state of Figure 1.1, for instance, sits at 166.667 qga and reads $+500$ m-qg; at the equator, one m-qg of reading equals one milliradian, that is, 0.159 qga (159 m-qga).

In code, since version 0.14.0, `qang.qga` does these conversions: `p0_from_qga(250)` gives 0.5, `milliqang_from_qga(1000/6)` gives +500 m-qg and `qga_from_milliqang(-1000)` gives 500 qga.

The azimuthal coordinate. The phase is read from two transverse measurements, $\varphi = \text{atan2}(qg_Y, qg_X)$, because $qg_X = \sin\theta\cos\varphi$ and $qg_Y = \sin\theta\sin\varphi$ and the ratio cancels $\sin\theta$. As a fraction of a turn, $qg_\varphi = (\varphi \bmod 2\pi)/2\pi \in [0, 1)$ is exactly the azimuthal angle in qga divided by 1000. It is undefined at the poles, where $qg_X = qg_Y = 0$.

Geometry. The Fubini–Study metric of a qubit and the Fisher information of the Z measurement carry a constant factor in qga:

$$ds^2 = \tfrac{1}{4}\left(d\theta^2 + \sin^2\theta\,d\varphi^2\right) = \frac{\pi^2}{10^6}\left(d\theta_{qga}^2 + \sin^2\theta\,d\varphi_{qga}^2\right), \qquad I_F(\theta) = 1\,\text{rad}^{-2} = 4\pi^2{\cdot}10^{-6}\,\text{qga}^{-2}$$

The second is the stabilization of the arcsine section in another unit: the standard deviation of the qang angle is $159.15/\sqrt{N}$ qga over the whole circle (5.03 qga with 1000 shots). The Berry curvature, $F = \frac{1}{2}\sin\theta\, d\theta \wedge d\varphi$, integrated over the sphere (θ spans 500 qga and φ 1000) gives 2π, that is, Chern number 1.

Corrections to the Foundation. Three claims of the source document are not adopted, because its own formulas contradict them:

- It assigns $p_0 = 0.5$ to 125 qga. Its own equation gives $p_0(125 \text{ qga}) = \cos^2(\pi/8) = 0.854$ (+707 m-qg); $p_0 = 0.5$ belongs to 250 qga.
- It gives curvature $K = 1$ to the metric $ds^2 = \frac{1}{4}\,d\theta^2$. That metric, with its azimuthal term, is the one of a sphere of radius $1/2$, whose Gaussian curvature is $K = 4$; $K = 1$ belongs to the Bloch sphere of radius 1.
- It suggests that a lattice in m-qga avoids barren plateaus. Changing the unit of the parameter multiplies every derivative by the constant $2\pi/1000$: it changes neither the ratio of gradient variance to its scale nor its decay with the number of qubits. The unit has no effect on plateaus. The qang tools are another matter: they do mitigate them in tested regimes, namely the qg local cost (Section 23.7) and the weight filter under T1 (Section 20.6).

1.4.1 Unverified proposal: energy and antimatter in qga

The Foundation extends the scale to relativistic kinematics. What follows is an **unverified proposal**: the identities are corrected, but the physical reading is unproven and the rest of the book does not use it.

Relativistic angle. Dividing $E^2 = (mc^2)^2 + (pc)^2$ by E^2 gives $\cos\theta_r = mc^2/E = 1/\gamma$ and $\sin\theta_r = pc/E = \beta$. The identity is exact, because $1/\gamma^2 + \beta^2 = 1$. With θ_r in qga, $E = mc^2/\cos(2\pi\theta_{r,qga}/1000)$: rest is at 0 qga and the ultra-relativistic limit ($E = pc$) at 250 qga. The angle θ_r is tied to the rapidity η by $\cos\theta_r = 1/\cosh\eta$ (the Gudermann function), with no Wick rotation needed.

Antimatter. The Dirac equation has two branches, $E_\pm = \pm\sqrt{(pc)^2 + (mc^2)^2} = \pm\, mc^2/\cos\theta_r$. The Foundation assigns the negative branch a shift of 500 qga, tied to CPT symmetry; that identification has no support and is not adopted.

Probabilities. The Foundation equates $P_{\text{rest}} = (mc^2/E)^2$ with $p_0(\theta_{qga})$. There is a factor 2 in the angle: $(mc^2/E)^2 = \cos^2\theta_r$, whereas $p_0 = \cos^2(\theta/2)$, so the equality holds only if the Bloch angle is $2\theta_r$. If instead the Bloch angle is taken equal to θ_r, there is a known result behind the document's formulas $\cos^2(\pi\,\text{qg}/1000)$ and $\sin^2(\pi\,\text{qg}/1000)$: in the free Dirac spinor the large component weighs $(E + mc^2)/2E = \cos^2(\theta_r/2)$ and the small one $\sin^2(\theta_r/2)$. But

those weights are not probabilities of "matter" and "antimatter", as the source document calls them: both belong to the same positive-energy solution.

Appendix G (Appendix G) develops this proposal. It finds an exact reading of $E = mc^2$, the expectation value of the Dirac matrix β, which equals mc^2/E and is therefore a qg_Z. It also shows that charge conjugation inverts that reading.

Proposed names. Table 1.2 lists the names the Foundation introduces and how the book treats them.

Table 1.2: The Foundation's names for energy, antimatter and the derived units, and how the book treats them.

Foundation name	definition	treatment in the book
relativistic angle θ_r	$\cos\theta_r = mc^2/E$, $\sin\theta_r = \beta$	exact identity; in qga, a proposal
$E(\text{qg}) = mc^2/\cos(2\pi\,\text{qg}/1000)$	total energy as a function of θ_r	exact, with θ_r in qga
$E_\pm$, matter and antimatter sheets	$\pm mc^2/\cos\theta_r$	Dirac branches; the 500 qga CPT shift is not adopted
P_{rest}, P_{kinetic}	$(mc^2/E)^2$, β^2	they add to 1; they are not p_0, p_1 of the same angle (factor 2)
P_{matter}, $P_{\text{antimatter}}$	$\cos^2$, $\sin^2$ of the half angle	weights of the spinor's large and small components, not antimatter probabilities
Q-Hertz (qHz)	qga per second	unit of angular frequency; a proposal
Q-Fisher (qFi)	Fisher information per qga^2	$4\pi^2 \cdot 10^{-6}$ qga^{-2} for a pure qubit
Q-Shannon (qg_S)	Shannon entropy of the outcome	it is qg_S; it is not the von Neumann entropy
Q-Curvature, Q-SNR, Q-Geodesic, dissipative deficit	qg_κ, qg_{SNR}, qg_L, qg_τ	defined and checked in Section 1.11
Q-Force, Q-Conductance, Q-Momentum	phase gradient, qga per volt, $(\pi/500)\hbar$	proposed names, not used in the book

1.5 Why measure in qg: the geometry–probability gap

The relation $p_0 = \cos^2(\theta/2)$ between the control parameter and the observed quantity is non-linear, and its derivative vanishes at the poles. This has concrete consequences. If a variational algorithm wants to adjust the state according to what it measures, the chain rule introduces the factor

$$\frac{dE}{dq} = -\frac{1}{\sin\theta}\,\frac{dE}{d\theta},$$

which diverges at $\theta = 0$ and $\theta = \pi$ (research notes, §1). And if θ is to be estimated from N shots, the variance of $\hat{p}_0$ changes with the state, so the uncertainty of each reading depends on where the qubit is (§5). Neither problem is exotic: both appear as soon as one reasons in terms of what the experiment delivers, and the probability equation is the tool to handle them.

1.6 Inversion and the branch sign

The angle is recovered with $\theta = \arccos(qg_Z)$. The inversion has a limit: $\cos\theta$ is even, so qg_Z does not distinguish θ from $-\theta$. As long as one works on $[0, \pi]$ there is no problem, but an optimizer that needs to cross a pole gets stuck. The remedy is to measure the sign of the sine as well: the pair (qg_Z, s) with $s = \mathrm{sign}(qg_X)$ is one-to-one on the whole circle (§65). In the H_2 molecule the variational optimum is at $\theta \approx -0.22$; every update that uses only qg_Z stops at the Hartree–Fock state, while the signed version reaches the exact energy.

1.7 The outcome entropy: qg_S

The polar bias says *where* the qubit points; it does not say how unpredictable the measurement is. That is the role of the **superposition spread**, the Shannon entropy of the outcome:

$$qg_S = H(p_0) = -p_0 \log_2 p_0 - (1 - p_0)\log_2(1 - p_0) \in [0, 1].$$

Since qg_Z and qg_S depend on the same population p_0, they are tied by an exact, branch-free identity (§7):

$$qg_S = H\left(\frac{1+qg_Z}{2}\right),$$

valid for every single-qubit state, pure or mixed, however it was prepared. Its derivative, $\frac{d\,qg_S}{d\,qg_Z} = \frac{1}{2}\log_2\frac{1-qg_Z}{1+qg_Z}$, diverges at the poles: it is the entropic counterpart of the gradient singularity. Near a pole, a minimal change of the bias produces an enormous relative change of the entropy.

1.8 Invertibility of qg_S and its error bars

Unlike qg_Z, the entropy is not invertible over the whole domain: $qg_S(\theta) = qg_S(\pi-\theta)$. The population p_0 decreases from 1 to 0 on $[0,\pi]$ and H rises up to $p = \frac{1}{2}$ and then falls, so qg_S is strictly increasing on $[0,\pi/2]$ and strictly decreasing on $[\pi/2,\pi]$ (Figure 1.5). Inversion requires choosing one of the two halves (§2).

Estimating qg_S from finite shots raises a further problem: the entropy is flat at its maximum, and the delta method gives a zero standard error exactly where $qg_S \approx 1$. The native qang solution is to build the Wilson interval for qg_Z and map it through the identity above; since H is unimodal, the image is a valid interval. Over 2000 repetitions that interval kept a coverage of 0.90–0.99 in every case tested, even where the Wald interval or the bootstrap fell below 0.40 (§51).

1.9 The qang angle as the arcsine transformation

The central statistical property of the unit follows from propagating shot noise. With $\widehat{qg}_Z = 2\hat{p}_0 - 1$,

$$\mathrm{Var}(\widehat{qg}_Z) = \frac{\sin^2\theta}{N}, \qquad \mathrm{Var}(\hat{\theta}) \approx \mathrm{Var}(\widehat{qg}_Z)\left(\frac{d\theta}{d\,qg_Z}\right)^2 = \frac{\sin^2\theta}{N}\cdot\frac{1}{\sin^2\theta} = \frac{1}{N}.$$

The $\sin^2\theta$ factors cancel exactly: **the uncertainty of the qang angle does not depend on the state** (§5), which is the horizontal dashed line of Figure 1.3. It is not a coincidence. Since $p_0 = \cos^2(\theta/2)$, the qang angle is, up to the reflection $\theta \to \pi-\theta$, the classical arcsine transformation of a proportion, $\phi = 2\arcsin\sqrt{p}$,

introduced by Bartlett (1947), Anscombe (1948) and Freeman and Tukey (1950) precisely to stabilize the variance (§125).

The claim was checked with exact binomial sums, without simulation: with $n = 100$, $(n + \frac{1}{2})\,\mathrm{Var}(\hat{\theta})$ lies between 1.000 and 1.001; the coverage of the interval $\hat{\theta} \pm z/\sqrt{n + 1/2}$ stays between 0.930 and 0.962 for $n \in \{30, 100, 1000\}$. The prediction failed in one case, and it is worth saying: with $np \approx 2.5$ both the angular and the Wald intervals lose coverage. The stabilization is asymptotic, and the regime of very few events requires the exact intervals of chapter 4.

1.10 Mixed states and multi-qubit registers

The definitions extend unchanged to mixed states and to registers:

- **Mixed states.** $qg_Z = \mathrm{Tr}(\rho\sigma_z)$ no longer determines the state. The classical mixture of $|0\rangle$ with weight 0.2 and $|1\rangle$ with weight 0.8 has $qg_Z = -0.6$ and von Neumann entropy 0.72, a combination no pure state reproduces (§3). The qang angle $\theta = \arccos qg_Z$ is still a function of p_0, so it keeps $\mathrm{Var}(\hat{\theta}) = 1/N$ here too: the probability does not know whether the state is pure. What loses the stabilization is the estimate of the Bloch polar angle, $\arccos(qg_Z/r)$, of a qubit of length $r < 1$: $\mathrm{Var} = \left[1 + \frac{1-r^2}{r^2 \sin^2\theta}\right]/N$, equal to the quantum Cramér–Rao bound $1/(Nr^2)$ on the equator (§45).
- **Registers.** Each qubit has its local profile, obtained by partial trace, and the register has a joint entropy normalized by the number of qubits (§4). The Bell state illustrates it: each qubit has $qg_Z = 0$ and a local entropy of 1 bit, but the normalized joint Shannon entropy of the outcomes is 0.5 (the von Neumann entropy of the register is 0). The register is globally pure and locally mixed; the local profiles are not enough to describe it, and chapter 5 returns to this point with the Bloch radius.

1.11 Derived quantities of the reading

The treatise's formulation (Monteverde 2026a) names, on top of this base, seven derived quantities. All of them come from the reading —from qg_Z, or from the full qg vector when needed—, so they add no information to p_0: they name combinations that the rest of the book uses. Table 1.3 collects them with their definition, what they are in the usual terms of statistics or physics, and how they are treated here. Appendix H (Appendix H) catalogues these and the book's other sub-measures with the decision each one improves and its evidence.

Table 1.3: Derived quantities of the formulation and how the book treats them.

quantity	definition	what it is	treatment in the book
Q-Hertz (qHz)	qga per second	angular frequency on the qga scale: 1 qHz $= 2\pi/1000$ rad/s	unit of convenience
Q-Fisher (qFi)	Fisher information per qga^2	$I_F = 4\pi^2 \cdot 10^{-6}$ qga^{-2} about the qang angle, for every state	exact (Section 1.4)
Q-Shannon (qg_S)	$H(p_0) = H\big((1+qg_Z)/2\big)$	Shannon entropy of the outcome	exact (Section 1.7)
Q-Curvature (qg_κ)	$-d^2 qg_S / d\,qg_Z^2 = 1/[\ln 2\,(1-qg_Z^2)]$	Fisher information about qg_Z, in bits	exact; in optimizers, no gain (Section 30.4)
Q-SNR (qg_{SNR})	$\lvert qg_Z\rvert/\sqrt{1-qg_Z^2} = \lvert\cot\theta\rvert$	signal-to-noise of one shot	exact; sets the shots needed to read a sign
Q-Geodesic (qg_L)	$\arccos\big(\sqrt{p_0 p_0'} + \sqrt{p_1 p_1'}\big) = \lvert\theta-\theta'\rvert/2$	Bhattacharyya angle between two readings	exact; it is the Fubini–Study distance on a common meridian
Dissipative deficit (qg_τ)	$\sqrt{1-r^2} = \sqrt{1-qg_X^2-qg_Y^2-qg_Z^2}$	$2\sqrt{\det\rho}$, the lack of purity	exact; measures loss of purity of any origin

The first three have already appeared. On the Q-Fisher, one precision: the formulation describes its constancy as that of a flat manifold, and it is not one. What is constant is the information about a single coordinate, the polar angle; the Bloch sphere with the Fubini–Study metric has Gaussian curvature 4 (Section 1.4).

Information curvature. Differentiating twice the identity $qg_S = H\big((1+qg_Z)/2\big)$ of Section 1.7,

$$\frac{d\,qg_S}{d\,qg_Z} = \frac{1}{2}\log_2\frac{1-qg_Z}{1+qg_Z}, \qquad qg_\kappa = -\frac{d^2 qg_S}{d\,qg_Z^2} = \frac{1}{\ln 2\,(1-qg_Z^2)}.$$

The denominator is the variance of one shot, so $qg_\kappa = I_F(qg_Z)/\ln 2$: the curvature of the entropy is the Fisher information that one shot carries about the polar bias, expressed in bits. It is $1/\ln 2 = 1.443$ on the equator and diverges at the poles, where the entropy is most sensitive. For a pure state, $1 - qg_Z^2 = \sin^2\theta$ and $qg_\kappa = 1/(\ln 2 \, \sin^2\theta)$. **Mind the notation:** qg_κ is not the pulse curvature $\kappa_{qg} = 1000\,|\cos\theta|$ of Section 1.2. That one is the second derivative of the reading with respect to the angle; this one, that of the entropy with respect to the reading. The formulation proposes using qg_κ to damp an optimizer's step near the poles. Scaling the step by the inverse Fisher information is the natural gradient (Amari 1998), and Chapter 30 measures it (Section 30.4): on the coordinate qg_Z that step is, to first order, gradient descent on θ, and it brings no speed-up of its own.

Signal-to-noise per shot. A shot is ± 1, with mean qg_Z and standard deviation $\sqrt{1 - qg_Z^2}$, for any state, pure or mixed. Their ratio and the number of shots a sign requires are

$$qg_{SNR} = \frac{|qg_Z|}{\sqrt{1 - qg_Z^2}} = |\cot\theta|, \qquad N_{\text{sign}} = \left(\frac{z}{qg_{SNR}}\right)^2,$$

with $\theta = \arccos qg_Z$, the qang angle. With N shots the signal-to-noise is $\sqrt{N}\, qg_{SNR}$, and N_{sign} is the number of shots for the sign of qg_Z to stand z standard deviations from zero: it is the stopping criterion the formulation proposes. At 60° (+500 m-qg) $qg_{SNR} = 0.577$ and 27 shots are enough for a 3σ sign; at 80° it takes 290, and at 89°, about 29 500. On the equator the signal vanishes and the criterion never ends. Like every formula of this chapter based on the normal approximation, it holds when N_{sign} is not too small; with few shots the exact binomial test is used.

Geodesic distance. Two readings (p_0, p_1) and (p_0', p_1') are at Wootters' statistical distance of Section 1.3, and for two outcomes that distance is

$$qg_L = \arccos\left(\sqrt{p_0 p_0'} + \sqrt{p_1 p_1'}\right) = \frac{|\theta - \theta'|}{2},$$

because $\sqrt{p_0 p_0'} + \sqrt{p_1 p_1'} = \cos\frac{\theta}{2}\cos\frac{\theta'}{2} + \sin\frac{\theta}{2}\sin\frac{\theta'}{2} = \cos\frac{\theta-\theta'}{2}$. It is half the difference of qang angles, that is, the distance between readings on the scale that stabilizes the variance. If both states are pure and share the meridian, qg_L equals the Fubini–Study distance $\arccos|\langle\psi|\psi'\rangle|$. With different azimuths it is a lower bound, because a measurement cannot tell two states apart better than their overlap does. Chapter 3 uses it between the states of a reflector (Section 3.4).

Dissipative deficit. The three readings qg_X, qg_Y and qg_Z form the qg vector **r**, of length $r \leq 1$, and

$$qg_\tau = \sqrt{1 - r^2} = 2\sqrt{\det \rho}, \qquad \operatorname{Tr} \rho^2 = 1 - \tfrac{1}{2}\, qg_\tau^2.$$

It is 0 for a pure state and 1 for the maximally mixed state. The 0.2/0.8 mixture of the previous section has $r = 0.6$ and $qg_\tau = 0.8$; each qubit of a Bell pair, $qg_\tau = 1$. The formulation associates it with T_1 and T_2 relaxation, which is one of its causes, but the quantity measures the loss of purity without telling its origin apart: an average over calibration errors produces it too (Chapter 2, Section 2.6). Unlike the previous ones, it is not a function of p_0: it requires measuring in three bases.

1.12 Scope and limits of the base

It is worth stating precisely what this base brings and what it does not, before building the rest of the book on it.

What it is not. The qang unit is not new physics. The polar bias is the expectation value of σ_z, the qang angle is the arcsine transformation of a proportion and the milli-qang is, mathematically, $1000\,\langle \sigma_z \rangle$. None of these readings holds information that the probability p_0 does not already have; the probability equation says so explicitly. Nor does it reduce the shots a given precision requires (Chapter 4).

What it is. A working scale, in the sense of the decibel: it puts preparation, gate error, shot noise and read-out in the same units, so that they add, compare and can be budgeted. The question that matters is then empirical: does reasoning in milli-qang predict what a real processor does?

First test on hardware (§168). It was answered with a pre-registered experiment on one qubit of IQM's superconducting Garnet processor: 26 circuits of 4000 shots, with four predictions fixed before the run. The result was mixed, and it is reported in full:

- **The cost of a gate is measurable and budgetable in m-qg.** Chaining 3 to 41 X gates, the shortfall of the reading (corrected for the readout errors measured at the poles) grows with a single slope: **1.64 m-qg per X gate**, with a χ^2 per degree of freedom of 0.81. After 41 gates it is 63 m-qg, eight standard deviations above zero (Figure 1.6, left). Linear rather than quadratic growth points to incoherent error, relaxation during the pulses, and not to a coherent over-rotation.

- **The scale makes visible an error that the readout model hides.** The affine readout predicted from the poles fits at the equator, but the four intermediate angles read between 16 and 55 m-qg low (Figure 1.6, right). The processor calibrates the angles 0, $\pi/2$ and π well; the others are over-rotated by $1.3°$ to $3.9°$. The prediction passed by the 4-of-5 rule, not cleanly.
- **One prediction failed.** The ten copies of the equator circuit scattered half as much as binomial shot noise allows, something independent shots cannot produce and which happens by chance with probability 0.013. A single run does not decide whether it was a fluke or a property of how the device returns batched circuits.
- **Another prediction was badly designed.** Fitting two parameters on four points left an extrapolation band so wide that any reading passed. It is recorded as a design flaw, not as a success.

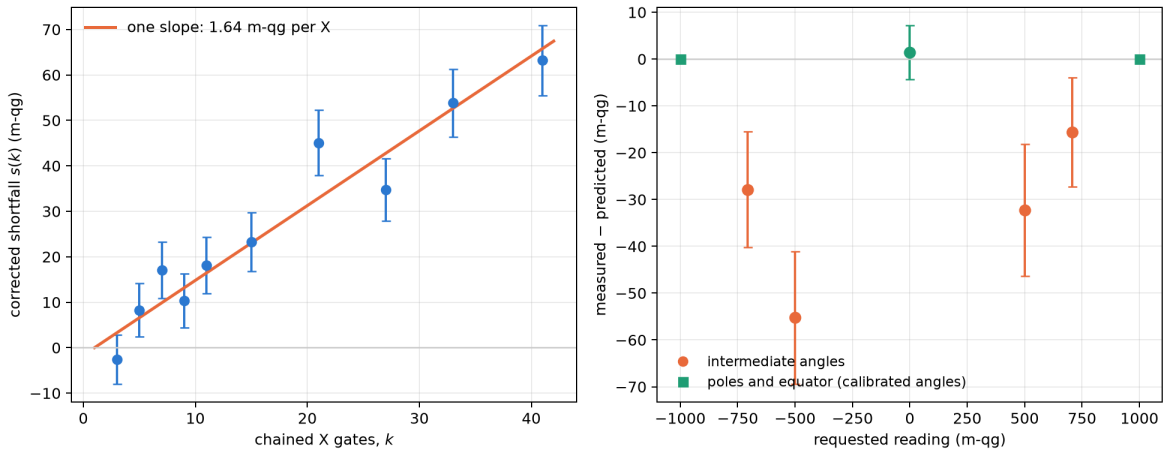


Figure 1.6: First test of the milli-qang on hardware (§168, IQM Garnet, one qubit, 4000 shots per circuit). Left: shortfall of the reading after k X gates, readout-corrected, with the single-slope line. Right: difference between measured and predicted by the affine readout; the poles are zero by construction, because they define the calibration.

The balance: on a real chip, the milli-qang turned the error of a gate into a budgetable number and exposed an angle error that the usual model does not see. But it was tested on one qubit, once, and shot noise did not describe the scatter of repeated circuits. Replication on more qubits and on another day is pending, and the book treats this base as solid mathematically and promising, but not yet consolidated, experimentally.

Second test on hardware (§170). On the same Garnet qubit, $E = mc^2$ was read as a reading, with the β qubit of Appendix G (Section G.6). The Dirac eigenstate

at $0.6c$ read $+759.9 \pm 14.9$ m-qg, where $1000/\gamma = +800$ was predicted, and its mirror on the negative branch read -809.4. The *Zitterbewegung* showed at the predicted frequency, 0.44% off. All five predictions passed. The low bias at intermediate angles appeared again, in the same direction as in §168, which points to a feature of the device rather than a fluke of one run. Both tests are on one qubit, and replication is still pending.

Open proposals. The author's working version adds three proposals to this basis, which the book records as **unverified**:

- A regularized variational energy, $E_{\text{total}} = \langle H \rangle + \lambda_\theta L_{\text{qang}}(\theta) + \lambda_\varphi L_{\text{qang}}(\varphi)$. Here the energy is the expectation value of a Hamiltonian, not the relativistic energy of the previous section. The form of the penalties L_{qang} and the values of the weights λ are not defined.
- A normalized radiation function, $R(\theta, \varphi) = N\,[\,p_1 + \alpha \cos\varphi\,]$, with the excited population $p_1 = \sin^2(\theta/2)$ as the emitting power. Neither N nor α is fixed, and it is not checked against an emission model.
- Mechanisms against barren plateaus (local costs in m-qg, minimum-curvature κ_{qg} compilation, initialization with the branch sign). Two qang tools are already tested on the book's networks. The **qg local cost**, $(1 - \overline{qg_Z})/2$, the idea of Cerezo et al. (2021) written as a reading, has 470 times more gradient variance than the global cost with 10 qubits and shallow circuits (Section 23.7). The **weight filter** returns, under T1, exactly the noiseless gradient, to 10^{-16}, against a noise-induced plateau (Section 20.6; pre-registered prediction, §112). Both have measured limits: with deep circuits the local cost also decays exponentially and, where it escapes the plateau, it is classically simulable; and the weight sector itself flattens the gradient as the register grows, with or without the filter: like a polynomial at weight 1, which is simulable, and exponentially at half filling. Chapter 23 gathers the three sources of a plateau and what qang does with each (Section 23.8). Compilation by κ_{qg} and initialization with the branch sign have not been tested against plateaus.

1.13 Without qang and with qang

Table 1.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
describe the outcome of a qubit	the probability p_0 (or $\langle Z \rangle$)	qg_Z, the angle θ, m and qg_S, all exact functions of p_0	**none in content**: the same information; qang gives it a working scale from -1000 to $+1000$
error bar of a reading	$\sqrt{p_0(1-p_0)/N}$, different for each state	the angle has variance $1/N$ for every state; in m-qg, $1000 \sin\theta/\sqrt{N}$	**convenience**: the variance-stabilized coordinate of classical statistics (the arcsine) is the gate parameter
compare a gate error with shot noise	infidelity on one side, binomial deviation on the other, in different units	both in m-qg (on Garnet, 1.64 m-qg per X gate against the run's shot noise)	**in favour**: one scale for the error budget
decide how many shots the sign of a reading needs	binomial test or normal approximation	$N = (z/qg_{SNR})^2$, with $qg_{SNR} = \lvert\cot\theta\rvert$	**convenience**: it is the normal approximation written with the reading
express a gate or phase angle	radians, with factors of π	qga: 1000 per turn (X = 500 qga, equator = 250 qga)	**convenience**: the same information without factors of π; it changes neither gradients nor plateaus
improve the device	—	—	**none**: the scale measures, it does not correct; prediction M2 of §168 failed

Balance. The base adds no information to p_0: it reorganizes it. What it gains is a uniform error bar in the angle and a common unit for gate errors and shot noise,

which are exactly the two things the rest of the book uses.

1.14 In code

The `qang.probability` module is this chapter's analysis tool: it reads counts and returns every reading, each with its interval.

```
import math
from qang import probability as pb

pb.readings_from_p0(0.75)      # theta = 1.047 (60°), qg_Z = 0.5, qg_S = 0.811, 500 m-qg

# 750 zeros in 1000 shots, against a target qg_Z = 0.5
r = pb.analyze(750, 1000, target_qg_z=0.5)
r.milliqang               # (500.0, 444.5, 551.7) m-qg: estimate and Wilson interval
r.theta                   # (1.047, 0.986, 1.110) rad = 60.0°, 56.5°, 63.6°
r.qg_s                            # (0.811, 0.768, 0.852)
r.deviation_milliqang             # 0.0 m-qg from the target

pb.angle_per_milliqang(math.pi / 2)      # 0.001 rad: one m-qg is one milliradian on the e
pb.pole_angle(1.0)                    # 0.0447 rad: one m-qg below the pole
pb.shots_for_resolution(10, math.pi / 2)  # 10000 shots to resolve 10 m-qg
pb.pulse_error_milliqang(math.pi, 0.05)   # 12.3 m-qg: X gate with a 5 % amplitude error
pb.statistical_angle([0.5, 0.3, 0.2], [0.4, 0.4, 0.2])   # 0.113 rad

# derived quantities
pb.quantum_snr(0.5)                # 0.577: signal-to-noise of one shot at 60°
pb.shots_for_sign(0.5)                    # 27 shots for a 3σ sign
pb.information_curvature(0.0)          # 1.443 = 1/ln 2, the minimum curvature (equator)
pb.dissipative_deficit(0.0, 0.0, -0.6)     # 0.8: the 0.2/0.8 mixture
```

The figures of this chapter are generated by `tools/make_cap01_figures.py`, which uses the same functions.

1.15 Summary

- **Definition.** The polar bias $qg_Z = \langle \sigma_z \rangle = 2p_0 - 1 = \cos\theta$ is bijective, directly measurable and valid for mixed states; the qang angle $\theta = \arccos qg_Z$ and the entropy $qg_S = H(p_0)$ complete the unit.

- **Milli-qang.** 1 qg $= 1000$ m-qg; the scale is linear in probability (1 m-qg $= \Delta p_0 = 0.0005$) and not linear in angle (1 mrad on the equator, 2.56° next to the pole). Resolving δ m-qg takes $(1000 \sin\theta/\delta)^2$ shots, and gate errors are budgeted in the same units.
- **Probability equation.** θ, qg_Z, qg_S and m are exact functions of a single probability p_0; with several outcomes, the qang angle generalizes to Wootters' statistical distance.
- **Angular scale.** The angular qang measures angles with 1000 qga per turn and is distinct from m-qg, which is a reading: $p_0 = \cos^2(\pi\theta_{qga}/1000)$ and $I_F = 4\pi^2 \cdot 10^{-6}$ qga^{-2}. The extension to relativistic energy and antimatter is an unverified proposal, with three errors of the source document corrected.
- **Statistics.** The qang angle is the arcsine transformation of a proportion: its variance is $1/N$ for every state, pure or mixed; what needs a correction in a mixed state is the estimate of the Bloch angle.
- **Derived quantities.** qg_κ is the Fisher information about qg_Z in bits (and not the pulse curvature κ_{qg}); $qg_{SNR} = |\cot\theta|$ sets the shots needed to read a sign; $qg_L = |\theta - \theta'|/2$ is Wootters' distance between two readings; $qg_\tau = \sqrt{1 - r^2}$ measures the lack of purity and is the only one that needs three bases.
- **Scope.** The unit is not new physics but a working scale. On real hardware (§168) it measured a cost of 1.64 m-qg per X gate, linear in the number of gates, and exposed reading shifts of 16–55 m-qg (over-rotations of 1.3°–3.9°); a prediction about shot noise failed and replication is pending.

Bibliographic notes

The Born rule (Born 1926) and the Bloch sphere (Bloch 1946) are the starting point; for states, measurements and channels in general, see Nielsen and Chuang (2010). Wootters' statistical distance (1981) and its relation to Fisher information are in Braunstein and Caves (1994); the Cramér–Rao bound, in Cramér (1946) and Rao (1945). The arcsine transformation that stabilizes the binomial variance comes from Bartlett (1947), Anscombe (1948) and Freeman and Tukey (1950), and the score interval is that of Wilson (1927). The binary entropy behind qg_S is that of Shannon (1948); the natural gradient, against which the information curvature qg_κ is compared, is that of Amari (1998). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c). The angular qga scale and the relativistic proposal come from the working document Monteverde (2026a); the free Dirac spinor and its components are in any text on relativistic quantum mechanics.

2 Native gates, closed algebra and the universal operator

Chapter 1 described the state of a qubit in qg units and reduced all its readings to a single probability, p_0, with the milli-qang as the working scale. This chapter does the same for operations: how a state is prepared from its probability or its m-qg reading, how any single-qubit gate is written as an algebraic function of qg_Z, why its action on the populations is the probability equation of Chapter 1, how the idea extends to systems with more levels, and which compilation rules follow when the error of a gate, measured in m-qg, depends on its angle.

2.1 The elementary gate R_{qg}

The basic operation is to prepare, from $|0\rangle$, the state with polar bias q: a rotation $R_y(\theta)$ with $\theta = \arccos q$. qang provides it as a native gate in the three most widely used development kits (§6):

- **Qiskit:** `RQangGate(q)` and `FullRQangGate(q, φ)`, the latter with the azimuthal phase.
- **Cirq:** `rqang_gate(q)`, built on `cirq.ry`, and `full_rqang_gate(q, φ)`, a `MatrixGate` with the same convention as Qiskit's `U` gate.
- **PennyLane:** the same pair, differentiable with respect to q.

By the probability equation (Section 1.3), the gate parameter can be set from any of the readings: for a probability p_0 use $q = 2p_0 - 1$, and for a reading of m m-qg, $q = m/1000$. Asking for $p_0 = 0.75$ or for $+500$ m-qg is asking for the same gate, $R_y(60^\circ)$. On the angular qga scale of Section 1.4 (1000 qga per turn), that gate is $R_y(166.667\ \text{qga})$: a gate angle is written without factors of π, with X at 500 qga and the equator at 250 qga.

All three prepare exactly the same state for the same qg value, which is checked both on the measured probabilities and on the state-vector fidelity. The gate accepts symbolic parameters (Qiskit `Parameter`, sympy symbols), so it can sit inside parameterized circuits.

A signed variant, `SignedRQangGate`, applies $R_y(s \cdot \arccos q)$ with $s = \pm 1$: it is the gate form of the branch sign of Chapter 1, and prepares the same state as the full gate with $\varphi = 0$ or π (§68).

2.2 The universal operator in qg

Every single-qubit gate is, up to a global phase, $U(\theta, \varphi, \lambda)$. With $q = \cos\theta$ the half-angles become algebraic,

$$\cos\frac{\theta}{2} = \sqrt{\frac{1+q}{2}}, \qquad \sin\frac{\theta}{2} = \sqrt{\frac{1-q}{2}},$$

and the universal gate is written without trigonometric functions of q:

$$U_3^{qg}(q, \varphi, \lambda) = \begin{pmatrix} \sqrt{\dfrac{1+q}{2}} & -e^{i\lambda}\sqrt{\dfrac{1-q}{2}} \\ e^{i\varphi}\sqrt{\dfrac{1-q}{2}} & e^{i(\varphi+\lambda)}\sqrt{\dfrac{1+q}{2}} \end{pmatrix}.$$

The amplitudes depend on the polar bias only, and the phases are decoupled into the factors $e^{i\varphi}$ and $e^{i\lambda}$. An immediate consequence is the **transition matrix**, the probabilities $|U_{ij}|^2$:

$$T(q) = \begin{pmatrix} \frac{1+q}{2} & \frac{1-q}{2} \\ \frac{1-q}{2} & \frac{1+q}{2} \end{pmatrix},$$

doubly stochastic, rational in q and independent of the phases. Seen from a measurement in the computational basis, and for inputs without coherences, every single-qubit gate is a binary symmetric channel with parameter $(1-q)/2$.

2.3 The gate as a probability equation

The transition matrix is Chapter 1's probability equation written for an operation. With $p_0 = (1+q)/2$,

$$T = \begin{pmatrix} p_0 & 1-p_0 \\ 1-p_0 & p_0 \end{pmatrix}, \qquad p_0 = \frac{1000+m}{2000} :$$

on an input without coherences ($|0\rangle$, $|1\rangle$ or a classical mixture of them), the action of any single-qubit gate on the populations is fixed by **a single number**, its reading in m-qg. The identity and Z sit at $+1000$ m-qg (they do not move populations), X at -1000 (it swaps them) and the Hadamard at 0 (it equalizes them). The phases φ and λ do not appear: they are invisible to a Z measurement and only show up when gates are composed, as Chapter 3 shows. On an input with coherences the full formula is $p_0' = T(q)\,p - \sqrt{1-q^2}\ \mathrm{Re}(e^{-i\lambda}\rho_{01})$: the Hadamard on $|+\rangle$ gives $p_0 = 1$, not the 0.5 that $T(0)$ would give.

The entropic reading has its version for operations too. The capacity of the binary symmetric channel with parameter $(1-q)/2$ is $1 - H\big((1-q)/2\big)$, and since H is symmetric, $H\big((1-q)/2\big) = H(p_0) = qg_S$:

$$C = 1 - qg_S.$$

A gate at a pole ($qg_S = 0$) carries one full bit per use; the Hadamard ($qg_S = 1$) carries nothing from the computational basis to itself.

The information curvature of Chapter 1 (Section 1.11) has a direct reading here. Since $C = 1 - qg_S$,

$$\frac{d^2C}{dq^2} = qg_\kappa = \frac{1}{\ln 2\,(1-q^2)} :$$

the capacity is convex in the gate's bias, with its minimum at the Hadamard, and it bends more and more near the poles. Moving the same amount away from a pole or from the equator does not cost the same: a gate at $+990$ m-qg, 10 m-qg below the pole, carries 0.955 bits, already 0.045 less than the identity; one at $+10$ m-qg, 10 m-qg above the equator, carries $7 \cdot 10^{-5}$ bits, almost the same as the Hadamard.

2.4 Verification against standard libraries

The formula was checked against Qiskit 2.6's `U` gate on 2000 random gates, with a largest difference of 1.8×10^{-16}, and likewise against Cirq, PennyLane and QuTiP (§166, `tools/verify_cap02.py`). The usual gates appear at simple values:

Table 2.1: Standard gates as points of U_3^{qg}.

gate	q	m-qg	θ (qga)	φ	λ
I	1	$+1000$	0	0	0
X	-1	-1000	500	0	π (500 qga)
Z	1	$+1000$	0	0	π (500 qga)
H	0	0	250	0	π (500 qga)

The table reads in qg: X is the only case with bias -1 (it flips the qubit), Z and the identity do not change the populations ($q = 1$) and differ only by a phase, and the Hadamard takes the qubit to the equator ($q = 0$).

2.5 Beyond the qubit: the Cartan subalgebra

For a d-level system (a qudit), the analogue of σ_z is not one matrix but the **Cartan subalgebra** of $\mathfrak{su}(d)$: the $d-1$ diagonal generalized Gell-Mann matrices,

$$\lambda_k = \sqrt{\frac{2}{k(k+1)}}\ \mathrm{diag}(\underbrace{1,\dots,1}_{k},-k,0,\dots,0), \qquad k = 1,\dots,d-1,$$

normalized so that $\mathrm{Tr}(\lambda_j\lambda_k) = 2\delta_{jk}$. Their expectation values are linear in the populations and determine them:

$$p = \frac{1}{d} + \frac{1}{2}\sum_{k=1}^{d-1}\langle\lambda_k\rangle\ \mathrm{diag}(\lambda_k).$$

For $d = 2$ a single matrix remains, $\lambda_1 = \sigma_z$, its expectation value is qg_Z and the formula reduces to $p_0 = \frac{1}{2} + \frac{1}{2}\,qg_Z$: the probability equation of Chapter 1. This is the natural generalization of the unit: a vector of $d-1$ qg values measurable in the computational basis, each expressible in m-qg, which fixes the d probabilities the way qg_Z fixes p_0 (§166).

Definition and usefulness should be kept apart. Whether a three-level readout gives qang an advantage was asked on a concrete case, a transmon leaking to the level $|2\rangle$ in a surface-code memory, and the answer was no: reading $|2\rangle$ as 1 is usually right, and declaring the qubit erased discards that information (§70). qang remains a qubit construction; the qudit generalization is an identity, not a demonstrated advantage.

2.6 The minimal pulse-area principle

On devices with fractional gates (rotations $R_X(\theta)$ or $R_{ZZ}(\theta)$ with a continuous angle), the pulse area grows with $|\theta|$, and so do the amplitude-miscalibration error and the relaxation during the pulse. In qg units, a random amplitude error of standard deviation σ damps the read-out as

$$qg_Z = \cos\theta \; e^{-\theta^2\sigma^2/2},$$

and a fixed one ε gives $\cos(\theta(1+\varepsilon))$, with first-order error $-1000\,\varepsilon\,\theta\sin\theta$ m-qg. The sensitivity to an absolute angle error, $d\,qg_Z/d\theta = -\sin\theta$, is largest at $\pi/2$ and vanishes at 0 and π; the fixed error is relative ($\varepsilon\theta$) and weighs most near 116° (§160). This is Chapter 1's milli-qang error budget (Figure 1.4), and the scale makes the gate error comparable with the shot noise of the read-out: with $\sigma = 0.05$, the shift of an X (12.3 m-qg) is about eight times the resolution of 10^4 shots at its reading, -987.7 m-qg, where $\sigma_m = 1000\sin\theta/\sqrt{N} = 1.6$ m-qg (on the equator it would be 10 m-qg). In terms of the curvature of Chapter 1, the random shift is, to second order, $\frac{1}{2}\theta^2\sigma^2\kappa_{qg}$: it depends on the area (θ^2) and on where the state ends ($\kappa_{qg} = 1000|\cos\theta|$), which vanishes on the equator.

Every gate has shorter representatives. $R_X(\theta) = -R_X(\theta - 2\pi)$ has period 2π up to a global phase (1000 qga), and Figure 2.1 shows what is gained: asking for $R_X(324^\circ) = R_X(900 \text{ qga})$ as is shifts the reading by -31.7 m-qg, while its representative $R_X(-36^\circ) = R_X(-100 \text{ qga})$, the same gate, shifts it by -0.4 m-qg.

The dissipative deficit of Chapter 1 (Section 1.11) completes the budget. Averaging over the amplitude error does not move the qg vector off its meridian; it shortens it:

$$\mathbf{r} = e^{-\theta^2\sigma^2/2}\,(0,\,-\sin\theta,\,\cos\theta), \qquad qg_\tau = \sqrt{1 - e^{-\theta^2\sigma^2}} \approx |\theta|\,\sigma.$$

Unlike the m-qg shift, which depends on where the state ends and vanishes on the equator, qg_τ depends only on the area. With $\sigma = 0.05$ it is 0.156 for X and 0.078 for $R_X(\pi/2)$, whose reading does not shift at all; for $R_X(324^\circ)$ it is 0.277 and for its representative $R_X(-36^\circ)$, 0.031. The m-qg reading sees the area only when the state ends far from the equator; qg_τ, which requires measuring in three bases, always sees it. Here the deficit does not come from relaxation but from averaging gates that differ from shot to shot: the formulation's name describes one of its causes, not the only one.

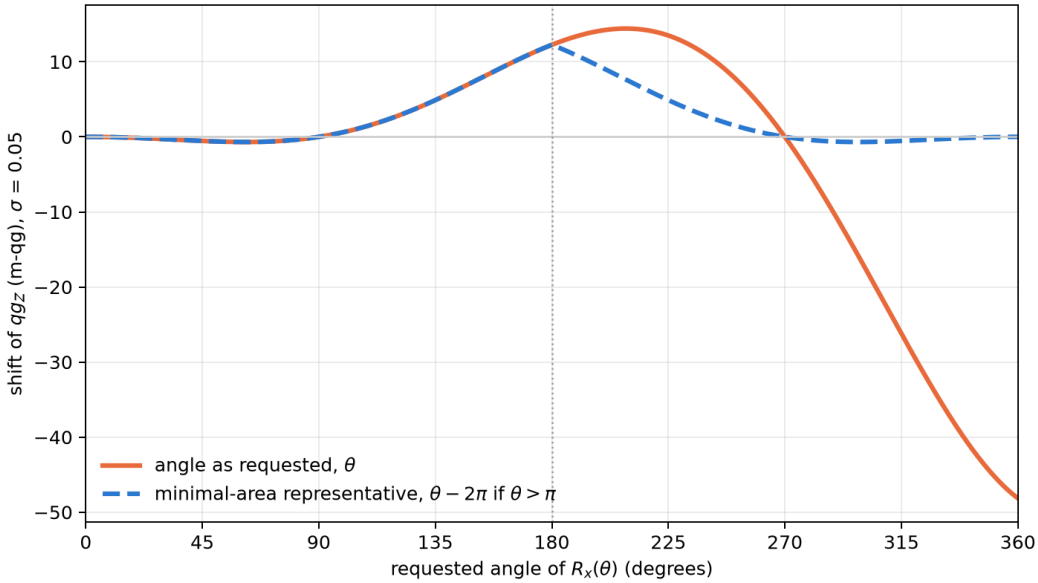


Figure 2.1: Milli-qang error of $R_x(\theta)$ under a random amplitude error $\sigma = 0.05$: the angle as requested against the minimal-area representative. Below 180° they coincide; above it, the short representative avoids the error that, at multiples of π, grows with the square of the angle.

More interesting: $R_{ZZ}(\theta+\pi) = R_{ZZ}(\theta)\,(-i\,Z\otimes Z)$, and $Z\otimes Z$ is two virtual $R_Z(\pi)$ rotations, free on superconducting processors. So R_{ZZ} has period π (500 qga) and its area never needs to exceed $\pi/2$ (250 qga). The function `minimal_area_circuit` rewrites a circuit with those representatives and preserves the unitary exactly, global phase included.

On a 4-qubit QAOA model with errors proportional to the area (150 instances), reducing the angles modulo 2π cut the total area 4.4× and the error (total variation distance) 4.1×; using the period π of R_{ZZ} as well cut it another 1.37× in the same ratio of medians. The recorded prediction was about the median of the per-instance ratios, which was 1.14× against 1.5×, and it failed: half of the area, that of the R_X rotations, was already minimal after the first reduction. The general result stands: **the error follows the area**. Two scope notes: on IBM processors ordinary single-qubit rotations are virtual R_Z plus fixed $\sqrt{X}$ pulses, of duration independent of the angle, so the rule applies to fractional gates; and the benefit depends on how strongly the error scales with the area on each platform.

The author's working version further proposes using this minimum-curvature compilation, together with a regularized variational energy and a radiation function tied to the excited population $p_1 = (1-q)/2$, against noise-induced barren plateaus. These are the open proposals of Chapter 1 (Section 1.12): **unverified**, and the book attributes no result to them.

2.7 An exposure screen before running

The symmetry filter (Chapters 20–21) discards the shots that left the conserved Hamming-weight sector. But a decay that happens while part of the state is outside the sector can rotate back in and pass the filter. That suggests a prior measure, computable without noise: the **exposure**, the population outside the sector right after each two-qubit gate (`qang.sectors.sector_exposure`).

Over 16 compilations of one step of an XXZ chain, the rank correlation between exposure and T1 leak was exactly 0.80, the recorded threshold, and the zero-exposure compilations removed at least 98% of the T1 error (§68):

Table 2.2: Exposure computed without noise against the measured leak.

compilation	exposure	T1 leak
3 CNOT	0.42	16–52%
Mølmer–Sørensen, ZZ by basis change	0.16–0.24	20–78%
Mølmer–Sørensen rotations	0.0006–0.047	0.4–17%

number-conserving	0	0.0–2.0%

The screen ranks compilations, not time steps: within one compilation the relation is weaker. Its practical value is choosing the compilation before spending shots.

2.8 One qubit order for every kit

A practical detail that causes silent errors: Qiskit writes qubit 0 at the right of the bitstring, while Cirq, PennyLane, Braket and pytket put it first. `qang.interop` converts any result to $\{(b_0, \dots, b_{n-1}) : \text{count}\}$ and computes the qg read-outs on that form: per-qubit qg_Z with its error, the filter and the register entropy (§130). A 3-qubit W state sampled with 20 000 shots in each kit gave qg_Z between 0.316 and 0.343 (ideal $1/3$), all shots in weight 1 and register entropy 0.528 (ideal $\log_2 3/3$), alike in all three. In milli-qang the readings run from 316 to 343 m-qg against 333 ideal; with 20 000 shots the resolution is $\sigma_m = 1000 \sin\theta/\sqrt{N} = 6.7$ m-qg, so the largest deviation (-17 m-qg) sits at about 2.6σ: it is shot noise, not a difference between kits.

2.9 Without qang and with qang

Table 2.3: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
prepare a state with a given probability	compute $\theta = 2 \arccos \sqrt{p_0}$ and apply $R_Y(\theta)$	the R_{qg} gate takes the target reading directly, in qg or m-qg	**convenience**: the physical gate is the same
describe a gate's action on populations	the full unitary, with phases	a doubly stochastic transition matrix set by one number in m-qg, with capacity $1 - qg_S$	**convenience** for populations, with inputs without coherences; **against** for phases, which the description loses (Chapter 3)

action	without qang	with qang	difference
budget the error of a fractional gate	process fidelity	error in m-qg: from -31.7 to -0.4 m-qg for $R_X(324^\circ)$ with the minimum-area representative	the improvement is the **minimum-area principle**'s; m-qg makes it comparable with shot noise
see the area of a pulse on the equator	process fidelity or tomography	$qg_\tau = \sqrt{1 - e^{-\theta^2\sigma^2}}$, which grows with the area even when the m-qg reading does not shift	**convenience**: it is the purity of the tomographic state, with a qg name
generalize to qudits	Gell-Mann matrices	the Cartan subalgebra as a qg identity	**none shown**: no advantage of reading more levels was found

Balance. For gates, qang is a compact way of writing what happens to the populations, and m-qg is the unit in which pulse error and shot noise are compared. The measured gain, the minimum area one, belongs to the physical principle, not to the unit.

2.10 In code

```
import math
from qang import algebra, probability, pulse_area

U = algebra.u3_qang(0.0, 0.0, math.pi)    # the Hadamard as the point (q, φ, λ) = (0, 0, π)
algebra.transition_matrix(0.3)          # [[0.65, 0.35], [0.35, 0.65]]: p0 = 0.65, +300 m
algebra.cartan_values([0.5, 0.3, 0.2])    # the two qg values of a qutrit

pulse_area.minimal_area_angle(3.5, math.pi)   # (0.358, 1): RZZ(3.5) = RZZ(0.358) · (Z ⊗
pulse_area.rx_qg_z_random_amplitude(math.pi, 0.05)  # -0.988 instead of -1
probability.pulse_error_milliqang(math.pi, 0.05)    # 12.3 m-qg of error on an X
```

```
th, s = math.radians(324), 0.05              # R_X(324°) with a 5 % amplitude error
r = math.exp(-(th * s) ** 2 / 2)
probability.dissipative_deficit(0.0, -r * math.sin(th), r * math.cos(th))  # 0.277; with
```

2.11 Summary

- Preparation by polar bias is native in Qiskit, Cirq and PennyLane, with the same convention and a signed variant.
- Every single-qubit gate is an algebraic function of $q = qg_Z$ and two phases; its transition matrix is doubly stochastic, does not depend on the phases and is the probability equation of Chapter 1: a single m-qg number fixes its action on the populations of an input without coherences, and the channel capacity is $1 - qg_S$, with curvature qg_κ: capacity is lost much faster moving away from a pole than from the equator.
- The Cartan subalgebra of $\mathfrak{su}(d)$ generalizes qg_Z to qudits as an identity; no advantage of a multi-level readout was shown.
- In fractional gates the error follows the pulse area; choosing the shortest representative reduces it (from -31.7 to -0.4 m-qg for $R_X(324^\circ)$), and the budget is stated in m-qg, the same scale as the shot noise. The dissipative deficit $qg_\tau \approx |\theta|\sigma$ depends only on the area and shows it also where the m-qg reading does not shift.
- The exposure outside the sector, computed without noise, ranks compilations by how much T1 the filter will let through.

Bibliographic notes

Universality of single-qubit gates plus an entangler is in Barenco et al. (1995) and Nielsen and Chuang (2010); the Cartan decomposition of $SU(2^n)$, in Khaneja and Glaser (2001). The kits the gate is checked against are Qiskit (Javadi-Abhari et al. 2024), PennyLane (Bergholm et al. 2018), Cirq (Cirq Developers 2026) and Amazon Braket (Amazon Web Services 2020). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

3 Reflections, Householder and Grover dynamics

Grover's algorithm is, at heart, a sequence of reflections. Written in qg units, each reflection is an algebraic matrix in the polar bias, the trajectory of the algorithm is a Chebyshev polynomial and the information each iteration brings has a closed form. Everything rests on the base of Chapter 1: the transition probability between two states is a probability like p_0, with its own qang angle and m-qg reading, and the bias of Grover's flag is read on the same scale. The derived quantities of Chapter 1 (Section 1.11) are read here on the algorithm: the geodesic distance qg_L between the states of a reflector, the signal-to-noise qg_{SNR} and the information curvature qg_κ along the trajectory, and the dissipative deficit qg_τ as a measure of the visibility that noise leaves to amplitude estimation. The chapter develops that reading and sets it against what happens when the gates are noisy, where the qg reading loses to the simplest tool.

3.1 Transition matrices without transcendental functions

Chapter 2 showed that every single-qubit gate has the transition matrix

$$T(q) = \begin{pmatrix} \frac{1+q}{2} & \frac{1-q}{2} \\ \frac{1-q}{2} & \frac{1+q}{2} \end{pmatrix},$$

doubly stochastic and rational in q. By construction it is *unistochastic*: it comes from a unitary. Its composition shows the difference between classical probability and quantum amplitude. Two classical channels with biases q_1 and q_2 compose by multiplication: $T(q_1)T(q_2) = T(q_1 q_2)$. Two successive quantum rotations instead give

$$q_{12} = \cos(\theta_1 + \theta_2) = q_1 q_2 - \sqrt{(1-q_1^2)(1-q_2^2)}.$$

The term $-\sqrt{(1-q_1^2)(1-q_2^2)}$ is the interference: the transverse part of each state, invisible in T, adding up in amplitude. With $q_1 = 0.3$ and $q_2 = -0.5$, the classical composition gives -0.15 and the quantum one -0.976: in milli-qang, -150 against -976, an interference of 826 m-qg, close to half the scale. The transition matrix describes each gate on its own, but not the sequence: that needs the phases. On the angular qga scale of Chapter 1 (Section 1.4) quantum composition is simply additive: $\theta_1 = 201.5$ qga, $\theta_2 = 333.3$ qga, the sum is 534.8 qga and $\cos(2\pi \cdot 534.8/1000) = -0.976$. What is not additive is the reading in m-qg.

3.2 The Householder reflector in qg space

A reflection with respect to the state $|\psi\rangle$ is the Householder operator $I - 2|\psi\rangle\langle\psi|$. For $|\psi(q)\rangle = \sqrt{\frac{1+q}{2}}|0\rangle + \sqrt{\frac{1-q}{2}}|1\rangle$ it is

$$H(q) = I - 2|\psi(q)\rangle\langle\psi(q)| = -q\,\sigma_z - \sqrt{1-q^2}\;\sigma_x,$$

a Hermitian, unitary matrix with $H(q)^2 = I$, whose coefficients are the state's own qg values: q on σ_z and its transverse part $\sqrt{1-q^2}$ on σ_x (§166). On the Bloch sphere it is the rotation by π (500 qga) about the axis of the state, together with an overall sign change.

3.3 The Grover diffuser as a qg reflection

The Grover diffuser reflects about the initial state: $D = 2|\psi\rangle\langle\psi| - I$. Comparing with the previous section,

$$D(q) = -H(q) = q\,\sigma_z + \sqrt{1-q^2}\;\sigma_x.$$

In the two-dimensional plane of the algorithm (the "good" state being searched for and the "bad" one, its complement), the oracle is the reflection about the bad axis and the diffuser the reflection about the initial state. The product of two reflections is a rotation by twice the angle between their axes: that is the whole of Grover's dynamics. In the Fubini–Study metric (Section 1.4), each iteration covers the same arc length along a great circle: $2t$ in Hilbert space, that is, $4t$ in the Bloch angle of the flag qubit of the Chebyshev section.

3.4 Transition probability between qg states

The fidelity of two pure states is written with their qg vectors $\mathbf{r} = (qg_X, qg_Y, qg_Z)$:

$$|\langle\psi_1|\psi_2\rangle|^2 = \frac{1 + \mathbf{r}_1 \cdot \mathbf{r}_2}{2}.$$

If the two states share a meridian (the same azimuthal phase), it depends on the polar biases alone:

$$P(q_1 \to q_2) = \frac{1 + q_1 q_2 + \sqrt{(1 - q_1^2)(1 - q_2^2)}}{2}.$$

The condition matters: with different azimuths the polar biases are not enough and the full transverse part is needed. The formula was checked against the inner product on 200 random pairs (§166).

The transition probability in the equation of Chapter 1. P is a probability, so it has its own qang readings. If γ is the angle between the two Bloch vectors, $\mathbf{r}_1 \cdot \mathbf{r}_2 = \cos\gamma$ and $P = \cos^2(\gamma/2)$: it is Born's rule with γ in place of θ. Therefore

$$2\arccos\sqrt{P} = \gamma, \qquad m_P = 2000\,P - 1000 = 1000\;\mathbf{r}_1 \cdot \mathbf{r}_2 \text{ m-qg} :$$

the qang angle of the transition probability is the angle between the states, and its milli-qang reading is a thousand times the dot product of their qg vectors. It is Chapter 1's Wootters statistical distance applied to two states (Section 1.3). With $q_1 = 0.3$ and $q_2 = -0.5$ on the same meridian, $P = 0.838$, $m_P = +676$ m-qg and $\gamma = 0.828$ rad, which is exactly $|\theta_1 - \theta_2| = |\arccos 0.3 - \arccos(-0.5)|$. In qga, $\gamma = 131.8$ qga and $P = \cos^2(\pi\gamma_{qga}/1000)$.

The geodesic distance. The Fubini–Study distance between the two states, the shortest path between them over the state space, is

$$qg_L = \arccos\sqrt{P} = \frac{\gamma}{2},$$

the quantity qg_L of Chapter 1: in the example, 0.414 rad (65.9 qga). On a common meridian it equals the Bhattacharyya angle of the two Z readings, so it is obtained from two single-qubit histograms without tomography. With different azimuths the Z reading underestimates it: two states with $q = 0.45$ and phases 0 and π have the same reading, and hence a zero Bhattacharyya angle in Z, but

$|\langle\psi_1|\psi_2\rangle| = 0.45$ and they are $qg_L = \arccos 0.45 = 1.104$ rad (175.7 qga) apart. The full transverse part is then needed.

3.5 The Grover trajectory as a Chebyshev polynomial

Let $a = \sin^2 t$ be the fraction of good states, read with a flag qubit whose initial bias is $qg_0 = 1 - 2a = \cos 2t$. Each iteration advances the amplitude angle by $2t$, so after k iterations the angle is $(2k+1)t$ and

$$qg_k = \cos\left((2k+1)\,2t\right) = T_{2k+1}(qg_0),$$

with T_m the Chebyshev polynomial of the first kind, $T_m(\cos x) = \cos mx$ (§37). By the probability equation, the success probability is $P_k = (1 - qg_k)/2$, and certain success is the pole at -1000 m-qg. For instance, with $a = 0.1$, $qg_0 = 0.8$ and the successive iterations give $+800$, -352, -997 and -206 m-qg: the bias crosses the pole at $k = 2$ ($P_2 = 0.9986$) and then moves away (Figure 3.1). In qga the step is uniform: $t = 51.2$ qga, the flag's Bloch angle starts at $2t = 102.4$ qga and advances $4t = 204.8$ qga per iteration; at $k = 2$ it reaches 512.1 qga, 12 qga past the pole (500 qga).

The Chebyshev structure yields for free the one-shot Fisher information at depth $m = 2k + 1$. Since $T'_m = mU_{m-1}$ and $1 - T_m^2 = (1 - q^2)U_{m-1}^2$,

$$F_m = \frac{m^2}{1 - qg_0^2} :$$

depth multiplies the qg information of Chapter 1 by m^2, uniformly along the fringe. It is the qg form of the known factor $4m^2$ for the angle (Suzuki et al. 2020). On the milli-qang scale, the resolution on qg_0 with N shots at depth m is $1000\sqrt{1 - qg_0^2}/(m\sqrt{N})$ m-qg: with $a = 0.1$ and 100 shots, 60 m-qg without iterating and 12 m-qg at depth $m = 5$. With respect to the flag's initial angle, the information is m^2 times that of Chapter 1 and does not depend on the state: $m^2 \cdot 4\pi^2 \cdot 10^{-6}$ qga^{-2}, with a resolution of $159.15/(m\sqrt{N})$ qga, that is, 15.9 qga without iterating and 3.2 qga at depth 5. In terms of the searched fraction, $1 - qg_0^2 = 4a(1 - a)$, so $F_m = m^2/[4a(1 - a)]$ and the resolution is $2000\sqrt{a(1-a)}/(m\sqrt{N})$ m-qg.

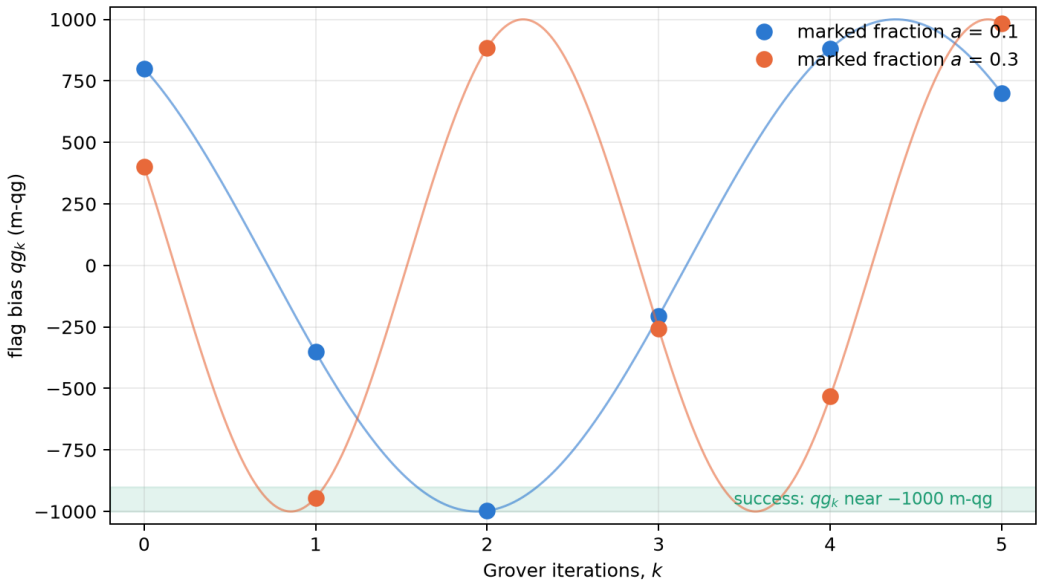


Figure 3.1: The Grover trajectory in milli-qang: the flag bias samples a cosine at the odd points $(2k+1)\,2t$. With $a = 0.1$ success arrives at $k = 2$; with $a = 0.3$ the angle is coarse, success arrives at $k = 1$ and by $k = 2$ the bias has overshot the pole.

3.5.1 The trajectory as a harmonic oscillation

The treatise's formulation reads $qg_k = \cos\big((2k+1)\,2t\big)$ as a sampled harmonic oscillator: the flag runs along a cosine at a constant angular step, $4t$ per iteration. If each iteration lasts Δt, that is an angular frequency of $4t_{qga}/\Delta t$ qga per second, that is, in qHz (Section 1.11); with $a = 0.1$, 204.8 qga per iteration. The picture has a useful consequence. At the pole the cosine is at its turning point, where the velocity $d\,qg/dk$ vanishes: there the success is insensitive, to first order, to an angle error. If the flag ends δ away from the pole, $P = \cos^2(\delta/2) \approx 1 - \delta^2/4$. The 12.1 qga of excess at $k = 2$ cost only 0.14 % of success ($P_2 = 0.9986$); 25 qga would cost 0.6 %, and 50 qga, 2.4 %. That is why the rounding of k^* matters little when a is small and a lot when the step is coarse. The geodesic distance to the solution says it in one figure: after $k = 2$, $qg_L(qg_2, -1) = 6.04$ qga, half the excess of Bloch angle.

The outcome entropy runs through the same cycle: qg_S is 0.469, 0.909, 0.016 and 0.969 bits at $k = 0, \ldots, 3$. It falls almost to zero at the optimum and rises again past it. Like the other derived quantities, it is a function of qg_k and cannot tell on which side of the pole the flag is.

Matter and antimatter: what the analogy allows. The treatise proposes reading the Grover rotation as a conversion channel between matter and antimatter, with the equator as a "neutral meson" and the inversion as charge conjugation. As a physical claim it is **not adopted**: the marked and unmarked subspaces of a search are not particle and antiparticle. What is exact is more modest. The oscillation of a neutral meson between its particle and antiparticle states is also a two-level rotation, and its mixing asymmetry $(N_{\text{unmixed}} - N_{\text{mixed}})/(N_{\text{unmixed}} + N_{\text{mixed}}) = \cos(\Delta m\, t)$, valid when the two widths are equal (Navas et al. 2024), is a qg_Z: it has qang angle $\Delta m\, t$ and a uniform error bar in that angle, like any reading of Chapter 1. Grover and the meson share the kinematics, a cosine, and nothing else: the meson also decays. Appendix G (Appendix G) gives the correct reading of charge conjugation, which flips the expectation value of the Dirac matrix β.

Two orthogonal states, such as the marked one and its complement, are $qg_L = \pi/2$ apart, that is, 250 qga, the maximum distance; an ideal unitary evolution preserves purity ($qg_\tau = 0$) along the whole arc.

3.6 Analytic limits: coarse angles and overshoot

The optimal number of iterations is $k^* = \text{round}(\frac{\pi}{4t} - \frac{1}{2})$; with t in qga, $k^* = \text{round}(125/t_{qga} - \frac{1}{2})$. When the searched fraction $a = M/N$ is small, t is small,

$k^* \approx \frac{\pi}{4}\sqrt{N/M}$ and the rounding hardly matters. When M/N is not negligible, the angle per iteration is *coarse*: the bias jumps in large steps along the curve $\cos((2k+1)2t)$ and can overshoot the pole. Iterating more does not help; it hurts, because the rotation keeps going. The qg reading makes this behaviour visible: qg_k samples a cosine at odd points. With $a = 0.3$, $k^* = 1$ gives -944 m-qg ($P = 0.972$) and one more iteration jumps to $+884$ m-qg ($P = 0.058$), the orange curve of Figure 3.1.

3.6.1 Signal-to-noise and curvature along the trajectory

The formulation proposes reading the trajectory with two derived quantities of Chapter 1: the per-shot signal-to-noise $qg_{SNR} = |qg_k|/\sqrt{1 - qg_k^2}$ and the information curvature $qg_\kappa = 1/[\ln 2\,(1 - qg_k^2)]$. Table 3.1 gives them for $a = 0.1$.

Table 3.1: The Grover trajectory with $a = 0.1$ and its derived quantities.

iteration k	depth m	qg_k (m-qg)	Bloch angle (qga)	success P_k	qg_{SNR}	qg_κ
0	1	$+800.0$	102.4	0.1000	1.33	4.01
1	3	-352.0	307.2	0.6760	0.38	1.65
2 (optimum)	5	-997.1	512.1	0.9986	13.15	250.8
3 (overshoot)	7	-206.4	716.9	0.6032	0.21	1.51

Near the pole both shoot up, and the formulation proposes using them as a pole-proximity detector to stop the iterations before the overshoot. Two precisions. First, they are functions of qg_k: a threshold on qg_{SNR} or on qg_κ is a threshold on $|qg_k|$, and it cannot tell whether the flag reached the pole or already passed it. At $k = 2$ it is 12 qga beyond the pole, and both take the same values as 12 qga before. What says whether another iteration pays is the formula for k^*, not the quantity. Second, reading them costs shots: measuring the flag destroys the state, so every k queried is a separate run. The use they do bring is that of Chapter 1: qg_{SNR} says how many shots are needed to confirm the sign of the final reading, $N = (3/qg_{SNR})^2$, that is, one shot at $k = 2$ and 64 at $k = 1$.

3.6.2 Grover with noise

With noisy gates the picture changes, and it is worth recording as it came out (§63). In searches on 4 to 6 qubits, with each multi-controlled Z compiled to 14, 36 and 84 CNOTs and depolarizing noise on every CNOT:

- **Noise moves the best iteration earlier and eats the speed-up.** With $n = 6$ and $p_2 = 0.005$ the best iteration drops from 4 to 1 and the cost per verified success rises from 6.1 to 27.6 oracle calls, against 32 for classical search.
- **The qg formula fails under gate noise.** If the unmarked outcomes were uniform, each qubit would have $qg_{Z,i} = \pm(NP-1)/(N-1)$, and the signs would give the answer and the magnitude the success probability without knowing it. Depolarizing noise is not uniform after the diffuser, and the label-free estimate was worse than the frequency of the most common outcome in every case.
- **Reading the answer from signs loses.** With $n = 5$, 100 shots find the marked item 66% of the time with the qg signs and 91% with the most frequent bitstring.

The honest balance: in Grover search the outcome histogram is a better estimator than the per-qubit qg reading. qg provides the Chebyshev picture and the information factor, not a better estimator. The author's working version further proposes a regularized variational energy for reflection circuits and a radiation function that would follow P_k during amplification; these are the open proposals of Chapter 1 (Section 1.12), **unverified**.

3.7 Amplitude estimation

Where the Chebyshev structure does pay off is in *amplitude estimation*: using increasing depths to estimate a with fewer oracle queries than direct sampling. With maximum-likelihood estimation (depths 0, 1, 2, 4, ..., 100 shots each), without noise, the error falls as queries$^{-0.93\ldots-1.12}$ against queries$^{-0.5}$ for Monte Carlo, and is 8.4–9.3 times lower at 26 200 queries (§37).

Noise sets limits. With depolarization $p = 0.01$ per iteration, a likelihood that ignores the noise ends worse than Monte Carlo ($6.8\cdot10^{-3}$ against $2.7\cdot10^{-3}$), while one that includes it keeps a 4× advantage. With $p = 0.05$, not even the noise-aware likelihood beats Monte Carlo: beyond a depth of order $1/p$ the iterations carry almost no signal.

The dissipative deficit of Chapter 1 places that limit on the qg scale. In the visibility model of the likelihood, every iteration shortens the flag's qg vector by a factor V; after k iterations the fringe has amplitude V^k and

$$qg_\tau = \sqrt{1 - V^{2k}}.$$

With $V = 1 - p$, $V^k \approx e^{-kp}$: towards $k \approx 1/p$ the fringe falls to $1/e$ of its amplitude and qg_τ exceeds 0.93, so further iterations add depth on an almost mixed state. With a visibility of 0.975 per iteration, the middle of the range fitted on IonQ (next paragraph), qg_τ is 0.63 after 10 iterations and 0.80 after 20. It is the same $1/p$ depth frontier, read as loss of purity.

On IonQ's calibrated noisy simulator (3 qubits, $a = 0.3$, 36 000 queries), the estimate was accurate, with an RMS error of 0.0005–0.0015 and a fitted visibility of 0.965–0.985 per iteration, and 7 to 30 times better than a realistic Monte Carlo, whose depth-zero circuit is biased by the device noise (§39).

3.8 Without qang and with qang

Table 3.2: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
transition probability between two states	$\lvert\langle\psi\rvert\phi\rangle\rvert^2$	$(1 + \mathbf{r}_1 \cdot \mathbf{r}_2)/2$, with qang angle equal to the angle between the states and reading $1000\,\mathbf{r}_1 \cdot \mathbf{r}_2$ m-qg	**none in value**; qang makes it geometric
Grover trajectory	$\sin^2((2k+1)\theta_0)$	$qg_k = T_{2k+1}(qg_0)$, success at -1000 m-qg	**equivalent**; the Chebyshev form shows that the m-qg resolution improves by a factor m with depth
distance between two states	$\arccos\lvert\langle\psi\rvert\phi\rangle\rvert$	$qg_L = \gamma/2$, which on a common meridian comes from two Z histograms	**convenience**; with different azimuths the Z reading is not enough

action	without qang	with qang	difference
decide when to stop iterating	k^* from the formula	threshold on qg_{SNR} or qg_κ	**none**: they are functions of $\lvert qg_k \rvert$ and do not see whether the pole was passed; the formula for k^* does
read the search with noise	bit-string histogram	per-qubit qg reading	**against**: the qg reading loses to the histogram
estimate an amplitude with noise	maximum likelihood on counts	the same, with the noise written into the qg likelihood	**in favour** up to a depth near $1/p$, not beyond

Balance. The chapter's identities are exact, but none gives a number that the usual amplitude does not. The only comparison with real noise in the chapter is unfavourable to the qg reading, and it is written down as such.

3.9 In code

```
from qang import algebra, probability

algebra.householder_qg(0.45)      # -0.45 Z - 0.893 X, a reflection: H @ H = I
algebra.grover_diffuser_qg(0.45)    # +0.45 Z + 0.893 X
P = algebra.transition_probability(0.3, -0.5)      # 0.838, same meridian
algebra.transition_probability(0.3, -0.5, 0.0, 1.2)  # different azimuths
probability.readings_from_p0(P)    # theta = 0.828 rad (the angle between the states), +6
algebra.geodesic_distance(0.3, -0.5)            # qg_L = 0.414 rad (65.9 qga) = gamma / 2

qg_2 = algebra.chebyshev_trajectory(0.8, 2)       # -0.9971 (-997 m-qg, P = 0.9986), a =
probability.quantum_snr(qg_2)                          # 13.15
probability.information_curvature(qg_2)                # 250.8
algebra.grover_fisher_information(0.8, 5)        # 69.4 per shot: 25 times that of m = 1
algebra.geodesic_distance(qg_2, -1.0)           # 0.038 rad (6.04 qga) from the solution
```

3.10 Summary

- The transition matrix of a gate is rational in q, but it does not compose like the gate: the difference between $q_1 q_2$ and $\cos(\theta_1 + \theta_2)$ is the interference.
- The Householder reflector and the Grover diffuser are $\mp(q\,\sigma_z + \sqrt{1-q^2}\,\sigma_x)$.
- The transition probability between two states is $(1+\mathbf{r}_1 \cdot \mathbf{r}_2)/2$; on a common meridian it depends on the polar biases alone. Read with the equation of Chapter 1, its qang angle is the angle between the states and its reading is $1000\,\mathbf{r}_1 \cdot \mathbf{r}_2$ m-qg. The geodesic distance is $qg_L = \gamma/2$.
- The Grover trajectory is $qg_k = T_{2k+1}(qg_0)$, with certain success at -1000 m-qg, and the Fisher information at depth m is $m^2/(1 - qg_0^2)$: the m-qg resolution improves by a factor m.
- qg_{SNR} and qg_κ shoot up near the pole, but as a stopping criterion they do not see whether the flag has already passed it; what they bring is the number of shots to confirm the sign of the final reading.
- The trajectory is a sampled cosine: at the pole it is at its turning point, so an excess of 12 qga costs only 0.14 % of success. The matter–antimatter reading of the rotation is not adopted; what is exact is that a neutral meson's mixing asymmetry, $\cos(\Delta m\, t)$, is a qg_Z.
- With noise, the per-qubit qg reading loses to the histogram in search; in amplitude estimation the advantage holds if the noise enters the likelihood, and ends near depth $1/p$, where the flag's dissipative deficit, $qg_\tau = \sqrt{1 - V^{2k}}$, exceeds 0.93.

Bibliographic notes

The reflector is that of Householder (1958). Grover search is Grover (1996), and amplitude amplification and estimation, Brassard et al. (2002); maximum-likelihood estimation without phase estimation is Suzuki et al. (2020). The transition probability between qg states is the distance of Wootters (1981). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

4 Sampling statistics, Bayesian inference and the Haar distribution

Every number a quantum processor delivers is an estimate made from a finite number of shots. This chapter gathers the statistical tools of the qg unit: how shot noise propagates, which prior is the natural one for an unknown qubit, how to build intervals that do not collapse near the poles, how to correct the bias of quadratic quantities and how many shots to plan for. All of them rest on the base of Chapter 1: what is estimated is always the probability p_0, and since the milli-qang is linear in it, every interval and every shot plan is stated directly in m-qg. It closes with the mathematical foundation that explains why qg_Z is the canonical read-out.

4.1 Propagating shot noise

With N shots and k_0 outcomes 0, the natural estimator is $\widehat{qg}_Z = 2k_0/N - 1$. It is unbiased and its variance is exactly

$$\mathrm{Var}(\widehat{qg}_Z) = \frac{1 - qg_Z^2}{N} = \frac{\sin^2\theta}{N},$$

which attains the Cramér–Rao bound for the parameter qg_Z: no unbiased estimator does better with the same information. In milli-qang it is the resolution $\sigma_m = 1000\sqrt{1-qg_Z^2}/\sqrt{N}$ of Chapter 1 (Figure 1.3). Propagated to the angle by the delta method, the variance becomes constant, $\mathrm{Var}(\hat{\theta}) \approx 1/N$, as seen in Chapter 1 (§5). The two statements are compatible: qg_Z is the *curved* coordinate, with Fisher information $1/(1-qg_Z^2)$ per shot, and θ the *flat* one, with information 1 (Chapter 23).

The delta method is a first-order approximation and fails where the derivative vanishes. Near a pole, $\widehat{qg}_Z$ is exactly ± 1 in a sizeable fraction of experiments, and the estimated standard error $\sqrt{(1-\widehat{qg}_Z^2)/N}$ becomes zero: the interval collapses exactly where the state is most extreme.

4.1.1 The azimuthal phase

The phase is estimated from the two transverse bases, $\hat{\varphi} = \operatorname{atan2}(\widehat{qg}_Y, \widehat{qg}_X)$ (Section 1.4). With N shots in each basis, the delta method gives

$$\operatorname{Var}(\hat{\varphi}) \approx \frac{1 - \frac{1}{2}\sin^2\theta\,\sin^2 2\varphi}{N\sin^2\theta} \;\leq\; \frac{1}{N\sin^2\theta}.$$

The bound on the right, which is the one the author's working version uses, is exact on the axes (φ a multiple of $90°$) and near the poles, but on the equator at $\varphi = 45°$ it doubles the true variance. The essential point stands: the phase becomes unobservable towards a pole, where $\sin\theta \to 0$. On the angular scale of Chapter 1, on the equator and on an axis, the deviation is $159.15/\sqrt{N}$ qga, the same as for the polar angle. The transverse reading $\widehat{qg}_X = \sin\theta\cos\varphi$, by contrast, does not diverge: its variance is $(1 - qg_X^2)/N \leq 1/N$, for the same reason as that of $\widehat{qg}_Z$.

4.2 The Haar distribution in qg space

What is known about a qubit before it is measured? If its state is "any state", the natural distribution is the Haar measure, the only one invariant under rotations. In qg units it takes a remarkable form: qg_Z **is uniform on** $[-1, 1]$. This is Archimedes' hat-box theorem: the area of a band of the sphere between two latitudes is proportional to the difference of heights, not of angles. It was checked on 200 000 Haar states (§15).

The statistical consequence is direct. A uniform prior on qg_Z is equivalent to $p_0 \sim \text{Beta}(1, 1)$, and with k_0 zeros in N shots the posterior is

$$p_0 \mid k_0 \;\sim\; \text{Beta}(k_0 + 1,\; N - k_0 + 1),$$

from which a credible interval for $qg_Z = 2p_0 - 1$ follows with a beta function alone (`bayes_qg_estimate`, implemented without external dependencies and checked against `scipy.stats.beta`).

The probability equation of Chapter 1 explains why this prior is so simple. Since $m = 2000\,p_0 - 1000$ is linear, a distribution uniform in qg_Z is uniform in milliqang and uniform in p_0: the three statements are one. In the angle, by contrast, the same measure has density $\frac{1}{2}\sin\theta$ and piles up at the equator (Figure 4.1). The "uninformative" prior depends on the scale it is written in, and Haar's is flat exactly on the scale that is measured.

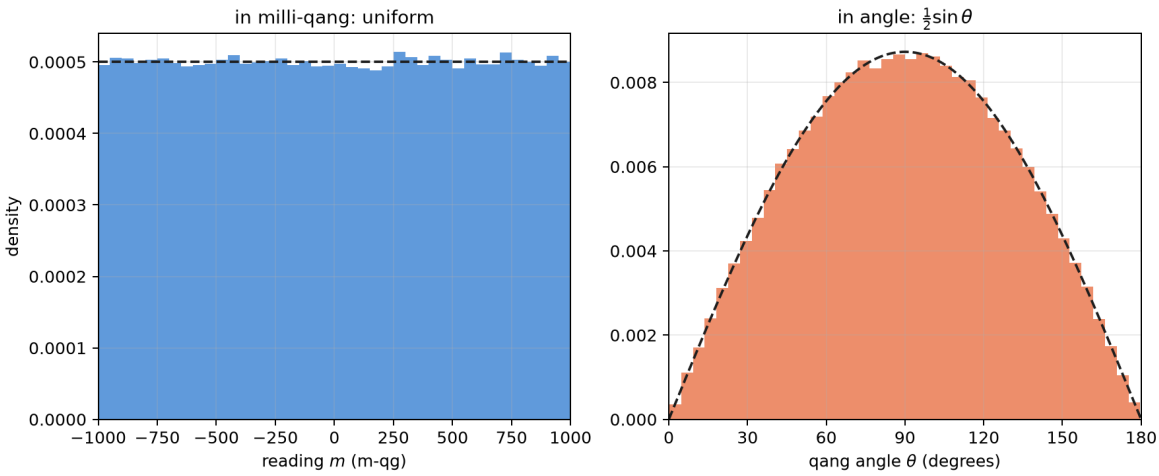


Figure 4.1: Two hundred thousand Haar states read in milli-qang (left) and in angle (right). The dashed line is the theoretical density: $1/2000$ per m-qg and $\frac{1}{2}\sin\theta$ in the angle.

4.3 Exact intervals near the poles

The comparison with 50 shots shows where the choice of interval matters (§15):

Table 4.1: Coverage of 95% intervals with 50 shots.

case	delta method	Wilson	Bayes–Haar
near a pole ($\theta = 0.08$): coverage of the 95% interval	7.6%	92.4%	92.4%
Haar-random states: coverage	89.8%	95.1%	94.9%

Near the pole the delta method covers the true value in fewer than one experiment in ten; Wilson and the Haar posterior recover the nominal coverage. On Haar states the posterior mean has a 4% lower mean squared error than the raw frequency.

Figure 4.2 shows the cause in milli-qang. With 50 shots, the width of the delta-method interval drops to zero at the ends of the scale: with 50 zeros in 50 shots it gives $[+1000, +1000]$ m-qg, an interval with no width. With 48 zeros

it gives $[+811, +1029]$ m-qg, which runs off the top of the scale. Wilson gives $[+731, +978]$ and Bayes–Haar $+885$ $[+731, +975]$ m-qg: widths that do not collapse and stay inside $[-1000, +1000]$.

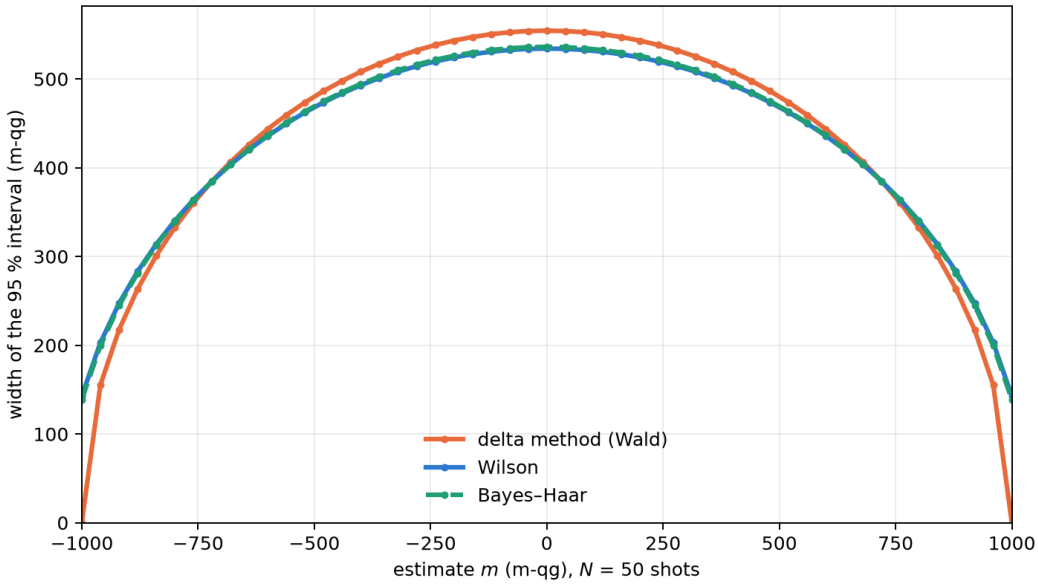


Figure 4.2: Width of the 95 % interval with 50 shots, in milli-qang, against the estimate. The delta method collapses at the poles; Wilson and Bayes–Haar do not.

The honest reading of the table has two parts. The Haar prior repairs the collapse of the delta method and has a clean geometric justification; but the Wilson interval, standard in statistics, achieves the same coverage, and the gain over good frequentist practice is in the point estimate and is modest. **The qg unit does not reduce the number of shots a given precision requires**; what it brings is making the cost explicit and choosing the interval well.

4.4 Second-order unbiased estimators

Quadratic quantities have a systematic bias. The plug-in estimator $\widehat{qg}^{\,2}$ overestimates qg^2 by $(1 - qg^2)/N$, because shot noise always adds variance. The correction is exact:

$$\widehat{qg^2} = \frac{N\,\widehat{qg}^{\,2} - 1}{N - 1}, \qquad \mathbb{E}\big[\widehat{qg^2}\big] = qg^2,$$

valid for $N \geq 2$ (`qg2_unbiased`). It can go negative when $qg \approx 0$, the price of being unbiased. With $N = 20$ and $qg = 0.6$, the plug-in estimator averages 0.390 and the corrected one 0.358, against the true value 0.36.

It matters because the Bloch radius of Chapter 5, $r^2 = qg_X^2 + qg_Y^2 + qg_Z^2$, is a sum of squares: without the correction each axis adds its bias, and an entangled state looks purer than it is (§97). The same correction enters the reading of the gravitational redshift with two clocks (§126, Chapter 8).

4.5 Planning the shots

The inverse relation, how many shots a given error needs, is the one used when preparing an experiment. `plan_shots` turns a target standard error of the filtered qg_Z into shots and estimated processor seconds, taking into account the fraction of shots the filter will keep (§143). For a target of 0.02 (20 m-qg) on two circuits per device it planned 3000–3400 shots, and over 30 repetitions the error of the worst qubit was 0.021–0.023 (21–23 m-qg): the plan held. For reference, Chapter 1's formula without the filter gives $N = (1000 \sin\theta/\delta)^2 = 2500$ shots for 20 m-qg on the equator; the plan asks for more because it discounts the shots the filter will discard. The simple rule is $N_{\text{run}} = N/(1 - P_{\text{discard}})$: with a 20 % discard, the 2500 useful shots ask for 3125.

The per-qubit prediction, however, was 1.4–1.6 times optimistic on average, and its analysis is instructive: the qubits that missed were those near a pole ($|qg| > 0.8$, with ratios of 2.4–3.1), whose qg_Z the noise pulls away from ± 1 and whose real variance is larger than the ideal one. The plan uses the worst qubit and therefore holds; the safe option, assuming no value of qg, is to plan for the worst case, $qg = 0$. The author's working version attributes that excess to the curvature $\kappa_{qg} = 1000\,|qg_Z|$ of Chapter 1, largest at the poles: a random pulse error shifts the reading towards the equator in proportion to κ_{qg}, and there the binomial variance is larger. It is an interpretation consistent with Chapter 1, but it was **not tested** on the data of §143.

4.6 Few-shot tomography and classical shadows

The full Bloch vector is (qg_X, qg_Y, qg_Z), one per measurement axis. With 10 shots per axis, six estimators were compared on 2000 states (§35). The qg version applies the Haar prior axis by axis, $\widehat{r}_a = (2k_a - N)/(N + 2)$, and projects onto the ball:

Table 4.2: Mean squared Bloch-vector error with 10 shots per axis.

ensemble	linear inversion	LI + projection	max. likelihood	**qg-Haar**	Jeffreys
Haar pure	0.203	0.157	0.150	0.165	0.159
uniform mixed	0.247	0.211	0.210	**0.185**	0.196
pure near a pole	0.199	0.159	0.109	0.163	0.159
strongly mixed ($\lvert r\rvert < 0.3$)	0.298	0.293	0.293	**0.208**	0.245

Linear inversion gives unphysical states in 52–88% of the runs on pure states. On mixed states the qg-Haar estimator is the best closed form (12–29% below maximum likelihood); on pure states its shrinkage towards the centre is a mistake, and it loses to maximum likelihood, by up to 50% near a pole. The right prior depends on which states are expected.

Would *classical shadows* (random Pauli measurements) do better than measuring the bases directly? For the quantities qang uses, no (§62). With exact per-shot variances on 6-qubit states, shadows need 3–41 times more shots for the same error: 3.0–4.3× for all the qg_Z, 5.8–8.1× for the register-mean weight witness, 9–12.6× for all the ZZ correlators and 41× for the filtered imbalance. They only match direct measurement when the target is all one- and two-qubit quantities at once. And the shot-level filter does not exist with shadows: only $3^{-6} = 0.14\%$ of the random shots have every qubit measured in Z.

4.7 The foundation: Riesz–Fréchet

A conceptual question remains: why qg_Z and not some other function of the populations? The operators of an n-qubit register form a finite-dimensional Hilbert space under the Hilbert–Schmidt product $\langle A, B\rangle = \mathrm{Tr}(A^\dagger B)$, and density matrices live in it. The Riesz–Fréchet representation theorem says that every continuous linear functional f on a Hilbert space has the form $f(x) = \langle y_f, x\rangle$ for a unique y_f; in finite dimension continuity is automatic and the theorem reduces to linear algebra (Appendix F).

The application is immediate. $\rho \mapsto qg_Z(\rho) = \mathrm{Tr}(\rho\,\sigma_z)$ is linear in ρ, so it has a unique representative, σ_z. Every linear read-out of a state (an expectation value, a population, a per-shot average) has that form, and qg_Z is the one whose

representative is the Z observable. By contrast, $qg_S(\rho) = H(\langle 0|\rho|0\rangle)$ is strictly concave in ρ: no operator represents it. Two properties already seen follow from this structural fact: linear functionals have unbiased sample-mean estimators and the entropy does not (Jensen's inequality), and the identity $qg_S = H((1 + qg_Z)/2)$ shows that qg_S carries no information about the state beyond what the representative of qg_Z already extracts.

The milli-qang inherits this property: $m(\rho) = 1000\ \mathrm{Tr}(\rho\,\sigma_z)$ is linear too, with representative $1000\,\sigma_z$, which is why its estimator is the shot average, without bias. It is the underlying reason Chapter 1 defined the unit on probability and not on angle: $\theta = \arccos(qg_Z)$ is not linear in ρ, has no representative and its direct estimator is biased. The same holds for the angle on the qga scale of Chapter 1: it is the same θ multiplied by $1000/2\pi$, so the angular scale changes the unit of the angle but does not make it linear in ρ.

Open proposals. The author's working version adds three proposals to this chapter, which the book records as **unverified**: watching the Fisher information $qg_F = 4000\ \mathrm{Var}(\sigma_z)$ so that an ansatz does not approach a global Haar 2-design (that the Haar measure over all of $SU(2^n)$ produces barren plateaus is a known result (McClean et al. 2018); that qg_F serves to watch it is not); shot-noise propagation formulas for the regularized energy and the radiation function of Chapter 1, which depend on penalties not yet defined; and that the natural gradient in qg canonically diagonalizes the Fubini–Study tensor, which is true only for the coordinates (θ, φ) of one qubit (Section 1.4).

4.8 The derived quantities in sampling

The treatise's formulation brings the derived quantities of Chapter 1 (Section 1.11) to sampling. All of them are estimated from the same counts, and each one inherits a statistical property of this chapter.

Two Fisher informations. A shot carries about the polar bias the information $F_{qg} = 1/(1-qg_Z^2)$, which diverges at the poles, and about the qang angle $F_\theta = 1$, constant. They are the same information in two coordinates, and the information curvature is the first one in bits: $qg_\kappa = F_{qg}/\ln 2$.

Signal-to-noise and curvature when all shots agree. The plug-in estimators of qg_{SNR} and qg_κ divide by $1-\widehat{qg}_Z^{\,2}$ and become infinite when the N shots coincide, which is exactly the case of Section 4.3. The Bayes–Haar posterior mean keeps them finite: with 50 out of 50, $qg_Z = 50/52 = 0.962$, $qg_{SNR} = 3.50$ and $qg_\kappa = 19.1$.

The geodesic distance between two samples. With N_1 and N_2 shots, each qang angle has variance $1/N$, so

$$\widehat{qg}_L = \frac{|\hat{\theta}_1 - \hat{\theta}_2|}{2}, \qquad \operatorname{Var}\left(\widehat{qg}_L\right) \approx \frac{1}{4}\left(\frac{1}{N_1} + \frac{1}{N_2}\right),$$

wherever the two states are. It is half of Cohen's h, the classical effect size for comparing two proportions, so the test is not qang's: qang gives it the name of a distance. With 750 and 700 zeros in 1000 shots each, $\widehat{qg}_L = 0.056$ rad (8.9 qga) with standard error 0.0224: a difference of 2.5σ. What Chapter 1 said holds: on a common meridian it is the Fubini–Study distance; with different azimuths, a lower bound.

The dissipative deficit from counts. $qg_\tau = \sqrt{1 - r^2}$ is a quadratic quantity, and it inherits the bias of Section 4.4 with the opposite sign. With N shots per axis, the plug-in radius satisfies

$$\mathbb{E}\left[\hat{r}^{\,2}\right] = r^2 + \frac{3 - r^2}{N},$$

so shot noise makes the state look purer than it is. A pure state measured with 100 shots per axis averages $\hat{r}^{\,2} = 1.02$, outside the sphere. A state with $r = 0.9$ ($qg_\tau = 0.436$) averages $\hat{r}^{\,2} = 0.832$ with the plug-in estimator and 0.810 with the corrected one, which is the true value. The square root, however, is not unbiased even with the correction: $\widehat{qg}_\tau$ averages 0.376 with the plug-in radius and 0.404 with the corrected one, both below 0.436, because the root is concave and is clipped at zero. The formulation proposes $\sqrt{\max(0, 1 - \hat{r}^{\,2})}$ as an unbiased estimator; the corrected one halves the bias, it does not remove it.

4.9 Without qang and with qang

Table 4.3: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
confidence interval near a pole	delta method: coverage below 10%	Bayes–Haar, $\mathrm{Beta}(1,1)$: uniform in m-qg, nominal coverage	**in favour of the method, not only of qang**: the Wilson interval, without qang, also recovers it; qang justifies the prior (Haar is uniform in m-qg)
quadratic quantities (radii, qg^2)	square the estimate: biased upward	$(N\widehat{qg}^{\,2}-1)/(N-1)$	the **standard unbiased estimator**; qang packages it and applies it in every chapter
how many shots to take	binomial variance of the state	$N=(1000\sin\theta/\delta)^2$ to resolve δ m-qg	**equivalent**; the m-qg formula is the one used for planning
compare two readings	Cohen's h	$qg_L=h/2$, with variance $\frac{1}{4}(1/N_1+1/N_2)$ over the whole circle	**none**: it is the same test, with the name of a distance
estimate the lack of purity	$1-\mathrm{Tr}\,\rho^2$ with tomography	qg_τ with unbiased squares	the bias of $\widehat{r}^{\,2}$ disappears; that of $\widehat{qg}_\tau$ is halved, **not removed**
measure qang quantities	classical shadows	direct Z measurement	3–41 times fewer shots, but it is a **measurement choice**, not an effect of the unit
justify the reading	—	Riesz–Fréchet: qg_Z is the canonical linear reading	**in favour**: explains why qg_Z is averaged and not the angle

Balance. The chapter's statistics are the classical statistics of proportions. qang contributes the choice of coordinate (uniform under Haar, linear by Riesz–Fréchet) that says which prior to use and what to average; the coverage improvement exists without qang too.

4.10 In code

```
import math
from qang import probability, proportions, statistics

statistics.bayes_qg_estimate(48, 50)    # Haar posterior: qg_Z 0.885, interval [0.731, 0
statistics.wilson_qg_estimate(48, 50)    # the frequentist alternative
statistics.qg2_unbiased(16, 20)       # qg^2 without the (1 - qg^2)/N bias
statistics.radius2_estimate([(77, 100), (50, 100), (86, 100)])  # unbiased r^2 and its e
proportions.cohens_h(0.75, 0.70) / 2    # qg_L between two readings: 0.056 rad

probability.analyze(48, 50).milliqang   # (920.0, 730.8, 977.9) m-qg with the Wilson int
probability.shots_for_resolution(20, math.pi / 2)  # 2500 shots for 20 m-qg on the equato
```

4.11 Summary

- $\widehat{qg}_Z$ is unbiased and attains the Cramér–Rao bound; the delta method collapses near the poles, where its m-qg interval loses its width or runs off the scale.
- Under the Haar measure qg_Z is uniform on $[-1, 1]$, that is, uniform in milliqang and in p_0; the matching prior is Beta$(1, 1)$ and gives exact intervals with a beta function.
- Wilson and Bayes–Haar recover the nominal coverage where the delta method covers less than 10%; neither reduces the shots required.
- Quadratic quantities need the correction $(N\widehat{qg}^{\,2} - 1)/(N - 1)$; without it, Bloch radii are overestimated.
- Of the derived quantities, qg_L between two samples is half of Cohen's h, with variance $\frac{1}{4}(1/N_1 + 1/N_2)$; qg_τ estimated from counts stays biased low even with the quadratic correction.
- Planning on the worst qubit meets the target; the qubits near a pole are the ones the noise pulls away most.
- For qang's quantities, measuring directly in Z costs 3–41 times fewer shots than classical shadows.

- qg_Z, and with it the milli-qang, is the canonical linear read-out by the Riesz–Fréchet theorem; qg_S and the angle are not linear and have no representative.

Bibliographic notes

The Cramér–Rao bound is in Cramér (1946) and Rao (1945), and the invariant prior, in Jeffreys (1946). The score interval is Wilson (1927). For sampling the Haar measure, see Mezzadri (2006); classical shadows are Huang et al. (2020). The Riesz–Fréchet representation theorem is Riesz (1907) and Fréchet (1907). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part II

II. Local geometry, quantum optics and geometric phases

5 The local Bloch radius as a measure of purity and entanglement

Part I treated the qubit as a point on the surface of the sphere: a pure state, a probability p_0 and its four readings. A qubit that belongs to a register is almost never pure: it is mixed by noise or by its entanglement with the others, and it sits *inside* the sphere. This chapter adds to the base of Chapter 1 the missing quantity, the **local radius** r, and shows that the probability equation still holds with a single change: the polar bias is no longer $\cos\theta$ but $r\cos\theta$. From the radius come an entanglement measure, a way to separate the direction of the state from its purity, and wave–particle duality, all read in qg units.

5.1 The probability equation with a radius

For a qubit in a mixed state, its Bloch vector (qg_X, qg_Y, qg_Z) has length $r \le 1$, and Born's rule in the Z basis reads

$$p_0 = \frac{1 + qg_Z}{2} = \frac{1 + r\cos\theta}{2}.$$

It is the probability equation of Chapter 1 with the radius inside: the readings $qg_Z = 2p_0 - 1$ and $m = 2000\,p_0 - 1000$ m-qg do not change, but they no longer fix the angle, because the same p_0 can come from a tilted pure state or from a mixed one closer to the pole. Each axis has its own probability, $p_0^{(a)}$ in the basis $a \in \{X, Y, Z\}$, and the squared radius is the sum of the three readings squared (§97):

$$r^2 = qg_X^2 + qg_Y^2 + qg_Z^2 = \frac{m_X^2 + m_Y^2 + m_Z^2}{10^6} = 2\,\mathrm{Tr}\,\rho^2 - 1.$$

The radius is linear in the purity $\mathrm{Tr}\,\rho^2$: it is 1 on the surface (pure state) and 0 at the centre (maximally mixed state). The sphere of radius r has area $4\pi r^2$, and

the **surface deficit** $4\pi(1-r^2)$ is the area the qubit has lost; the library works with $1-r^2$, which carries the same information.

5.2 The deficit as an entanglement measure

If the whole register is in a pure state, all the radius a qubit lacks is due to its entanglement with the rest. Then $1-r_q^2 = 4\det\rho_q$ is the *one-tangle* of qubit q, which for two qubits is the squared concurrence (Coffman et al. 2000), and its mean over the qubits is the global entanglement of Meyer and Wallach (2002):

$$Q = \frac{1}{n}\sum_q \left(1-r_q^2\right).$$

Two examples set the scale. In the Bell state each qubit sits at the centre of its sphere, $r=0$, and $Q=1$. In the 3-qubit W state each has $r^2 = 1/9$ and $Q = 0.889$. If the global state is mixed, the deficit mixes entanglement and noise and cannot separate them: the radius measures entanglement only when the register is known, or guaranteed, to be pure. Reading r^2 takes three measurement settings for the whole register (every qubit in X, in Y and in Z), not three per qubit.

5.3 The radius under gates and algorithms

A single-qubit gate rotates the sphere and preserves the radius, also for mixed qubits. SWAP exchanges the radii of two qubits. CNOT, CZ, iSWAP, Toffoli and Fredkin can take a qubit of a product input to the centre ($r=0$, from inputs of the type $|{+}0\rangle$ or $|{+}{+}\rangle$) and can also bring it back to the surface; on computational-basis states they create no deficit, because they act as classical permutations (§97). In algorithms, the radius of the read-out qubits says where the answer is:

Table 5.1: The radius of the read-out qubits in standard algorithms (§97).

algorithm	r^2 of the read-out qubits	reading
Bernstein–Vazirani, constant Deutsch–Jozsa	1 on every qubit	product output: the local qg_Z is the whole answer
Simon ($s=110$)	0 on every qubit	the answer lies entirely in the correlations

algorithm	r^2 of the read-out qubits	reading
Grover ($n = 5$), $k = 0 \dots 4$	$1 \to 0.82 \to 0.66 \to 0.84 \to 0.998$	entangles and disentangles; back to 1 at the optimum
counting ($N = 16$, $M = 3$)	$(1 - 2M/N)^2 = 0.39$	deficit $4(M/N)(1 - M/N) = 0.61$
QAOA, ring of 4, $p = 1$	0.25	the cost lives in ZZ

5.4 Definite weight: the radius from the Z basis

In a state of definite Hamming weight every local qg_X and qg_Y is zero, so $r_q = |qg_Z^{(q)}|$ and the radius is read from Z shots alone. With the probability equation, the deficit becomes

$$1 - r^2 = 1 - \left(\frac{m}{1000}\right)^2 = 4\,p_0\,(1 - p_0)\,:$$

four times the binomial variance of the outcome, the same $p_0(1 - p_0)$ that governs the shot noise of Chapter 1. Figure 5.1 shows it on the milli-qang scale. The deficit is zero at the poles, largest on the equator and very sensitive near a pole: a qubit at 990 m-qg has a deficit of 0.02, and one at 800 m-qg, which is what 10 % of misplaced shots on a basis state gives, already reads 0.36.

This is where the symmetry filter (Chapters 20–21) and the radius meet. Under equal T1 on every qubit the filtered state is the noiseless pure state, so the filtered deficit is an exact entanglement measure. The predictions were committed before running (§98); on RBS circuits of 4 and 6 qubits:

Table 5.2: The deficit as an entanglement measure under noise, with and without the qg filter (§98).

noise	exact error of the deficit, with qang	without qang	filter MSE lower
equal T1	$< 3 \cdot 10^{-15}$	0.061–0.370	30 of 30
unequal T1	0.013–0.045	0.070–0.342	30 of 30
dephasing 0.01	0.031–0.307	0.080–0.352	24 of 30

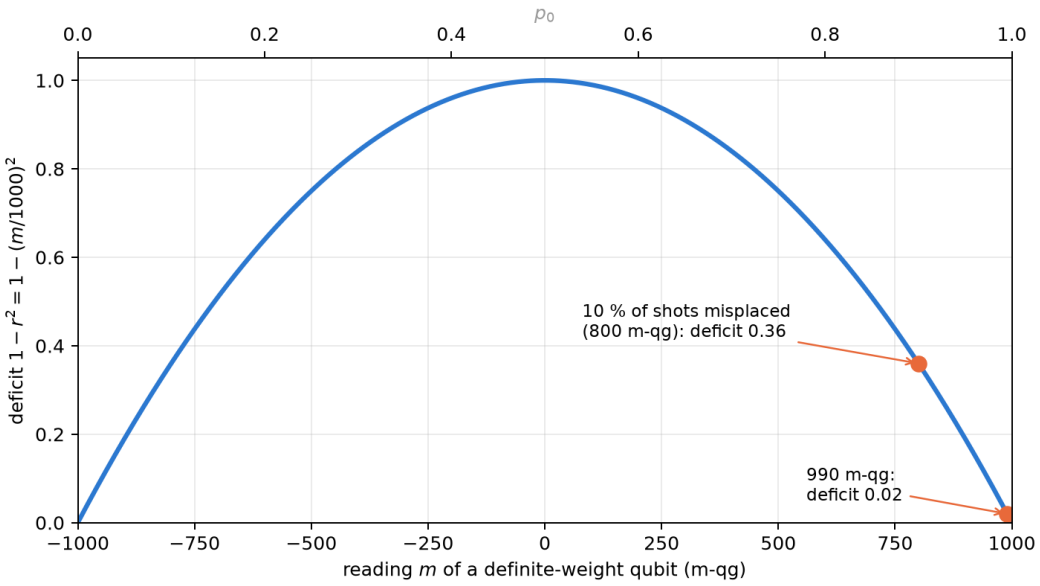


Figure 5.1: The deficit of a definite-weight qubit against its milli-qang reading (bottom) and its probability (top). Near a pole, a few misplaced shots produce a large deficit.

On product states the filtered deficit is exactly 0 under all three noise models, while the raw read-out shows 0.60–0.95 on every excited qubit from decay alone: false entanglement. One prediction failed: the raw error was expected to be smaller at half filling, and it was larger (0.177 against 0.074), because the first-order cancellation holds only at exactly $p = 1/2$.

On noise models of real devices (three from IBM and two from IonQ, §100), the filter cut the error of the deficit 3.2–4.0× where the deficit is large, but it did not work as a test of "no entanglement": on a product state it removed only 15–33 % of the false deficit. The cause is the one in Figure 5.1: errors inside the sector leave 10–21 % of misplaced shots, and near the pole that already reads as 0.36. A small deficit needs an echo calibration before it is interpreted.

5.5 Direction and radius apart

The polar bias $qg_Z = r\cos\theta$ mixes the direction of the vector with its length. The usual reading $\theta = \arccos(qg_Z)$ drifts towards the equator whenever noise shortens the vector. With the radius the two separate (§102):

$$\theta = \operatorname{atan2}\left(\sqrt{qg_X^2 + qg_Y^2},\, qg_Z\right) = \arccos\frac{qg_Z}{r}.$$

A state at 30° shortened to $r = 0.8$ by depolarizing noise has $qg_Z = 0.693$: $\arccos(qg_Z)$ reads 46.1°, and the separated angle reads the true 30°. Figure 5.2 shows the effect for several radii: the naive angle flattens towards 90°, and the separated one does not depend on r.

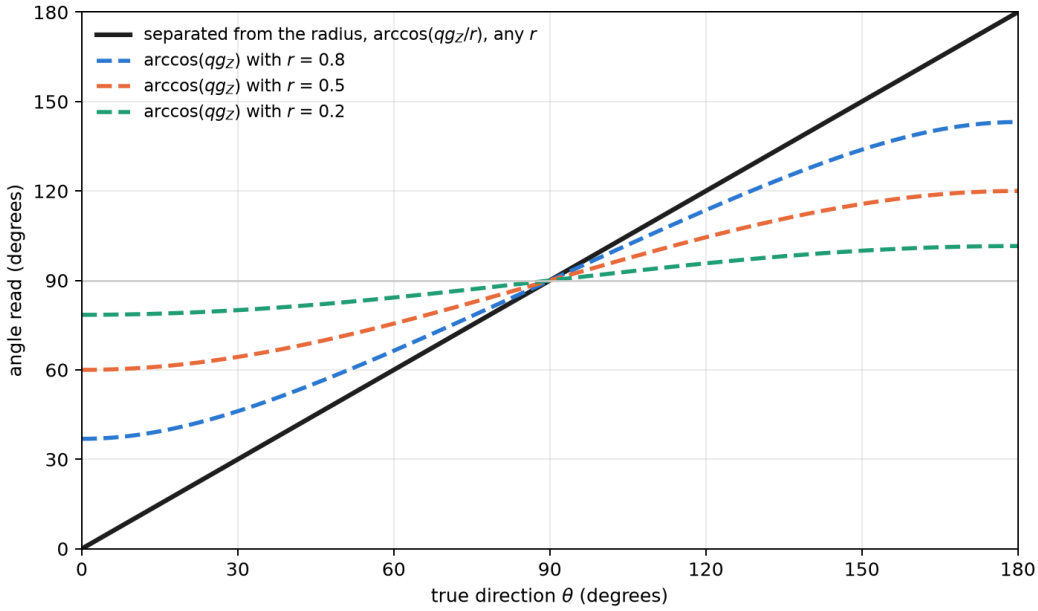


Figure 5.2: The angle read against the true direction. The dashed curves are $\arccos(qg_Z)$ for shortened vectors; the straight line is the angle separated from the radius, exact for any r under depolarizing noise.

The separation has a cost and does not always win. With the same total number of shots split over three bases:

- **Depolarizing:** the separated angle is exact and $\arccos(qg_Z)$ errs by up to 0.70 rad near the poles; with shots, the root-mean-square error drops 5.8×.
- **Dephasing:** this is where it loses. Dephasing does not shorten qg_Z, so $\arccos(qg_Z)$ is exact, and the separated angle is pulled 0.03–0.12 rad towards the poles.
- **No noise:** measuring three bases costs 1.6× in root-mean-square error.
- **Devices:** on the IBM models the vector shrinks almost uniformly and the separated angle is 5–9× more accurate. On IonQ, $\arccos(qg_Z)$ won at first,

because the compiler had removed pairs of identical CX gates; in the native gate set the separated angle wins there too, 8.7–9.0× (§103).

The practical rule: if the measured radius is close to 1, use $\arccos(qg_Z)$ with all shots in Z; if it is clearly below 1 and the noise is not pure dephasing, use the separated angle. On the angular scale of Chapter 1 (Section 1.4), the example above is 83.3 qga true against 128.1 qga read.

The azimuthal phase does not depend on the radius. Since $qg_X = r\sin\theta\cos\varphi$ and $qg_Y = r\sin\theta\sin\varphi$, the ratio cancels r and $\sin\theta$, and $\varphi = \operatorname{atan2}(qg_Y, qg_X)$ is the same for the shortened vector. That holds whenever the noise shortens both transverse components equally: depolarizing noise, dephasing and T1 relaxation preserve the phase. A channel that treats X and Y differently shifts it: with a bit flip of probability 0.15, which keeps qg_X and shortens qg_Y, a phase of 0.70 rad reads 0.53 rad.

Open proposals. The author's working version adds a purity penalty $\lambda_r(1-r^2)$ to the regularized energy of Chapter 1, damps the azimuthal part of the radiation function with r and, against noise-induced barren plateaus (Wang et al. 2021), proposes watching the radius layer by layer, discarding layers with $r < 0.3$ and compiling with minimum curvature κ_{qg}, with a 35 to 60% reduction of exposure time. That noise produces plateaus is a known result; the three proposals and those figures **have not been verified**.

5.6 Wave–particle duality and entropic uncertainty

The radius also orders a classical result of the foundations. The **predictability** $P = |qg_Z| = |m|/1000$ measures how much is known about the path (particle) and the **visibility** $V = \sqrt{qg_X^2 + qg_Y^2}$ measures the contrast of the interference (wave). Their relation is exactly the radius (Greenberger and Yasin 1988; Englert 1996):

$$P^2 + V^2 = r^2 \leq 1.$$

The duality inequality is, in qg units, the statement that the state lies inside the sphere, and it is an equality for pure states (Figure 5.3). A state at 60° with $r = 0.9$ has $P = 0.45$, $V = 0.779$ and $P^2 + V^2 = 0.81$. It was checked on 2000 random states with a largest deviation of $2.2 \cdot 10^{-16}$ (§127).

The entropic version uses Chapter 1's spread qg_S in two bases. The Maassen–Uffink relation, with the correction of Berta and co-workers for mixed states, says

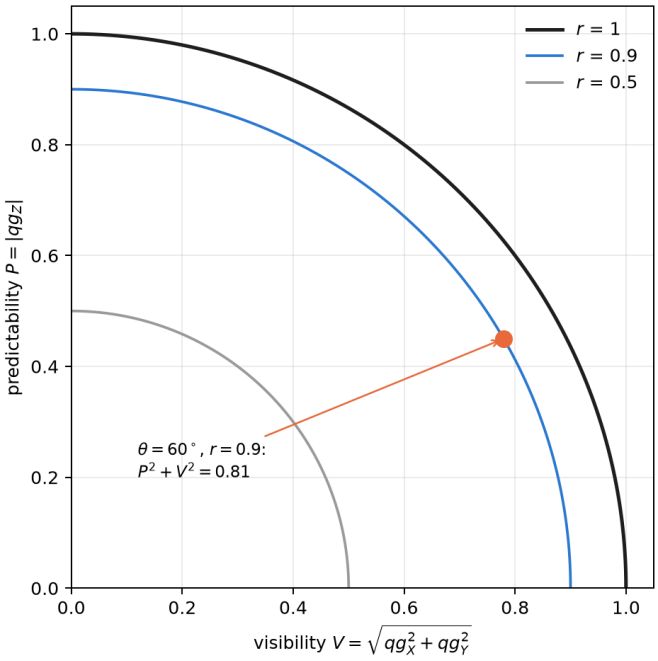


Figure 5.3: Wave–particle duality in qg units: each state is a point of the quarter disc, at distance r from the origin. Pure states lie on the arc $r = 1$.

$$qg_S(Z) + qg_S(X) \;\geq\; 1 + S(\rho), \qquad S(\rho) = H\left(\frac{1+r}{2}\right).$$

For the state above the sum is 1.349 and the bound 1.286. The bound is tight: the state $|0\rangle$ meets it with equality. That gives a practical warning. With plain (uncorrected) entropy estimates, a pure state seems to violate the relation in 93 % of runs at 100 shots, because the estimated entropy is biased downwards; with the Miller–Madow correction, in 37 %. A test of the bound needs corrected estimators or intervals, like those of Chapter 4.

5.7 Estimating the radius from shots

Since r^2 is a sum of squares, it inherits the bias of the quadratic estimators of Chapter 4 (Section 4.4): each axis adds $(1-qg^2)/N$, and an entangled state looks purer than it is. The exact correction $(N\widehat{qg}^{\,2}-1)/(N-1)$ is applied axis by axis (`radius2_estimate`), and in a definite-weight state the Z axis is enough.

5.8 Without qang and with qang

Table 5.3: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
measure local purity	tomography in three bases and $\mathrm{Tr}\,\rho^2$	$r^2 = qg_X^2 + qg_Y^2 + qg_Z^2$; at definite weight, $1 - r^2 = 4p_0(1-p_0)$ from **a single basis**	**in favour**: one basis instead of three, thanks to the probability equation together with the definite weight
measure the entanglement of a pure register	one-tangle and Meyer–Wallach by definition	the deficit $1 - r^2$, averaged	**none in value**; with a mixed global state the deficit mixes entanglement and noise, like the one-tangle

action	without qang	with qang	difference
direction of a noisy state	$\theta = \arccos qg_Z$	$\theta = \arccos(qg_Z/r)$	**in favour** with depolarizing noise; **against** with pure dephasing
wave–particle duality	visibility and predictability separately	$P^2 + V^2 = r^2$	**reformulation**: the same relation

Balance. The practical difference is reading purity at definite weight with a single measurement basis. The rest are known identities written on the qang scale, and the direction–radius split wins or loses depending on the noise.

5.9 In code

```
import math
import numpy as np
from qang import complementarity, formulation, statistics

bell = np.array([1, 0, 0, 1]) / math.sqrt(2)
formulation.radius_profile(bell, 2)       # [0.0, 0.0]: each qubit at the centre
w = np.zeros(8); w[[1, 2, 4]] = 1 / math.sqrt(3)
formulation.meyer_wallach(w, 3)          # 0.889: global entanglement of the W state
formulation.weight_sector_radius([0.99, 0.8])   # [0.0199, 0.36]: deficit from qg_Z (990

formulation.direction_from_qg(0.4, 0.0, 0.6928) # (0.524, 0.0, 0.8): θ = 30° separated f
complementarity.duality_sum(0.779, 0.0, 0.45)    # 0.81 = P² + V² = r²
complementarity.entropic_sum(0.45, 0.779)        # 1.35 ≥ bound 1.29
```

The figures are generated by `tools/make_cap05_figures.py`.

5.10 Summary

- For a mixed state, Chapter 1's probability equation still holds with $qg_Z = r\cos\theta$; the local radius $r^2 = qg_X^2 + qg_Y^2 + qg_Z^2 = 2\,\mathrm{Tr}\,\rho^2 - 1$ is the missing reading.

- If the register is pure, the deficit $1-r^2$ is the one-tangle and its mean is the Meyer–Wallach measure; with a mixed global state, it mixes entanglement and noise.
- At definite weight, $1-r^2 = 4p_0(1-p_0)$ is read in Z alone, and with the filter under equal T1 it is exact; near a pole a few misplaced shots inflate it, and on devices a small deficit needs calibration.
- $\theta = \arccos(qg_Z/r)$ separates direction from purity: it wins under depolarizing noise and loses under pure dephasing.
- Wave–particle duality is $P^2 + V^2 = r^2$, and entropic uncertainty bounds $qg_S(Z) + qg_S(X)$ by $1 + S(\rho)$.

Bibliographic notes

The one-tangle and the concurrence are from Coffman et al. (2000), and global entanglement, from Meyer and Wallach (2002). Wave–particle duality is Greenberger and Yasin (1988) and Englert (1996); entropic uncertainty, Maassen and Uffink (1988) and, with quantum memory, Berta et al. (2010). The Deutsch–Jozsa, Bernstein–Vazirani, Simon and Grover algorithms are Deutsch and Jozsa (1992), Bernstein and Vazirani (1997), Simon (1997) and Grover (1996). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

6 Optical polarization and Stokes parameters in qg units

The polarization of a photon is a qubit, and classical polarization optics is, unchanged, the physics of a qubit written with other names. This chapter reads that optics with the base of Chapter 1. The key piece is that a polarizer is a two-outcome measurement, pass or block, and its transmission is exactly the probability p_0 of the qang equation. From there follow Malus's law, the Stokes parameters, the degree of polarization and photon-counting polarimetry, all in qg and in milli-qang (§83).

6.1 Malus's law as a probability equation

Light linearly polarized at an angle α goes through a polarizer set at ϑ. With $\Delta = \vartheta - \alpha$, the transmitted fraction is Malus's law, which for a single photon is a probability:

$$p_0 = \cos^2 \Delta .$$

The comparison with Born's rule in Chapter 1, $p_0 = \cos^2(\theta/2)$, is immediate: **the qang angle is twice the physical angle**, $\theta = 2\Delta$, and the milli-qang reading is $m = 1000 \cos 2\Delta$. That is why two orthogonal polarizations, H and V, 90° apart in the laboratory, are opposite points of the sphere, 180° apart (Figure 6.1). An analyser turned 30° from the light transmits $p_0 = 0.75$: it is the state at $\theta = 60^\circ$ and $+500$ m-qg of Chapter 1's qang circle (Figure 1.1).

For light with any polarization, pure or partial, the transmission through a linear polarizer at ϑ is

$$p_0 = \frac{1 + q g_Z \cos 2\vartheta + q g_X \sin 2\vartheta}{2} = \frac{1 + \mathbf{r} \cdot \mathbf{a}}{2}, \qquad \mathbf{a} = (\sin 2\vartheta,\ 0,\ \cos 2\vartheta),$$

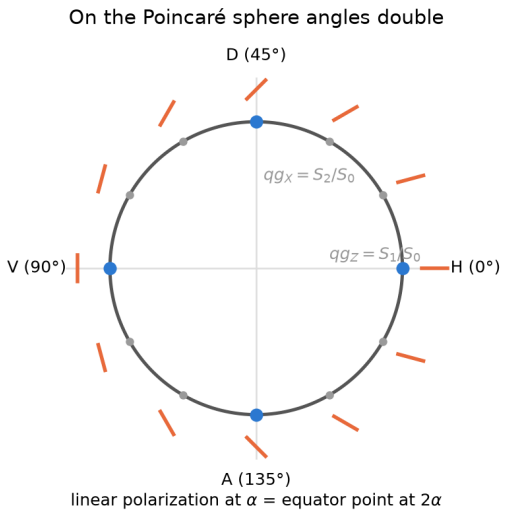


Figure 6.1: Angle doubling. Each orange stroke is the physical direction of linear polarization at α; its point on the equator of the sphere is at 2α. H and V, perpendicular in the laboratory, sit at opposite poles of the qg_Z axis.

with $\mathbf{r} = (qg_X, qg_Y, qg_Z)$ the qg vector of the light and $\mathbf{a}$ that of the state the analyser lets through. It is Chapter 3's transition probability between two qg vectors (Section 3.4), and its milli-qang reading is $m = 1000\,\mathbf{r}\cdot\mathbf{a}$. Figure 6.2 shows four cases: H light ($\cos^2\vartheta$), D light (the same curve shifted $45°$), unpolarized light (always $\frac{1}{2}$, that is, 0 m-qg) and partially polarized light with $qg_Z = 0.6$ and $qg_X = 0.3$. The last transmits 0.80 at $0°$, a maximum of 0.835 at $13.3°$ and a minimum of 0.165 at $103.3°$; the difference between maximum and minimum, 0.671, is its degree of polarization.

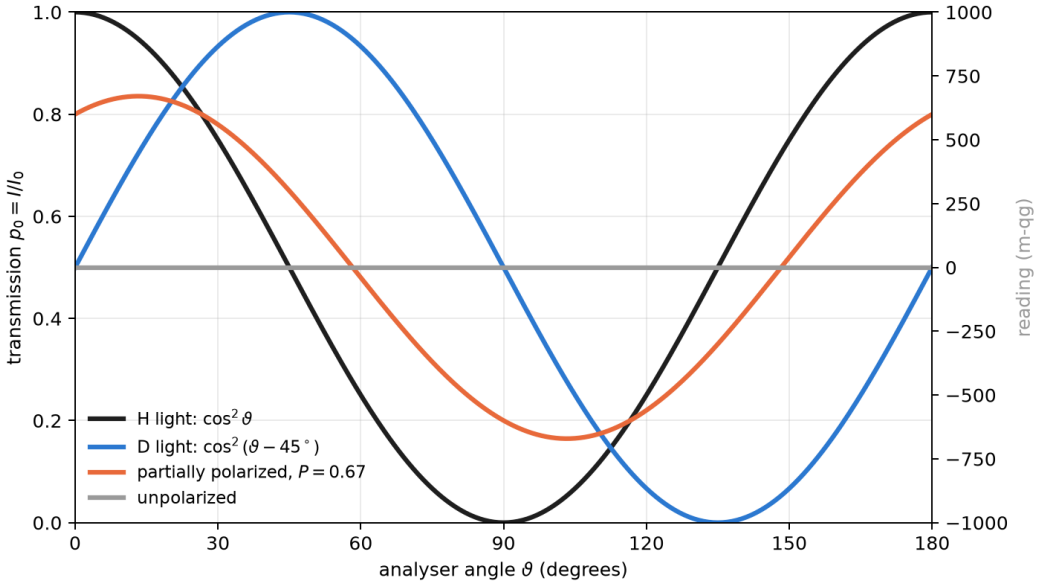


Figure 6.2: Malus's law as a probability equation: transmission of a linear polarizer against its angle, with its milli-qang reading on the right.

6.2 The Stokes parameters are qg values

The Poincaré sphere is the Bloch sphere. With $|H\rangle = |0\rangle$ and $|V\rangle = |1\rangle$, the normalized Stokes parameters are the three qg values, each the difference of the two probabilities of an analyser:

Table 6.1: The normalized Stokes parameters as qg values (§83).

qg value	Stokes	analyser	reading
qg_Z	S_1/S_0	H / V	$p_H - p_V$

qg value	Stokes	analyser	reading
qg_X	S_2/S_0	D / A, at $\pm 45^\circ$	$p_D - p_A$
qg_Y	S_3/S_0	R / L, circular	$p_R - p_L$

The convention for right circular light varies between texts; here $R = (|H\rangle + i|V\rangle)/\sqrt{2}$, so $S_3 > 0$ means $qg_Y > 0$ (the library takes the opposite convention with `circular_sign=-1`). Since each parameter is $2p_0 - 1$ in its own basis, each is expressed in milli-qang with the equation of Chapter 1.

The polarization ellipse is read from the angles of the qg vector, halved: its orientation is $\psi = \frac{1}{2}\,\mathrm{atan2}(qg_X, qg_Z)$ and its ellipticity $\chi = \frac{1}{2}\arcsin(qg_Y/P)$. For the partial light of the example, $\psi = 13.3^\circ$: the angle where Figure 6.2 has its maximum.

6.3 Degree of polarization and purity

Partially polarized light is a mixed state. Its degree of polarization is the length of the qg vector, the radius of Chapter 5 (Section 5.1):

$$P = \sqrt{qg_X^2 + qg_Y^2 + qg_Z^2} = r, \qquad \mathrm{Tr}\,\rho^2 = \frac{1 + P^2}{2}.$$

Unpolarized light sits at the centre of the sphere. An ideal depolarizer has Mueller matrix $\mathrm{diag}(1, d, d, d)$: it shortens the vector by a factor d and does not change its direction. It is the case in which Chapter 5 separates direction and radius (Section 5.5): with $d = 0.7$, the light of the example goes from $P = 0.671$ to $P = 0.470$, and the orientation of its ellipse stays at 13.3°. The polar bias, however, falls from 0.60 to 0.42, and the naive reading $\arccos(qg_Z)$ goes from 53.1° to 65.2°: a false widening of 12° due to noise. The separated angle $\arccos(qg_Z/P) = 26.57^\circ$ (73.8 qga on the scale of Chapter 1) does not change.

Complementarity. The predictability in the H/V basis, $P_Z = |qg_Z|$, and the fringe visibility in the transverse bases, $V = \sqrt{qg_X^2 + qg_Y^2}$, satisfy exactly $P_Z^2 + V^2 = P^2 \le 1$. The Greenberger–Yasin–Englert inequality is here the condition that the Stokes vector lie inside the Poincaré sphere, and the Maassen–Uffink entropic bound corrected for purity (Berta et al. 2010) reads $qg_S(Z) + qg_S(Y) \ge 1 + H\big((1+P)/2\big)$: it is the duality of Chapter 5 (Section 5.6) read in polarimetry.

6.4 Lossless elements: the Mueller matrix is the gate rule

A Mueller matrix acts on (S_0, S_1, S_2, S_3); in the order (I, Z, X, Y) it is the Pauli transfer matrix. For a lossless element, with a unitary Jones matrix U, it is exactly qang's gate rule:

$$qg'_P = qg_{U^\dagger P U}.$$

Wave plates are single-qubit gates, and Chapter 2 applies as is (Section 2.3): a half-wave plate with its axis at 22.5° takes H to D ($qg_X = 1$), because it turns the physical polarization by 45° and the point on the sphere by 90°; a quarter-wave plate at 45° takes H to a circular polarization ($qg_Y = -1$ with the convention above). The polarizer, by contrast, is not unitary: it is the measurement. All these identities are checked numerically in `tests/test_polarization.py`.

6.5 Photon-counting polarimetry

With weak or single-photon sources, the Stokes parameters are estimated by counting: n_H, n_V in the H/V basis, n_D, n_A in the diagonal one and n_R, n_L in the circular one. Each pair gives a qg value with its interval, just like a qubit measurement (Chapter 4). With 1000 photons per basis and counts $(800, 200)$, $(650, 350)$ and $(500, 500)$:

$$qg_Z = 0.599\,[0.548,\, 0.647], \quad qg_X = 0.299\,[0.240,\, 0.358], \quad qg_Y = 0.000\,[-0.062,\, 0.062],$$

95 % Bayes–Haar intervals. In milli-qang, qg_Z is $+599$ m-qg with a resolution of $\sigma_m = 1000 \sin\theta/\sqrt{N} = 25$ m-qg.

The degree of polarization needs more care. P is the square root of a sum of squares, and each square inherits the bias $(1-qg^2)/N$ of Chapter 4 (Section 4.4): with few photons, unpolarized light looks polarized. Figure 6.3 quantifies it with a simulation from this chapter (4000 repetitions per point). With unpolarized light, the plug-in estimator averages $\hat{P} \approx 0.50$ with 10 photons per basis, 0.16 with 100 and 0.05 with 1000. Using the unbiased squares brings it down about three times (≈ 0.18, 0.055 and 0.017), but does not remove it: the square root and the clipping at zero leave a residual bias. With light of $P = 0.5$, the plug-in estimator overestimates at few photons and the unbiased one slightly underestimates; from about 200 photons per basis both match the true value.

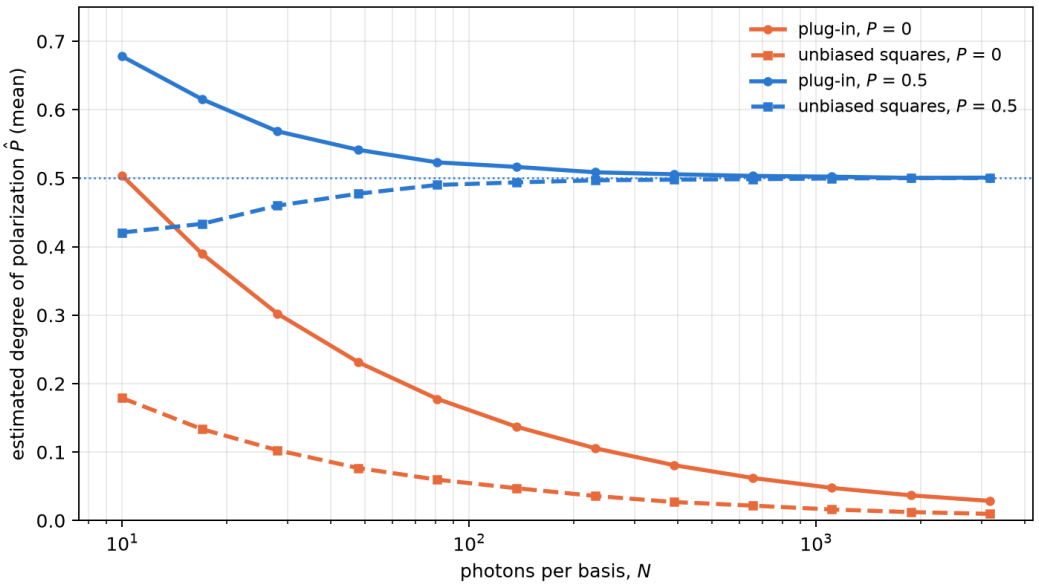


Figure 6.3: Bias of the degree of polarization estimated by counting, against the photons per basis. Solid lines: plug-in estimator; dashed: with the unbiased squares of Chapter 4. The horizontal dotted lines mark the true value.

The practical consequence: to claim that a source is polarized from few photons, one needs an interval for P or a threshold calibrated with unpolarized light, not the point value.

Open proposals. The author's working version carries the regularized energy of Chapter 1 over to programmable optics, with a penalty $\lambda_r(1 - P^2)$ against depolarization in the interferometer, and a radiation function whose azimuthal part is damped by P. As in the earlier chapters, the penalties are not defined and neither proposal has been verified.

6.6 Without qang and with qang

Table 6.2: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
transmission of a polarizer	Malus's law, $I = I_0 \cos^2 \Delta$	$p_0 = \cos^2 \Delta$ with $\theta = 2\Delta$; reading $1000\,\mathbf{r} \cdot \mathbf{a}$ m-qg	**none in value**; qang ties it to the transition probability of Chapter 3
describe the polarization	normalized Stokes parameters S_1, S_2, S_3	they are qg_Z, qg_X, qg_Y	**identity**: the same numbers under another name
degree of polarization with few photons	$\sqrt{S_1^2 + S_2^2 + S_3^2}$: biased upward	the unbiased squares of Chapter 4	**in favour**: the bias drops about three times, without vanishing
wave plates	Mueller matrices	the gate rule $qg'_P = qg_{U^\dagger PU}$	**reformulation**

Balance. Polarization optics already works with qg values under another name (Stokes). The measurable difference is statistical: correcting the squares reduces the bias of the degree of polarization with few photons.

6.7 In code

```
import math
import numpy as np
from qang import polarization as pz

light = {"Z": 0.6, "X": 0.3, "Y": 0.0}
pz.malus_intensity(light, 0.0)                 # 0.80: transmission at 0°
pz.malus_intensity(light, math.radians(13.28))   # 0.835: the maximum, at the orientatio

H = np.array([1, 0])
pz.jones_to_qg(pz.half_wave_plate(math.radians(22.5)) @ H)  # {'X': 1.0, ...}: H → D
pz.jones_to_qg(pz.quarter_wave_plate(math.radians(45)) @ H) # {'Y': -1.0, ...}: H → cir

c = pz.qg_from_counts(800, 200, 650, 350, 500, 500, method="bayes")
c["Z"], c["X"], c["P"]                     # 0.6, 0.3, 0.671 and the intervals in c["interv
```

The figures are generated by `tools/make_cap06_figures.py`.

6.8 Summary

- A polarizer is a two-outcome measurement: its transmission is $p_0 = \cos^2\Delta$, Chapter 1's probability equation with the qang angle equal to twice the physical one, $\theta = 2\Delta$.
- For any light, Malus's law is the transition probability $(1 + \mathbf{r}\cdot\mathbf{a})/2$ between the light and the analyser; its reading is $1000\,\mathbf{r}\cdot\mathbf{a}$ m-qg.
- The normalized Stokes parameters are qg_Z, qg_X and qg_Y; the ellipse is read from half the angles of the qg vector.
- The degree of polarization is the radius of Chapter 5; a depolarizer shortens it without turning the ellipse.
- Wave plates are gates: the lossless Mueller matrix is the rule $qg'_P = qg_{U^\dagger P U}$.
- From photon counts each parameter has its interval; the degree of polarization is biased upwards at few photons, and the unbiased squares cut that bias about three times without removing it.

Bibliographic notes

The polarization parameters are those of Stokes (1852). Malus's law, the Poincaré sphere and the Mueller calculus follow the classical treatment of Born et al. (1999). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

7 Geometric phases and Stokes' theorem on the sphere

A phase is not a probability. Chapter 2 showed that the phases of a gate do not appear in its transition matrix: a Z measurement does not see them. This chapter deals with a phase that does carry physical meaning, the **geometric phase** a qubit picks up when carried around a closed loop, and shows that the base of Chapter 1 reaches it twice. First, its value: for a loop at constant height, the phase in turns *is* the probability p_0. Then, its measurement: an interferometer turns it into a probability, and the qang angle of that probability is the phase (§84).

7.1 The Berry connection and curvature

If a state $|\psi(\mathbf{R})\rangle$ depends on parameters $\mathbf{R}$ and is carried slowly around a closed path, besides the dynamical phase it picks up a phase that depends only on the geometry of the path. Its density is the **Berry connection**, $\mathbf{A} = i\langle\psi|\nabla\psi\rangle$, and its curl is the **curvature**. For a qubit, with the state described by its qg vector on the sphere, the curvature is a monopole of strength $\frac{1}{2}$ at the centre: on the unit sphere its normal component is $\frac{1}{2}$ everywhere.

7.2 Stokes on the sphere

Stokes' theorem turns the line integral of the connection around the loop C into the flux of the curvature through the surface it encloses. Since the curvature is uniform, the flux is $\frac{1}{2}$ times the area, that is, the solid angle $\Omega(C)$:

$$\gamma(C) = -\frac{\Omega(C)}{2} \pmod{2\pi}.$$

The phase is gauge invariant (it does not depend on the phase chosen for each state along the path) and changes sign when the loop is run backwards.

7.3 Loops at constant height: the phase is the probability

A loop around the Z axis at constant polar bias qg_Z encloses the cap above it, of solid angle $\Omega = 2\pi(1 - qg_Z)$. In turns, the library's qg_Φ unit,

$$\varphi = \frac{\gamma}{2\pi} = \frac{qg_Z - 1}{2} = -p_1 \equiv p_0 \pmod 1, \qquad \varphi = \frac{1000 + m}{2000} \text{ turns.}$$

The geometric phase of a loop at constant height, measured in turns, is the probability p_0 of its height, and with the milli-qang it is a linear function of the reading. The reason is the same that made the Haar distribution uniform in Chapter 4 (Section 4.2): by Archimedes' theorem, the fraction of the sphere in the cap above the height qg_Z is $(1 - qg_Z)/2 = p_1$. The solid angle is $4\pi p_1$ and the phase $-2\pi p_1$.

Some values: a loop at $+500$ m-qg ($p_0 = 0.75$) picks up 0.75 turns, that is, $-\frac{1}{4}$ of a turn; one on the equator picks up half a turn, π; one hugging the pole $|0\rangle$, nothing. Figure 7.1 shows the line together with the three independent computations of the next section.

On the angular scale of Chapter 1 (Section 1.4), which measures in thousandths of a turn, the same sentence reads with no conversion: **the phase of a loop at constant height is $1000\,p_0$ qga** (modulo 1000). The loop at $+500$ m-qg picks up 750 qga, the equator's 500 qga and the octant of the last section -125 qga (875 qga). The reading in m-qg and the phase in qga are then two linear scales of the same probability: $\varphi_{qga} = (1000 + m)/2$.

The equator deserves a note. There the phase is π, and $+\pi$ and $-\pi$ are the same angle. A numerical computation can return either one depending on rounding, and a comparison that ignores this fails although the result is right. That is exactly what happened to one of the library's tests, now fixed to compare on the circle.

7.4 Three independent computations

The library computes the phase in three ways that share no code and checks that they agree (14 tests):

1. **Overlap product** (Pancharatnam): $\gamma = -\arg \prod_k \langle \psi_k | \psi_{k+1} \rangle$ over the discretized loop, gauge invariant.
2. **Solid angle**, as a fan of spherical triangles (Van Oosterom and Strackee's formula). It agrees with the overlap product to 10^{-10}.

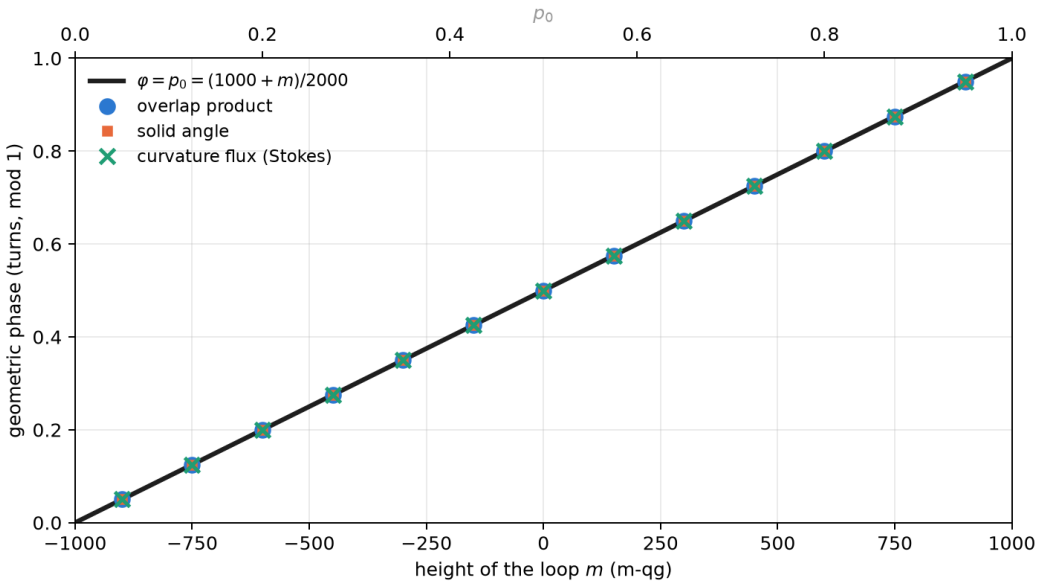


Figure 7.1: The geometric phase of a loop at constant height against its milli-qang reading (bottom) and its probability (top). The line is $\varphi = p_0$; the symbols are the three independent computations.

3. **Curvature flux** through the cap, integrated numerically. It agrees with the line integral to 10^{-5}: Stokes' theorem checked in numbers.

7.5 Measuring the phase: the interferometer and the probability equation

A Z measurement does not see the phase. Measuring it takes an **interferometer**: a control qubit in $|+\rangle$ decides whether the system runs the loop, and at the end the control is measured in the X basis (Hadamard test). With the dynamical phase cancelled, for instance by an echo, the probability of reading 0 on the control is

$$p_0 = \frac{1 + \cos\gamma}{2} = \cos^2\frac{\gamma}{2} \,:$$

Born's rule of Chapter 1 with γ in place of θ. Therefore **the qang angle of the measured probability is the geometric phase**, $\hat{\gamma} = 2\arccos\sqrt{\hat{p}_0}$, and its milli-qang reading is $m = 1000\cos\gamma$.

Everything the base says about that angle applies to the phase. In particular, being the arcsine transformation (Section 1.9), its uncertainty is $1/\sqrt{N}$ **for any value of the phase**. Figure 7.2 checks it with a simulation from this chapter: with 1000 shots per point, the standard deviation of $\hat{\gamma}$ stays around 32 mrad (about $1.8°$) from $10°$ to $170°$. At 0 and π it is zero because the outcome is deterministic, always 0 or always 1, and there the exact intervals of Chapter 4 are needed.

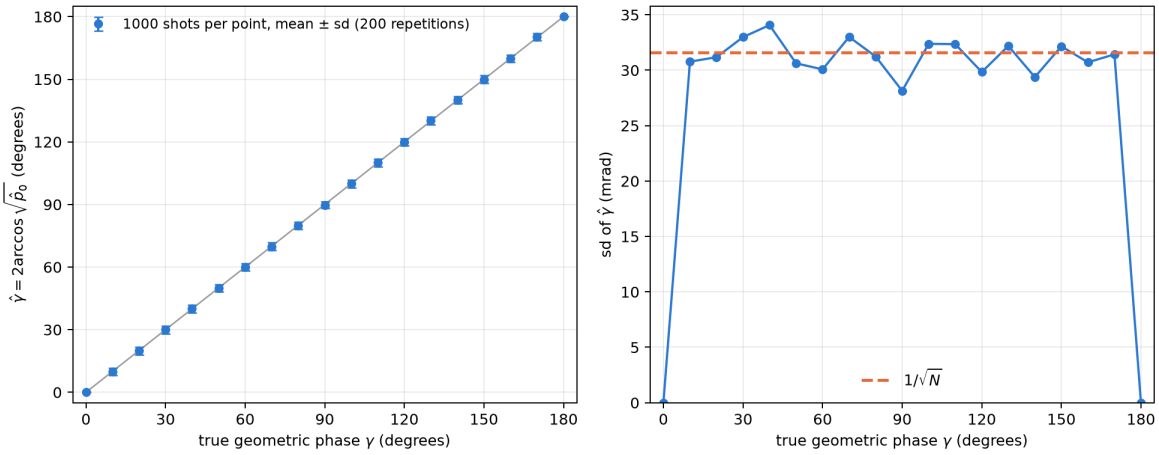


Figure 7.2: Interferometric measurement of the geometric phase with the probability equation. Left: the phase estimated as the qang angle of the measured probability, against the true one. Right: its standard deviation, equal to $1/\sqrt{N}$ over the whole range.

There is a limit, the same as in Chapter 1. Since $\cos\gamma$ is even, a single probability gives $|\gamma|$ in $[0, \pi]$, not its sign. The remedy is the branch sign (Section 1.6): measure the other quadrature too, shifting the reference phase by a quarter turn, which gives $\sin\gamma$. The two probabilities together fix the phase on the whole circle.

7.6 Pancharatnam's phase and the octant

For polarized light the same phase is Pancharatnam's on the Poincaré sphere, and Chapter 6 already set up the coordinates (Section 6.1). The circuit H → D → R → H runs around an octant of the sphere: three quarter arcs of great circles enclosing one eighth of its area, $\Omega = 4\pi/8 = \pi/2$. The phase is $\gamma = -\pi/4$, one eighth of a turn backwards (0.875 turns mod 1), and the library reproduces it from the three qg vectors.

7.7 Without qang and with qang

Table 7.1: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
geometric phase of a constant-height loop	$-\frac{1}{2}$ of the solid angle, $\Omega = 2\pi(1-\cos\theta)$	the phase in turns is $p_0 = (1000+m)/2000$	**equivalent**; with qang it is read from a single measured probability
measure the phase in an interferometer	invert $\cos^2(\gamma/2)$ with the delta method	the qang angle is the phase, with uncertainty $1/\sqrt{N}$ for any value	**same result** in the interior: the delta method also gives $1/\sqrt{N}$; qang gives it without a derivative. At 0 and π both give zero width and the intervals of Chapter 4 are needed
the sign of the phase	second quadrature	second quadrature	**none**: the limit is the same

Balance. The Berry phase is the same with or without qang; what the scale adds is reading it from a single probability and having its error bar without deriving anything. It adds no precision.

7.8 In code

```
import math
import numpy as np
from qang import geometric

geometric.cone_phase_turns(0.5)        # 0.75 turns = p0 at +500 m-qg
geometric.cone_phase_turns(0.0)        # 0.5: the equator, half a turn
```

```
# discretized loop at qg_Z = 0.8: the phase from overlaps, in turns
r = math.sqrt(1 - 0.8 ** 2)
loop = [np.array([r * math.cos(a), r * math.sin(a), 0.8]) for a in np.linspace(0, 2 * math
geometric.geometric_phase_qg(loop).phi         # 0.9 (= p0 = 0.9)

octant = [np.array([0, 0, 1.0]), np.array([1.0, 0, 0]), np.array([0, 1.0, 0])]  # H, D, R
geometric.geometric_phase_qg(octant).phi        # 0.875: -1/8 of a turn
```

The figures are generated by `tools/make_cap07_figures.py`.

7.9 Summary

- The geometric phase of a closed loop is minus half the solid angle it encloses: Stokes' theorem with the Berry curvature, a monopole of strength $\frac{1}{2}$.
- For a loop at constant height, the phase in turns is p_0 (mod 1), that is, $(1000 + m)/2000$: the same Archimedean reason that makes the Haar measure uniform in m-qg.
- Three independent computations (overlaps, solid angle, curvature flux) agree to 10^{-10} and 10^{-5}.
- An interferometer turns the phase into the probability $\cos^2(\gamma/2)$; its qang angle is the phase, with uncertainty $1/\sqrt{N}$ for any value, and the sign needs a second quadrature.
- The polarization circuit H → D → R → H encloses an octant and picks up $-\pi/4$.

Bibliographic notes

The geometric phase is that of Berry (1984), anticipated in optics by Pancharatnam (1956). Its extension to non-adiabatic cyclic evolutions is Aharonov and Anandan (1987), and the general formulation with the Pancharatnam connection, Samuel and Bhandari (1988). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

8 Fisher information, clocks, gravity and metrology

Metrology is the discipline of measuring a parameter with the smallest possible uncertainty, and its central tool is the Fisher information: how much each shot teaches about the parameter. This chapter writes it on the base of Chapter 1. The probability equation $p_0 = \cos^2(\theta/2)$ fixes how much information each reading carries, and from it follow the resolution in milli-qang, the operating point of an interferometer, the comparison of two atomic clocks one metre apart in height, the advantage (and the limits) of entanglement, and the thermometry of a qubit.

8.1 From the base: the information in one shot

A shot in the Z basis is a coin with probability p_0. Its Fisher information about the angle is

$$F_\theta = \frac{(dp_0/d\theta)^2}{p_0(1-p_0)} = \frac{\sin^2\theta/4}{\sin^2\theta/4} = 1,$$

constant over the whole sphere; about the reading $qg_Z = \cos\theta$, by contrast, it is $F_{qg} = 1/(1-qg_Z^2)$, growing towards the poles. **The qang angle is the flat coordinate of the information and qg_Z the curved one** (Figure 8.1). The two descriptions are the same physics: the $1/N$ variance of the angle in Chapter 1 is the Cramér–Rao bound in the flat coordinate (§167). On the angular scale of Chapter 1 (Section 1.4), the same constant information is $F_{qga} = 4\pi^2 \cdot 10^{-6}$ qga^{-2} per shot: it is the Foundation's Q-Fisher, and its Cramér–Rao bound is the deviation of $159.15/\sqrt{N}$ qga over the whole circle.

Expressed in milli-qang, the bound gives the resolution of a reading:

$$\sigma_m = \frac{1000\,|\sin\theta|}{\sqrt{N}} \text{ m-qg}, \qquad N = \left(\frac{1000\sin\theta}{\delta}\right)^2 \text{ shots to resolve } \delta \text{ m-qg}.$$

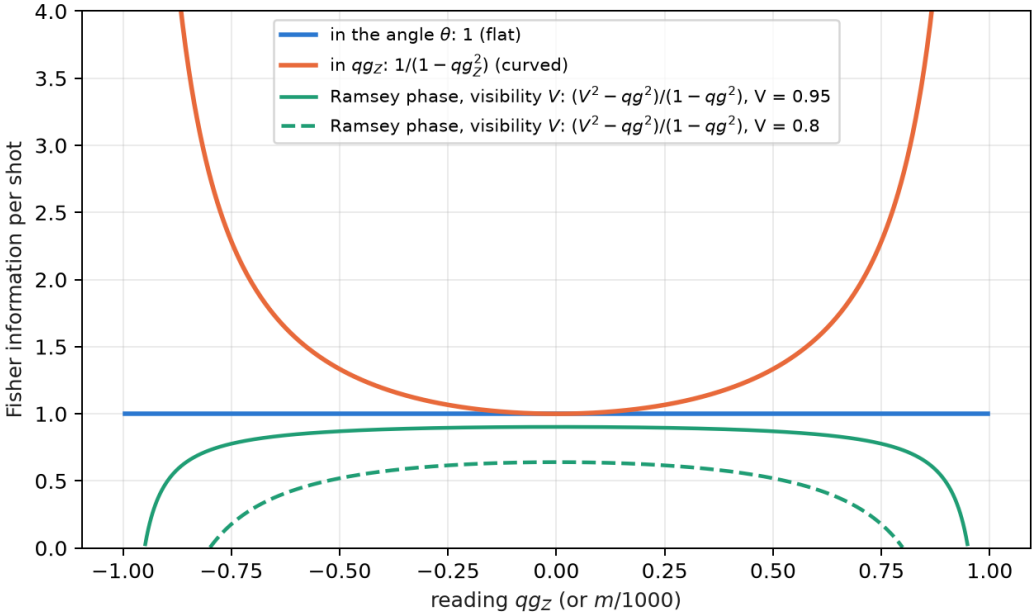


Figure 8.1: Fisher information per shot. In the angle it is constant; in qg_Z it grows towards the poles. In a Ramsey interferometer with visibility $V < 1$ (green) the information about the phase falls to zero at the extreme fringes and peaks in the middle.

On the equator, resolving 10 m-qg takes 10^4 shots and resolving 1 m-qg, 10^6; at 60° from the pole, 5 m-qg cost 30 000. Near a pole the reading is more precise in m-qg, but one m-qg there corresponds to a larger angle: a state δ m-qg from the pole is $\arccos(1-\delta/1000) \approx \sqrt{2\delta/1000}$ radians away, 2.56° for 1 m-qg.

8.2 Information geometry and the natural gradient

The many-parameter version of the Fisher information is the Fubini–Study metric, and it governs the **natural gradient**: descent that measures its steps in distance between states, not between parameters. In a real ansatz of R_Y rotations and CZ gates, every state of the circuit is real, $\langle Y_k \rangle = 0$ before each rotation and every diagonal entry of the metric is exactly $\frac{1}{4}$: the angles are flat one by one, and all the curvature sits in the off-diagonal terms, which within a layer are correlators $\langle Y_i Y_j \rangle / 4$, measurable with one circuit per layer (§159).

Two practical consequences, checked on a 4-qubit Ising chain with 20 seeds: descent in qg coordinates **never converges** (0 of 20), because the curved coordinate distorts the steps; and the natural gradient with the qg_{YY} block metric is 1.3 times faster than descent in the angle, and 4.6 times with the full metric. The gain comes from the metric, not from changing coordinates: the natural gradient in qg is Euclidean descent in θ. The rule is to optimize in the angle and read out in qg.

8.3 Two clocks and gravity

An atomic clock measures a phase, and two clocks at different heights tick at different rates: the gravitational redshift. Read with the same laser, they give $qg_{Z,1} = \cos(\varphi + \delta)$ and $qg_{Z,2} = \cos(\varphi - \delta)$, where φ is a random common phase (the laser noise) and 2δ the difference sought. Even with φ unknown, the pairs of readings lie on the ellipse

$$x^2 + y^2 - 2xy\,c = 1 - c^2, \qquad c = \cos 2\delta,$$

so the difference is read from the *shape* of the cloud, without controlling the common phase (Figure 8.2). For two strontium clocks 1 m apart, $2\delta = 2\pi\nu\, gh\, t/c^2 = 0.294$ rad in one second.

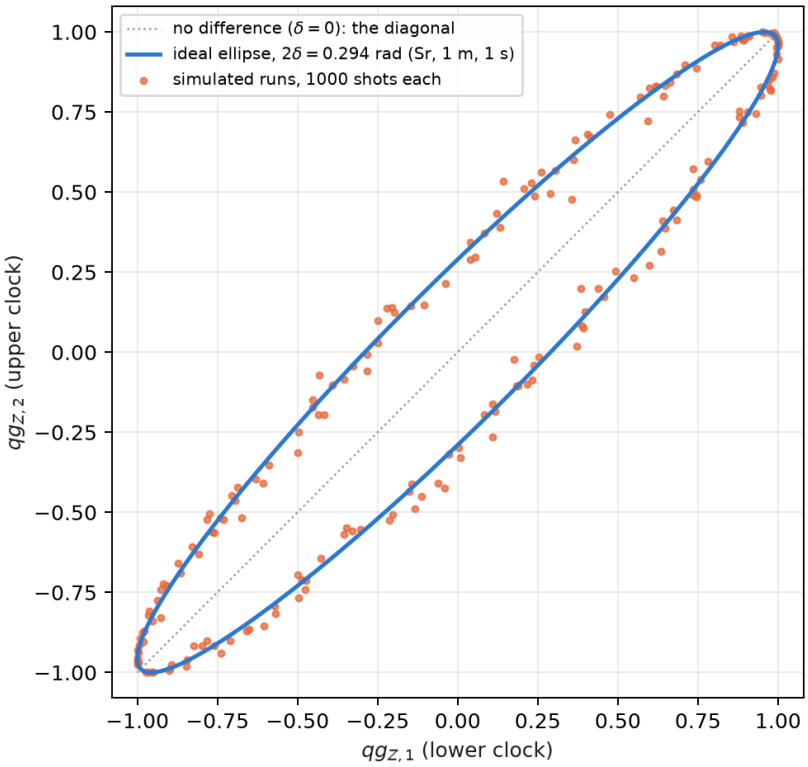


Figure 8.2: Readings of two clocks with a random common phase. The height difference opens the diagonal into an ellipse; its shape gives the phase difference without knowing the common phase.

Fitting the ellipse needs the squares of the readings, and there the unbiased estimator of qg^2 from Chapter 4 comes in: without the correction, the shot-noise bias flattens the ellipse. All four recorded predictions passed (§126):

Table 8.1: Clock comparison by ellipse fitting, simulated.

check	result
relative bias, 1000 shots, 200 runs, four values of δ	+0.09%, +0.12%, +0.05%, −0.07%
spread of the ellipse fit against the moment estimator ($\delta = 0.15$)	0.0022 against 0.0075
bias with 100 shots, without and with the correction	+11.8% and +0.9%
strontium, 1 m, 1 s	$\delta = 0.1471$; estimate 0.1472 ± 0.0022 (68σ)

In this idealized model the spread corresponds to about 1.5 cm of height per 200 runs. The limits are stated: no laser decoherence during the Ramsey time, contrast 1 and binomial projection noise.

8.4 Ramsey sensing: where to operate

A Ramsey interferometer turns a phase into a probability: $qg = a + bV\cos(\varphi - \alpha)$, with visibility V, operating point α and the read-out map (a, b) of Chapter 27. With ideal read-out, the information per shot about the phase is

$$F = \frac{V^2 - qg^2}{1 - qg^2}.$$

With $V = 1$ it is flat, $F = 1$ everywhere: the $1/(1 - qg^2)$ gain at the poles is exactly cancelled by the Jacobian. With $V < 1$ the optimum is the middle of the fringe, $qg = 0$, and the information vanishes at the bright fringe (Figure 8.1). This is the known practice of operating at quadrature, written in qg (§36). Three measured nuances:

- **Read-out asymmetry moves the optimum, but it does not matter:** moving to the new optimum gains 0.17% with the calibrated read-out, and 1.6% even with a 10% read-out error.

- **With few shots a prior helps:** with 10 shots, the Bayesian posterior mean is about 2 times better than direct inversion, because the prior range of the phase carries information.
- **With visibility 0.9, operating at the pole gets worse with more shots** (the scaled error goes from 1.09 to 2.50 between 10 and 1000 shots): small phases become invisible there. The middle of the fringe tracks its bound (1.11–1.24, with bound 1.12).

8.5 GHZ metrology: when entanglement wins

With N qubits a frequency ω can be measured in two ways: N independent Ramsey clocks, or a GHZ state, whose parity $X_1 \cdots X_N$ is a single qg value of the whole register, $qg_P = V_N \cos(N\omega t)$, oscillating N times faster. The advantage is paid for in visibility, which falls as the N-th power of that of one qubit. At mid-fringe the information per shot is V^2, and per unit time (§58):

$$\text{independent: } N\,t\,e^{-2t/T_2}(1-2e)^2, \qquad \text{GHZ: } N^2\,t\,e^{-2Nt/T_2}(1-2e)^{2N},$$

with e the read-out error per qubit. Optimizing the interrogation time:

Table 8.2: Gain in the Fisher-information rate, GHZ over independent qubits.

	Markovian, $e =$			Gaussian, $e =$		
N	$e = 0$	0.002	$e = 0.01$	$e = 0$	0.002	$e = 0.01$
2	1.000	0.992	0.960	1.414	1.403	1.358
10	1.000	0.930	0.695	3.162	2.942	2.198
30	1.000	0.793	0.310	5.477	4.341	1.697
100	1.000	0.452	0.018	10.00	4.522	0.183

The table sums up three known results, here in a single qg reading. With Markovian dephasing there is no gain for any N (Huelga et al. 1997), and with read-out error independent qubits are strictly better. With Gaussian, slow dephasing the gain grows as $\sqrt{N}$, the Zeno-type limit (Matsuzaki et al. 2011; Chin et al. 2012). And the read-out caps it: the parity carries $(1-2e)^N$, so there is an optimal size of the GHZ state, beyond which adding qubits makes the measurement worse (Figure 8.3).

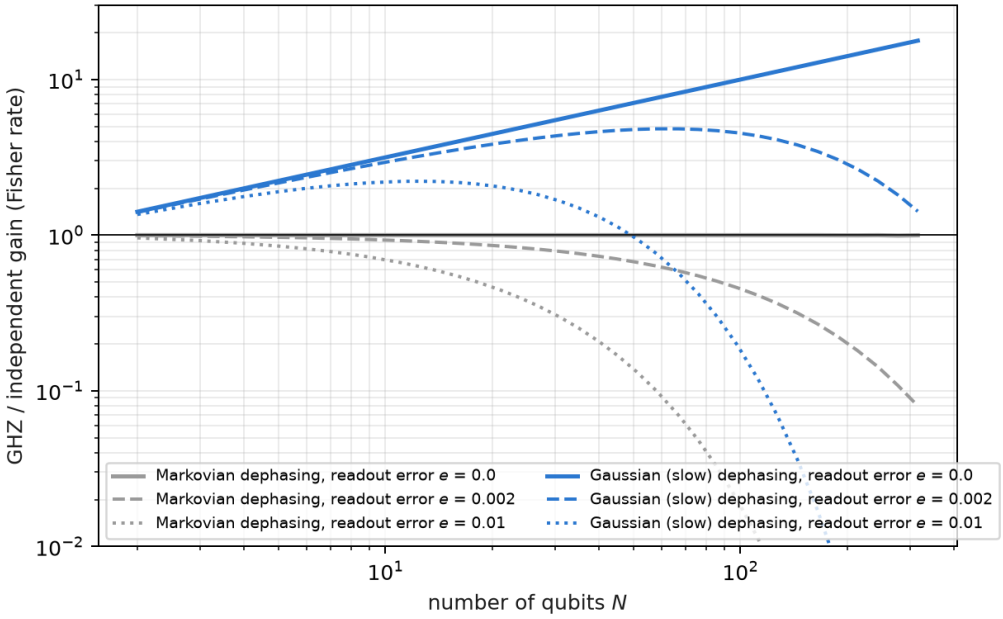


Figure 8.3: The gain of the GHZ state by type of dephasing and read-out error. With Markovian dephasing there is no gain for any N; with Gaussian dephasing it grows as √N, until the read-out error, raised to the N, eats it.

8.6 Thermal states: $qg_Z = \tanh(\beta h)$

A qubit with $H = -hZ$ in equilibrium at temperature T (with $k_B = 1$) is in the Gibbs state, and its polar bias is

$$qg_Z = \tanh(\beta h), \qquad \beta = 1/T.$$

The whole thermodynamics of the qubit is then a function of a single reading (§25):

Table 8.3: Thermodynamics of a qubit in qg units.

quantity	in qg
energy	$U = -h\, qg$
entropy (bits)	$S = qg_S(qg)$, exact because ρ is diagonal in Z
free energy	$F = -T \ln\!\left(2/\sqrt{1 - qg^2}\right)$
heat capacity	$C/k_B = \operatorname{artanh}(qg)^2\,(1 - qg^2)$
preparation angle	$\theta = \arccos(qg) = \pi/2 - \mathrm{gd}(\beta h)$ (Gudermannian)

Thermometry. A Z measurement is an energy measurement, and therefore the optimal one on a Gibbs state; its Fisher information about T is C/T^2. The best relative precision is at the maximum of the heat capacity:

$$qg \cdot \operatorname{artanh}(qg) = 1 \quad \Rightarrow \quad qg^* = 0.8336, \qquad \beta h = 1.1997,$$

the peak of the Schottky anomaly, the known optimal operating point of a two-level thermometer (Correa et al. 2015), here as a one-line condition on the reading: $m^* = +833.6$ m-qg, an angle of 33.53° (93.2 qga). The section is an exact rewriting of textbook thermodynamics, checked against $e^{-\beta H}$ to $5 \cdot 10^{-12}$; no advantage was expected or found.

8.7 Without qang and with qang

Table 8.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
resolution of a reading	variance of $\langle Z\rangle$, $(1-\langle Z\rangle^2)/N$, different at each point	information 1 per shot in the angle; $\sigma_m = 1000\lvert\sin\theta\rvert/\sqrt{N}$	it is the usual **Cramér–Rao bound**; the angle makes it uniform
optimize a variational circuit	descent in the gate angles	descent in qg coordinates: **does not converge** (0 of 20); natural gradient with qg_{YY} correlators: 1.3× to 4.6×	**against** using qg_Z as a coordinate; the gain is the **metric**'s, not qang's
compare two clocks	ellipse fit with naive squares: bias +11.8% at 100 shots	with unbiased qg^2: +0.9%	**in favour** of the estimator; it is the standard correction, applied on the qg scale
optimal thermometry point	maximum of the heat capacity (Schottky)	$qg{\cdot}\operatorname{artanh}(qg) = 1$, $qg^* = 0.8336$	**same point**; qang gives it as a one-line condition on the reading
GHZ gain	Fisher rates	the same rates, with a qg parity	**none**: the three results are known

Balance. In metrology qang organizes: it puts the Cramér–Rao bound in a flat coordinate and gives one-line conditions. The single result against the qg coordinate is explicit: for optimization, the angle and the metric, not qg_Z.

8.8 In code

```
from qang import probability

r = probability.analyze(750, 1000)  # 750 zeros in 1000 shots, in every qang unit
# p0, qg_Z, m = 500.0 [444.5, 551.7] m-qg, theta = 60.0° [56.5°, 63.6°], qg_S, with interv
```

8.9 Summary

- Per shot, the information is 1 in the angle and $1/(1-qg_Z^2)$ in qg_Z; the resolution in m-qg is $1000|\sin\theta|/\sqrt{N}$, the Cramér–Rao bound.
- The natural-gradient metric in a real ansatz has diagonal $\frac{1}{4}$ and qg_{YY} correlators; optimizing in qg coordinates does not converge, and the gain comes from the metric.
- Two clocks with a random common phase trace an ellipse; its shape gives the gravitational redshift at 1 m with 68σ in the model, if the bias of the squares is corrected.
- A Ramsey interferometer with visibility $V<1$ is operated at mid-fringe; its information is $(V^2-qg^2)/(1-qg^2)$.
- The GHZ state gains nothing with Markovian dephasing, gains $\sqrt{N}$ with slow dephasing and loses it to the read-out error raised to the N.
- A thermal qubit has $qg_Z=\tanh(\beta h)$; its best thermometry point is $qg^*=0.8336$, the Schottky peak.

Bibliographic notes

The information geometry of a quantum state is Braunstein and Caves (1994), and the natural gradient, Amari (1998). The separated-field method is Ramsey (1950); for quantum metrology in general, see Giovannetti et al. (2004). The optical clocks that measure gravitational time dilation are Chou et al. (2010). The GHZ results under dephasing are Huelga et al. (1997), Matsuzaki et al. (2011) and Chin et al. (2012), and optimal thermometry with a two-level system, Correa et al. (2015). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part III

III. Neutrino physics, oscillations and quantum simulation (`qang.neutrino`)

9 Quantum kinematics of flavour oscillation in vacuum

A neutrino is born and detected in a flavour state (ν_e, ν_μ or ν_τ), but it travels in mass states. Since the mass states pick up different phases, the flavour changes on the way: this is flavour oscillation, and in vacuum it is exactly a two- or three-level rotation. This chapter writes it on the base of Chapter 1. An oscillation probability is a qang probability like any other, with its angle, its milli-qang reading and its error bar; two-flavour oscillation is a precession on a cone of the sphere; and three flavours fit on three qubits with a circuit that conserves the weight, so the qg filter applies. The chapter ends with the project's first real-hardware data point, taken on IQM Garnet.

Everything in this chapter is in vacuum. A 1300 km baseline crosses the Earth and has sizeable matter effects; those are the subject of Chapter 10.

9.1 From the base: the oscillation probability in m-qg

With two flavours on one qubit, $|0\rangle = \nu_e$ and $|1\rangle = \nu_\mu$, the survival probability is $p_0 = 1 - P$, with $P = P(\nu_e \to \nu_\mu)$. The probability equation of Chapter 1 gives all its readings:

$$qg_Z = 1 - 2P, \qquad m = 1000\,(1 - 2P) \text{ m-qg}, \qquad \vartheta = \arccos qg_Z = 2 \arcsin \sqrt{P},$$

and its shot resolution, $\sigma_m = 1000\,|\sin\vartheta|/\sqrt{N} = 2000\sqrt{P(1-P)}/\sqrt{N}$ m-qg. A 1% appearance probability is 20 m-qg; a probability difference of 0.025 is 50 m-qg. With three flavours on three qubits (below) each qubit carries the probability of one flavour and is read the same way: $m_\beta = 1000\,(1 - 2P_\beta)$.

The scale adds no physics: it is the same probability. What it adds is a common ruler for comparing the signal with shot noise and, as shown next, a coordinate in which the mixing angle is read directly.

9.2 Oscillation as a rotation on a cone

In the flavour basis, going to the mass basis is a rotation $R_Y(-2\theta)$, with θ the mixing angle; in the mass basis the two states acquire the relative phase

$$\Delta = \frac{\Delta m^2 L}{2E} = 2 \times 1.267 \, \frac{\Delta m^2[\mathrm{eV}^2]\, L[\mathrm{km}]}{E[\mathrm{GeV}]},$$

which is an $R_Z(\Delta)$; and returning to flavour is $R_Y(2\theta)$. The one-qubit circuit is then $R_Y(2\theta)\, R_Z(\Delta)\, R_Y(-2\theta)$, and its reading is (§151)

$$qg_Z = \cos^2 2\theta + \sin^2 2\theta \, \cos\Delta, \qquad P(\nu_e \to \nu_\mu) = \sin^2 2\theta \, \sin^2 \frac{\Delta}{2}.$$

The first formula is the **spherical law of cosines**. The mass axis is 2θ from the pole; the Bloch vector, which starts at the pole, precesses about that axis on a cone of half-angle 2θ, and after turning by Δ it is an arc distance $\vartheta = \arccos qg_Z$ from the pole. It is the spherical triangle with two sides 2θ and included angle Δ. The qang angle of the reading is the distance travelled on the sphere (Figure 9.1).

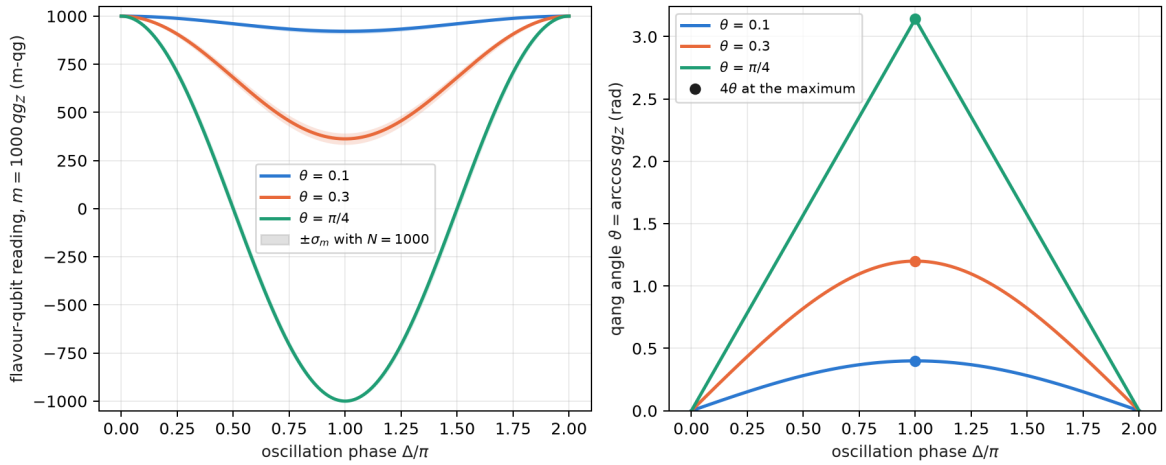


Figure 9.1: Two flavours on one qubit. Left, the m-qg reading as a function of the phase, with the shot-noise band for 1000 shots. Right, the qang angle $\vartheta = \arccos qg_Z$: at the oscillation maximum ($\Delta = \pi$) it is exactly 4θ.

9.3 The mixing angle is a quarter of the qang angle

At the oscillation maximum, $\Delta = \pi$, the law of cosines reduces to $qg_Z = \cos^2 2\theta - \sin^2 2\theta = \cos 4\theta$, and therefore

$$\theta_{\text{mix}} = \frac{\vartheta}{4} = \frac{\arccos qg_Z}{4} \qquad (\theta \le \pi/4).$$

The mixing angle is read as a quarter of a qang angle, without inverting any $\sin^2$. It must be said precisely what this is and what it is not. It is exact only at the maximum: at any other phase the relation between ϑ and θ is not linear, and the trigonometry is still there. And it is not new physics: it is the arcsine transformation of a proportion from Chapter 4 (§125), which at the maximum coincides with the geometry of the cone.

The condition $\theta \le \pi/4$ is not a detail either. For $\theta > \pi/4$ the angle 4θ exceeds π and arccos returns $2\pi - 4\theta$, that is, $\pi/2 - \theta$. Since $\sin^2 2\theta$ is symmetric under $\theta \to \pi/2 - \theta$, no two-flavour measurement can tell the two values apart: this is the **octant degeneracy**, the same one the atmospheric angle θ_{23} has (49° in the values used here; the reading at the maximum would give 41°). The qg scale makes it visible; it does not resolve it.

9.4 Error bars that do not depend on the angle

The qang angle inherits the stabilized variance of the arcsine transformation: its shot-noise standard deviation is $1/\sqrt{N + \frac{1}{2}}$, the same over the whole sphere (Chapter 4, §125). Divided by four:

$$\sigma_\theta = \frac{1}{4\sqrt{N + \frac{1}{2}}},$$

which does not depend on θ. With 1000 shots it is 0.0079 rad (0.45°), that is, 1.26 qga on the angular scale of Chapter 1 (Section 1.4), where the relation is just as simple: $\theta_{\text{mix},qga} = \vartheta_{qga}/4$, and the condition $\theta \le \pi/4$ reads $\theta \le 125$ qga. Read in milli-qang, the same measurement does depend on the point: for $\theta = 0.1$ the maximum is at $m = 1000 \cos 0.4 = 921$ m-qg with $\sigma_m = 12$ m-qg. The pre-registered prediction (N1) held over 2000 binomial repetitions:

Table 9.1: The mixing angle read as $\arccos(qg_Z)/4$ at the maximum, 1000 shots, 2000 repetitions.

θ (rad)	0.1	0.3	0.5	0.7
mean of $\hat{\theta}$	0.0999	0.2999	0.5000	0.7006
std of $\hat{\theta}$ (theory 0.0079)	0.0079	0.0077	0.0078	0.0081

9.5 Three flavours on three qubits

With three flavours, flavour state k is the weight-1 state with only qubit k excited: $\nu_e = |100\rangle$, $\nu_\mu = |010\rangle$, $\nu_\tau = |001\rangle$ (in qubit order 0, 1, 2). The PMNS matrix in its standard form,

$$U = R_{23}\, U_\delta\, R_{13}\, U_\delta^\dagger\, R_{12}, \qquad U_\delta = \mathrm{diag}(1, 1, e^{i\delta_{CP}}),$$

is three real rotations and two phases. Each rotation R_{ij} is an RBS gate (an $XX+YY$) between qubits i and j; each phase is a phase gate on one qubit; and propagation, in the mass basis, is one phase gate per qubit. **Every gate conserves the Hamming weight**, so a shot with weight other than 1 is an error and can be discarded: this is the qg symmetry filter of Chapter 21 (§77). For antineutrinos, $\delta_{CP} \to -\delta_{CP}$ is enough. The circuit reproduces the analytic vacuum probabilities to 10^{-15}.

Figure 9.2 shows the case of the prediction: an initial ν_μ, $L = 1300$ km (a DUNE-like baseline), $\delta_{CP} = -\pi/2$ and NuFIT-like mixing values in normal ordering. The first oscillation maximum is at 2.63 GeV, where survival drops to $P(\nu_\mu \to \nu_\mu) = 0.013$, that is, $m_\mu = 974$ m-qg.

9.6 Noise suppression with the qg filter

The compiled three-flavour circuit has 18 two-qubit gates on IBM. It was run on **simulated noise models** of three devices (FakeBrisbane with ECR gates, FakeTorino and FakeFez with CZ), at 8 energies between 0.8 and 5 GeV with 4000 shots. The error is the mean of $|\hat{P} - P|$ over the three flavours and eight energies (§151):

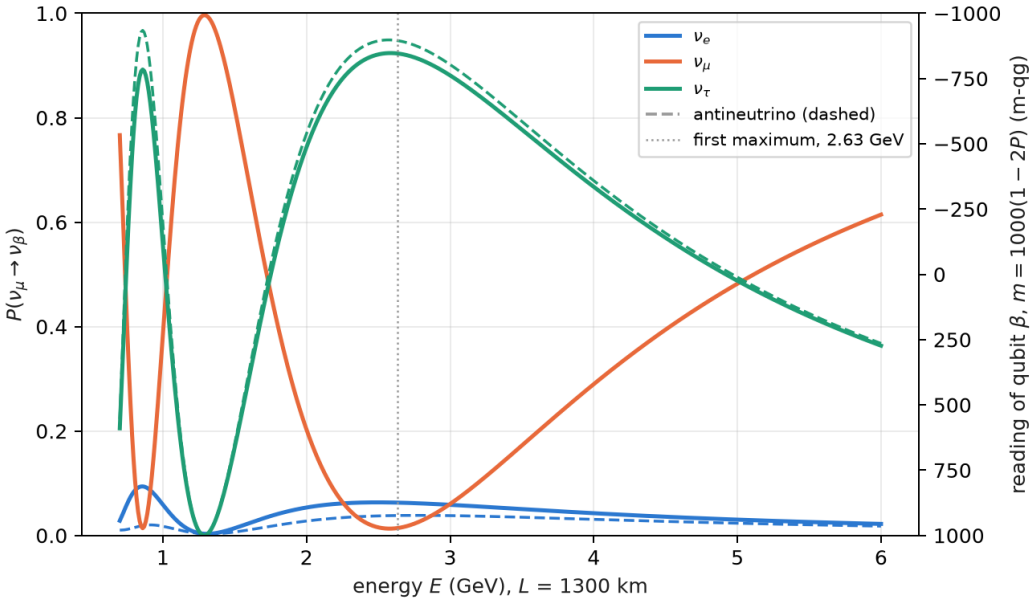


Figure 9.2: Three-flavour probabilities for an initial ν_μ at 1300 km, in vacuum; on the right, the m-qg reading of each flavour's qubit. Dashed lines are antineutrinos: ν_μ survival is identical (CPT), ν_e appearance is not (CP).

Table 9.2: Mean error of the three-flavour probabilities, simulated noise models.

simulated device	raw	readout correction	qg filter	filter + readout	kept	raw / filter + readout
FakeBrisbane	0.048	0.038	0.014	0.013	85–87%	3.8×
FakeTorino	0.034	0.028	0.010	0.010	90–92%	3.4×
FakeFez	0.025	0.024	0.009	0.009	93–94%	3.0×

The error falls between 3.0 and 3.8 times, and almost all of the gain is the filter, not the readout correction. All four pre-registered predictions held, including the position of the minimum of $P(\nu_\mu \to \nu_\mu)$ at 2.6 GeV on all three devices.

On real hardware (§165). The same circuit was run on IQM Garnet, a 20-qubit superconducting processor, compiled to its native gates with 12 CZ; 8 energies, 300 shots each, with predictions recorded before the job. All four held, and Figure 9.3 shows the data:

Table 9.3: The three-flavour circuit on IQM Garnet, 8 × 300 shots, no readout correction.

check	result
filtered error below raw	0.039 < 0.062
ratio raw / filtered ≥ 1.5	1.59 (narrowly)
filtered minimum of $P(\nu_\mu \to \nu_\mu)$ at 2.6 GeV	yes; **raw puts it at 0.8 GeV**
kept fraction ≥ 0.5	0.88

Three readings of this data point. On real hardware **the filter puts the oscillation maximum back where it is**. The chip did better than its simulated model (0.062 raw against 0.126) and so the filter gained less (1.6× against 2.2× in the rehearsal): fewer shots leave the sector, so there is less to remove. And what remains is **inside the sector**, where the filter cannot see: at the maximum, $P(\nu_\mu \to \nu_\mu)$ reads 0.10 against an exact 0.013 (about 175 m-qg apart), and the ν_τ error barely moves (0.062 → 0.052). It is one run on one day, with 300 shots per circuit.

9.7 The CP asymmetry in vacuum

CP violation shows up as a difference between neutrinos and antineutrinos:

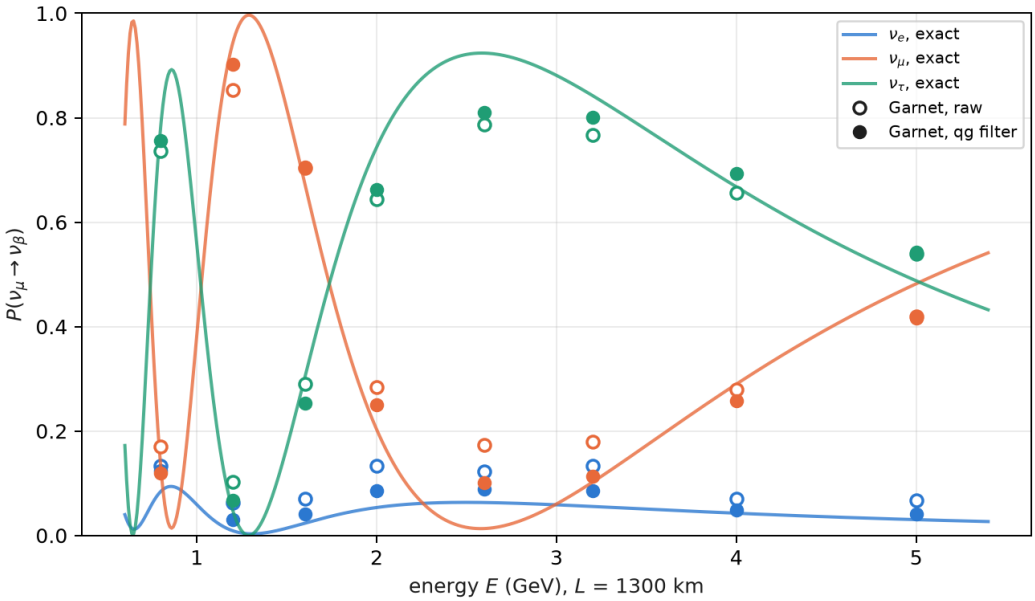


Figure 9.3: The three-flavour circuit on IQM Garnet. Curves are exact; open circles are raw readings, filled ones filtered. Without the filter, ν_μ survival (orange) is lowest at 0.8 GeV; with it, at 2.6 GeV, where the oscillation maximum is.

$$A_{CP} = P(\nu_\mu \to \nu_e) - P(\bar{\nu}_\mu \to \bar{\nu}_e).$$

At 2.6 GeV with $\delta_{CP} = -\pi/2$ it is 0.0250 (0.0630 against 0.0380), that is, **50 m-qg** on the ν_e qubit. With 20 000 shots per circuit, the shot noise of the difference is about 4.4 m-qg: the signal is about 11σ. Survival, on the other hand, cannot tell them apart: in vacuum CPT imposes $P(\nu_\mu \to \nu_\mu) = P(\bar{\nu}_\mu \to \bar{\nu}_\mu)$ (the orange curves of Figure 9.2 coincide).

Table 9.4: CP asymmetry at 2.6 GeV, exact 0.0250 (50 m-qg), 20 000 shots, simulated noise models.

simulated device	raw	filter + readout
FakeBrisbane	0.024 (48 m-qg)	0.026 (51 m-qg)
FakeTorino	0.017 (35 m-qg)	0.022 (44 m-qg)
FakeFez	0.020 (41 m-qg)	0.022 (44 m-qg)

The asymmetry survives even without the filter. It is the difference of two nearly identical circuits (only the sign of two phases changes), so the shift that noise produces in common cancels in the subtraction, as in the fidelity classifier of §149. What does not cancel is the contraction, which shrinks both probabilities and the difference with them: that is why FakeTorino's raw value reads 35 m-qg and not 50. The filter halves the error on the CZ devices.

9.8 What the link between neutrinos and qang is, and is not

It is modelling: the oscillation is a rotation, and the qang angle reads the mixing angle directly at the maximum, with a known error bar. It does not remove the trigonometry in general, and it does not speed up variational methods for dense neutrino gases: those are many-body problems (Chapter 10). The limits are stated: vacuum only, fixed parameters, simulated device noise except for the single IQM Garnet run.

9.9 Without qang and with qang

Table 9.5: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read the mixing angle at the maximum	$\theta = \frac{1}{2}\arcsin\sqrt{P}$ and the delta method	$\theta = \arccos(qg_Z)/4$, $\sigma_\theta = 1/(4\sqrt{N+\frac{1}{2}})$	**same leading order**: the delta method also gives $1/(4\sqrt{N})$; qang gets it without a derivative, with the $+\frac{1}{2}$ correction for few shots; at exactly $P = 0$ or 1 both need the intervals of Chapter 4
see the octant degeneracy	symmetry of $\sin^2 2\theta$	the reading at the maximum returns $\pi/2 - \theta$	**none**: the scale makes it visible, it does not resolve it
compare the CP signal with noise	$\Delta P = 0.025$ against σ_P	50 m-qg against 4.4 m-qg	**unit**: the ratio (11σ) is the same
reduce the error of the three-flavour circuit	readout correction: little effect	weight filter: 3.0–3.8× (simulated), 1.6× on Garnet	the gain is the **symmetry filter**'s, not the qang scale's

Balance. The geometry of the cone and the reading of the mixing angle are qang's own contribution, with an error bar that need not be derived. The improvement against noise, the chapter's largest number, belongs to the filter.

9.10 In code

```
import math
from qang import neutrino as nu

# two flavours: the reading and the mixing angle at the maximum
q = nu.two_flavour_qg_z(theta=0.1, delta=math.pi)  # cos(0.4) = 0.921 -> 921 m-qg
nu.mixing_angle_from_maximum(q)                    # 0.1
nu.mixing_angle_std(1000)                          # 0.0079 rad
```

```
# three flavours: exact probabilities, weight-1 circuit and CP asymmetry
P = nu.three_flavour_probabilities(L=1300, E=2.6, delta_cp=-math.pi / 2)  # P[mu] = [0.0
qc = nu.three_flavour_circuit("mu", L=1300, E=2.6, delta_cp=-math.pi / 2)
nu.cp_asymmetry(L=1300, E=2.6, delta_cp=-math.pi / 2)              # 0.0250
```

9.11 Summary

- An oscillation probability P is read as $m = 1000(1 - 2P)$ m-qg, with angle $2 \arcsin \sqrt{P}$ and resolution $2000\sqrt{P(1-P)/N}$ m-qg.
- With two flavours, $qg_Z = \cos^2 2\theta + \sin^2 2\theta \cos\Delta$ is the spherical law of cosines of the precession on a cone of half-angle 2θ.
- At the maximum, $\theta = \arccos(qg_Z)/4$ with standard deviation $1/(4\sqrt{N + \frac{1}{2}})$, independent of θ; for $\theta > \pi/4$ the reading returns $\pi/2 - \theta$ (octant degeneracy).
- Three flavours are three qubits in the weight-1 sector; PMNS is three RBS gates and two phases, and the qg filter lowers the error 3.0–3.8 times on simulated models.
- On IQM Garnet the filter gains 1.6× and puts the oscillation maximum back at 2.6 GeV; the error that remains is inside the sector.
- The CP asymmetry at 2.6 GeV is 50 m-qg and survives without the filter, because the common shift cancels in the subtraction.

Bibliographic notes

The lepton mixing matrix comes from Pontecorvo (1968) and Maki et al. (1962); the oscillation parameters used are of the type of the NuFIT global fit (Esteban et al. 2020) and the *Review of Particle Physics* (Navas et al. 2024). The RBS gate is that of the orthogonal networks of Landman et al. (2022). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

10 Matter effects, solar neutrinos and collective oscillations

Chapter 9 left the neutrino in vacuum. But a 1300 km beam crosses the Earth, solar neutrinos are born at the centre of the Sun, and in a supernova neutrinos are so dense that they interact with each other. In all three cases the medium adds a term to the flavour Hamiltonian, and the chapter's question is the same in all three: **which number is still conserved**, because that decides whether the qg filter applies, and in which basis. On the base of Chapter 1, each effect of the medium is read as a shift in milli-qang, on the same scale as shot noise.

10.1 From the base: the medium as a shift in m-qg

Each flavour (or mass) qubit is read as in Chapter 9, $m = 1000\,(1 - 2P)$ m-qg, so an effect of the medium that changes a probability by δP is a shift of $2000\,\delta P$ m-qg, directly comparable with the resolution $\sigma_m = 2000\sqrt{P(1-P)/N}$. Three examples from this chapter, on that scale:

Table 10.1: Effects of the medium as shifts in milli-qang.

effect	δP	in m-qg
matter on $P(\nu_\mu \to \nu_e)$, 1300 km, 2.6 GeV, normal ordering	0.063 → 0.087	48
fake asymmetry from matter at $\delta_{CP} = 0$	0.042	84
solar MSW drop of P_{ee}, 0.5 → 10 MeV, read on the ν_2 qubit	p_2: 0.35 → 0.93	1173

With 1000 shots and $P \approx 0.07$ the resolution is about 16 m-qg: the Earth's effect shows at 3σ in a single circuit, and the solar drop spans most of the scale.

10.2 The MSW effect in the Earth

Electrons in the Earth forward-scatter ν_e and add to the flavour Hamiltonian a potential $A = 2\sqrt{2}\,G_F N_e E$, which in oscillation units is

$$A = 1.52 \times 10^{-4}\, Y_e\, \rho[\mathrm{g/cm^3}]\, E[\mathrm{GeV}]\ \mathrm{eV}^2,$$

with the opposite sign for antineutrinos (Mikheyev–Smirnov–Wolfenstein). In the crust ($\rho = 2.85$ g/cm³, $Y_e = 0.5$) and at 2.6 GeV, $A = 5.6 \times 10^{-4}$ eV², 22% of Δm^2_{31}: not a small correction.

The potential acts on ν_e only, that is, it is a phase on qubit 0. The circuit alternates the exact vacuum step of Chapter 9 with that phase (second-order, or Strang, Trotter), and **every gate still conserves the weight**: the qg filter applies just as in vacuum (§152). With $L = 1300$ km and between 0.8 and 5 GeV:

Table 10.2: The MSW effect as a Trotter circuit, simulated noise models.

check	result
M1 noiseless Trotter error, $r = 2$ / 4 (needed < 0.005 / < 0.001)	0.0052 / 0.0011 — **failed** narrowly
M2 matter-only asymmetry at $\delta_{CP} = 0$, 2.6 GeV	0.042 (the vacuum CP signal was 0.025)
M3 best error raw / filter + readout	Brisbane 0.046 / 0.017, Torino 0.027 / 0.009, Fez 0.023 / 0.014 (median 2.8×)
M4 optimal steps with filter $\geq$ without	$r^* = 1$ everywhere (holds vacuously)
M5 $A_{CP}(\delta_{CP} = 0)$, filter + readout, $r = 2$	0.035, 0.035, 0.029 (exact 0.042; Trotter $r = 2$: 0.033)

Three lessons. **Matter fakes CP violation**: at $\delta_{CP} = 0$, where vacuum gives no asymmetry at all, the Earth produces 84 m-qg of difference between ν and $\bar{\nu}$, more than the 50 m-qg of Chapter 9's CP signal; that is why long-baseline experiments must model matter, and the filtered circuit sees it. **There is no depth trade-off**: a single Trotter step already has less error (0.012) than the noise, so $r^* = 1$ with or without the filter, and M4 holds only because both are 1. And **M5's gap to the exact value is Trotter error, not noise**: the filtered values sit within 0.004 of the $r = 2$ Trotter value.

10.3 The mass ordering read through matter

Vacuum cannot tell the sign of Δm^2_{31}; matter can, because what matters is the sign of the product $A\,\Delta m^2_{31}$. In normal ordering ($\Delta m^2_{31} > 0$) matter enhances $\nu_\mu \to \nu_e$ and suppresses the antineutrino channel; in inverted ordering, the reverse (Figure 10.1). With the exact values at 2.6 GeV:

Table 10.3: $P(\nu_\mu \to \nu_e)$ at 1300 km and 2.6 GeV, $\delta_{CP} = -\pi/2$.

	ν, vacuum	ν, matter	$\bar{\nu}$, vacuum	$\bar{\nu}$, matter
normal	0.063	0.087	0.038	0.021
inverted	0.063	0.041	0.038	0.058

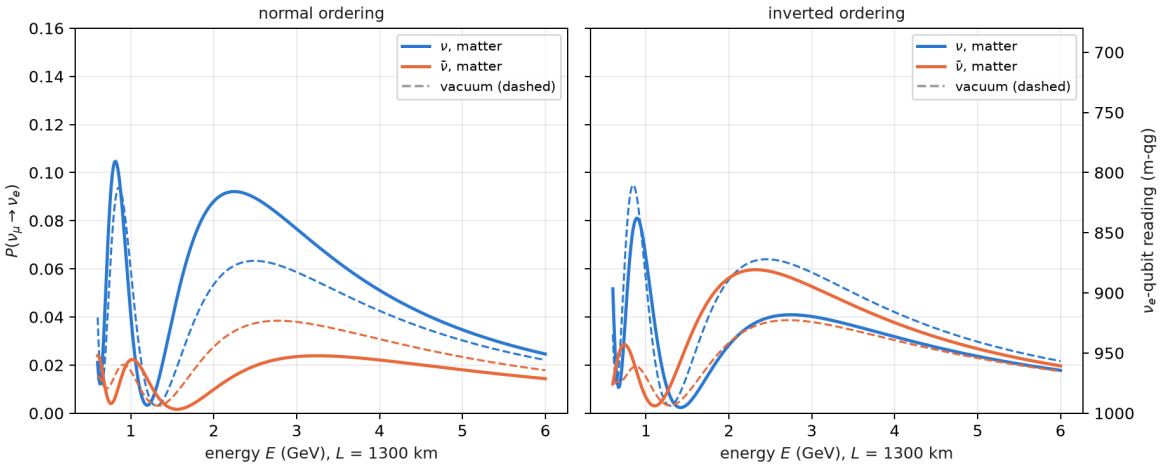


Figure 10.1: $P(\nu_\mu \to \nu_e)$ at 1300 km in matter (solid) and in vacuum (dashed), for neutrinos and antineutrinos, in both orderings. On the right, the ν_e-qubit reading in m-qg. The sign of the matter shift flips with the ordering.

On the simulated noise models, with the $r = 2$ Trotter circuit and the filter, **the ordering is read correctly on all three devices**, with the right sign of the $\nu/\bar{\nu}$ asymmetry (§163). The filtered values sit 0.010–0.033 above the noiseless-circuit reference (prediction O3 failed): noise inside the sector moves probability toward the other flavours. But the bias is common, and the difference between orderings survives: 0.034–0.036 against 0.037 in the reference.

10.4 Solar neutrinos: the whole Sun is one rotation

A ν_e is born at the solar core, where $(\rho Y_e) \approx 100$ g/cm³, and travels out to vacuum. At low energy matter barely matters and it reaches the Earth with the averaged vacuum probability; above the resonance (about 2 MeV at the core), the central density turns the matter mixing angle θ_m almost to 90°, the neutrino adiabatically follows the heavy eigenstate and leaves the Sun as ν_2. The phases are lost on the way to the Earth, so what arrives is a population in the mass basis, and the survival probability (§156)

$$P_{ee} = c_{13}^4 \left(\cos^2\theta_{12}\, p_1 + \sin^2\theta_{12}\, p_2\right) + s_{13}^4$$

is **linear in qg_Z of the ν_2 qubit**. It falls from 0.54 to 0.30 between 0.5 and 15 MeV (Figure 10.2); on the ν_2 qubit that is going from $+304$ to -868 m-qg, more than half of the scale.

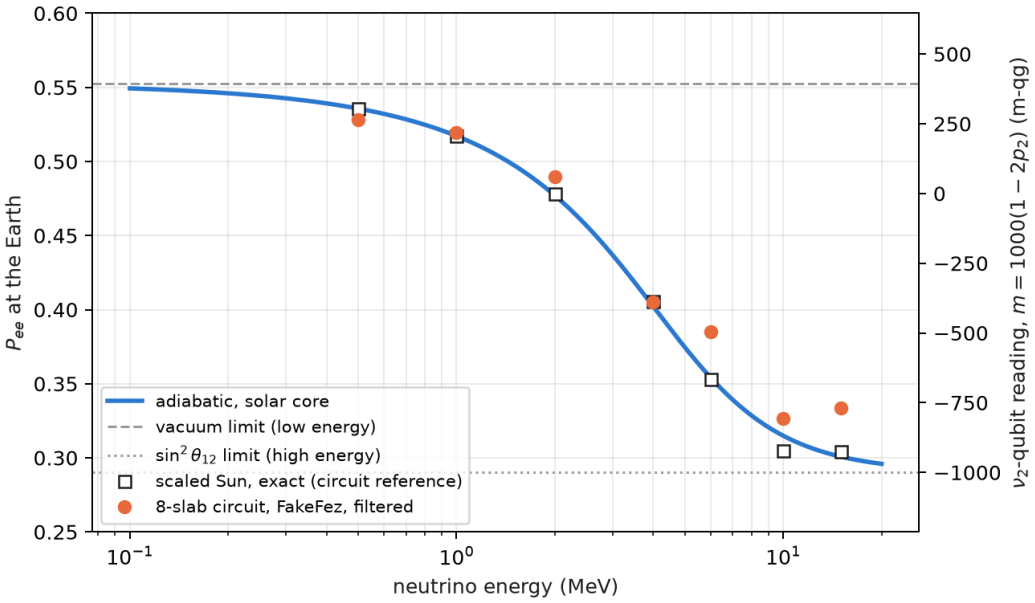


Figure 10.2: Survival probability of solar ν_e. The curve is the adiabatic conversion from the core; the squares, the exact reference of the scaled Sun the circuit uses; the circles, the 8-slab circuit on the FakeFez noise model with the filter. On the right, the same reading in m-qg of the ν_2 qubit.

The circuit uses two effective flavours (ν_3 decoupled; it agrees with three flavours to 0.001) and a scaled Sun: an exponential profile with a 100 km scale height, cut

into r slabs uniform in θ_m. Each slab is one exact 2×2 propagator, $R_Z \cdot \mathrm{RBS} \cdot R_Z$ on two qubits in the weight-1 sector, and the filter applies. All six pre-registered predictions held:

Table 10.4: Solar neutrinos on simulated noise models.

read-out (mean P_{ee} error against the noiseless circuit)	Brisbane	Torino	Fez
slabs $r = 8$, raw / filter + readout	0.025 / 0.011	0.018 / 0.008	0.027 / 0.015
compiled (a single RBS), raw / filter + readout	0.008 / 0.002	0.005 / 0.002	0.006 / 0.003
one qubit, raw / readout	0.004 / 0.002	0.031 / 0.003	0.004 / 0.002
MSW drop $P_{ee}(0.5) - P_{ee}(10\ \mathrm{MeV})$, $r = 8$ filtered (exact 0.231)	0.218	0.218	0.202

The filter gains 1.9–2.3× on the slab circuit, as in the Earth. But the important result is another one: **the physics makes the circuit unnecessary**. Because the phases are lost before detection, the observable is a single population, and the compiled circuit (the whole Sun in the one RBS that loads p_2) beats every slab circuit even without the filter; with two flavours, one qubit with readout correction is as good as anything. The adiabatic answer is classical, $\sin^2\theta_m$ at the core. A quantum computer could add something only where coherence survives: day–night regeneration in the Earth, or the collective effects of the next section.

Open: the real Sun. The profile is scaled (the real one has a scale height of about 6×10^4 km and is far more adiabatic, with the same adiabatic answer), and the slab discretization error (0.009 at $r = 8$) is of the size of the filtered noise. The comparison with a Standard Solar Model density profile has not been done: it stays a proposal, not a result.

10.5 Collective oscillations: what is conserved is the mass number

In a supernova the neutrino density is so high that neutrinos scatter off each other. With N two-flavour neutrinos, one qubit each, in the single-angle approximation, the Hamiltonian in the mass basis is

$$H_m = -\sum_i \frac{\omega_i}{2} Z_i + \frac{\mu}{2N} \sum_{i<j} \sigma_i \cdot \sigma_j,$$

with $\omega_i = \Delta m^2 / 2E_i$ and μ the neutrino–neutrino coupling. Section §151 had named these gases a natural target for the filter because "the total flavour number is conserved". **That holds only for the neutrino–neutrino term.** The product $\sigma_i \cdot \sigma_j$ is the same in every basis, but the vacuum term is diagonal only in the mass basis: what is conserved is **the number of** ν_2, not the flavour number (§154). Figure 10.3 shows it with the exact evolution of four neutrinos: the number of ν_2 is constant and the number of ν_x moves, a little at the experiment's angle and more at a large angle.

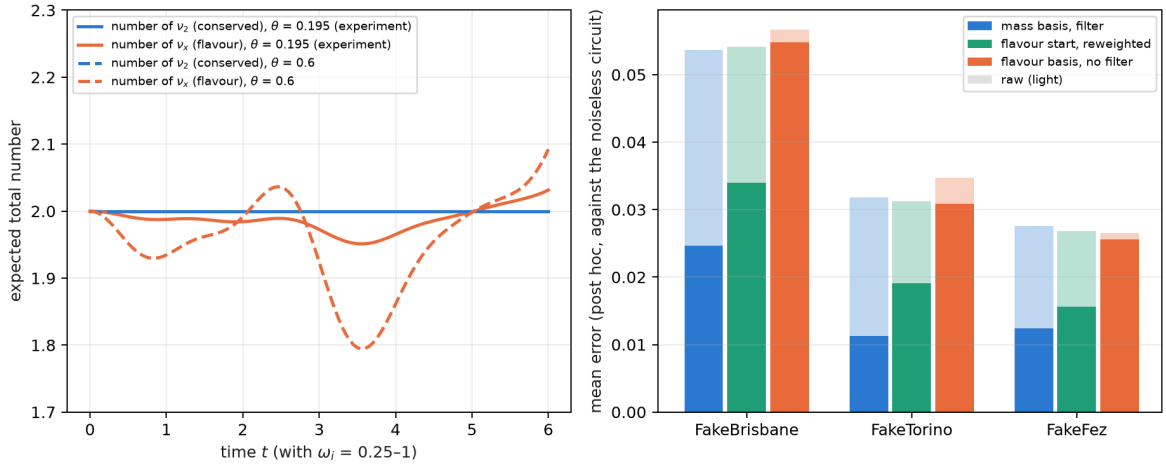


Figure 10.3: Left: exact evolution of four neutrinos (ω = 0.25–1, μ = 4, initial $\nu_e\nu_e\nu_x\nu_x$). The number of ν_2 is conserved; the flavour number is not. Right: the error of the three read-outs on simulated noise models (post hoc analysis, not pre-registered), raw in light and corrected in dark.

A flavour initial state is a superposition of mass sectors with known, conserved probabilities (for $\nu_e\nu_e\nu_x\nu_x$ at $\theta = 0.195$: 0.863 in the two-ν_2 sector and 0.067

in each neighbour). The function `reweight_sectors` rescales each sector to its probability; when the distribution is a delta, it is the filter.

All seven pre-registered predictions failed, for one reason: the noiseless Trotter error (0.095) was larger than the whole noisy error against the exact evolution (0.074–0.088), so every read-out was scored mostly on Trotter error. A post hoc analysis, not pre-registered, scored against the noiseless circuit, gives what the predictions meant to measure:

Table 10.5: Collective oscillations, $N = 4$, post hoc analysis against the noiseless circuit.

simulated device	mass basis, raw / filter	flavour start in mass basis, raw / reweighted	flavour basis, raw / readout
FakeBrisbane	0.054 / 0.025	0.054 / 0.034	0.057 / 0.055
FakeTorino	0.032 / 0.011	0.031 / 0.019	0.035 / 0.031
FakeFez	0.028 / 0.012	0.027 / 0.016	0.027 / 0.026

The filter on the conserved mass number gains 2.2–2.8×; the known distribution over sectors, 1.6–1.7×; and in the flavour basis, where nothing is conserved, there is no gain beyond readout correction. The lesson for the next pre-registration is on record: score against the noiseless circuit when the Trotter error is not known to be small, or measure that error before fixing the predictions.

10.6 Without qang and with qang

Table 10.6: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
quantify an effect of the medium	probability change δP	shift of $2000\,\delta P$ m-qg	**unit**: lets the Earth, the fake asymmetry and the Sun be tabulated on one scale with the noise
read solar survival	P_{ee} linear in the population p_2	linear in qg_Z of the ν_2 qubit	**equivalent**

action	without qang	with qang	difference
simulate the Earth with noise	unfiltered Trotter circuit	with the weight filter: 2.8× in the median	the **filter**'s, not qang's
collective oscillations	flavour-basis reading: no gain	filter in the mass basis: 2.2–2.8× (post hoc)	the **conserved number**'s; qang does not decide which number is conserved

Balance. In this chapter qang is the presentation scale. Everything that improves the simulation comes from identifying the conserved number and filtering on it, and the table says so.

10.7 In code

```
import math
from qang import neutrino as nu
from qang.sectors import reweight_sectors

# the Earth: exact matter probabilities and the weight-preserving Trotter circuit
P = nu.matter_probabilities(L=1300, E=2.6, delta_cp=-math.pi / 2, rho=2.85)      # P[mu,
qc = nu.matter_trotter_circuit("mu", L=1300, E=2.6, delta_cp=-math.pi / 2, rho=2.85, ste
nu.cp_asymmetry(L=1300, E=2.6, delta_cp=0.0, rho=2.85)                          # 0.042: matt
Pio = nu.matter_probabilities(1300, 2.6, -math.pi / 2, 2.85, params=nu.NUFIT_IO)   # P[m

# the Sun: the adiabatic answer and the compiled circuit (a single RBS)
nu.solar_pee_adiabatic(0.010)                                          # 0.315 at 10 MeV
qc_sun = nu.solar_circuit(0.010, h=100.0, X=800.0, r=8, compiled=True)

# collective oscillations: the conserved mass sectors of a flavour state
nu.mass_sector_probabilities("eexx", theta=0.195)          # [0.001, 0.067, 0.863, 0.06
```

10.8 Summary

- An effect of the medium that changes a probability by δP is a shift of $2000\,\delta P$ m-qg, comparable with shot noise on the same scale.

- The Earth adds a phase on the ν_e qubit; the Trotter circuit conserves the weight, the filter gains 2.8× in the median, and a single step is enough.
- Matter fakes CP violation (84 m-qg at $\delta_{CP} = 0$) and reveals the mass ordering: the sign of the shift flips between normal and inverted.
- In the Sun, P_{ee} is linear in qg_Z of the ν_2 qubit and falls by 1173 m-qg; since the phases are lost, the whole Sun is one rotation, and the answer is classical.
- In collective oscillations the number of ν_2 is conserved, not the flavour number: the filter works in the mass basis (2.2–2.8×, post hoc) and not in the flavour basis; all seven pre-registered predictions failed on Trotter error.
- The comparison with a realistic solar density profile is still open.

Bibliographic notes

The matter effect is Wolfenstein (1978) and Mikheyev and Smirnov (1985); the parameters, Esteban et al. (2020) and Navas et al. (2024). For collective neutrino oscillations, see the review Duan et al. (2010). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

11 Majorana phases, neutrinoless double beta decay and sterile neutrinos

Chapters 9 and 10 read what an oscillation circuit sees: mixing angles, the CP phase, the mass ordering through matter. This chapter takes two open questions of neutrino physics that have a clear answer in qg terms, one negative and one positive. Majorana phases change no oscillation probability, and therefore no qang reading: the compiler deletes them. A sterile neutrino, on the other hand, is one more qubit in the weight-1 sector, and the filter treats it like any other flavour. At the end, closing Part III, the four quantum numbers of the atom written with the same rule.

11.1 From the base: what changes no probability changes no reading

All the readings of Chapter 1 are functions of probabilities: $qg_Z = 2p_0 - 1$, $m = 1000\, qg_Z$, the angle $\arccos qg_Z$ and the entropy qg_S. So a physical parameter that changes no measurement probability is invisible to all of them at once, without exception and with any number of shots. The question "does a circuit see it?" reduces to "does it change any probability?", and is answered before running anything.

Conversely, a parameter that does change a probability by δP shows up as $2000\,\delta P$ m-qg on the qubit of the corresponding flavour, as in Chapters 9 and 10. The sterile neutrino of this chapter is of that kind: with the hypothetical parameters used below, it lowers ν_μ survival by 0.085, that is, 171 m-qg.

11.2 Majorana phases are invisible to oscillations

If the neutrino is its own antiparticle (a Majorana particle), the mixing matrix has two more phases,

$$U_M = U\ \mathrm{diag}\big(1,\ e^{i\alpha_{21}/2},\ e^{i\alpha_{31}/2}\big)\,.$$

The added matrix is diagonal in the mass basis, and so is propagation. The two commute, and in the evolution operator $U_M\, D\, U_M^\dagger$ the phase cancels against its inverse: $U_M D U_M^\dagger = U D U^\dagger$. No oscillation probability depends on α_{21} or α_{31}, in vacuum or in matter (in matter the Hamiltonian is $U_M\,\mathrm{diag}(0, \Delta m^2_{21}, \Delta m^2_{31})\, U_M^\dagger + V$, and the cancellation already happens in the first term, before the potential is added). Numerically the largest difference among the nine probabilities with and without the phases is 4×10^{-16}, machine rounding (§163).

In the three-flavour circuit of Chapter 9 the phases are phase gates in the mass basis, inserted on either side of propagation (`majorana_circuit`). What is instructive is what the transpiler does: on all three simulated devices, **the circuit compiled with phases $(\alpha_{21}, \alpha_{31}) = (\pi/2,\, 3\pi/4)$ is identical to the one compiled without them** (on FakeFez, 18 CZ, 48 R_Z and 47 $\sqrt{X}$ in both). The optimizer merges each phase gate with its inverse and deletes them. It is not an "analytic elimination" designed by anyone, but the direct consequence of the operator not changing: a compiler that simplifies consecutive phases finds the identity on its own.

11.3 Majorana phases in neutrinoless double beta decay

Where the phases do show is in a process that violates lepton number: neutrinoless double beta decay, $(A, Z) \to (A, Z+2) + 2e^-$, possible only if the neutrino is Majorana. Its rate is proportional to the square of the effective mass

$$m_{\beta\beta} = \left| \sum_i U_{ei}^2\, m_i \right|,$$

and since U_{ei}^2 enters and not $|U_{ei}|^2$, the phases do not cancel: they can add the three contributions or make them interfere destructively. With a massless lightest neutrino, scanning the two phases gives (§163)

Table 11.1: Effective mass of neutrinoless double beta decay with a massless lightest neutrino, scanning the Majorana phases.

ordering	masses (meV)	$m_{\beta\beta}$ (meV)
normal	0; 8.6; 50.1	1.4 to 3.7
inverted	49.2; 49.9; 0	18.7 to 48.3

Figure 11.1 shows both bands as a function of the lightest mass. In inverted ordering the band has a floor of about 19 meV, which is the target of the next generation of experiments. In normal ordering, for a lightest mass between about 2 and 7 meV, the phases can make $m_{\beta\beta}$ vanish completely: a cancellation that no double beta experiment could tell apart from a Dirac neutrino.

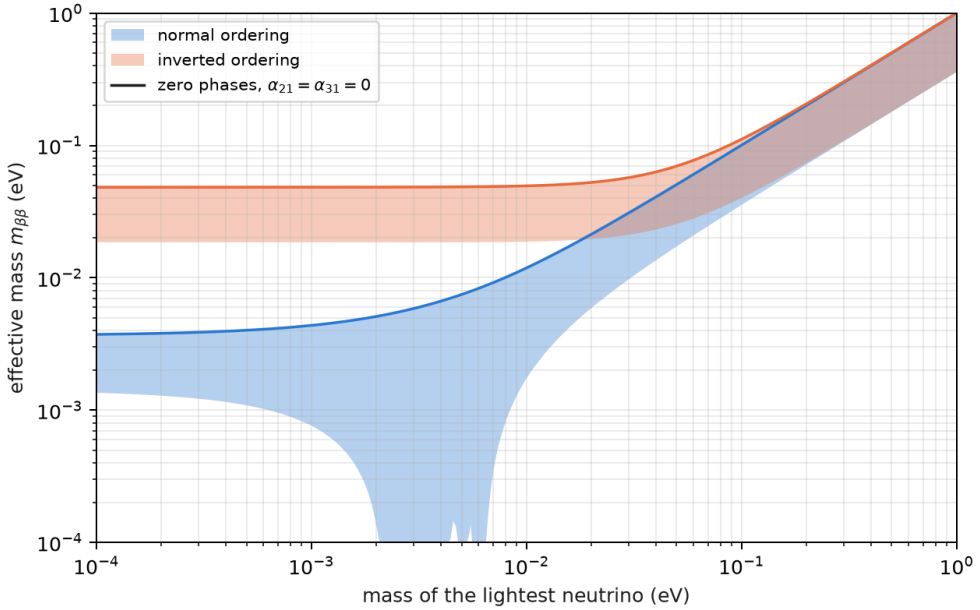


Figure 11.1: The effective mass $m_{\beta\beta}$ as a function of the lightest neutrino mass, scanning the two Majorana phases (bands), and with zero phases (lines). In normal ordering the phases can cancel it between about 2 and 7 meV.

The qg reading of this section is just the previous section said the other way round: an oscillation circuit cannot measure what this figure plots. What the scale does do is separate the two questions cleanly, the oscillation one (probabilities, qang readings) and the mass one (an amplitude with phases, outside any reading of a flavour qubit).

11.4 The sterile neutrino: a fourth qubit in the weight-1 sector

A sterile neutrino ν_s is a fourth flavour that does not interact through the weak force, but mixes with the other three. In the 3+1 model there is a fourth mass state with $\Delta m^2_{41} \sim 1$ eV², much larger than the other two, and the real mixing is a sequence of six rotations,

$$U = R_{34}\, R_{24}\, R_{14}\, R_{23}\, R_{13}\, R_{12}$$

(without CP phases, to keep it simple). In the circuit each flavour is a qubit with a single 1: the sterile one is simply **one more qubit in the weight-1 sector** (`sterile_circuit`). Each rotation is an RBS and propagation is one phase gate per qubit; everything conserves the weight and the filter applies unchanged.

At short distances, where the three-flavour oscillations have not started yet, survival is effectively two-level,

$$P(\nu_\mu \to \nu_\mu) \approx 1 - 4|U_{\mu 4}|^2 \left(1 - |U_{\mu 4}|^2\right) \sin^2 \frac{\Delta_{41}}{2}.$$

With the hypothetical values $\theta_{14} = \theta_{24} = 0.15$ and $\Delta m^2_{41} = 1$ eV², $|U_{\mu 4}|^2 = 0.022$ and the amplitude is 0.085; at $L = 0.5$ km the first minimum falls at 0.40 GeV (Figure 11.2). It is the same case as Chapter 9, with the same geometry: a precession cone with an effective mixing angle, now between ν_μ and ν_s.

On the simulated noise models, all three pre-registered predictions for the sterile neutrino held:

Table 11.2: The 3+1 sterile neutrino on simulated noise models, 4 qubits, 8000 shots.

check	result
S1 filtered error below raw at all 18 (energy, device) combinations	yes
S2 median raw / filtered ratio ≥ 2	5.5
S3 filtered minimum of $P(\nu_\mu \to \nu_\mu)$ at 0.4 GeV on all three devices	yes

At the minimum (exact 0.915), FakeBrisbane reads 0.678 raw and 0.879 filtered; FakeTorino, 0.749 and 0.890; FakeFez, 0.792 and 0.890. **The gain is larger**

than with three flavours (5.5× against 3.0–3.8×) because the circuit has four qubits: there are more ways for a shot to leave the weight-1 sector, and the filter discards them all.

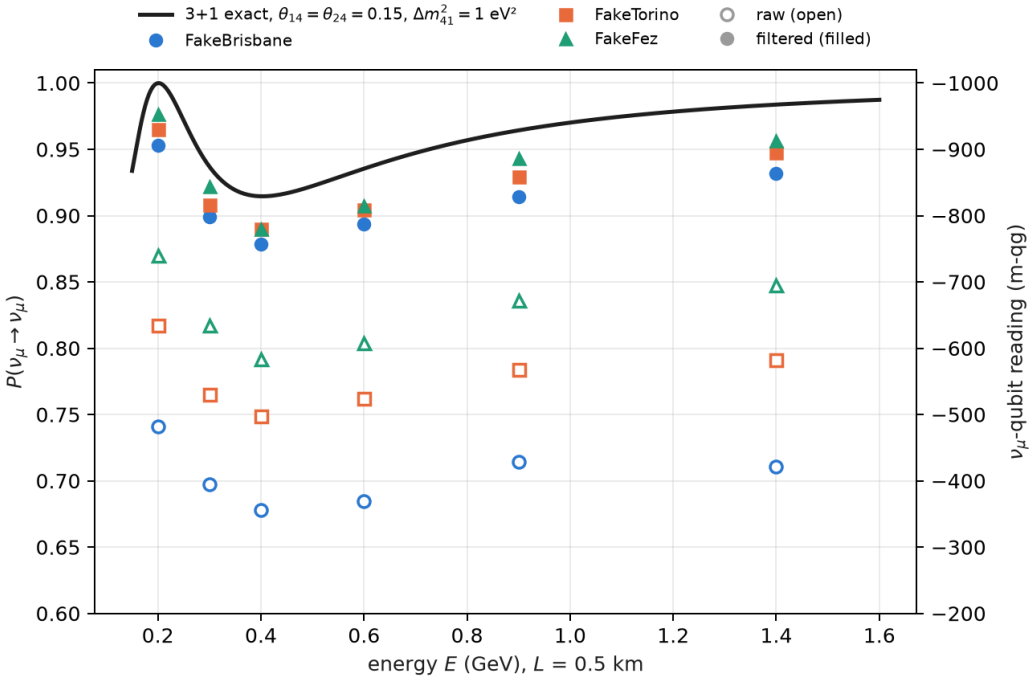


Figure 11.2: ν_μ survival at 0.5 km with a hypothetical sterile neutrino. The curve is exact; open markers are raw readings, filled ones filtered, on three simulated noise models. On the right, the ν_μ-qubit reading in m-qg.

What remains is, again, inside the sector: the filtered values still sit 0.025 to 0.036 below the exact value at the minimum, because noise that moves a 1 from one qubit to another conserves the weight and the filter cannot see it.

Open: the data. The mixing angles in this section are hypothetical, chosen so that the signal is visible; they are not a fit. Phenomenological limits and the comparison of 3+1 models with short-baseline experiments have not been done: they stay a proposal.

11.5 The four quantum numbers in qg terms

Chapter 10 ended on the question of which number is conserved. The same question, applied to the atom, gives a qg reading of the four quantum numbers (§155).

It is a translation, not new physics.

- **The magnetic number.** A register of n qubits in its symmetric subspace is a spin $j = n/2$, and in each shot the magnetic number is read from the **Hamming weight**: $m = n/2 - w$. The weight filter is the projection onto one value of m, and $qg_Z = \langle J_z \rangle / j$ for any angular momentum.
- m_s is one qubit, with $qg_Z = 2m_s$. m_l of an l orbital is the weight of a symmetric register of $2l$ qubits: the three complex p orbitals are $|00\rangle$, $(|01\rangle + |10\rangle)/\sqrt{2}$ and $|11\rangle$. The real p_x and p_y mix $m_l = \pm 1$, which is why they have $qg_Z = 0$ like p_z, but with entropy 0.63 instead of zero.
- l is the permutation symmetry, and is read with qg values of SWAP: for two qubits, $\langle J^2 \rangle = 1 + qg_{\text{SWAP}}$.
- n: by hydrogen's hidden SO(4) symmetry, shell n is two spins $(n-1)/2$, that is, two symmetric registers of $n-1$ qubits; m_l is $(n-1)$ minus the total weight of the $2(n-1)$ qubits, and the weight filter fixes it.

The qang angle of an eigenstate $|l, m\rangle$ is $\arccos(m/l)$. The textbook vector-model angle is $\arccos(m/\sqrt{l(l+1)})$, and always lies closer to the equator (Figure 11.3), because $|l, m\rangle$ has a transverse spread $\langle J_x^2 + J_y^2 \rangle = l(l+1) - m^2 > 0$:

Table 11.3: The qang angle of $|l, m\rangle$ against the vector-model angle.

l, m	qang angle	vector model	$\langle J_x^2 + J_y^2 \rangle$
1, +1	0°	45.0°	1
2, +2	0°	35.3°	2
2, +1	60.0°	65.9°	5
3, +3	0°	30.0°	3
3, +1	70.5°	73.2°	11

qg_Z measures where $\langle \mathbf{J} \rangle$ points; the vector model draws a vector of length $\sqrt{l(l+1)}$. The two agree only at $m = 0$. In a demonstration (not pre-registered) on the FakeTorino noise model, p_z ($m_l = 0$) leaks to $m_l = \pm 1$ in 19% of shots on the default qubit layout and in 2.2% on the chain chosen by `ibm.best_chain`; in both cases the weight filter returns $P(m_l = 0) = 1$.

11.6 Without qang and with qang

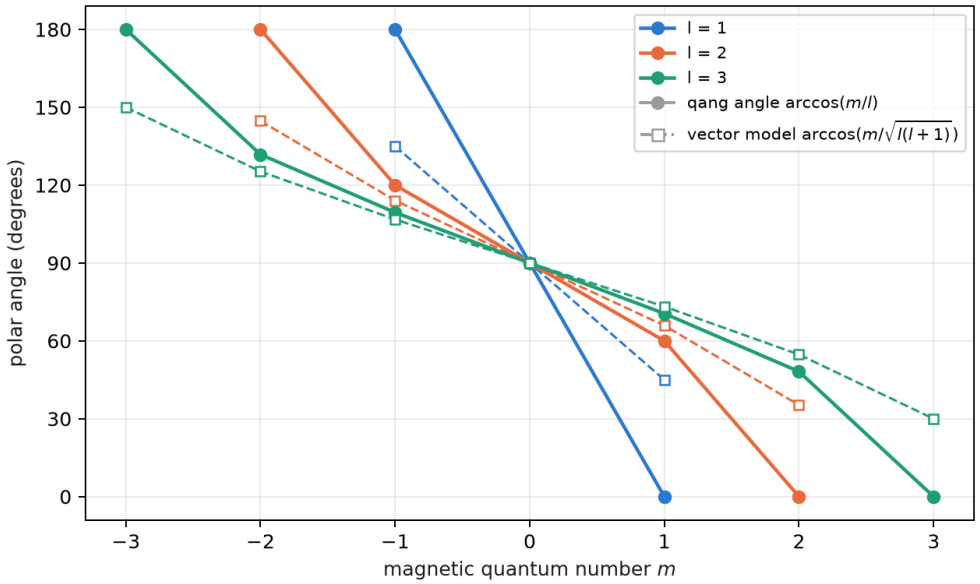


Figure 11.3: Polar angle of the eigenstates $|l, m\rangle$: the qang angle $\arccos(m/l)$ (circles) and the vector-model one (squares). They agree only at $m = 0$; the qang angle is the direction of $\langle \mathbf{J} \rangle$, the vector model draws a vector of length $\sqrt{l(l+1)}$.

Table 11.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
decide whether a parameter is observable	check whether any probability depends on it	"what changes no probability changes no reading"	**equivalent**; the qang rule says it for all readings at once
compute $m_{\beta\beta}$	$\lvert\sum U_{ei}^2 m_i\rvert$	—	**qang does not apply**: it is an amplitude with phases, not a probability
quantify the sterile signal	$\delta P = 0.085$	171 m-qg	**unit**
reduce the error with four flavours	raw reading	weight filter: 5.5×	the **filter**'s
compare the qang angle with the vector model	$\arccos(m/\sqrt{l(l+1)})$	$\arccos(m/l)$	**in favour** of qang as a reading: it is the direction of $\langle\mathbf{J}\rangle$, which is what a qg_Z measures; the vector model draws something else

Balance. The chapter has one case where qang has nothing to say ($m_{\beta\beta}$), and it is worth having it written. Its own contribution is interpretive: what a reading sees and does not see, and which angle a qg_Z really measures.

11.7 In code

```
import math
from qang import neutrino as nu
from qang import angular

# Majorana: no probability changes, in vacuum or in matter
nu.majorana_invisible(1300, 2.6, -math.pi / 2, math.pi / 2, 3 * math.pi / 4, rho=2.85)  #
nu.m_bb(0.0, a21=0.0, a31=0.0, ordering="IO")                                            # 0.048 eV
```

```
# the sterile neutrino: a fourth qubit in the weight-1 sector
P = nu.sterile_probabilities(L=0.5, E=0.4, dm41=1.0, theta14=0.15, theta24=0.15)      #
qc = nu.sterile_circuit("mu", L=0.5, E=0.4, dm41=1.0, theta14=0.15, theta24=0.15)

# the quantum numbers: the qang angle and the vector-model angle
math.degrees(angular.qang_angle_of_eigenstate(2, 1)), math.degrees(angular.vector_mode
```

11.8 Summary

- Every qang reading is a function of probabilities: a parameter that changes no probability is invisible to all of them at once.
- Majorana phases cancel in $U_M D U_M^\dagger$ (difference 4×10^{-16}); the transpiler outputs the same circuit with and without them.
- Neutrinoless double beta decay does see them: $m_{\beta\beta}$ ranges from 1.4 to 3.7 meV (normal) and from 18.7 to 48.3 meV (inverted) with a massless lightest neutrino, and can vanish in normal ordering.
- A sterile neutrino is a fourth qubit in the weight-1 sector; with hypothetical mixing, the filter gains 5.5× and places the minimum at 0.4 GeV. Limits from real data are still open.
- The four quantum numbers are read in qg: m from the Hamming weight, l from SWAP, n from the register size; the qang angle $\arccos(m/l)$ is not the vector-model one.

Bibliographic notes

The Majorana neutrino is Majorana (1937); the Majorana phases and their invisibility in oscillations, Schechter and Valle (1980) and Bilenky et al. (1980). For neutrinoless double beta decay, see Dolinski et al. (2019); the signal that motivated the sterile neutrino is Aguilar et al. (2001). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part IV

IV. Quantum teleportation in optical networks and entanglement distribution (qang.qpon)

12 Optical teleportation and the Bell-pair limit

Part IV takes the qang scale to quantum communication networks. This chapter deals with teleporting a qubit through a passive optical network: a central office shares an entangled state with a user through a fibre and a passive $1{:}N$ splitter, and the photon arrives with low probability. The question is which state to share. The usual answer, a Bell pair, stops being optimal as soon as the channel loses photons, and the optimal state is fixed by a single number: the qg_Z reading of one of its qubits. The chapter closes with the qg reading of Bell pairs in a network, which is what Chapter 13 uses as a predictor of the fidelity.

12.1 From the base: the fidelity is a probability

The fidelity of teleporting an input state is the probability that Bob, after undoing the preparation, measures 0: it is a p_0 of Chapter 1, with its reading $m = 1000(2p_0 - 1)$ m-qg and its error bar. The average fidelity over the Bloch sphere is obtained exactly with six inputs, the cardinal states $|0\rangle$, $|1\rangle$, $|{\pm}\rangle$, $|\pm i\rangle$, because they form a spherical 3-design.

The shared state is $\cos\alpha|00\rangle + \sin\alpha|11\rangle$, and each of its qubits has

$$qg_Z = \cos^2\alpha - \sin^2\alpha = \cos 2\alpha, \qquad \alpha = \frac{\arccos qg_Z}{2}.$$

The resource angle is half the qang angle of either of its qubits. The Bell pair, $\alpha = \pi/4$, is the equator: $qg_Z = 0$, $m = 0$. A less entangled resource leans toward $|00\rangle$ ($m > 0$); a "more than maximally entangled" one, in the sense of this section, leans toward $|11\rangle$ ($m < 0$). The whole design rule of this chapter is an m-qg reading of a single qubit. On the angular qga scale of Chapter 1 (Section 1.4), the Bell pair has $\alpha = 125$ qga and its qubit sits on the equator, at 250 qga.

12.2 Teleportation with a Bell pair

In the protocol of Bennett et al. (1993), Alice measures her input qubit and her half of the resource in the Bell basis, sends Bob two classical bits, and Bob applies one of four Pauli corrections. With a perfect Bell pair the fidelity is 1. With an imperfect resource, the average fidelity depends only on its singlet fraction $f = \langle\Phi^+|\rho|\Phi^+\rangle$:

$$F = \frac{2f + 1}{3},$$

and teleportation beats every classical strategy (measure and resend) only if $F > 2/3$, that is, $f > 1/2$. In the circuit the Bell measurement is in deferred form (CX, H, CX, CZ), with no mid-circuit measurements.

12.3 A channel that loses photons

In a downlink passive optical network (§157) the resource's second qubit is a photon present or absent (single rail) that travels through a fibre of transmission $T = 10^{-0.02L}$ (L in km, 0.2 dB/km) and reaches a $1{:}N$ splitter. For Bob's mode both are beam splitters, that is, RBS gates that conserve the photon number, and the photon arrives with probability

$$\eta = \frac{T}{N}.$$

A lost photon leaves Bob in $|0\rangle$: the channel is amplitude damping, which does not treat the two branches of the resource alike. The $|00\rangle$ branch has nothing to lose; the $|11\rangle$ branch loses its photon with probability $1 - \eta$. So it pays to put less weight on the fragile branch, and the optimum moves away from $\alpha = \pi/4$.

12.4 The design rule

The singlet fraction of the resource after the channel is

$$f = \frac{(\cos\alpha + \sqrt{\eta}\,\sin\alpha)^2}{2\,(\cos^2\alpha + \kappa\sin^2\alpha)},$$

with $\kappa = 1$ in the non-assisted protocol, in which only Alice and Bob act. In the **assisted** protocol the other $N-1$ users report whether the photon reached them, and those shots are discarded: it is a **sector filter** on the other users' register, and the only loss left unheralded is the fibre loss, with $\kappa = \eta + 1 - T$. Maximizing f gives

$$\tan\alpha^* = \frac{\sqrt{\eta}}{\kappa}, \qquad F^* = \frac{2+\eta/\kappa}{3}, \qquad qg_Z^* = \cos 2\alpha^* = \frac{\kappa^2-\eta}{\kappa^2+\eta}.$$

The original master index wrote the rule as $\sqrt{\eta/\kappa}$; the correct form is $\sqrt{\eta}/\kappa$, which is what the code verifies. The whole rule is the third formula: **one resource qubit measured in Z fixes α.** On the four links of the experiment:

Table 12.1: The design rule on the four links; in bold, the cases where the Bell pair does not beat the classical limit.

link	protocol	η	α^* (rad)	F^*	m^* (m-qg)	Bell pair	kept at α^*
0 km, 1:4	non-assisted	0.250	0.464	0.750	+600	0.708	1
0 km, 1:4	assisted	0.250	1.107	1.000	−600	0.933	0.40
0 km, 1:8	non-assisted	0.125	0.340	0.708	+778	**0.639**	1
0 km, 1:8	assisted	0.125	1.231	1.000	−778	0.876	0.22
20 km, 1:4	non-assisted	0.100	0.306	0.700	+819	**0.622**	1
20 km, 1:4	assisted	0.100	0.423	0.714	+663	0.672	0.95
20 km, 1:8	non-assisted	0.050	0.219	0.683	+905	**0.583**	1
20 km, 1:8	assisted	0.050	0.330	0.692	+790	0.635	0.96

Three readings of the table. **Without assistance, the Bell pair may be useless**: on a 1:8 splitter, or at 20 km, it falls below 2/3, and only the optimal resource, leaning toward $|00\rangle$, beats the classical limit. **With assistance and no fibre, the optimum has $F^* = 1$** at $\tan\alpha^* = \sqrt{N} > 1$: the resource leans further toward $|11\rangle$ than a Bell pair, and the splitter's filtering brings it back to exactly a Bell pair;

the price is the kept fraction. In that case $\kappa = \eta$ and $qg_Z^* = -(1-\eta)/(1+\eta)$, exactly the opposite reading of the non-assisted case: -600 against $+600$ m-qg at $N = 4$. **With fibre**, the unheralded loss dominates again and the assisted optimum returns to the $|00\rangle$ side (Figure 12.1).

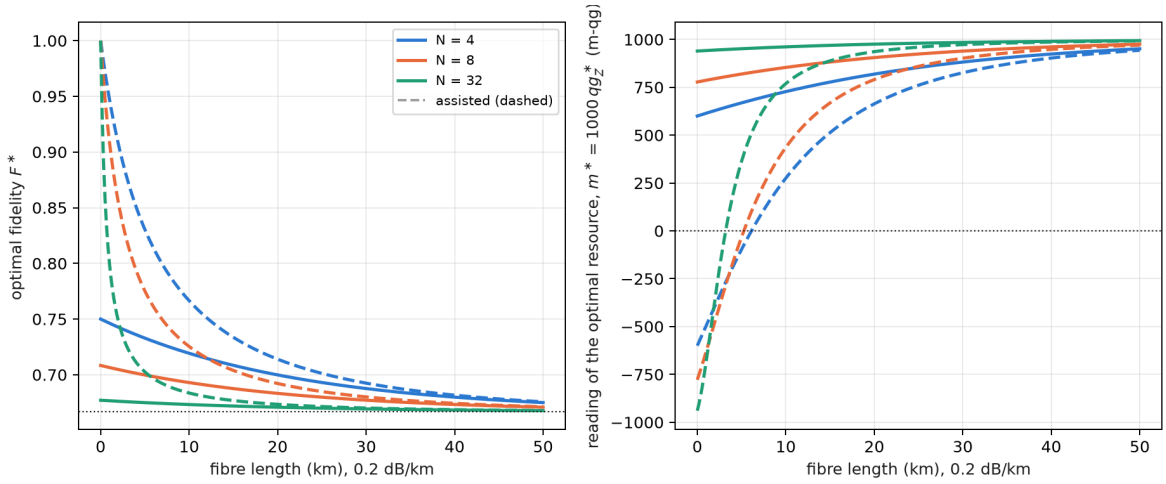


Figure 12.1: Left: optimal fidelity as a function of fibre length for three splitters; dashed lines are the assisted protocol. Right: the reading of the optimal resource in m-qg. Without fibre the assisted optimum is on the $|11\rangle$ side ($m^* < 0$); a few kilometres of fibre bring it back to the $|00\rangle$ side.

It is the discrete-variable version of a continuous-variable result: in a passive optical network, non-assisted teleportation has a finite optimal squeezing, and assisted teleportation one that goes to infinity at short distance (Yaghoobianzadeh and Salehi 2026). The non-maximally entangled resource is the analogue of finite squeezing.

12.5 The rule on simulated devices

The circuit has five qubits (input, Alice, Bob, the fibre's environment and the other users), with the fibre and the splitter as RBS gates. It was run with 14 resource angles plus $\pi/4$, six inputs and 4000 shots, on three **simulated noise models** (§157). Five of the six pre-registered predictions held:

Table 12.2: The design rule on three simulated noise models.

check	result
Q1 best non-assisted angle within 0.1 rad of α^*	12 of 12
Q2 0 km, 1:8: Bell pair below 2/3, best angle above	0.624–0.627 and 0.690–0.701 on all 3
Q3 assisted beats non-assisted	12 of 12
Q4 assisted optimum above $\pi/4$ at 0 km	1.0–1.2 rad on all 3
Q5 $\lvert F_{\text{best}} - F^* \rvert < 0.05$	**failed**: 19 of 24; the 5 misses are assisted at 0 km ($F^* = 1$, measured 0.897–0.957)
Q6 kept fraction within 0.03 of $\cos^2\alpha + \kappa \sin^2\alpha$	12 of 12 (largest gap 0.028)

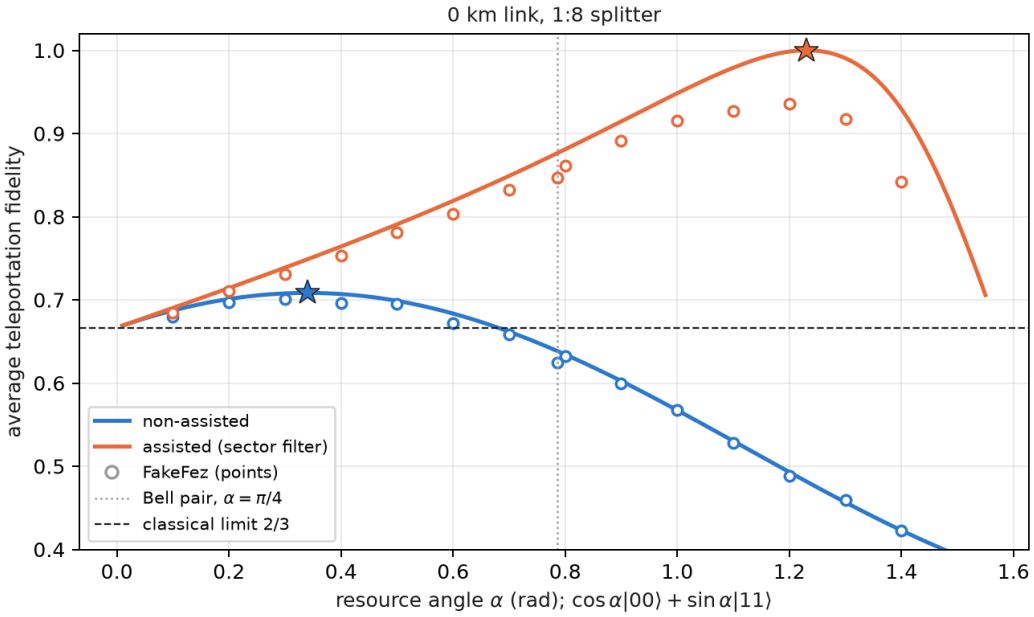


Figure 12.2: Average fidelity as a function of the resource angle, 0 km link with a 1:8 splitter. Exact curves, circles from the FakeFez noise model, stars at the optimum. The Bell pair (dotted line) falls below the classical limit without assistance; the optimal resource stays above it.

The design rule survives device noise (Figure 12.2): on a 1:8 splitter only

the optimal resource beats the classical 2/3. And the qang scale, here, **adds nothing against the noise**: the teleportation circuit (H, CX) does not conserve the weight, so the only filter is the other users' heralding. That is why the Q5 misses are where F^* is near 1: gate noise weighs more the higher the fidelity one is after. Chapter 13 asks what changes when every gate conserves the photon number.

The limits are stated: downlink only, single rail, simulated device noise, and no detector inefficiency or dark counts beyond the readout errors of the models.

12.6 Bell pairs in a network, read in qg

Before teleporting it pays to know how good the shared pair is. A pair meant to be $|\Phi^+\rangle$ is described by five qg readings: three two-qubit correlators, qg_{XX}, qg_{YY}, qg_{ZZ}, and the polar biases of each side, $m_A = \langle Z_A \rangle$ and $m_B = \langle Z_B \rangle$. Three measurement settings are enough (§57). The fidelity with $|\Phi^+\rangle$ is exact for any state,

$$F_{\Phi^+} = \frac{1 + qg_{XX} - qg_{YY} + qg_{ZZ}}{4},$$

and with it the teleportation fidelity the pair will give, $F = (2F_{\Phi^+} + 1)/3 = (3 + qg_{XX} - qg_{YY} + qg_{ZZ})/6$, without running the teleportation (Chapter 13 puts it to the test). The weights of the four Bell states say which error dominates: Φ^- is a phase flip, Ψ^+ a bit flip, Ψ^- both. With noise $p = 0.1$ on each half:

Table 12.3: qg signatures of noise on a Bell pair, $p = 0.1$ on each half.

noise	qg_{XX}	qg_{YY}	qg_{ZZ}	$m_A = m_B$	F_{Φ^+}	dominant error
dephasing	+0.64	−0.64	**+1.00**	0	0.820	Φ^-
X flip	**+1.00**	−0.64	+0.64	0	0.820	Ψ^+
Y flip	+0.64	**−1.00**	+0.64	0	0.820	Ψ^-
depolarizing	+0.81	−0.81	+0.81	0	0.858	all equal
amplitude damping	+0.90	−0.90	+0.82	**+0.10**	0.905	Ψ^+

Each Pauli noise leaves one correlator untouched, and only damping (the T1 relaxation of a memory) moves the local biases (Figure 12.3). That reading decides how to **distil**. The DEJMPS protocol combines the Φ^+ weight with the error

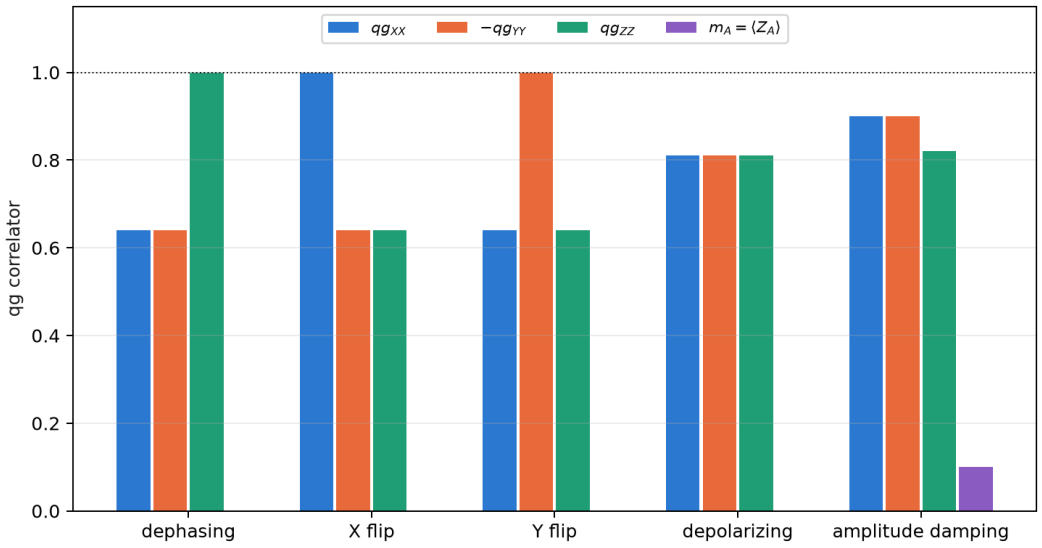


Figure 12.3: The signatures of the table: each Pauli noise leaves one correlator at its ideal value (the dotted line), and only amplitude damping moves the local polar biases.

sitting in a fixed slot, and local rotations choose which error goes there. The textbook order can make things worse: with Y-flip noise, one round lowers the fidelity from 0.820 to 0.705, and with damping it gains almost nothing (0.828). The rule "smallest estimated error in that slot", read from the three correlators with 200 shots per setting, reaches the best output in all three cases: 0.954, 0.954 and 0.915 (against an optimum of 0.920).

12.7 Without qang and with qang

Table 12.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
choose the optimal resource	optimize the Schmidt coefficient α	one reading: $qg_Z^* = (\kappa^2 - \eta)/(\kappa^2 + \eta)$ on one qubit; the sign of m^* says which side of the Bell pair it is on	**in favour** for design and verification: it is calibrated by measuring one qubit in Z
decide whether the Bell pair works	compute $F(\pi/4)$	the Bell pair is $m = 0$; the optimum, m^*	**equivalent** in value; qang places it on the scale
reduce device noise	—	the filter does not apply: teleportation does not conserve the weight	**none**
diagnose a Bell pair	full tomography	three qg correlators and two biases; the signature picks the distillation (0.705 → 0.954 with Y flips)	**in favour**: three settings instead of nine, and a direct decision rule

Balance. This is the chapter where the qang reading looks most like an engineering tool: one number sets the resource and three correlators diagnose the pair. Against noise it contributes nothing, and that is said.

12.8 In code

```
import math
from qang import qpon

eta, kappa = qpon.link(length_km=0, n_users=8, assisted=False)  # (0.125, 1.0)
qpon.optimal_angle(eta, kappa)                                  # 0.340 rad
qpon.optimal_fidelity(eta, kappa)                                  # 0.708
```

```
qpon.average_fidelity(math.pi / 4, eta, kappa)              # 0.639: the Bell pair, below
qpon.optimal_qg_z(eta, kappa)                             # 0.778 -> +778 m-qg
qpon.resource_angle_from_qg_z(0.778)                        # 0.340: the resource, from one

qc = qpon.teleport_circuit(alpha=0.340, transmission=1.0, n_users=8, input_index=2)  #
qpon.fidelity_from_correlators(xx=0.81, yy=-0.81, zz=0.81)     # 0.905 = (2 F_Phi+ + 1)/
```

12.9 Summary

- The teleportation fidelity is a probability; the resource $\cos\alpha|00\rangle + \sin\alpha|11\rangle$ has $qg_Z = \cos 2\alpha$ on each qubit, and the Bell pair is $m = 0$.
- On a channel that loses photons ($\eta = T/N$) the optimum is not the Bell pair: $\tan\alpha^* = \sqrt{\eta}/\kappa$, $F^* = (2 + \eta/\kappa)/3$, $qg_Z^* = (\kappa^2 - \eta)/(\kappa^2 + \eta)$.
- Without assistance, on a 1:8 splitter or at 20 km the Bell pair falls below 2/3 and the optimal resource stays above; with assistance and no fibre, $F^* = 1$ with a resource on the $|11\rangle$ side.
- The rule survives simulated noise (five of six predictions); the qg filter does not help here because teleportation does not conserve the weight.
- A Bell pair is read with three correlators and two biases; each Pauli noise leaves one untouched, and that signature picks the right distillation round.

Bibliographic notes

The teleportation protocol is Bennett et al. (1993), and the Bell pairs, those of Bell (1964). The continuous-variable result in passive optical networks is Yaghoobian-zadeh and Salehi (2026). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

13 Dual rail, the sector filter and long-distance fidelity

Chapter 12 ended on a limitation: the qg filter did not help against device noise, because the teleportation circuit does not conserve the weight. This chapter asks the two questions that follow. The first is whether the fidelity a resource will give can be read before teleporting, with few measurements; the answer is yes, with three qg correlators. The second is what happens if every gate conserves the photon number; the answer is dual rail, in which the filter heralds every loss, and which at 20 km beats the noiseless single-rail optimum until false heralds set a floor. The chapter closes Part IV with entanglement swapping along a chain, read with the same correlators.

13.1 From the base: three qg readings predict a fidelity

Each two-qubit correlator is a qg reading like those of Chapter 1: the parity of two bits measured in a basis, $qg_{PP} = 1 - 2P(\text{odd parity})$, with its m-qg reading and its error bar $\sqrt{(1 - qg^2)/N}$. Chapter 12 showed that the fidelity with $|\Phi^+\rangle$ is linear in three of them; so is the average teleportation fidelity (§158):

$$F = \frac{3 + qg_{XX} - qg_{YY} + qg_{ZZ}}{6}.$$

It is a **linear read-out of the resource alone**, without the Bell measurement: three circuits with 4–5 two-qubit gates, against 10–11 for the teleportation. Being linear, its error bar combines the three: with N shots per correlator, $\sigma_F = \sqrt{\sum(1 - qg^2)/N}/6$, about 0.0025 with $N = 10\,000$ near $F \approx 0.7$.

13.2 The fidelity from three correlators

On the four links of Chapter 12, with both protocols and on three **simulated noise models**, the fidelity predicted by the correlators was compared with

the teleported one. In all 24 combinations the difference is at most 0.029 (Figure 13.1): **three short circuits predict the teleported fidelity**.

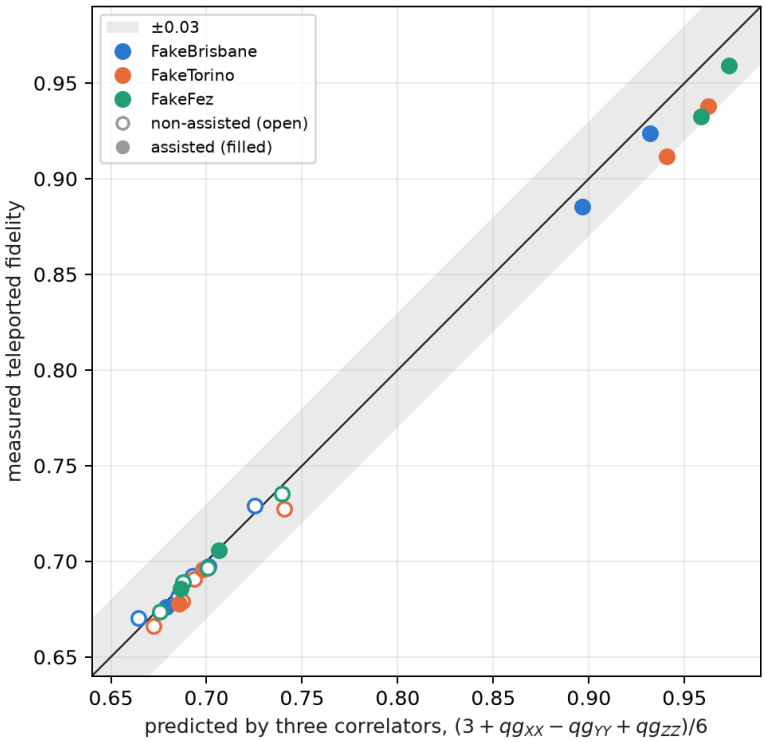


Figure 13.1: Measured teleported fidelity against the one predicted by three correlators, 4 links × 2 protocols × 3 simulated noise models. All 24 fall within ±0.03 of the diagonal.

The pre-registered predictions wanted more: that the difference would isolate the noise of the Bell measurement. That failed. The prediction fell below the noiseless optimum in 5 of 24 cases (A1), was not always above the teleported one (A2, 3 of 24, all within two shot errors), and the difference was not always larger with assistance (A3, 2 of 12). The reason is that the three correlator circuits read two data qubits and the herald, so they carry almost all the noise the teleportation sees. **They work as a predictor, not as a decomposition.**

13.3 Dual rail: every loss is heralded

In single rail a qubit is a photon present or absent, and losing the photon turns a $|1\rangle$ into a valid $|0\rangle$: the loss is invisible. In **dual rail** each qubit is one photon in

two modes, $|0_L\rangle = |10\rangle$ and $|1_L\rangle = |01\rangle$. The logical gates are built to conserve the photon number: the logical CNOT is a controlled SWAP of the target's modes, the logical CZ a CZ between modes, and one-qubit gates act on the subspace $\{|10\rangle, |01\rangle\}$. Every loss leaves Bob's pair with weight 0, and the qg filter (weight 1 on each pair) **heralds every loss, fibre included, without help from the other users**.

Without gate noise, the result is clean:

$$F_{\text{filtered}} = 1 \text{ at any distance}, \qquad \text{kept} = \eta, \qquad F_{\text{raw}} = \frac{1+\eta}{2},$$

where the raw value assumes Bob guesses when no photon arrives, and falls below the single-rail optimum $(2+\eta)/3$. The filter turns loss into rate: the fidelity does not fall with distance; the number of useful pairs does. The price is the circuit: 65–66 two-qubit gates, six times those of single rail.

13.4 What dual rail measures on simulated devices

With the Bell pair ($\alpha = \pi/4$), six inputs and 10 000 shots (§158), only one of the five dual-rail predictions held (B4: the raw value is below single rail in all 12 combinations, 0.516–0.595). The other four failed, mostly for the same reason:

Table 13.1: Dual rail with the qg filter on three simulated noise models; links in the order (0, 4), (0, 8), (20, 4), (20, 8).

simulated device	filtered, (0 km, 1:4) → (20 km, 1:8)	kept / η	20 km: filtered vs noiseless single-rail F^*
FakeBrisbane	0.786; 0.720; 0.709; 0.650	0.91 → 1.98	0.709 / 0.714 (1:4); 0.650 / 0.692 (1:8)
FakeTorino	0.847; 0.790; 0.783; 0.727	0.97 → 2.14	0.783 / 0.714; 0.727 / 0.692
FakeFez	0.869; 0.814; 0.791; 0.710	0.98 → 1.74	0.791 / 0.714; 0.710 / 0.692

False heralds set a floor (B1, B2 and B5 failed). When the photon is lost, a gate or readout error can still leave Bob's pair with weight 1: at 20 km, $(\text{kept} - \eta)/(1-\eta)$ is 3 to 6%. Those shots carry a fidelity of about 1/2, and as η falls they

weigh more among the kept ones. With a false-herald probability p_f, the filtered fidelity is

$$F_{\text{filtered}} = \frac{\eta\, F_0 + (1-\eta)\, p_f/2}{\eta + (1-\eta)\, p_f},$$

which falls with distance even though the channel adds no error (Figure 13.2). It is the dark-count floor of a real heralded link, produced here by gate and readout errors.

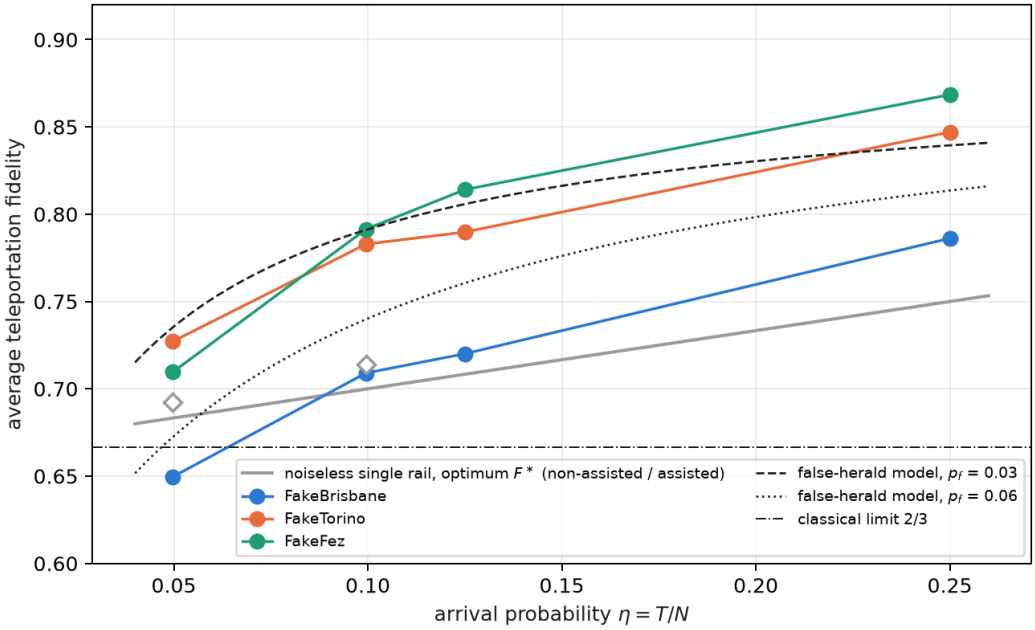


Figure 13.2: Filtered dual-rail fidelity as a function of the arrival probability, on three simulated noise models, against the noiseless single-rail optimum (grey line; the diamonds are the assisted one at 20 km). The black curves are the false-herald model with $F_0 = 0.87$, an illustration, not a fit.

Heralding still wins at 20 km. In a post hoc analysis, not pre-registered, on FakeTorino and FakeFez the filtered dual rail beats the **noiseless** single-rail optimum on both splitters: 0.783 and 0.791 against 0.714 at 1:4; 0.727 and 0.710 against 0.692 at 1:8. It does so with six times the two-qubit gates. The raw value loses everywhere (B4), so the gain is the filter's. B3, the pre-registered version of that comparison (against the noisy teleported fidelity), failed only on FakeBrisbane at (20 km, 1:8).

The comparison with the continuous-variable work is qualitative: there the fidelity is deterministic and falls with distance and with N; here the filter keeps the fidelity high and pays in rate (kept $\approx \eta$), until false heralds, not loss, set the limit.

13.5 Entanglement swapping along a chain

To go further than one link, a network chains pairs: an intermediate node makes a Bell measurement between its two halves, and the ends become entangled (entanglement swapping). This section is new in the book: it is exact density-matrix algebra, verified in the figure script, with no device run.

For pairs diagonal in the Bell basis, which are what the Pauli noises of Chapter 12 leave, the result in qg is a one-line rule: **the correlators multiply**. After swapping two pairs, with the Pauli correction that belongs to each outcome,

$$qg_{PP}^{(1\to3)} = qg_{PP}^{(1\to2)} \cdot qg_{PP}^{(2\to3)}, \qquad P \in \{X, Y, Z\},$$

and along a chain of n equal links each correlator is raised to the n-th power. With depolarized pairs of parameter w ($qg_{XX} = qg_{ZZ} = w$, $qg_{YY} = -w$), the teleportation fidelity of the final pair is $F = (1 + w^n)/2$, and it beats the classical limit while $w^n > 1/3$, that is, up to $n \leq \ln 3 / \ln(1/w)$: 21 links at $w = 0.95$, 10 at $w = 0.9$ and 4 at $w = 0.8$ (Figure 13.3).

The correlator reading also says which error accumulates. With dephasing, $qg_{ZZ} = 1$ on every link and therefore along the whole chain, while qg_{XX} and qg_{YY} decay as 0.64^n: the final pair keeps its Z correlations and loses its phase ones, and its fidelity tends to 2/3 from above. It is the same signature as in the noise table of Chapter 12, propagated along the chain, and it is what decides where to distil.

The link with dual rail is direct: in a chain of dual-rail pairs, every intermediate Bell measurement also conserves the photon number, and the filter heralds the losses of all links at once. That experiment has not been done.

13.6 What remains open

Three topics of the index have no support in the repository yet, and stay proposals, not results:

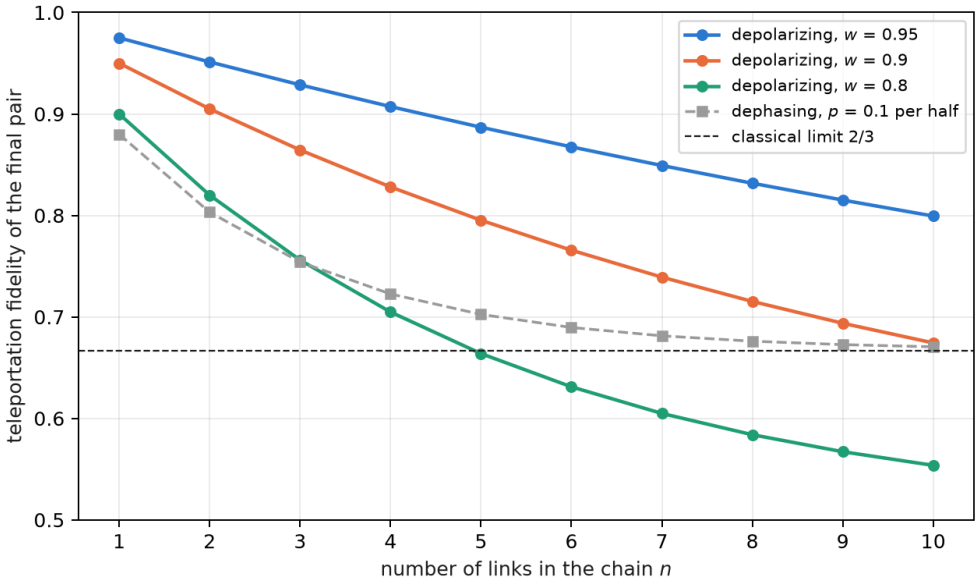


Figure 13.3: Teleportation fidelity of the final pair of a chain of n links without distillation. With depolarizing noise it falls as $(1+w^n)/2$ and crosses the classical limit; with dephasing, qg_{ZZ} stays at 1 and the fidelity tends to 2/3 without crossing it.

- routing and passive quantum switching in multi-carrier metropolitan topologies;
- a dual-rail repeater chain with the filter on simulated devices, which would test this section's multiplication rule against real gate noise;
- gate teleportation and its reading in qg correlators.

13.7 Without qang and with qang

Table 13.2: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
predict the teleportation fidelity	teleport and measure (10–11 two-qubit gates)	three correlators of the resource (4–5 gates): largest error 0.029 over 24 cases	**in favour** in circuit cost; the formula is known, qang makes it a linear reading with an error bar
herald losses	single rail: the loss is invisible	dual rail with the weight filter: every loss is heralded	the **filter and the encoding**'s, not the scale's
propagate noise along a chain of swaps	compose channels	the qg correlators multiply; $F = (1 + w^n)/2$	**in favour** in clarity: a one-line rule and the signature of which error accumulates

Balance. The qg correlators are the part of qang that pays most in networks: they predict and propagate with linear rules. The chapter's large gain, dual rail, belongs to the encoding and the filter.

13.8 In code

```
from qang import qpon
```

```
# the teleportation fidelity, from three correlators of the resource
qpon.fidelity_from_correlators(xx=0.84, yy=-0.81, zz=0.73)      # 0.897
qc = qpon.resource_circuit(alpha=1.231, transmission=1.0, n_users=8, basis="X")

# dual rail: noiseless closed forms and the 8-qubit circuit
qpon.dual_rail_fidelities(alpha=0.7854, eta=0.05)          # filtered 1.0, raw 0.525,
qc_dual = qpon.dual_rail_teleport_circuit(alpha=0.7854, eta=0.05, input_index=2)
```

13.9 Summary

- The teleportation fidelity is linear in three qg correlators of the resource, $F = (3 + qg_{XX} - qg_{YY} + qg_{ZZ})/6$; in 24 simulated cases it predicts the teleported one to 0.029, although it does not isolate the noise of the Bell measurement.
- In dual rail every gate conserves the photon number and the filter heralds every loss: without noise, $F = 1$ at any distance with rate η.
- With noise, false heralds (3–6% at 20 km) set a floor and the filtered fidelity falls with distance; even so, on two of three models it beats the noiseless single-rail optimum at 20 km (post hoc analysis).
- Along a chain of swaps, the correlators of Bell-diagonal pairs multiply; with depolarization $F = (1 + w^n)/2$, and with dephasing qg_{ZZ} survives whole.
- Routing, repeaters on devices and gate teleportation remain open.

Bibliographic notes

Dual-rail encoding in linear optics is that of Knill et al. (2001); entanglement swapping, Żukowski et al. (1993), and quantum repeaters, Briegel et al. (1998). The underlying teleportation is Bennett et al. (1993). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part V

V. Relativity, time loops and black-hole evaporation (`qang.postselection`)

14 Postselected closed timelike curves (P-CTC) and the resolution of paradoxes

Part V deals with two information-theory models that use the vocabulary of relativity: time loops and black-hole evaporation. It is worth saying from the start: **they are models of information theory, not statements about spacetime** (§164). Both are teleportation with postselection on a Bell outcome, that is, a filter on the Bell sector, and that is why they belong in this book. This chapter treats the time loop in its two known formulations, the postselected one (P-CTC) and Deutsch's (D-CTC), and what a qg reading tells apart between them.

14.1 From the base: the success probability has qang angle φ

In the simplest model the qubit that "travels to the past" undergoes a rotation $R_X(\varphi)$, and the postselection succeeds with probability

$$P_{\text{success}} = \cos^2 \frac{\varphi}{2}.$$

It is the probability equation of Chapter 1 with $p_0 = P_{\text{success}}$, so its qang angle is exactly the rotation, $\arccos(2P - 1) = \varphi$, its reading is $m = 1000 \cos\varphi$, and the success rate measured over N attempts gives φ with uncertainty $1/\sqrt{N}$. The grandfather paradox, $\varphi = \pi$, sits at the south pole: $P = 0$, $m = -1000$.

14.2 Two ways to close a time loop

Deutsch (D-CTC, 1991). (Deutsch 1991) The system in the loop must be self-consistent: its state ρ_C is a fixed point of the evolution, $\rho_C = \mathrm{Tr}_S[U(\rho_S \otimes \rho_C)U^\dagger]$. At least one fixed point always exists, so there is never a paradox, but there may be

several, and an extra rule is needed to choose; the usual one takes the maximum-entropy fixed point.

Postselection (P-CTC; Bennett and Svetlichny, Lloyd and co-workers, 2011). (Svetlichny 2011; Lloyd et al. 2011) The qubit sent to the past, C, is half of a Bell pair (C, A). After C interacts with the system S, (C, A) is projected back onto the Bell state: it is a teleportation of C to its own past, with the Bell measurement forced to its trivial outcome. The effective evolution of S is

$$\rho_S \mapsto \frac{\mathrm{Tr}_C(U)\,\rho_S\,\mathrm{Tr}_C(U)^\dagger}{\mathrm{Tr}[\ldots]},$$

which is not linear because of the normalization. In the circuit the projection is a Bell measurement keeping the shots that read 00: a filter on the Bell sector.

14.3 The success law and the grandfather paradox

With $R_X(\varphi)$ on C alone, $\mathrm{Tr}_C(U) = \mathrm{Tr}\,R_X(\varphi)/2 = \cos(\varphi/2)$, and the success is $\cos^2(\varphi/2)$. At $\varphi = \pi$ the qubit flips its own past: it is the grandfather paradox, and in P-CTC **it has probability zero**. It is not resolved: it does not happen. In D-CTC the same case has the maximally mixed state as its fixed point, and it happens with probability 1, like a coin.

On three **simulated noise models**, with 8000 shots (§164), the success law holds within 0.05 except at $\varphi = 0$ on two devices (C1 failed: 0.950 and 0.931 against 1), and the paradox is suppressed to the noise floor on two of three (C2 failed on FakeTorino):

Table 14.1: The P-CTC success law at the extremes, simulated noise models.

simulated device	success at $\varphi = 0$ (exact 1)	success at $\varphi = \pi$, the paradox (exact 0)
FakeBrisbane	0.950	0.019
FakeTorino	0.931	0.056
FakeFez	0.983	0.009

Figure 14.1 shows something specific to the angle reading. In the interior, the qang angle of the measured rate reproduces φ to a few hundredths of a radian; at the poles it deviates by up to half a radian. It is the non-linearity of the milli-qang in angle from Chapter 1: next to the pole, 1 m-qg is 2.56°, and the devices'

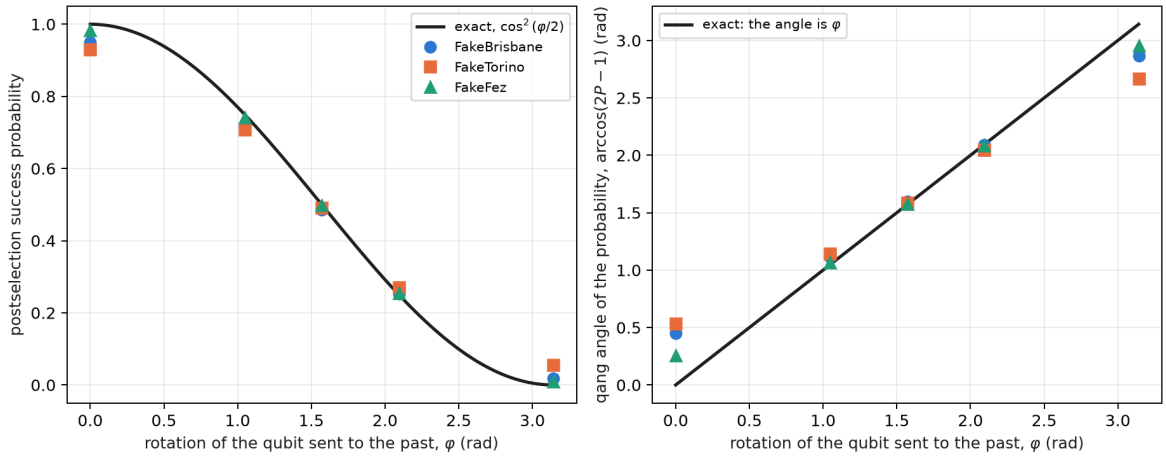


Figure 14.1: Left: postselection success probability against the rotation, exact and on three simulated noise models. Right: its qang angle, which is the rotation itself; the largest deviations are at the poles, where a small probability error is a large angle error.

34–138 m-qg of error at $\varphi = 0$ become 0.26–0.53 rad, that is, 41.6–84.6 qga on the angular scale of Chapter 1 (Section 1.4). On that scale the success rate is $\cos^2(\pi\varphi_{qga}/1000)$, and the grandfather paradox sits at $\varphi = 500$ qga.

14.4 The rotation controlled by the system

If it is the system S that controls the rotation of C, the partial trace is an operator on S:

$$\mathrm{Tr}_C(U) = \mathrm{diag}\big(1,\ \cos\tfrac{\varphi}{2}\big)\,,$$

and the branch of S that would produce the paradox is damped by $c = \cos(\varphi/2)$. From $|+\rangle$, the postselected state of S has

$$qg_Z(S) = \frac{1-c^2}{1+c^2}, \qquad qg_X(S) = \frac{2c}{1+c^2},$$

with success $(1+c^2)/2$. At $\varphi = \pi$, $qg_Z(S) = 1$: the system is forced into the branch that does not flip the past. On the simulated models FakeFez follows the

curve within 0.061 over the whole range, FakeBrisbane reads 0.903 at $\varphi = \pi$, and FakeTorino loses most of the suppression (0.600): prediction C3 failed.

14.5 What a qg reading tells apart between D-CTC and P-CTC

This section is new in the book: it is exact density-matrix algebra, computed in the figure script, with no device run. The same controlled circuit is taken, with S in $|+\rangle$, and the output of S is computed in both formulations.

In D-CTC the channel on C is $\rho_C \mapsto \frac{1}{2}\rho_C + \frac{1}{2}R_X\rho_C R_X^\dagger$, and every state that commutes with $R_X(\varphi)$ is a fixed point. With the maximum-entropy rule, $\rho_C = I/2$, and the output of S is

$$qg_Z(S) = 0, \qquad qg_X(S) = \cos\frac{\varphi}{2}, \qquad r = \cos\frac{\varphi}{2}.$$

The two formulations separate in different readings (Figure 14.2):

Table 14.2: The system S after the controlled loop, in both formulations.

φ	P-CTC: m_Z, m_X (m-qg)	D-CTC: m_Z, m_X (m-qg)	P-CTC success
0	0, 1000	0, 1000	1
$\pi/3$	143, 990	0, 866	0.875
$\pi/2$	333, 943	0, 707	0.750
$2\pi/3$	600, 800	0, 500	0.625
π	1000, 0	0, 0	0.500

The difference is qualitative. **P-CTC changes the populations** of S (the paradoxical branch loses weight) and leaves a pure state; **D-CTC does not touch the populations** and only removes coherence. A Z reading of S already decides between the two models, and the simulated devices, with all their noise, sit on the P-CTC side, as they must, because the circuit is the P-CTC one.

Deutsch's ambiguity can be read too. At $\varphi = \pi/2$ the state $|+\rangle$ is also a fixed point, and it gives the same $qg_Z(S) = 0$ and the same $qg_X(S) = 0.707$, but $qg_Y(S) = -0.707$ and radius 1, against $qg_Y = 0$ and radius 0.707 with $I/2$. The two D-CTC solutions differ in the local radius of Chapter 5: the maximum-entropy rule picks the fixed point that leaves S most mixed.

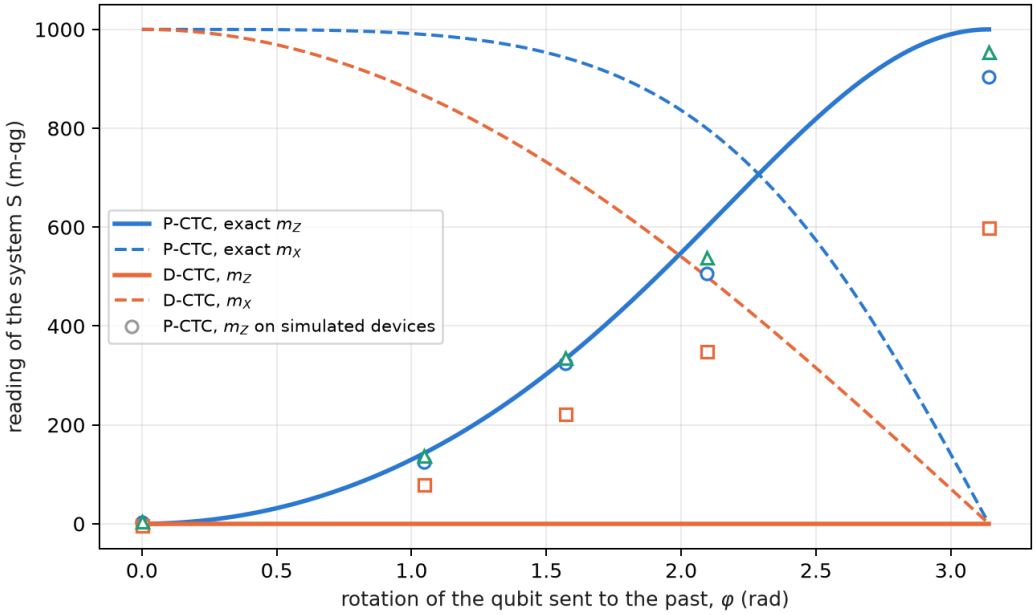


Figure 14.2: Readings of the system S after the controlled loop. P-CTC (blue) moves the population, m_Z, toward the paradox-free branch; D-CTC (orange) leaves $m_Z = 0$ and only decoheres, lowering m_X. The markers are the P-CTC on three simulated noise models.

14.6 The cost of postselection

"Time travel" in this model costs shots. Each useful shot needs, on average, $1/P_{\text{success}}$ shots: 2 at $\varphi = \pi/2$, 4 at $2\pi/3$, 400 at a tenth of a radian from the paradox (Figure 14.3). With the controlled rotation the cost never exceeds 2, because half of the superposition of S pays nothing. It is the same accounting as the sector filter of the previous chapters, with one difference: here the discarded fraction is not noise but the model's physics, and it cannot be reduced.

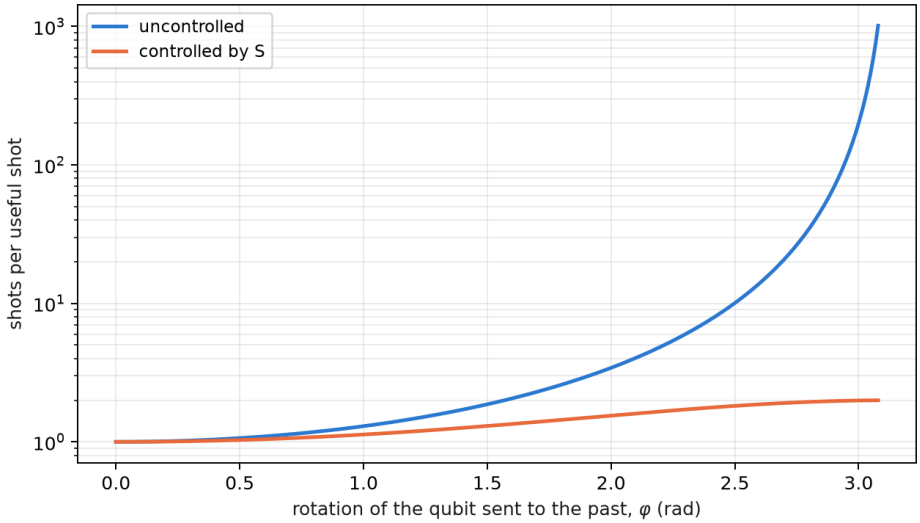


Figure 14.3: Shots needed per useful postselected shot, uncontrolled and with the rotation controlled by S.

14.7 What remains open

Two topics of the index have no support yet and stay proposals: the non-linearity of P-CTCs applied to discriminating non-orthogonal states (the circuit and the prediction are missing), and a device run of the D-CTC against P-CTC comparison, which would require preparing Deutsch's fixed point and not only postselecting.

14.8 Without qang and with qang

Table 14.3: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read the loop's rotation	invert $\cos^2(\varphi/2)$	the qang angle of the success rate is φ, with $\sigma = 1/\sqrt{N}$	**equivalent** in the interior; at the poles the angle reading **amplifies** the error (34–138 m-qg become 0.26–0.53 rad), which counts **against**
verify the suppression of the paradox	success probability and the state of S by tomography	$m_Z(S)$: 903, 600 and 955 m-qg against 1000	**unit**; the suppression is the postselection's
tell D-CTC from P-CTC	compare density matrices	P-CTC moves m_Z; D-CTC does not and lowers m_X	**in favour**: a single Z reading decides the model
tell apart the D-CTC fixed points	entropy of each solution	the local radius of Chapter 5: 0.707 against 1	**in favour**: the ambiguity becomes a reading
reduce noise	—	the postselection is the filter; the qang scale adds nothing	**none**

Balance. In this chapter qang pays as a diagnostic tool: with one or two readings it says which loop model is operating and which of Deutsch's solutions was chosen. The success law is the probability equation of Chapter 1, and with it come its limits next to the poles.

14.9 In code

```
import math
from qang import postselection as ps
```

```
ps.pctc_success(2 * math.pi / 3)             # 0.25: postselection success
math.acos(2 * ps.pctc_success(2.0) - 1)       # 2.0: the qang angle is the rotation
ps.pctc_controlled_qg_z(math.pi)             # 1.0: the paradoxical branch, suppressed
qc = ps.pctc_circuit(math.pi, controlled=True)  # S, C, A; success = C and A read 0
```

14.10 Summary

- The time loops of this chapter are information models: teleportation with postselection on a Bell outcome.
- In P-CTC the success probability is $\cos^2(\varphi/2)$ and its qang angle is φ; the grandfather paradox has probability zero and, on two of three simulated models, sits at the noise floor.
- With the controlled rotation, $qg_Z(S) = (1 - c^2)/(1 + c^2)$: the paradoxical branch is suppressed; FakeTorino loses most of the suppression.
- D-CTC does not change the populations of S and only removes coherence; a Z reading separates the two models, and the local radius separates Deutsch's fixed points.
- Postselection costs $1/P$ shots per useful shot; it is the model's physics, not noise, and it cannot be reduced.

Bibliographic notes

The self-consistent model of closed timelike curves is Deutsch (1991); the postselection model and its experimental test, Lloyd et al. (2011), and the comparison between the two, Svetlichny (2011). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

15 Black-hole thermodynamics, quantum chaos and the Yoshida–Kitaev decoder

This is the second model of Part V, and the warning of chapter 14 holds again: **it is an information-theoretic model, not a statement about real black holes** (§164). A black hole of a few qubits, a unitary that scrambles information, and a teleportation postselected on a Bell outcome: the same filter on the Bell sector as in the previous chapter, now set to recover a message that fell into the horizon.

15.1 From the base: success 1/4 and the equator of the fidelity

The protocol produces two numbers, and both are probabilities, so the probability equation of chapter 1 reads them directly. The postselection succeeds with probability $P = 1/4$: its reading is $m_P = 1000(2P - 1) = -500$ and its qang angle $\arccos(-1/2) = 2\pi/3$. The fidelity F of the recovered pair with a Bell pair is the probability that the Bell measurement gives its trivial outcome, and it reads the same way, $m_F = 1000(2F - 1)$.

The scale has a concrete virtue here: **the threshold that matters falls on the equator**. A Bell fidelity above $1/2$ certifies entanglement, and $F = 1/2$ is $m_F = 0$. Without scrambling the pair comes out maximally mixed, $F = 1/4$, and reads -500; with scrambling, $F = 1$, and it reads 1000. Decoding is crossing the equator.

15.2 The information paradox and complementarity

Hawking (1976) showed that a black hole radiates with a thermal spectrum, and that the radiation does not seem to depend on what fell in. If the hole evaporates completely, a pure state would end as a mixed one: evolution would not be unitary. Complementarity (Susskind et al. 1993) and, above all, Page's argument (1993)

led to the opposite position, the majority one today: the information comes out in the radiation, encoded in fine correlations.

Hayden and Preskill (2007) asked when it comes out. Their answer, for a black hole past the middle of its life whose dynamics scrambles information fast, is "almost at once": a message of a few qubits thrown into the hole can be rebuilt from just a few more qubits of radiation, given the early radiation. Yoshida and Kitaev (2017) gave an explicit decoder, and it is the one run in this chapter.

15.3 Scrambling: a 3-qubit micro black hole

The toy black hole has 3 qubits: the message A and two qubits B_1, B_2 that were already inside. Its dynamics is a unitary U on the three, and "U scrambles information" has an operational definition: **every single-qubit Pauli operator ends up, after U, acting on all three qubits**. For a Clifford this is checked with no approximation, by evolving the 9 Paulis X, Y, Z of each qubit. `find_scrambler` takes the first random Clifford (seed 164) that meets the condition, and its 9 Paulis have weight 3.

Scrambling is the quantum version of chaos: a local error (a Pauli) becomes, after the dynamics, an error of the whole system, and no measurement of one qubit detects it. So it is no surprise that a black hole "hides" the message, and it is a surprise that the message can be recovered.

15.4 The Hayden–Preskill protocol and the decoder

The circuit has 8 qubits. The message A is entangled with a reference R, the one who throws it; the black hole B is entangled with the early radiation B', held by the outside observer. After U on (A, B_1, B_2), one qubit D (the output of B_2) is radiated. The observer holds D and B', and prepares a copy of the experiment: a pair (A', R') and the conjugate unitary U^* on (A', B_1', B_2'). Then (D, D') **is postselected on the Bell state**. On success, the entanglement of A with R has moved to R': the message was teleported out of the interior.

For a scrambler, the postselection succeeds with probability $1/d_A^2 = 1/4$ and the fidelity of the pair (R, R') is 1. Without scrambling ($U = I$), the postselection almost always succeeds, because D carries nothing of A, and the fidelity is that of a random pair, $1/4$. The postselection is, as in chapter 14, a filter on the Bell sector of (D, D').

15.5 The exact law $P \cdot F = 1/4$

This section is new in the book. In Yoshida and Kitaev's probabilistic decoder, the success and the fidelity are not independent:

$$P\,F = \frac{1}{d_A^2} = \frac{1}{4}.$$

The figure script checks it on exact states over 60 random Cliffords and 25 Haar unitaries of 3 qubits, and at every size of the next section: on 3 qubits the largest deviation is 3×10^{-16}. On the chapter 1 scale the law is a hyperbola,

$$(1000 + m_P)(1000 + m_F) = 10^6, \qquad \text{or} \qquad \cos\frac{\theta_P}{2}\,\cos\frac{\theta_F}{2} = \frac{1}{2},$$

with a practical consequence: **without noise, the success already tells the fidelity**. A success of $1/4$ ($m_P = -500$) forces $F = 1$; a success of 1 ($m_P = 1000$) forces $F = 1/4$. Among random 3-qubit Cliffords, 27 of 60 decode perfectly.

With noise the law breaks, and the direction in which it breaks is what makes it useful. Full noise takes (D, D') to the mixed state, whose success is *also* $1/4$, with $F = 1/4$: the success alone does not tell a perfect decoder from one destroyed by noise. The product does: $4PF$ is 1 for the first and $1/4$ for the second.

15.6 On simulated noise models

With 8000 shots on three **simulated noise models** (§164), the circuit with the scrambler, 90 two-qubit gates after transpilation, gives:

Table 15.1: The Yoshida–Kitaev decoder on three simulated noise models.

simulated device	success P (m_P)	fidelity F (m_F)	$4PF$	no scrambling: F (m_F)
exact	0.250 (−500)	1 (1000)	1	0.250 (−500)
FakeBrisbane	0.260 (−481)	0.547 (93)	0.568	0.248 (−504)
FakeTorino	0.248 (−504)	0.578 (156)	0.573	0.250 (−501)
FakeFez	0.246 (−508)	0.742 (484)	0.731	0.250 (−500)

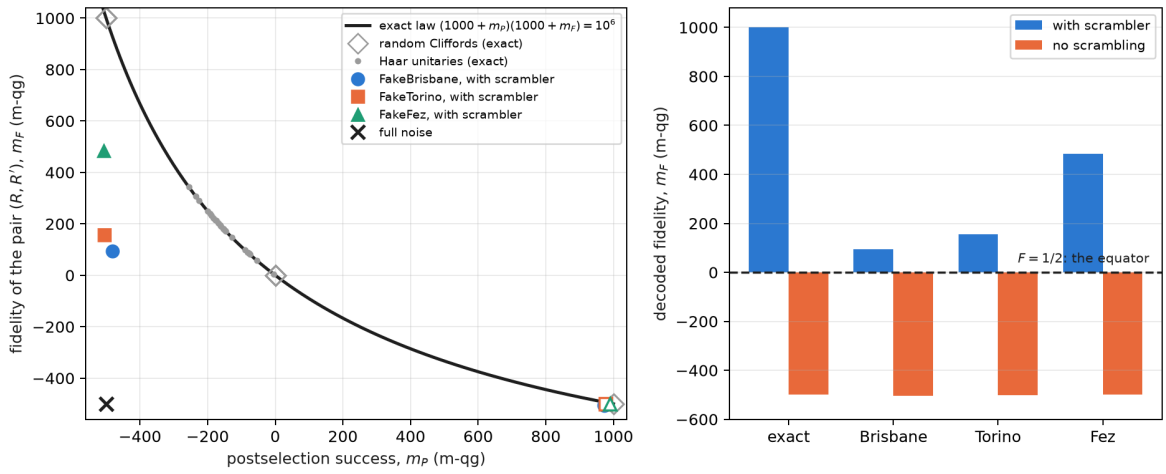


Figure 15.1: Left: the success–fidelity plane in m-qg. The curve is the exact law $PF = 1/4$; Cliffords and Haar unitaries fall on it; the devices with the scrambler (filled markers) fall below it, towards the full-noise point; without scrambling (open) they stay on it, at $m_P \approx 1000$. Right: the decoded fidelity in m-qg; the three devices with the scrambler cross the equator.

The three runs with the scrambler cross the equator (H3), and the success stays within 0.05 of $1/4$ (H2). But Figure 15.1 shows why H2 is no evidence of decoding: the success of the three devices is the one full noise would give. What sets FakeFez apart from the other two is the fidelity, and $4PF$ sums it up in one number: 0.73 against 0.57. Prediction H1 (a fidelity advantage above 0.3 over the unscrambled case on all three) missed by 0.001 on FakeBrisbane.

15.7 Larger black holes

This section is new as well: exact states, gate counts and a simulation with the FakeFez noise model, no device run. The circuit is generalised to an n-qubit black hole with k radiated qubits: $2n + 2$ qubits in all, 16 for $n = 7$.

Two results organise the topic. The first: **with a single radiated qubit, a generic black hole is not enough**. For a Haar unitary the average success is $P \approx \frac{1}{4} + 4^{-k} - 4^{-k-1}$, and by the law $F = 1/(4P)$: with $k = 1$, $F \approx 4/7 = 0.57$, barely above the equator, whether $n = 3$, 5 or 7. One has to radiate a little more than the message: with $k = 2$, $F \approx 0.84$; with $k = 3$, $F \approx 0.95$ (Figure 15.2, left). It is the quantitative version of Hayden and Preskill's "a few more qubits". The

3-qubit Clifford scrambler, with $F = 1$ at $k = 1$, is a special case, not the typical one.

The second: **the gate cost separates the two kinds of scrambler**. A Clifford scrambler is transpiled as a Clifford circuit: 35, 126 and 252 two-qubit gates for $n = 3$, 5 and 7. A Haar unitary is synthesised as a generic unitary, and its cost grows as 4^n: 92, 1908 and 32 680 (Figure 15.2, right). There is also a lesson for the §164 run: the same 3-qubit scrambler, handed to the transpiler as a matrix, was synthesised with 90 gates; handed over as a Clifford, with 35, and on the FakeFez model the fidelity rises from 0.742 to 0.887 (a new simulation with another seed, not the saved run).

Table 15.2: The decoder with perfect Clifford scramblers of 3, 5 and 7 qubits, exact and on the FakeFez noise model.

black hole	radiated qubits	two-qubit gates (FakeFez)	success on FakeFez (exact 0.25)	F on FakeFez (m_F; exact 1)	$4PF$ on FakeFez
3 qubits	1	35	0.253	0.887 (773)	0.899
5 qubits	2	126	0.191	0.802 (605)	0.613
7 qubits	3	252	0.105	0.882 (764)	0.370

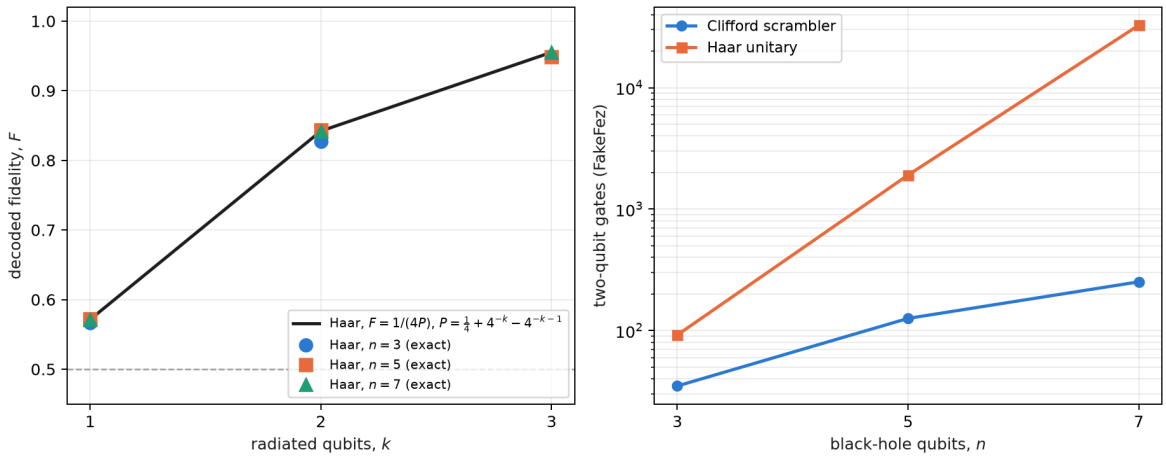


Figure 15.2: Left: decoded fidelity of Haar unitaries against the number of radiated qubits, exact for $n = 3, 5, 7$ and the large-n formula. Right: two-qubit gates of the decoder transpiled for FakeFez, with a Clifford scrambler and with a Haar unitary.

Table 15.2 shows how far the filter reaches. The fidelity stays between 0.80 and 0.89 even with 252 gates, and not because noise does not act: the postselection on k Bell pairs discards the shots in which noise touched (D, D'), and with more pairs it discards more. The success falls from 0.253 to 0.105 against the exact $1/4$ (with full noise it would be 4^{-k}, $1/64$ for $k = 3$). The fidelity alone would say the 7-qubit decoder is as good as the 3-qubit one; $4PF$ says it delivers 0.37 of what it should, because it pays in shots what the fidelity does not show. It is the rule of the sector filters of earlier chapters: the filter improves what is kept at the cost of the rate, and $4PF$ puts both in one number. With noise the law now breaks on the other side: on 3 qubits the fidelity drops, on 7 the success does.

15.8 A toy Page curve

Page (1993) predicted how the entanglement of the radiation of a unitarily evaporating black hole grows: it rises while the radiation is the smaller part, peaks when half has been radiated, and falls to zero at the end, because the total state is pure. An exact toy: a Haar state of 8 qubits, with the first k as the radiation.

The index proposed reading this curve with qg_S. **It does not work**, and it is worth saying why. qg_S is the Shannon entropy of a Z reading, and a random state looks random in any basis: the joint reading of the k radiated qubits has entropy close to k through the whole evaporation, including when the radiation is already nearly pure (Figure 15.3, triangles). It is the limit `qang.circuits` already documents with graph states as the counterexample: qg_S depends on the basis and does not measure entanglement.

What does work uses the same qg readings, squared. The purity of the radiation is the sum of the squares of its Pauli correlators,

$$\operatorname{Tr}\rho^2 = \frac{1}{2^k} \sum_{P \in \{I,X,Y,Z\}^k} qg_P^2,$$

and $S_2 = -\log_2 \operatorname{Tr}\rho^2$, the Rényi entropy of order 2, is a lower bound on the von Neumann entropy that follows the Page curve: at $k = 4$, the half, $S_2 = 2.98$ bits against $S = 3.24$, and both return to zero at $k = 8$, while the Shannon entropy of the Z reading reaches 7.46. The cost is the unfavourable side: $4^k - 1$ correlators, 255 for $k = 4$. (The script computes it from the correlators up to $k = 4$; beyond that it uses $\operatorname{Tr}\rho^2$ directly, the same quantity.)

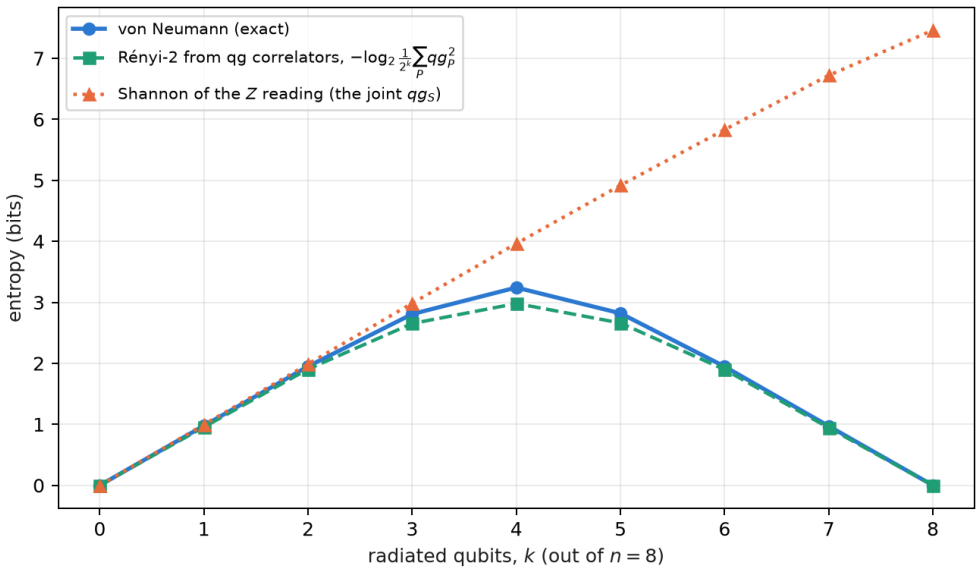


Figure 15.3: A toy Page curve, 8 qubits. The von Neumann entropy of the radiation and the Rényi-2 entropy computed from the qg correlators rise and fall; the Shannon entropy of the Z reading keeps rising and does not see the curve.

15.9 Traversable wormholes

The Yoshida–Kitaev decoder is, as a circuit, the same protocol that the holographic literature calls teleportation through a traversable wormhole (Gao et al. 2017; Maldacena et al. 2017): two entangled systems, a scrambling dynamics on each (one side conjugated) and a coupling between them that lets a qubit through. The ER = EPR conjecture (Maldacena and Susskind 2013) proposes that such entangled pairs are, in a dual gravitational theory, an Einstein–Rosen bridge. The circuit decides none of that; it does let the result be read with the tools of this book.

The fidelity of the recovered pair is linear in three qg correlators, as in chapter 13:

$$F = \frac{1 + qg_{XX} - qg_{YY} + qg_{ZZ}}{4}.$$

On the exact state, with the scrambler, $(qg_{XX}, qg_{YY}, qg_{ZZ}) = (1, -1, 1)$ and $F = 1$; without scrambling, $(0, 0, 0)$ and $F = 1/4$. The experiment of Jafferis et al.

(2022) on a Google processor used a learned 7-qubit model to mimic this physics, and its reading as a wormhole was disputed (Kobrin et al. 2023). This book does not run that experiment; the correlator reading stays a tool, not evidence.

15.10 What remains open

None of the new results of this chapter (the law $PF = 1/4$ as a diagnostic, the 5- and 7-qubit black holes, the Page curve from correlators) has a device run. A real-hardware run of the 3-qubit decoder, with 90 two-qubit gates, would be the next step and costs credits; it stays outside this book for now.

15.11 Without qang and with qang

Table 15.3: The actions of the chapter without the qang scale and with it.

action	without qang	with qang	difference
decide whether the message was recovered	compare F with 1/2	the sign of m_F: decoding is crossing the equator	**unit**: the threshold sits at zero, but it is the same threshold
judge the quality of the decoder	P and F separately	the law $(1000 + m_P)(1000 + m_F) = 10^6$ and its summary $4PF$	**in favour**, with a caveat: the law is Yoshida and Kitaev's, not qang's; the scale turns it into a readable hyperbola and shows that $m_P = -500$ proves nothing
predict how much radiation is needed	the Haar formula	the same formula, $F = 1/(4P)$	**none**
the Page curve	von Neumann entropy by tomography	qg_S in Z	**against**: qg_S does not see the curve

action	without qang	with qang	difference
the Page curve, again	tomography	Rényi-2 from the sum of qg_P^2	**equivalent**: it is Pauli tomography written in qg, with its $4^k - 1$ cost
the wormhole fidelity	Bell measurement of the pair (one setting)	three qg correlators	**none** in cost, since the Bell measurement already gives F; **in favour** as a diagnostic: the three correlators tell which component fails
reduce noise	—	the postselection is the filter; the scale adds nothing	**none**

Balance. The clearest contribution of qang in this chapter is one of reading: the decoder's exact law becomes a curve in the plane of the two readings, and in that plane one sees at once that the right success is no evidence of decoding. Where the index promised more, the Page curve with qg_S, the result is negative, and what replaces it is Pauli tomography, not a property of the scale.

15.12 In code

```
from qiskit.quantum_info import Statevector
from qang import postselection as ps

scr = ps.find_scrambler()                       # 3-qubit Clifford, seed 164
qc = ps.yk_decoder_circuit(scr, measure=False)  # 8 qubits: R, A, B1, B2, B1', B2', A', R
probs = Statevector(qc).probabilities_dict()
ps.yk_from_counts({k: round(1e6 * v) for k, v in probs.items()})
# {'success': 0.25, 'fidelity': 1.0}: P F = 1/4
```

15.13 Summary

- The Yoshida–Kitaev decoder is teleportation postselected on the Bell sector of a radiated qubit and its copy: a model of information, not of gravity.
- On the chapter 1 scale, the success $1/4$ is $m_P = -500$ and the fidelity that certifies recovery, $F > 1/2$, is $m_F > 0$: decoding is crossing the equator.
- The exact law $PF = 1/4$ is a hyperbola in m-qg; with noise, the success stays at $1/4$ but the fidelity drops, and $4PF$ (0.57–0.73 on three simulated models) measures the quality of the decoder.
- With a single radiated qubit, a generic black hole gives $F \approx 0.57$; 2 or 3 radiated qubits are needed for $F \approx 0.84$–0.95.
- A Clifford scrambler costs 35, 126 and 252 two-qubit gates on 3, 5 and 7 qubits; with 7, the fidelity on FakeFez stays at 0.88 but the success falls to 0.105 and $4PF$ to 0.37: the filter keeps the quality and charges in rate.
- qg_S does not see the Page curve; the Rényi-2 entropy from the qg correlators does, at the cost of $4^k - 1$ readings.

Bibliographic notes

The information paradox starts from Hawking (1976); complementarity is Susskind et al. (1993), and the Page curve, Page (1993). The recovery protocol is Hayden and Preskill (2007), and its decoder, Yoshida and Kitaev (2017). Traversable wormholes are Gao et al. (2017) and Maldacena et al. (2017), and ER = EPR, Maldacena and Susskind (2013); the experiment is Jafferis et al. (2022), and the discussion of its interpretation, Kobrin et al. (2023). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part VI

VI. Quantum cybersecurity, random-number generation and post-quantum frontiers

16 Demarcation: quantum cryptography against post-quantum cryptography (PQC)

Part VI is about security. Before reading quantum key distribution (Chapter 17) and certified random numbers (Chapter 18) in qg, it is worth drawing the border. There are two answers to the quantum computer that share the word "quantum" and little else. **Quantum cryptography** (QKD, QRNG) uses qubits: its security rests on a physical measurement, and so it enters the probability equation of Chapter 1. **Post-quantum cryptography** (PQC) is classical software that runs on ordinary computers: its security rests on the hardness of a mathematical problem, and there is no qubit to read. This chapter says where the line runs, with numbers, and what qang reads on each side.

16.1 From the base: which probabilities exist on each side

Any probability p_0 enters the equation of Chapter 1: it has qang angle $\theta = 2\arccos\sqrt{p_0}$ and reading $m = 1000(1-2p_1)$ m-qg. The question is not whether it can be written down, but whether the reading says anything.

- **In QKD**, the probability that decides security is the quantum bit error rate (QBER) of a measurement in one basis. It is the probability of a qubit, and its reading is the one of Chapter 17: $Q = (1-|qg|)/2$, that is, $m = 1000\,|qg| = 1000(1-2Q)$.
- **In PQC**, the only probability in the scheme is that decryption fails, δ. For ML-KEM, the FIPS 203 standard bounds it by $2^{-138.8}$, $2^{-164.8}$ and $2^{-174.8}$ for the three levels. Written in m-qg it is $m = 1000 - 3.3\times10^{-39}$ at best: it sits against the pole, and there is no Hermitian operator, measurement basis or Bloch sphere behind that number.

The security of PQC, moreover, is not that probability: it is the hardness of a lattice problem (module learning with errors, MLWE), which is not the frequency

of anything. That is the demarcation in one line: **qang reads probabilities of measurements; PQC has no measurements**.

16.2 The nature of PQC

NIST's standards of August 2024 are three: ML-KEM (FIPS 203, formerly Kyber), a lattice key-encapsulation mechanism; ML-DSA (FIPS 204, formerly Dilithium), lattice signatures; and SLH-DSA (FIPS 205, formerly SPHINCS+), signatures built only on hash functions. In 2025 HQC, based on error-correcting codes, was also chosen as a second key mechanism. All of them are integer arithmetic: polynomials modulo $q = 3329$ in ML-KEM, hash trees in SLH-DSA. None of them prepares a state, applies a gate or measures.

In ML-KEM, decrypting is subtracting and rounding: the message survives if an accumulated noise (the sum of many products of small coefficients) stays within $\pm q/4$. If not, decryption fails. That is the δ above.

16.3 Decryption failure is counted, not measured on a qubit

To see what qang would do with δ, the figure script builds a toy scheme with the same mechanism: noise $e = \sum_{i<2n} a_i b_i + c$ with a_i, b_i, c from ML-KEM's centred binomial distribution, $n = 256$, $q = 3329$, without the standard's compression (which in ML-KEM accounts for most of the failure), and with the width η as the knob. The failure probability is computed exactly (a convolution modulo q) and counted over 200,000 simulated encryptions.

Table 16.1: Decryption failure of the toy scheme, exact and counted.

η	σ/q	exact failure	counted	reading (m-qg, 95% CI)
16	0.054	5.1×10^{-6}	1 of 200,000	1000.0 [999.9, 1000.0]
20	0.068	2.5×10^{-4}	60	999.4 [999.2, 999.5]
24	0.082	2.2×10^{-3}	437	995.6 [995.2, 996.0]

η	σ/q	exact failure	counted	reading (m-qg, 95% CI)
32	0.109	2.2×10^{-2}	4360	956.4 [955.1, 957.7]
40	0.136	6.6×10^{-2}	13,072	869.3 [867.1, 871.4]

The count agrees with the exact sum, and `qang.probability.analyze` reads it in m-qg with its interval. But that reading is the Wilson proportion in another unit: there is no angle of a state behind it. And no count reaches where real PQC lives (Figure 16.1, right). Seeing a single ML-KEM-512 failure takes on the order of $1/\delta \approx 6 \times 10^{41}$ decryptions; with 10^{12} trials, about the most one can count, the floor sits 30 orders of magnitude above. The shot rule of Chapter 1 says the same in m-qg: next to the pole, $\sigma_m \approx 2000\sqrt{\delta/N}$, and to bring it below the offset $2000\,\delta$ one needs $N > 1/\delta$. **PQC's δ is proved, not measured.**

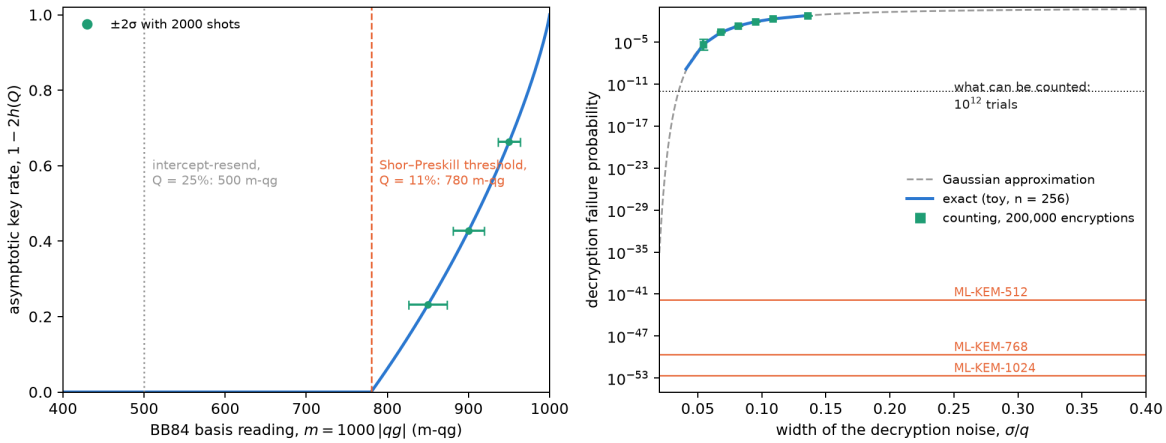


Figure 16.1: Left: the security frontier of BB84 is a line in m-qg; the Shor–Preskill asymptotic key rate vanishes at 780 m-qg (Q = 11%), and with 2000 shots the reading has ±2σ of 14–24 m-qg. Right: the decryption failure of a toy ML-KEM-like scheme (n = 256, no compression), exact, Gaussian and counted; the orange lines are the FIPS 203 bounds, 30 to 41 orders of magnitude below what can be counted.

On the other side of the line (Figure 16.1, left), the quantity that decides BB84's security sits in the middle of the scale. The Shor–Preskill asymptotic rate, $1 - 2h(Q)$, vanishes at $Q = 11.0\%$, which is $m = 780$ m-qg, a qang angle

of $2\arcsin\sqrt{Q} = 0.676$ rad (38.7°, or 107.6 qga on the scale of Chapter 1); intercept-resend gives $Q = 25\%$, $m = 500$. With 2000 shots, σ_m is 11.8, 9.7 and 7.0 m-qg at 850, 900 and 950 m-qg: the margin to the threshold is resolved comfortably. It is the probability of a qubit, readable and with an error bar; Chapter 17 uses it to separate the eavesdropper from the noise.

16.4 The quantum attack vector: Grover

The quantum computer is quantum, and its attacks do enter the base. Grover's attack on an n-bit key has, after k iterations, failure probability

$$p_{\text{fail}} = \cos^2\big((2k+1)\,\beta\big), \qquad \beta = \arcsin 2^{-n/2},$$

which is the probability equation with qang angle $2(2k+1)\beta$: **each iteration advances the qang angle by 4β, in a straight line**. This was checked with the circuits of `qang.algorithms` for $n = 4$, 6, 8 and 10: the angle of the exact probability matches the law to 10^{-10} rad (Figure 16.2). In m-qg, the failure goes from +1000 to −1000 along a cosine and crosses the equator (success 1/2) halfway.

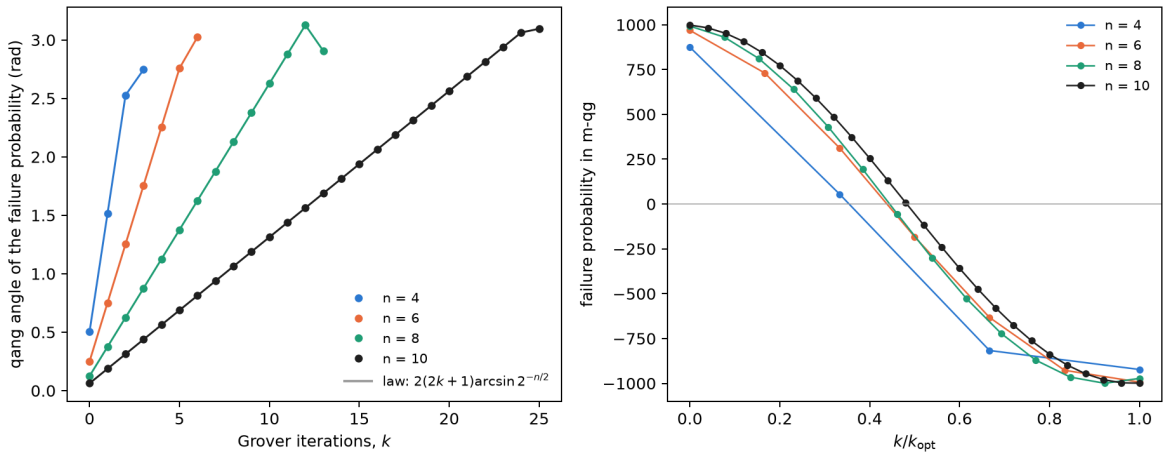


Figure 16.2: Left: the qang angle of the probability that Grover fails grows linearly with the iterations, $4\arcsin 2^{-n/2}$ per step (dots: exact circuits; lines: the law). Right: the same probability in m-qg against the fraction of the optimal iterations.

For AES-128 the step is $4\arcsin 2^{-64} = 2.2\times10^{-19}$ rad per iteration, and reaching the pole takes 1.4×10^{19} serial iterations, each a full AES circuit; for AES-256,

2.7×10^{38}. PQC's answer to this attack is to double the symmetric key length, and the geometry explains it: the angle advances as $2^{-n/2}$, so Grover halves the security bits and nothing more.

16.5 The quantum attack vector: Shor and phase estimation

The attack that breaks RSA and elliptic curves, and that PQC sidesteps, is Shor's algorithm: phase estimation (QPE) on modular multiplication. The notes (§13) show that QPE recovers exact phases, and the book's index said that "the eigenstate register projects onto $qg_Z = -1$ without resolving the phase". That holds for the QPE of an eigenstate, and the exact computation of this chapter says how far a qg reading goes in Shor's case (Figure 16.3).

QPE of an eigenstate. With the T gate on $|1\rangle$ ($\varphi = 1/8$), the target register stays at $m_Z = -1000$ and carries no information: the phase lives in the counting register. There, before the inverse QFT, qubit k has $m_X = 1000\cos(2\pi 2^k\varphi)$: 707, 0 and −1000 m-qg. Each qubit carries one bit of the phase in its qang angle, which is Kitaev's iterative phase estimation.

Shor for $N = 15$. The work register is not in an eigenstate but in a superposition of the r eigenstates, and the counting qubits end up entangled with it. With four counting qubits:

Table 16.2: Single-qubit readings of the counting register in Shor's algorithm (exact state).

case	local radius of each counting qubit, before (m-qg)	m_Z of each qubit, after	outcomes
$a = 7$ ($r = 4$)	0, 0, 1000, 1000	1000, 1000, 0, 0	0, 4, 8, 12
$a = 11$ ($r = 2$)	0, 1000, 1000, 1000	1000, 1000, 1000, 0	0, 8
$N = 21$, $a = 2$ ($r = 6$), six qubits	0, 0, 0, 0, 0, 0	0, 250, 312, 328, 332, 0	peaks at 0, 11, 21, 32, 43, 53

For $N = 15$ the order can be read qubit by qubit: the number of qubits with zero radius before the QFT, or with $m_Z = 0$ after it, is $\log_2 r$. But that is an accident of $N = 15$, where r is a power of 2. With $N = 21$ and $r = 6$, every local radius is

zero and the m_Z after do not say 6: the order is in the joint distribution, and is extracted with continued fractions. The work register does not help either: for $a = 7$ its qubits read $m_Z = \pm 500$, the mixture of 1, 7, 4 and 13.

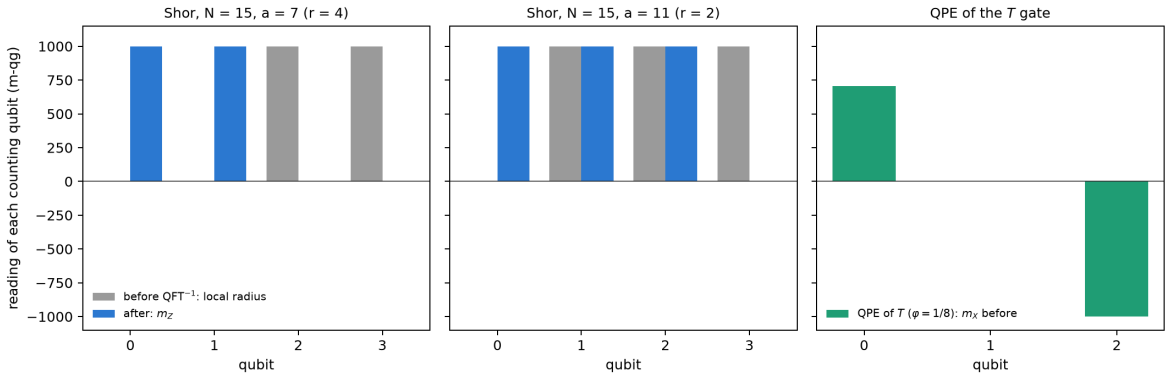


Figure 16.3: Single-qubit readings of the counting register. Grey: local radius before the inverse QFT; blue: m_Z after. For N = 15 the number of zeros is $\log_2 r$; for the QPE of T, each qubit carries one bit of the phase in m_X.

The conclusion holds for the whole chapter: single-qubit readings see the phase when there is an eigenstate, and lose it when the information is in the correlations. It is limitation 3 of the notes (§14) in another form.

16.6 The demarcation

Table 16.3: Where the line runs.

	quantum cryptography (QKD, QRNG)	post-quantum cryptography (PQC)	quantum attack (Grover, Shor)
substrate	qubits, photons	integers, polynomials, hashes	qubits
security rests on	a physical measurement	the hardness of MLWE, of codes or of the hash	—
the key probability	QBER, in the middle of the scale (780 m-qg)	$\delta \leq 2^{-138.8}$, against the pole	Grover's failure, Shor's success

	quantum cryptography (QKD, QRNG)	post-quantum cryptography (PQC)	quantum attack (Grover, Shor)
does qang read it?	**yes**: Chapters 17 and 18	**no**: there is no qubit; δ can only be counted, and cannot be	**partly**: Grover yes; Shor per qubit only with an eigenstate

qang is a physical-layer tool: it reads what comes out of a measurement on a qubit. It says nothing about the hardness of a lattice, and a module claiming to "measure in qg" the security of ML-KEM would be a category error.

16.7 What remains open

There is no `qang.crypto` module, and this chapter does not propose one. The project's cryptography lives in `examples/` (`bb84_qg_eve_vs_noise.py`, `bb84_finite_key_qg.py`, `bb84_attacks_beyond_ir_qg.py`, `qrng_qg_certified.py`, `ionq_sim_bb84.py`, `ionq_sim_qrng.py`) and in `manuscript/crypto`. Two questions remain open: whether the per-qubit reading of iterative QPE (one recycled counting qubit) saves shots against QPE with a QFT when there is noise, and a run of Shor for $N = 15$ on a device, read per qubit; neither has a circuit or a written prediction yet.

16.8 Without qang and with qang

Table 16.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read BB84's security	QBER and $1 - 2h(Q)$	threshold at 780 m-qg, σ_m 7–12 m-qg with 2000 shots	**unit**; the detail is in Chapter 17
estimate PQC's decryption failure	counting with a Wilson interval	the same proportion in m-qg	**none**: it is the same statistic, and there is no state behind it

action	without qang	with qang	difference
ML-KEM's failure	mathematical proof ($\leq 2^{-138.8}$)	$1000 - 3 \times 10^{-39}$ m-qg, with no possible resolution	**none**; qang does not apply
describe Grover's attack	$\sin^2((2k+1)\beta)$	the qang angle advances 4β per iteration, in a straight line	**equivalent**: it is the standard Grover geometry under another name, though it reads at a glance
read the phase in the QPE of an eigenstate	histogram of the register after the QFT	m_X of each counting qubit: one bit per qubit	**in favour**, modestly: iterative QPE read on the Chapter 1 scale
read the order in Shor	continued fractions on the histogram	local radii and m_Z per qubit	**against** except for r a power of 2: with $N = 21$ everything reads zero

Balance. This is the chapter where qang contributes least, and its value lies in saying so. On the PQC side there is nothing to read; on the attack side, Grover is a probability equation and reads cleanly, and Shor shows the known limit of single-qubit readings. On the quantum side, the QBER is a qang reading, and the next two chapters use it.

16.9 In code

```
import math
from qang import algorithms as al
from qang import probability as pb

pb.milliqang_from_p0(1 - 0.11)             # 780: the BB84 threshold in m-qg
pb.analyze(200_000 - 437, 200_000).milliqang    # (995.6, 995.2, 996.0): a counted failu
p = 1 - al.grover_success_probability(8, 1, 5)
math.acos(2 * p - 1)                    # 2(2*5+1) asin(2**-4) = 1.376: Grover's angle
```

16.10 Summary

- Quantum cryptography has measurements and enters the probability equation; PQC is classical arithmetic and does not.
- BB84's security is a line in m-qg: the key rate vanishes at 780 m-qg, in the middle of the scale.
- ML-KEM's only probability, the decryption failure, is 3×10^{-39} m-qg from the pole: it is proved, not measured; counting it is impossible, and reading it in m-qg is the same proportion in another unit.
- Grover's attack advances the qang angle by $4 \arcsin 2^{-n/2}$ per iteration, in a straight line; hence doubling the key is enough.
- In the QPE of an eigenstate each counting qubit carries one bit of the phase in m_X; in Shor the information is in the correlations, and single-qubit readings see the order only when it is a power of 2.

Bibliographic notes

The post-quantum standards are FIPS 203, 204 and 205 (National Institute of Standards and Technology 2024b, 2024a, 2024c), and the selection of HQC is in Alagic et al. (2025). The quantum attack vectors are Grover (Grover 1996) and Shor (Shor 1997), the latter built on the phase estimation of Kitaev (1995). BB84 is Bennett and Brassard (2014) (reprint of the 1984 original). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

17 Quantum key distribution (QKD) and BB84 in qg space

Chapter 16 left BB84's QBER on the side of what qang reads: it is the probability of a qubit, and the security threshold sits in the middle of the scale, at 780 m-qg. This chapter opens it up. The QBER is a single number, and a real link has four: each state Alice sends reaches Bob with its own error probability. On the scale of Chapter 1 those four readings say where the image of the Bloch sphere ended up after the channel, and the shape of that image tells apart two causes the QBER confuses: a memory that degrades and an eavesdropper who intercepts. The chapter says how far that distinction goes, with finite key blocks, with subtler attacks and with trapped-ion noise, and where it does not go.

17.1 From the base: four readings and the image of the sphere

In BB84 Alice sends $|0\rangle$ or $|1\rangle$ (Z basis) or $|+\rangle$ or $|-\rangle$ (X basis); the qubit waits in a memory or travels through a channel, and Bob measures. After sifting only the rounds in which Bob measured in Alice's basis remain. For each sent state, Bob's reading in its own basis is a reading of Chapter 1:

$$m(0) = 1000\,(1 - 2e_{0\to1}), \qquad m(1) = -1000\,(1 - 2e_{1\to0}),$$

and likewise for $|\pm\rangle$ in X. An ideal channel gives ±1000. A real channel maps the sphere to an ellipsoid, and in each basis the two readings give its **centre** and its **semi-axis**:

$$c_Z = \frac{m(0) + m(1)}{2}, \qquad r_Z = \frac{m(0) - m(1)}{2}, \qquad Q_Z = \frac{1}{2}\Big(1 - \frac{r_Z}{1000}\Big).$$

The QBER is exactly the semi-axis, with no approximation. The index's formula, $Q = (1-|qg|)/2$, is the symmetric case, in which the two readings of a basis have the same modulus; in general $|m(0)| \neq |m(1)|$, and the difference is the centre.

The asymmetry $A_Z = qg(0) + qg(1)$ of the notes (§40) is $2c_Z$: the T1 witness of §24.

The two causes move the image in different directions. T1 relaxation of the memory is affine, $z \mapsto (1-\gamma)z+\gamma$: it leaves $|0\rangle$ at the pole and lifts $|1\rangle$, so that c_Z goes up and r_Z goes down **by the same amount**. An eavesdropper who intercepts and resends a fraction f of the qubits, measuring in a random basis, shrinks the sphere before the memory: r_Z goes down and c_Z **does not move**, exactly, because the affine memory puts the centre at a place that depends only on γ. In the (c_Z, r_Z) plane the drift is a line of slope -1 and the eavesdropper a vertical (Figure 17.1, left).

The channel of the notes: a memory with $\gamma_0 = 0.02$ (T1), dephasing $p_0 = 0.01$ and an asymmetric detector ($e_{01} = 0.005$, $e_{10} = 0.01$, calibrated as in §31).

Table 17.1: The four BB84 readings in m-qg (the centre in X is $c_X = 5.0$ in every case, from the detector). The secret fraction is Shor–Preskill's asymptotic one, $1 - h(Q_Z) - h(Q_X)$.

scenario	$m(0)$	$m(1)$	c_Z	r_Z	r_X	Q_Z	Q_X	secret fraction
baseline	990.0	−940.6	24.7	965.3	955.6	1.74%	2.22%	0.720
T1 drift, $\gamma = 0.04$	990.0	−901.2	44.4	945.6	945.8	2.72%	2.71%	0.640
T1 drift, $\gamma = 0.06$	990.0	−861.8	64.1	925.9	935.9	3.70%	3.21%	0.567
intercept-resend, $f = 0.05$	965.9	−916.5	24.7	941.2	931.7	2.94%	3.41%	0.594
intercept-resend, $f = 0.10$	941.7	−892.3	24.7	917.0	907.8	4.15%	4.61%	0.481

In a window of N compared bits, $N/4$ per state, the shot noise of c_Z and of r_Z is 16.4, 8.2 and 3.7 m-qg for $N = 500$, 2000 and 10,000. The memory going from

$\gamma = 0.02$ to 0.04 moves the centre by 19.7 m-qg: with 2000 bits, 2.4 standard deviations.

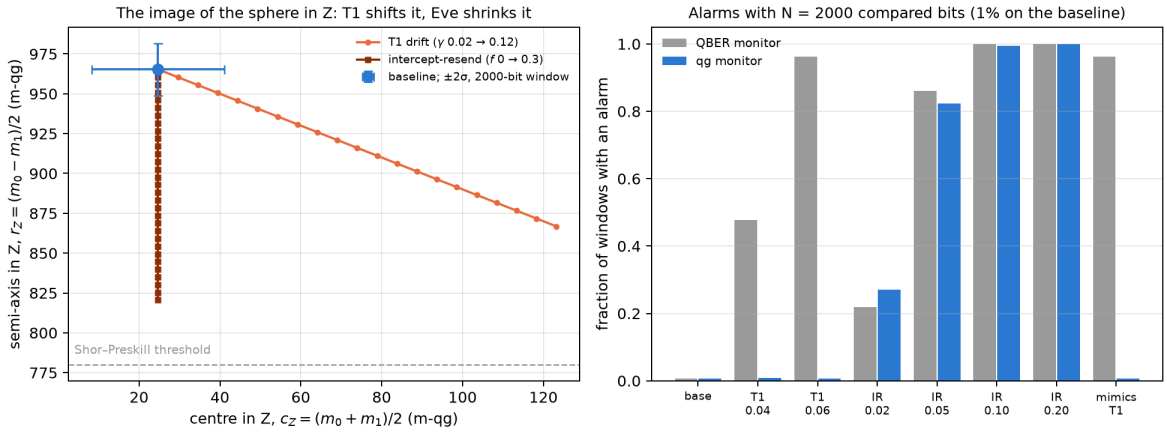


Figure 17.1: Left: the image of the sphere in the Z basis, centre against semi-axis (m-qg). T1 drift runs along a line of slope -1; intercept-resend lowers the semi-axis without moving the centre; the bar is ±2σ with a 2000-bit window. Right: alarms of the two monitors with N = 2000 compared bits, tuned to 1% false alarms on the baseline.

17.2 The monitor: drift or eavesdropper

The notes (§40) compare two monitors on windows of N sifted and compared bits, both tuned to 1% false alarms on the baseline:

- **QBER monitor**: a one-sided test on the total error count;
- **qg monitor**: a generalized likelihood ratio on the four error counts, where "attack" ($f > 0$) must beat both "baseline" and "T1 drift" (γ free).

Table 17.2: Fraction of windows with an alarm (1000 windows per scenario).

scenario	QBER, $N = 500$	qg, $N = 500$	QBER, $N = 2000$	qg, $N = 2000$	QBER, $N = 10^4$	qg, $N = 10^4$
baseline	0.007	0.004	0.007	0.006	0.010	0.007
T1 drift, 0.04	0.147	0.009	0.478	**0.009**	0.992	0.000

scenario	QBER, $N = 500$	qg, $N = 500$	QBER, $N = 2000$	qg, $N = 2000$	QBER, $N = 10^4$	qg, $N = 10^4$
T1 drift, 0.06	0.463	0.009	0.962	**0.006**	1.000	0.000
intercept, $f = 0.02$	0.072	0.065	0.218	0.271	0.847	0.848
intercept, $f = 0.05$	0.337	0.338	0.861	0.823	1.000	1.000
intercept, $f = 0.10$	0.821	0.730	1.000	0.994	1.000	1.000
T1-mimicking attack	0.463	0.009	**0.962**	0.006	1.000	0.000

Three things can be read from the table.

- **Detection is the same; false alarms are not.** The qg monitor catches the naive eavesdropper as well as the QBER monitor, at every N and f, within the sampling spread. When the memory degrades, the QBER monitor fires in 48% and 96% of windows ($N = 2000$) and the qg monitor stays at 0.9% and 0.6%, at the baseline level. The reduction in false alarms holds for T1 drift, which is in the qg monitor's model, and not for any degradation whatsoever.
- **The price, by design.** An eavesdropper who couples the qubit to her ancilla through an amplitude-damping interaction produces exactly the counts of a T1 drift, because in the (c_Z, r_Z) plane she runs along the same line. The qg monitor misses her (0.6%, the same figure as the drift, because the counts are identical), and the QBER monitor does not (96%). The two are complementary: QBER says "something changed"; qg says "symmetric (attack-like)" or "T1-like (drift, or an eavesdropper hiding as one)".
- **No security gain.** The secret fraction falls the same way in both cases (0.720 → 0.567 with drift to 0.06; 0.720 → 0.594 with $f = 0.05$). A security proof must charge every error to Eve, and the qg reading does not change that.

The setting is a link with matter qubits, or an on-chip demonstration, where the qubit waits with T1 and T2 before Bob measures. In a photon-polarization link the noise is nearly unital, the centre does not move, and the distinction does not exist.

17.3 Finite blocks

A real link cuts the key into blocks. The notes (§43) use the finite-key bound of Tomamichel et al. (2012) on the same channel:

$$\ell = n\big[1-h(Q_{\mathrm{tol}}+\mu)\big]-1.1\,n\,h(Q_Z)-\log_2\frac{2}{\varepsilon_{\mathrm{sec}}^2\varepsilon_{\mathrm{cor}}},\qquad \mu=\sqrt{\frac{n+k}{nk}\,\frac{k+1}{k}\,\ln\frac{2}{\varepsilon_{\mathrm{sec}}}},$$

with $\varepsilon = 10^{-10}$, n key bits in Z and k test bits in X. The protocol aborts if the test error exceeds Q_{tol}; k and Q_{tol} are optimised for each n.

Table 17.3: Secret bits per signal with a finite key.

scenario	minimum n	$n = 10^4$	10^5	10^6	10^7	asymptotic
baseline	$1.2 \cdot 10^3$	0.120	0.268	0.407	0.516	0.707
T1 drift, 0.04	$1.6 \cdot 10^3$	0.091	0.223	0.348	0.446	0.622
T1 drift, 0.06	$2.0 \cdot 10^3$	0.068	0.182	0.294	0.383	0.544
intercept, $f = 0.05$	$1.7 \cdot 10^3$	0.079	0.201	0.318	0.410	0.575

A block of 10^4 bits keeps 17% of the asymptotic rate; 10^5, 38%; 10^6, 58%; 10^7, 73%. Below about 10^3 bits there is no key. Drift and attack cost key alike.

Where qg enters, at no key cost. After error correction and its check, Bob knows Alice's bits: he has the four error counts, that is the four readings, over the n key bits, without disclosing anything. Announcing an alarm bit costs at most one key bit, and the bits of an aborted block are discarded anyway. Two nested likelihood ratios on the joint model (γ and f free) give an "attack-like" flag (symmetric error beyond any T1) and a "drift-like" flag (T1 beyond the baseline, attack or not), each with 1% false alarms.

With the protocol tuned on the baseline for $n = 10^5$ ($k = 7609$, $Q_{\mathrm{tol}} = 2.67\%$, that is, abort if the semi-axis in X falls below 946.6 m-qg), 400 blocks per scenario:

Table 17.4: Protocol aborts and the qg diagnosis of the same blocks ($n = 10^5$).

scenario	protocol aborts	qg: attack-like	qg: drift-like
baseline	0.010	0.020	0.015
T1 drift, 0.03	0.100	0.005	**1.000**
T1 drift, 0.04	0.557	0.005	**1.000**
T1 drift, 0.06	0.995	0.007	**1.000**
intercept, $f = 0.02$	0.495	**1.000**	0.043
intercept, $f = 0.05$	1.000	**1.000**	0.040
drift 0.04 + intercept 0.02	0.995	**1.000**	**1.000**
T1-mimicking attack	0.995	0.013	1.000

- **An abort does not say why; qg does.** With 10^5 key bits the attribution is nearly exact: every drifting block, aborted or not, comes out as drift and not as an attack, and every intercepted block comes out as an attack, including the half that did not abort. With $n = 10^4$ the flags are already at 0.97–1.00 (0.81 for drift 0.03).
- **An attack hidden under drift is caught.** Intercepting with $f = 0.02$ on a memory with $\gamma = 0.04$ raises both flags in every block. The one-parameter test of the previous section misses it: it gives 0.5% attack flags and calls 99% "drift". The joint model fixes a real weakness of §40.
- **Early warning.** At $\gamma = 0.03$ only 10% of blocks abort, but the drift flag is up in all of them: the memory can be serviced before the link stops. A Q_Z trend would also warn; what qg adds is the attribution.

17.4 Beyond intercept-resend

The flags know two error families. Seen from Bob, every individual attack is a Pauli channel (p_x, p_y, p_z), and on the image of the sphere a Pauli channel **shrinks the axes without moving the centre**: the semi-axis in Z is multiplied by $1 - 2(p_x + p_y)$ and the one in X by $1 - 2(p_z + p_y)$. That explains the table of the notes (§49):

Table 17.5: Individual attacks as Pauli channels. The identities hold to 10^{-16}.

attack	(p_x, p_y, p_z)	what it shrinks	same counts as
intercept in both bases, fraction f	$(f/4, 0, f/4)$	Z and X	—

attack	(p_x, p_y, p_z)	what it shrinks	same counts as
optimal phase-covariant cloner, disturbance D	$(0, D, 0)$	Z and X	intercept with $f = 4D$
Z-only intercept, fraction f	$(0, 0, f/2)$	X only	**T2 drift**, $(1 - 2p') = (1 - 2p_0)(1 - f)$
X-only intercept, fraction f	$(f/2, 0, 0)$	Z only	no natural noise

T2 dephasing also shrinks only the X axis. Hence the third family: with $n = 10^5$, $k = 7609$ and 400 blocks, the model adds a dephasing family and three one-sided nested flags: attack-like, T1-like ($\gamma > \gamma_0$) and T2-like ($p > p_0$).

Table 17.6: Flags with three error families ($n = 10^5$, 400 blocks).

scenario	attack-like	T1-like	T2-like	two-family model: attack
baseline	0.020	0.005	0.033	0.020
T1 drift, 0.04	0.018	**1.00**	0.033	0.015
T2 drift, $p = 0.02$ / 0.03	0.007 / 0.020	0.01	**1.00**	**0.075 / 0.268**
intercept, $f = 0.05$	**1.00**	0.055	0.028	1.00
cloner, $D = 0.0125$ / 0.025	**1.00**	0.055	0.02–0.04	1.00
Z-only, $f = 0.02$ / 0.04	0.013	0.02–0.03	**1.00**	0.06 / 0.27
X-only, $f = 0.04$	**1.00**	0.048	0.000	1.00
T2 drift + cloner	**1.00**	0.043	**0.995**	1.00

- **Partial cloning is caught.** The optimal cloner comes out as an attack in every block, like intercepting with $f = 4D$, because Bob sees the same counts. What no error statistic shows is that the cloner gives Eve more information

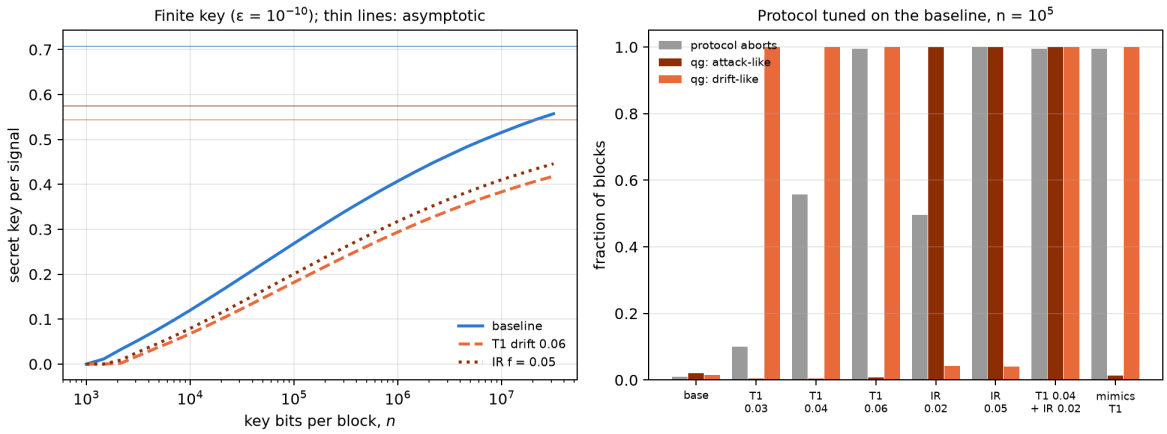


Figure 17.2: Left: finite-key rate against block size for the baseline, T1 drift to 0.06 and an eavesdropper with f = 0.05; thin lines are the asymptotic rates. Right: with n = 10^5, what fraction of blocks aborts and what fraction raises each qg flag.

per error; the key rate, which charges every error to Eve, already accounts for it.

- **The two-family model had a false alarm.** Without the dephasing family, T2 drift raised the attack flag in 7.5% and 27% of blocks. With the family, the false alarm drops to 0.7–2% and the verdict is T2 in every block.
- **A second blind spot, by construction.** A Z-only intercept shrinks only the X axis, like T2: it comes out as T2 drift and not as an attack. It is the twin of the T1-mimicking attack: an eavesdropper hiding in the error type the hardware makes by itself. An X-only intercept has no natural twin, and is always caught.
- **Mixtures are resolved:** T2 drift plus a cloner raises both flags.

17.5 On IonQ's noisy simulator

So far the channel was analytic. In the notes (§44) every part of BB84 is a circuit run on the aria-1 and forte-1 noise models of IonQ's simulator (free; nothing was sent to a QPU), and the model behind the flags is fitted, not given.

- **The memory** is amplitude damping built with an ancilla, $\mathrm{CRY}(2\arcsin\sqrt{\gamma})$ followed by CX, with a healthy $\gamma_0 = 0.02$ and drifts from 0.03 to 0.06.
- **Eve** measures and resends (CX to her ancilla in Z, H·CX·H in X), never reading the ancilla; a fraction f of the rounds is taken from those circuits.

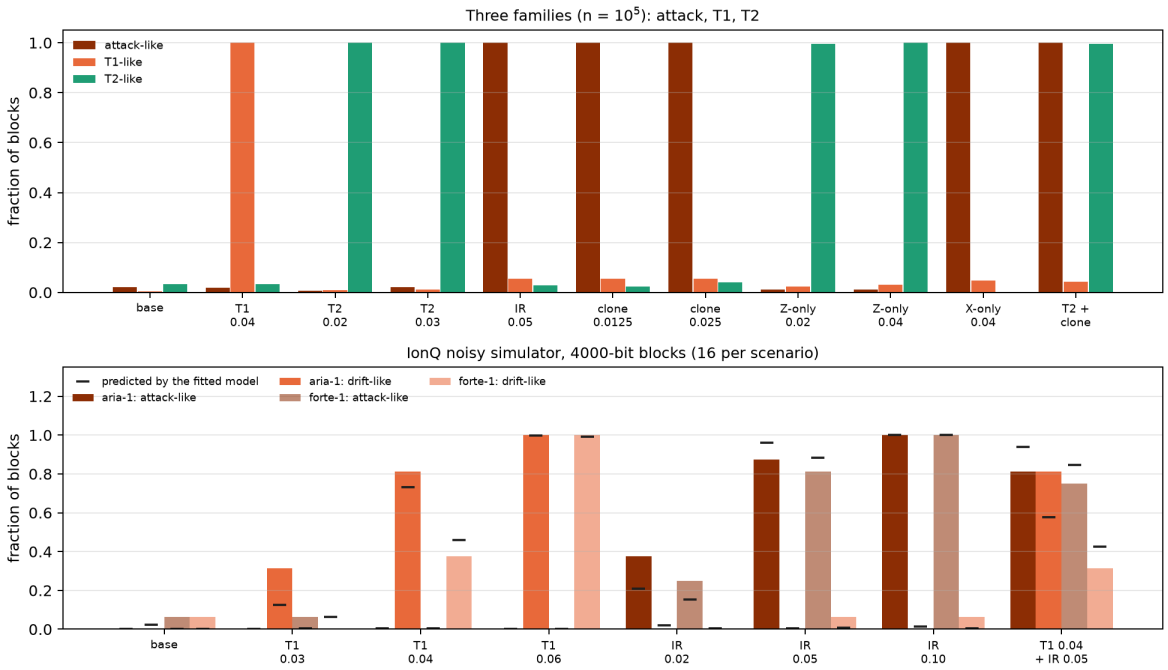


Figure 17.3: Top: flags of the three-family model with n = 10^5. T2 drift and the Z-only intercept get the same verdict, because they shrink the same axis. Bottom: flags on IonQ's noisy simulator (aria-1 and forte-1), observed in 16 blocks of 4000 bits; the black ticks are what the fitted model predicts.

- **The runs:** 32 circuits in native gates, 4 copies of each (12 qubits), 2000 shots, 4 repetitions per noise model. Each copy is a block of $n = k = 4000$ bits.
- **The flag model** is the joint one (γ, f) plus two symmetric device flip rates, ε_Z and ε_X, fitted leaving out each repetition, with thresholds from a parametric bootstrap at 1%.

The device adds a nearly symmetric error of 2.3–2.4% per bit (aria-1) and 2.8–3.1% (forte-1), mostly from the two MS gates of the memory circuit. On the image of the sphere that is a shrinking, not a shift. The designed healthy memory gives $c_Z = 20$, $r_Z = 980$ and $r_X = 990$ m-qg; the simulator gives $c_Z = 18.5$ and 20.6 m-qg (aria-1, forte-1), with the semi-axes 46–62 m-qg lower; with $\gamma = 0.06$ the centre rises to 60.5 and 56.5 m-qg, where the design says 60. The device noise leaves the centre in place, which is what the T1 witness needs to keep working. The fit recovers the designed γ_0 (0.020–0.0225) in every fold.

Table 17.7: Flags on IonQ's simulator, observed / predicted by the fitted model (16 blocks per scenario).

scenario	aria-1 QBER	aria-1 attack	aria-1 drift	forte-1 QBER	forte-1 attack	forte-1 drift
baseline	0.00/0.01	0.00/0.00	0.00/0.02	0.00/0.01	0.06/0.00	0.06/0.00
drift 0.03	0.56/0.36	0.00/0.00	0.31/0.12	0.06/0.20	0.06/0.01	0.00/0.06
drift 0.04	0.94/0.93	0.00/0.01	0.81/0.73	0.38/0.75	0.00/0.01	0.38/0.46
drift 0.06	1.00/1.00	0.00/0.00	**1.00**/1.00	1.00/1.00	0.00/0.00	**1.00**/0.99
intercept, $f = 0.02$	0.62/0.58	0.38/0.21	0.00/0.02	0.31/0.38	0.25/0.15	0.00/0.01
intercept, $f = 0.05$	1.00/1.00	**0.88**/0.96	0.00/0.01	1.00/0.99	**0.81**/0.88	0.06/0.01
intercept, $f = 0.10$	1.00/1.00	**1.00**/1.00	0.00/0.01	1.00/1.00	**1.00**/1.00	0.06/0.01
drift 0.04 + $f = 0.05$	1.00/1.00	0.81/0.94	0.81/0.58	1.00/1.00	0.75/0.84	0.31/0.42

- **The attribution transfers.** A drifting memory is never taken for an attack on aria-1 (0 of 48 blocks), and once in 48 on forte-1. An attack is taken for drift in at most 1 of 16 blocks. Large changes (drift 0.06, $f = 0.10$) are attributed in every block.
- **Small blocks, noisy device.** With 4000 bits and 2–3% of its own noise, small changes are seen only in part: 31–81% for drift 0.03–0.04 and 81–88% for $f = 0.05$. The QBER monitor sees changes as often or more often, but cannot say which kind.

- **The model predicts which flag rises.** Almost every observed rate falls within the ±0.12 spread of 16 blocks; the largest misses are near threshold. Extrapolated with the fitted noise to $n = 10^5$, the regime of the finite-block section, every scenario is attributed correctly in 99–100% of blocks.

These are vendor noise models on a simulator, not hardware; the memory and the attack are designed circuits.

17.6 What remains open

There is no hardware run: the Forte-1 plan (notes, §46) was rejected by the provider before the job was accepted, and the IQM credits at hand do not stretch to blocks of 10^5 bits. Collective attacks, detector attacks and photon links, where the noise does not move the centre, are not modelled. Nor is there a `qang.crypto` module (Chapter 16): everything lives in `examples/bb84_*.py` and `examples/ionq_sim_bb84.py`.

17.7 Without qang and with qang

Table 17.8: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
decide whether there is a key	QBER and $1 - h(Q_Z) - h(Q_X)$, or the finite-key bound	the semi-axis r in m-qg; abort if $r_X < 946.6$	**unit**: it is the same computation
describe the channel	four error rates	centre and semi-axis of the image of the sphere in each basis	**in favour**, modest: T1 shifts the centre, the eavesdropper shrinks, T2 shrinks X only; the geometry says which model is needed

action	without qang	with qang	difference
tell drift from eavesdropper	total QBER: cannot	likelihood ratio on the four counts	**in favour**, but the gain belongs to the likelihood model on per-state counts; a standard asymmetric analysis would do the same without the word qg
false alarms from T1 drift ($N = 2000$)	48% and 96%	0.9% and 0.6%	**in favour**, at a price: the T1-mimicking attack gets through (0.6% against 96%)
security	every error is charged to Eve	the same	**none**: the key does not grow
attack hidden under drift	only an abort	both flags, 100% of blocks ($n = 10^5$)	**in favour** of the joint model
noise of a real device (simulated)	QBER jumps with everything	attribution with 0–1 errors in 48 blocks	**in favour**, still partial with 4000-bit blocks

Balance. The contribution that belongs to the Chapter 1 base is the geometry: reading the four probabilities as the centre and semi-axes of the image of the sphere says at once which error families can be told apart (those that move different things) and which cannot (those that shrink the same axis). The detection power comes from the statistics on per-state counts, which does not need qang. And security does not change: this is an operational diagnostic, not a proof.

17.8 In code

```
import sys
sys.path.insert(0, "examples")
import bb84_qg_eve_vs_noise as B
```

```
e = B.error_probabilities(gamma=0.04)        # error rates of 0, 1, +, -
m = 1000 * B.qg_per_state(e)            # [990.0, -901.2, 950.8, -940.8] m-qg
c_z, r_z = (m[0] + m[1]) / 2, (m[0] - m[1]) / 2
c_z, r_z, (1 - r_z / 1000) / 2              # 44.4, 945.6 and Q_Z = 0.0272
```

17.9 Summary

- In BB84 each sent state has its reading; in each basis, the two readings give the centre and the semi-axis of the image of the sphere, and the QBER is exactly the semi-axis: $Q = (1 - r/1000)/2$.
- T1 drift shifts the centre and shrinks by the same amount; an intercepting eavesdropper shrinks without moving the centre. Hence the qg monitor stays quiet under drift (0.6–0.9% against 48–96%) and sees the eavesdropper as well as the QBER does.
- By construction, an eavesdropper who mimics T1, or who intercepts only in Z (and looks like T2), comes out as drift.
- With finite blocks of 10^5 bits the attribution is nearly exact, at no key cost, and an attack hidden under drift raises both flags; IonQ's simulator confirms it with 4000-bit blocks, in part.
- The secret key does not grow: the qg reading is a diagnostic, not security.

Bibliographic notes

The BB84 protocol is Bennett and Brassard (2014) (reprint of the 1984 original), and the security proof that ties the QBER to the key, Shor and Preskill (2000). The finite-key bound is Tomamichel et al. (2012). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

18 Certified quantum random numbers (QRNG) and channel auditing

A one-qubit quantum random number generator (QRNG) is the simplest experiment in the book: prepare $|+\rangle$, measure in Z, write down the bit. The result is a fair coin, and that is what the usual analysis measures. But a fair coin is not the same as a private coin. If the qubit lost coherence, the output still comes out 50/50, and yet someone holding the environment can know every bit in advance. This chapter reads that difference on the scale of Chapter 1. The output lives in one reading and privacy in the other two, and what the adversary knows is one more reading: the one the state lacks to reach the sphere. The chapter adds the readout calibration needed to avoid overestimating, and the same certifier against a designed leak on IonQ's simulator.

18.1 From the base: output and privacy are different readings

The source qubit has three readings, m_X, m_Y and m_Z. The QRNG output is the measurement in Z, and the usual analysis looks only at its bias: with $p_0 = (1000 + m_Z)/2000$ it certifies

$$H_{\text{naive}} = -\log_2 \max(p_0, p_1) = -\log_2 \frac{1000 + |m_Z|}{2000}$$

bits per shot, 1 bit when $m_Z = 0$. But $m_Z = 0$ is given by $|+\rangle$ and by the completely mixed state alike, and in the second case nothing is private: the environment holds the other half of the purification and knows what came out.

With a trusted measurement and an adversary (Eve) holding the purification, Eve's two conditional states are pure, and their overlap depends only on the transverse radius $r_\perp = \sqrt{qg_X^2 + qg_Y^2}$. By Helstrom, her guessing probability is (notes, §41)

$$P_{\text{guess}} = \frac{1 + \sqrt{1 - r_\perp^2}}{2}, \qquad H_{\min}(Z|E) = -\log_2 P_{\text{guess}}.$$

On the scale of Chapter 1, that probability is one more reading. Its value in m-qg is

$$m_E = 1000\sqrt{1 - r_\perp^2} = \sqrt{1000^2 - m_X^2 - m_Y^2},$$

what the transverse radius lacks to reach the sphere, and its qang angle is $\arcsin r_\perp$. The formula does not depend on m_Z. A pure state on the equator has $m_\perp = 1000$ and leaves Eve at $m_E = 0$: she guesses with probability 1/2, and the bit is private. A state that lost coherence has $m_\perp < 1000$, and the length it lost is what Eve knows (Figure 18.1, left).

Table 18.1: The private randomness of a qubit measured in Z, as a function of its transverse radius.

$m_\perp$ (m-qg)	1000	990	950	900	800	500	200
what Eve knows, m_E (m-qg)	0	141	312	436	600	866	980
private bits per shot	1.000	0.810	0.608	0.478	0.322	0.100	0.015

The shape of the formula has a practical consequence. Ten m-qg less of transverse radius, from 1000 to 990, give Eve 141 m-qg and take 0.19 bits per shot from the source. It is the pole of Chapter 1 seen from the other side: near $m_\perp = 1000$, the square root has infinite slope.

18.2 The output-bias analysis fails; the qg reading does not

The qg certifier estimates m_X and m_Y with test rounds (rotations that bring X and Y to the measurement basis) and uses a 3σ lower bound on $r_\perp$. The notes (§41)

compare four estimators with 10^4 test rounds per setting, 300 runs per scenario and the asymmetric readout of §31 ($e_{01} = 0.015$, $e_{10} = 0.04$):

- **naive**: the bias of the output;
- **qg, one rotation**: m_X and m_Y measured with one rotation each;
- **qg, ± pairs**: each axis with both rotations, so that the readout offset cancels;
- **qg calibrated**: the ± pairs, divided by the readout scale factor b measured with the protocol of §31.

Table 18.2: Certified bits per shot: mean estimate and, in parentheses, the fraction of runs above the truth (unsafe).

scenario	$m_\perp$	truth	naive	qg, one rotation	qg, ± pairs	**qg calibrated**
ideal $\|+\rangle$	1000	1.000	0.965 (0)	0.578 (0)	0.512 (0)	0.681 (0)
T2: $V = 0.8$	800	0.322	0.964 (**1.00**)	0.269 (0)	0.245 (0)	**0.286 (0)**
T2: $V = 0.5$	500	0.100	0.965 (**1.00**)	0.087 (0)	0.077 (0)	**0.087 (0)**
T2: $V = 0.2$	200	0.015	0.965 (**1.00**)	0.013 (**0.07**)	0.009 (0)	**0.010 (0)**
thermal $p = 0.05$, $V = 0.9$	810	0.334	0.965 (**1.00**)	0.279 (0)	0.254 (0)	**0.297 (0)**
strongly asymmetric readout, $e_{10} = 0.10$, $V = 0.95$	950	0.608	0.876 (**1.00**)	0.498 (0)	0.341 (0)	**0.513 (0)**

- **The usual estimator is unsafe whenever the state is not pure.** At $V = 0.2$ it certifies 0.97 bits per shot where only 0.015 are private, in 100% of runs. Dephasing and thermal mixing leave the output unbiased.
- **The qg estimator with ± pairs is safe in every run**, because the readout offset a cancels. With one rotation, the offset inflates $r_\perp$ and overestimates in 7% of runs at $V = 0.2$. So the "100% against 0%" safety holds for the paired estimator, not for any qg reading whatsoever.
- **Calibrating b with §31 recovers 10 to 50% more certified bits** (0.286 against 0.245 at $V = 0.8$; 0.513 against 0.341 with the strongly asymmetric readout), without losing safety.

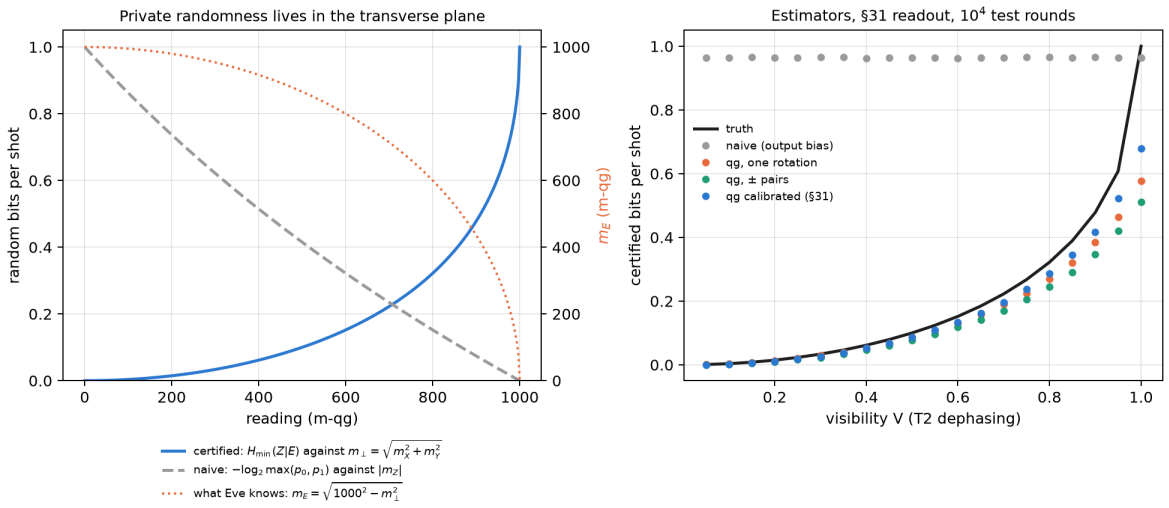


Figure 18.1: Left: random bits per shot. The naive estimator reads $|m_Z|$; the certified one reads the transverse radius $m_\perp$; the dotted curve is what Eve knows, $m_E = \sqrt{1000^2 - m_\perp^2}$ (right axis). Right: the four estimators against visibility, with the readout of §31 and 10^4 test rounds.

18.3 Near the pole certification is shot-hungry

A 3σ bound with n rounds per setting lowers the radius by $3000/\sqrt{n}$ m-qg, and near the pole that is expensive. For a perfect source, without readout error:

Table 18.3: Certified bits with a 3σ bound, closed form.

test rounds	10^3	10^4	10^5	10^6	10^7
bound on $m_\perp$ (m-qg)	905	970	991	997	999
what is conceded to Eve, m_E (m-qg)	425	243	137	77	44
certified bits (truth: 1)	0.489	0.686	0.814	0.892	0.939

test rounds	10^3	10^4	10^5	10^6	10^7
with $m_\perp = 950$ (truth: 0.608)	0.397	0.523	0.578	0.598	0.605

With the §31 readout and the calibrated certifier, the simulation gives 0.68, 0.81 and 0.89 bits with 10^4, 10^5 and 10^6 rounds, the same. Away from the pole the bound converges fast: with $m_\perp = 950$, 10^5 rounds already give 95% of the truth (Figure 18.2, left). A nearly perfect source is the most expensive to certify.

18.4 Readout calibration: temperature or detector error

The calibrated certifier divides by the readout scale factor b, and the question is where b comes from. On the scale of Chapter 1 the readout is affine (notes, §31):

$$m_{\text{measured}} = 1000\,a + b\,m, \qquad a = e_{10} - e_{01}, \quad b = 1 - e_{01} - e_{10}.$$

A real qubit starts with thermal population p, that is with $m_{\text{eq}} = 1000(1-2p) = 1000\tanh(hf/2k_BT_{\text{eff}})$. The standard calibration measures "$|0\rangle$" and "$|1\rangle$", which are really the thermal state and X on it: **two readings for three unknowns** (a, b, m_{eq}). It assumes $m_{\text{eq}} = 1000$ and charges the thermal population to the readout: it reports $e_{01} + p\,b$ and $e_{10} + p\,b$, exactly.

Table 18.4: The standard calibration against the qg protocol of §31 ($e_{01} = 0.015$, $e_{10} = 0.04$). The standard column is exact; the qg column is the median of 100 repetitions from §31.

p	T_{eff} (5 GHz)	m_{eq} (m-qg)	standard e_{01}	qg protocol e_{01} (§31)
0	0 mK	1000	0.0150	0.0149
0.01	52 mK	980	0.0245	**0.0151**
0.03	69 mK	940	0.0434	**0.0148**
0.08	98 mK	840	0.0906	**0.0152**

The qg protocol adds the third equation. It measures, reads again and uses the covariance of the two readings, $V = b^2(1 - qg_{\text{eq}}^2)$, which separates b from m_{eq} in

closed form. According to the notes, it recovers the temperature to ±1 mK from 52 mK up. Chapter 27 develops it. What matters here is its consequence for the QRNG: **calibrating with the standard suite makes the certifier unsafe**. With a thermal source ($V = 0.9$), dividing by the standard b, which is $b(1 - 2p)$, returns $r_\perp$ without the thermal contraction and certifies too much.

Table 18.5: Thermal source with $V = 0.9$, 10^4 test rounds, 300 runs: the certifier with each calibration.

p	truth	calibrated with the standard b	calibrated with the true b
0.01	0.443	0.418 (0.01 unsafe)	0.390 (0)
0.03	0.383	0.414 (**1.00** unsafe)	0.339 (0)
0.05	0.334	0.413 (**1.00** unsafe)	0.297 (0)

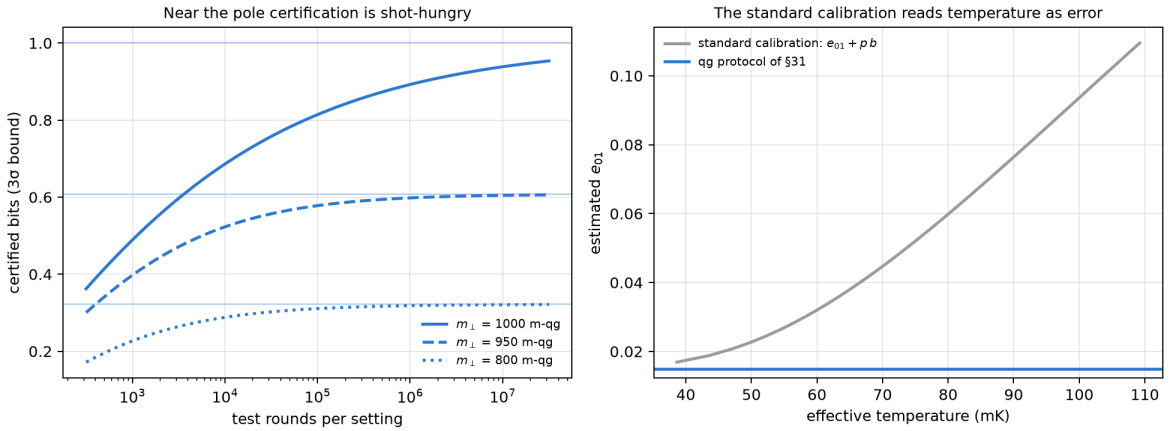


Figure 18.2: Left: certified bits with a 3σ bound against the number of test rounds, for three transverse radii; the thin lines are the truth. Right: the e_{01} reported by the standard calibration grows with the effective temperature, $e_{01} + p\,b$; the qg protocol of §31 keeps it at its value.

18.5 On IonQ's noisy simulator, with a leak

In the notes (§42) the leak is real: an ion q is the source and a second ion e plays the environment. An $\mathrm{MS}(0, 0, \theta)$ interaction lets e learn the X value of q, and the output is read in X. It is still a fair coin, but the private randomness falls to the formula above with $r = |\cos 2\pi\theta|$ (X and Z exchanged). Test rounds measure ±Z and ±Y, and two circuits calibrate the readout. Everything is in native gates (GPI, GPI2, MS), with 5 pairs per circuit, 10^4 samples per setting and three repetitions per noise model, on the free simulator. Nothing was sent to a QPU.

Table 18.6: Certified bits per shot, mean of 3 runs (aria-1 / forte-1 where they differ; $m_\perp$ from aria-1).

θ (turns)	truth	naive	qg, one rotation	**qg, ± pairs**	measured / ideal $m_\perp$
0 (no leak)	1.000	0.98	0.69	0.73	1000 / 1000
0.04	0.680	0.99 / 0.98	0.53	**0.56 / 0.55**	954 / 969
0.08	0.433	0.98 / 0.97	0.37 / 0.36	**0.38 / 0.37**	866 / 876
0.12	0.248	1.00 / 0.98	0.22 / 0.21	**0.22**	721 / 729

- **The naive estimator does not see the leak:** 0.97–1.00 bits per shot in every run, above the truth in 100% of runs with $\theta > 0$. At $\theta = 0.12$ it certifies four times the private randomness.
- **The qg estimator is safe in all 24 cells** (run × θ) and captures 81–90% of the truth when there is a leak, and 73% without one, because of the pole of the previous section.
- **Calibration adds nothing here.** The simulator's readout is almost perfect ($b = 0.9998$–1.0000, $|a| \leq 2 \cdot 10^{-4}$). The 10–50% gain needs hardware-like readout errors, and this is a limit of the noise model that only a QPU can check.
- **Gate noise is counted as lost privacy:** the measured $m_\perp$ sits 1–2% below the ideal, which is the MS error. That is as it should be, because the certifier cannot know who took that length.

18.6 What remains open

There is no hardware run. The missing test is a QRNG on a device with real readout errors, where the §31 calibration should show its 10–50% gain. It is a device-dependent QRNG (trusted measurement, independent and identical rounds), not

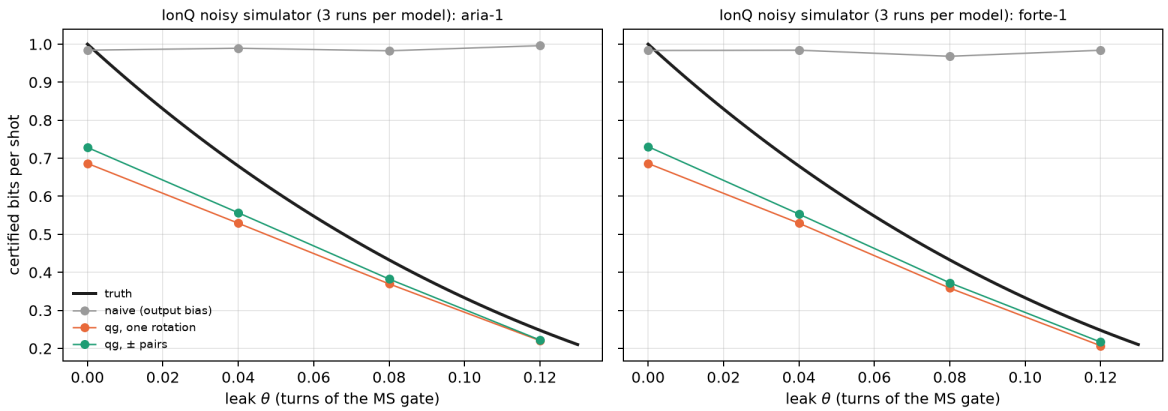


Figure 18.3: Certified bits per shot against the strength of the leak on IonQ's simulator (aria-1 and forte-1). The naive estimator stays near 1; the qg estimators always stay below the truth.

device-independent randomness. And the leak of §42 is a designed interaction, not an adversary's best strategy.

18.7 Without qang and with qang

Table 18.7: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
certify bits from the output bias	$-\log_2 \max(p_0, p_1)$	the same in m_Z	**none**: both fail alike, 100% unsafe under dephasing
certify private bits	Helstrom's $H_{\min}(Z \mid E)$ with the purification	$m_E = \sqrt{1000^2 - m_\perp^2}$: what Eve knows is the missing length	**equivalent** in the computation (the formula is standard); **in favour** in the reading: privacy is in the transverse plane and is measured with two rotations

action	without qang	with qang	difference
how many test rounds are needed	error propagation on $H_{\min}$	the pole of Chapter 1: 10 m-qg of radius give Eve 141 m-qg	**in favour**, modest: it shows at a glance why the perfect source is the most expensive
calibrate the readout	two circuits; temperature is read as error	affine readout and a third equation (§31)	**in favour**: it avoids an unsafe certifier (100% of runs with $p \geq 0.03$) and recovers 10–50% of bits
detect a leak (simulated IonQ)	the bias does not see it	safe in 24 of 24 cells, 81–90% of the truth	**in favour** of the transverse estimator, which is standard in device-dependent QRNGs

Balance. The entropy bound is the usual one, and a QRNG expert already measures transverse coherence. What the Chapter 1 base adds is the reading: what the adversary knows is a reading in m-qg, m_E, the length the state lacks to reach the sphere. Two things follow that can be seen at a glance: the drop of certification near the pole and the danger of calibrating the readout with two circuits. The second is the chapter's most concrete gain, and it is owed to the §31 protocol, not to the unit.

18.8 In code

```
import math
import sys
sys.path.insert(0, "examples")
import qrng_qg_certified as Q

m_x, m_y = 800.0, 0.0                          # transverse readings (m-qg)
m_perp = math.hypot(m_x, m_y)
```

```
m_e = math.sqrt(1000**2 - m_perp**2)        # 600: what Eve knows
-math.log2((1000 + m_e) / 2000)             # 0.322 private bits per shot
Q.h_min_certified(m_perp / 1000)            # 0.322, the same computation
```

18.9 Summary

- The output of a one-qubit QRNG is the reading m_Z; its privacy is the transverse radius $m_\perp$. What an adversary holding the purification knows is one more reading, $m_E = \sqrt{1000^2 - m_\perp^2}$.
- The output-bias analysis certifies almost 1 bit per shot even when the state is almost mixed: unsafe in 100% of runs with dephasing or thermal mixing. The qg estimator with ± pairs is safe in all of them.
- Near the pole certification is expensive: 10 m-qg of radius are worth 0.19 bits, and a perfect source gives 0.69 certified bits with 10^4 rounds and 0.89 with 10^6.
- The standard calibration reads temperature as readout error, and using it makes the certifier unsafe. The §31 protocol separates them and recovers 10–50% of bits.
- On IonQ's simulator, with a designed leak, the naive estimator does not see it and the qg one is safe in every cell.

Bibliographic notes

Optimal state discrimination is Helstrom (1976), and the operational meaning of the min-entropy, Konig et al. (2009). For quantum random number generators, see the reviews Herrero-Collantes and Garcia-Escartin (2017) and Ma et al. (2016). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

19 Quantum batteries and networks

A quantum battery is a system that stores energy to deliver it later as work. The smallest one is a qubit with $H = \omega|1\rangle\langle 1|$: charging it means bringing it towards $|1\rangle$, and discharging it means returning it to $|0\rangle$ with a unitary operation. The practical question is not how much energy the qubit has, but how much can be taken out, which is its **ergotropy**. This chapter reads it on the scale of Chapter 1. Energy, ergotropy and the part nobody can take out are three distances between readings. With that, the chapter answers how to charge a battery that will wait with T1 and T2, how much it costs to certify its charge from shots, and where the charge of a qubit register sits when it lives in the correlations. The last section turns to networks: a Bell pair stored in two memories is another thing stored with T1 and T2, and its useful time is read the same way.

19.1 From the base: three distances between readings

A qubit with readings (m_X, m_Y, m_Z) and radius $m_r = \sqrt{m_X^2 + m_Y^2 + m_Z^2}$ has energy $\omega(1000 - m_Z)/2000$: the distance of its Z reading from the north pole, which is the ground state. Its ergotropy, the maximum work a unitary operation extracts, is (notes, §56)

$$W = \omega\,\frac{m_r - m_Z}{2000},$$

and what is left unextracted, the energy of the passive state, is $\omega(1000{-}m_r)/2000$. **The energy splits between how far the state sits below the top of its own sphere, $m_r{-}m_Z$, and how much radius the sphere lost, $1000{-}m_r$.** A unitary rotates the state without changing its radius, so the best it can do is bring it to the top of its sphere, $m_Z = m_r$. The part lost to mixing is out of reach of any unitary.

For a pure state, $m_r = 1000$ and $W = \omega\sin^2(\theta/2) = \omega\,p_1$: **the ergotropy is the probability equation of Chapter 1**, with the qang angle as the charging angle (checked to $2 \cdot 10^{-16}$). On the angular qga scale (Section 1.4), $W = \omega\sin^2(\pi\theta_{qga}/1000)$: the battery is empty at 0 qga, half charged at 250 qga and full

at 500 qga. The notes also split $W = W_{\text{inc}} + W_{\text{coh}}$, with $W_{\text{inc}} = \omega \max(0, -q g_Z)$ (population inversion) and $W_{\text{coh}} = \frac{\omega}{2}(r - |q g_Z|)$ (coherence). The formulas match the general eigenvalue expression to $3 \cdot 10^{-16}$ on 200 random states.

Table 19.1: The energy of a qubit in units of ω, split according to its readings.

state	m_Z	m_r	energy	extractable	locked by mixing
pure, inverted	−1000	1000	1.000	1.000	0
pure, on the equator	0	1000	0.500	0.500	0
pure, $\theta = 2\pi/3$	−500	1000	0.750	0.750	0
mixed, down	−500	500	0.750	0.500	0.250
mixed, on the equator	0	500	0.500	0.250	0.250
inverted and aged, $t = 0.8\,T_1$	+101	101	0.449	**0**	0.449

The last row is a **passive** battery: it has almost half a unit of energy and delivers nothing, because it sits at the top of its sphere ($m_Z = m_r$). This happens to every inverted battery when its Z reading crosses the equator upwards.

19.2 Storage with T1 and T2: inversion dies at the equator

A battery charged at angle θ ages as

$$m_Z(t) = 1000\big[1 - (1 - \cos\theta)\, e^{-t/T_1}\big], \qquad m_\perp(t) = 1000 \sin\theta\, e^{-t/T_2}.$$

The inverted battery ($\theta = \pi$) has no coherence and stores only inversion: $m_Z = 1000(1 - 2e^{-t/T_1})$, which crosses the equator at $t = T_1 \ln 2$. There its ergotropy is zero, and it does not come back. A tilted battery stores less at first, but its coherence keeps it below the top of its sphere after the crossing, as long as T_2 allows.

Table 19.2: Best charging angle and stored ergotropy (best / inverted / equator, in units of ω).

T_2/T_1	crossover t_c	$t = 0.5\,T_1$	$t = T_1$	$t = 2\,T_1$
2	$0.288\,T_1 = \ln(4/3)$	0.70π: 0.297 / 0.213 / 0.240	0.58π: 0.129 / 0 / 0.122	0.038 / 0 / 0.038

T_2/T_1	crossover t_c	$t = 0.5\,T_1$	$t = T_1$	$t = 2\,T_1$
1	$0.405\,T_1 = \ln(3/2)$	0.78π: 0.234 / 0.213 / 0.165	0.59π: 0.054 / 0 / 0.050	0.005 / 0 / 0.005
0.5	$0.618\,T_1$	π: 0.213 / 0.213 / 0.073	0.60π: 0.008 / 0 / 0.007	≈ 0
0.2	$0.692\,T_1$	π: 0.213 / 0.213 / 0.004	≈ 0	≈ 0

- **The charging rule.** If the battery will be used before t_c, charge it fully ($\theta = \pi$). Otherwise tilt it, to about 0.6π at $t = T_1$. The crossover solves $x^{2T_1/T_2 - 1} = 2(2x - 1)$ with $x = e^{-t_c/T_1}$, which comes from a small-tilt analysis and matches the numerical optimum to 10^{-6}; for $T_2 = 2T_1$ and $T_2 = T_1$ it gives the closed forms $\ln(4/3)$ and $\ln(3/2)$.
- **The crossover tends to the equator.** As T_2 shrinks, t_c tends to $T_1 \ln 2 = 0.693\,T_1$, the moment the Z reading of the inverted battery crosses the equator. Without coherence there is nothing to tilt.
- **Coherence buys time only when T_2 is comparable to T_1.** At $t = T_1$ a tilted battery stores $0.13\,\omega$ ($T_2 = 2T_1$) or $0.05\,\omega$ ($T_2 = T_1$), where the inverted one stores nothing. At $T_2 = 0.2\,T_1$ it buys nothing. In trapped ions, where T1 hardly exists, always charge fully.

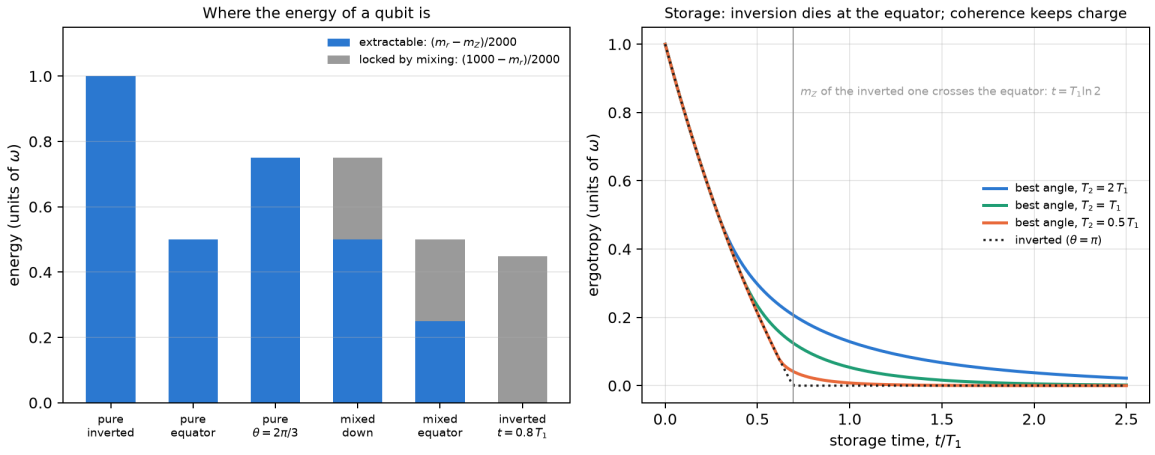


Figure 19.1: Left: the energy of six states, split between the extractable part, $(m_r - m_Z)/2000$, and the part locked by mixing, $(1000 - m_r)/2000$; the mixed states have radius 500 m-qg. Right: the ergotropy stored with the best charging angle for three ratios T_2/T_1, and that of the inverted battery, which dies when its m_Z crosses the equator.

19.3 Certifying the charge from shots

Ergotropy is estimated from shots on the three axes, and the formula has a trap. The estimated radius comes out inflated, because the noise of the three axes adds in quadrature: $E[\hat{r}^2] = r^2 + 3\sigma^2$. And a passive battery sits right at the top of its sphere, so any inflation of the radius makes it look charged.

Table 19.3: Estimated ergotropy (units of ω), mean of 4000 runs; in parentheses, the fraction of runs above the truth.

shots per axis	100	300	1000	3000	10,000
inverted, $t = 0.5\,T_1$ (truth 0.213): plug-in	0.237 (0.60)	0.220 (0.54)	0.215 (0.52)	0.214 (0.52)	0.213 (0.50)
inverted: 3σ lower bound	0.005	0.025	0.086	0.139	0.172
passive, $t = 0.8\,T_1$ (truth 0): plug-in	0.042 (0.99)	0.016 (1.00)	0.005 (1.00)	0.002 (1.00)	0.001 (1.00)
passive: 3σ lower bound	0	0	0	0	0

- **The plug-in estimate overestimates in more than half of the runs**, and reports charge in a passive battery in 99–100% of them.
- **The 3σ bound is never unsafe, but it is expensive:** with 1000 shots per axis it certifies 0.086 of $0.213\,\omega$, and with 10^4 still only 81%. A nearly discharged battery cannot be certified with any reasonable number of shots.

It is the same problem as the QRNG of Chapter 18: a quantity that depends on the radius, read next to an edge.

19.4 Registers: charge locked in the correlations

In a register of qubits the charge can live in the correlations. The notes compare two batteries of 4 qubits with 2 excitations, both with energy 2ω, aged with local

noise ($T_2 = T_1$):

- the **Dicke state** $D(4,2)$, the even superposition of the six registers with two ones;
- a **product** of four qubits on the equator, with the same energy.

Read qubit by qubit, they are opposites. In the Dicke state each qubit reads $(m_X, m_Y, m_Z) = (0,0,0)$: local radius zero, the centre of the sphere, a passive qubit. In the product each qubit reads $(1000, 0, 0)$: radius 1000, ergotropy $\omega/2$. Yet the mean Z reading of the register is identical in both at all times (0, 95, 259, 503 and 777 m-qg for $t = 0$, 0.1, 0.3, 0.7 and 1.5 T_1), and does not tell them apart.

Table 19.4: 4-qubit register with 2 excitations (ergotropy in units of ω).

t/T_1	0	0.1	0.3	0.7	1.5
Dicke: global ergotropy	2.000	1.489	0.710	0.049	0
Dicke: sum of the local ones	0	0	0	0	0
product: global = sum of the local ones	2.000	1.629	1.051	0.407	0.063
mean m_Z of the register (m-qg)	0	95	259	503	777

- **The Dicke charge is locked.** All of it is extractable with a global operation, and none with one-qubit operations: the sum of the local ergotropies is zero. The product's charge is all local.
- **The local radius tells them apart; the mean reading does not.** The sum of the local ergotropies, which comes from the three readings of each qubit, separates one battery from the other.
- **Locked charge is more fragile** under local noise: 0.71 against $1.05\,\omega$ at $0.3\,T_1$, and 0.05 against $0.41\,\omega$ at $0.7\,T_1$.

Local radius zero is not new in the book. It is what the qubits of a Bell pair in a network read (Chapter 12, §57), and the counting qubits of Shor with $N = 21$

before the QFT (Chapter 16). In all three cases what matters is in the correlations, and a one-qubit reading can only say that it is not in that qubit.

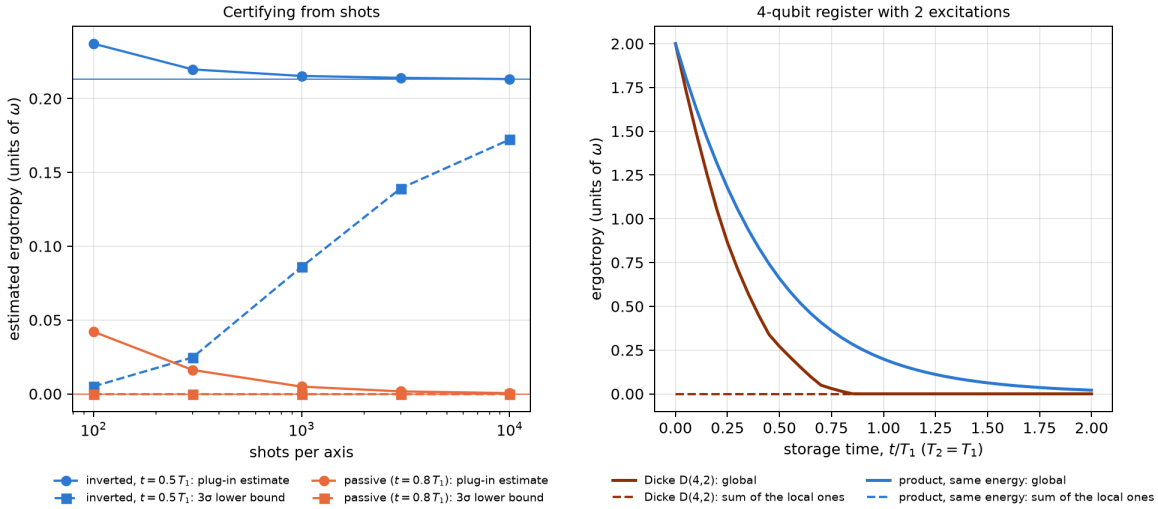


Figure 19.2: Left: plug-in estimate and 3σ lower bound of the ergotropy against shots per axis, for an inverted battery at $t = 0.5\,T_1$ and a passive one at $t = 0.8\,T_1$; the thin lines are the truth. Right: global ergotropy and sum of the local ones for the Dicke register and the product with the same energy; for the product the two curves coincide.

19.5 Networks: a Bell pair stored in two memories

A quantum network also stores something in memories with T1 and T2: the Bell pairs it distributes among its nodes. Chapter 12 read a pair in qg (three correlators and two local biases) and used that reading to choose how to distil it. The question of this section is still open, and it is the battery's question: how long a pair can wait in memory before it is no longer useful. For $|\Phi^+\rangle$ with both halves stored for a time t, the readings have closed forms (notes, §57; they match the density matrix to $5 \cdot 10^{-16}$):

$$c_x = -c_y = e^{-2t/T_2}, \qquad c_z = 1-2g+2g^2, \qquad m_A = m_B = 1000\,g, \qquad g = 1-e^{-t/T_1}.$$

The local biases m_A and m_B rise from zero: it is the T1 witness on a shared pair, the same reading that shifted the centre of the image of the sphere in BB84

(Chapter 17). The pair is good for a BBM92 key while $1 - h(Q_X) - h(Q_Z) > 0$, with $Q_X = (1 - c_x)/2$ and $Q_Z = (1 - c_z)/2$.

Table 19.5: How long a stored Bell pair can wait, for a BBM92 key (positive secret fraction per input pair).

T_2/T_1	2	1	0.5	0.2
cutoff without distillation (T_1)	0.184	0.129	0.088	0.051
readings at the cutoff: $1000\,c_x$ / $1000\,c_z$ / m_A (m-qg)	832 / 720 / 168	773 / 787 / 121	702 / 845 / 85	598 / 905 / 50
cutoff with one DEJMPS round (T_1)	0.395	0.232	0.144	0.076
distilling pays after (T_1)	0.085	0.083	0.056	0.032

- **A pair lasts much less than a battery.** The inverted battery keeps charge up to $0.69\,T_1$; the pair stops giving key at 0.05–$0.18\,T_1$, because the key needs correlators near the Shor–Preskill threshold of Chapter 16: 780 m-qg when the error is equal in both bases, and here shared between them (832 and 720 m-qg at $T_2 = 2T_1$). What is stored is a correlation, and the key needs it nearly intact.
- **One distillation round almost doubles the useful time** when T1 limits it (0.184 → 0.395 at $T_2 = 2T_1$). It costs half the pairs, so with fresh pairs it lowers the key per input pair. Distil only after the crossover in the last row, and that crossover is read from the aged correlators.
- **Both memories follow the same underlying rule.** In the battery, coherence (T_2) buys time after the inversion crossing. In the pair, distillation buys time after its own crossover. In both cases the crossover comes from the readings.

19.6 What remains open

Everything is exact or simulated: there is no hardware run. Missing are a real battery charged at different angles on a device with $T_2 \approx T_1$, where the charging

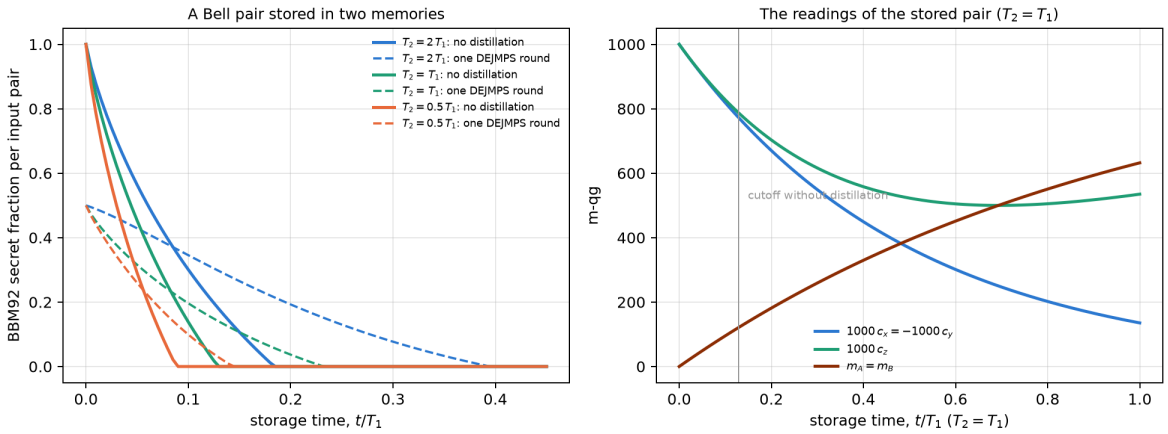


Figure 19.3: Left: BBM92 secret fraction per input pair of a stored Bell pair, without distillation and with one DEJMPS round, for three ratios T_2/T_1. Right: the readings of the stored pair with $T_2 = T_1$; the local biases $m_A = m_B$ rise from zero (the T1 witness), and the grey line is the cutoff without distillation.

rule can be tested, and the preparation of a Dicke $D(4, 2)$ with a measurement of its local ergotropy. On the network side, distillation is computed on the twirled (Bell-diagonal) state with perfect local gates, and the cutoff assumes equal memories at the two nodes.

19.7 Without qang and with qang

Table 19.6: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
compute the ergotropy	eigenvalues of ρ and the passive energy	$(m_r - m_Z)/2000$	**equivalent**: it is the standard one-qubit formula written in readings
know how much energy nobody can take out	energy of the passive state	$(1000 - m_r)/2000$, the lost radius	**in favour**, modest: the split is read at a glance

action	without qang	with qang	difference
ergotropy of a pure state	$\omega \sin^2(\theta/2)$	the probability equation of Chapter 1, $\omega\, p_1$	**equivalent**; it shows the Chapter 1 base is the base of the charge
decide how to charge	numerical optimisation of the angle	closed crossover $x^{2T_1/T_2-1} = 2(2x-1)$; inversion dies at the equator crossing	**in favour**: a closed rule and a picture (the equator); the rule is a contribution of the notes, not of the unit
certify the charge	plug-in estimate	the passive battery sits at the top of its sphere, and the inflated radius makes it "charged"	**in favour** as a diagnosis; the cure, the 3σ bound, is standard statistics
tell local from locked charge	global and local ergotropy by tomography	local radius of each qubit; the mean reading is useless	**equivalent** in the computation, **in favour** in the reading: it is the same local radius zero of Chapters 12 and 16
decide how long to store a Bell pair	fidelity and key rate of the aged state	closed-form correlators and biases; cutoff and distillation crossover read from them	**equivalent** in the computation; cutoffs are known, the closed form and the crossover are contributions of the notes

Balance. The formulas are the standard ones for the ergotropy of a qubit, and batteries with coherence-assisted storage or locked charge are already in the literature. What the Chapter 1 base adds is the reading: the energy splits between how far the state sits below the top of its sphere and how much radius it lost. With that reading one sees why the inverted battery dies at the equator, why the plug-in estimate sees charge where there is none, and why a Dicke register looks discharged qubit by qubit. The charging rule with its closed crossover is a

contribution of the notes.

19.8 In code

```
import math
import sys
sys.path.insert(0, "examples")
import battery_ergotropy_qg as BT

m_x, m_y, m_z = 500.0, 0.0, 0.0             # mixed, on the equator (m-qg)
m_r = math.sqrt(m_x**2 + m_y**2 + m_z**2)
(1000 - m_z) / 2000, (m_r - m_z) / 2000, (1000 - m_r) / 2000  # energy 0.5, extractable 0.2
BT.ergotropy(0.5, 0.0, 0.0)                  # 0.25, the same computation
BT.crossover_time_analytic(1.0), math.log(1.5)  # 0.405 and 0.405: the crossover with T2
```

19.9 Summary

- The energy of a qubit splits between the extractable part, $(m_r - m_Z)/2000$, and the part locked by mixing, $(1000 - m_r)/2000$; for a pure state the ergotropy is the probability equation of Chapter 1, $\omega\, p_1$.
- An inverted battery dies when its Z reading crosses the equator, at $t = T_1 \ln 2$. A tilted battery keeps charge after that if T_2 is comparable to T_1, and the charging rule has a closed crossover: $\ln(4/3)$ for $T_2 = 2T_1$, $\ln(3/2)$ for $T_2 = T_1$.
- The plug-in estimate of the ergotropy overestimates, and reports charge in a passive battery in almost every run; the 3σ bound is safe but expensive.
- In a Dicke register all the charge is in the correlations: each qubit reads the centre of the sphere, the sum of the local ergotropies is zero, and the mean reading does not tell it apart from a product with the same energy.
- A stored Bell pair gives key only up to 0.05–$0.18\, T_1$, much less than a battery; one distillation round almost doubles that time, and pays only after a crossover read from the correlators.

Bibliographic notes

Passive states are Pusz and Woronowicz (1978), and ergotropy, Allahverdyan et al. (2004). The entanglement advantage in extractable work is Alicki and Fannes (2013); for quantum batteries in general, see Campaioli et al. (2024). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part VII

VII. Quantum machine learning (QML) and the symmetry filter

20 Weight-conserving quantum neural networks (WeightQNN)

Part VII is about quantum machine learning, and it starts with the network that fits the base of Chapter 1 best. A quantum neural network (QNN) that conserves the Hamming weight (the number of ones in the register) lives in a fixed sector, and that has two consequences. The first is good: the qg filter, which discards the shots that left the sector, makes the readout exact under T1 relaxation. The second is honest: such a network can be simulated on an ordinary computer, and cannot give a quantum advantage. This chapter reads the network in m-qg, checks its structural properties with code, and summarises the classification experiments of the notes (§75–§114), all with predictions committed before running, with and without qang.

20.1 From the base: a sector and its readings

The RBS gate (*reconfigurable beam splitter*) rotates the pair $|01\rangle, |10\rangle$ by an angle θ and leaves $|00\rangle$ and $|11\rangle$ alone: it is a Givens rotation, and it does not change the number of ones. In the weight-1 sector of n qubits the state is a unit vector $u \in \mathbb{R}^n$, a layer of RBS gates is an orthogonal matrix, and each qubit has the reading of Chapter 1:

$$m_Z(q) = 1000\,(1 - 2u_q^2),$$

because u_q^2 is the probability that the excitation is on qubit q. The notes' network E (§76) loads the data as amplitudes, $v = (x_1, \ldots, x_4, 1)/|(x, 1)|$ on 5 qubits (the constant keeps the norm of x), applies 3 layers of RBS gates (9 sublayers, 15 parameters) and reads $z = \sum_q c_q\, qg_Z(q) + b$. Network W does the same in the weight-2 sector.

The filter under T1. T1 relaxation takes a weight-k state to weight $k-1$ and, if every qubit has the same T1, damps all weight-k amplitudes by the same factor. Keeping only the weight-k shots then returns the noiseless distribution, exactly, for any strength (notes, §76 F1 and §78 F4). The price is known before running:

the filter keeps a fraction $K = (1-\gamma)^{k\,d}$ of the shots, with d the number of sublayers, whatever the data or the parameters (§77).

Table 20.1: The filter under equal T1 ($\gamma = 0.08$ per qubit and per sublayer, 9 sublayers), in m-qg, over 40 random inputs and parameters. The last two columns are the filter's mean bias with unequal T1 (spread 0.5) and with dephasing 0.01 per sublayer, which the filter does not correct.

model	max $\lvert$filtered − noiseless$\rvert$	mean $\lvert$raw − noiseless$\rvert$	simulated K	$(1-\gamma)^{kd}$	bias with unequal T1	bias with dephasing
E, 5 qubits, weight 1	$6 \cdot 10^{-13}$	211	0.4722	0.4722	33	64
W pairs, 5 qubits, weight 2	$8 \cdot 10^{-13}$	426	0.2229	0.2229	45	75
W dual with qq_{ZZ}, weight 2	$8 \cdot 10^{-13}$	506	0.2229	0.2229	31	89
E, 8 qubits, weight 1	$7 \cdot 10^{-13}$	132	0.4722	0.4722	29	48
W ring, 8 qubits, weight 2	$5 \cdot 10^{-13}$	264	0.2229	0.2229	38	35

On the scale of Chapter 1 one sees what the noise does without the filter. A shot that decayed has one fewer one, and its qubits read $+1000$: **raw readings drift towards the north pole**. In the example of Figure 20.1 (weight 2, dual encoding), the five noiseless readings lie between 157 and 260 m-qg, the raw ones under T1 between 543 and 664, and the filtered ones return exactly to 157–260. The filter does not correct what does not change the weight: unequal T1 and dephasing leave biases of 30 to 90 m-qg.

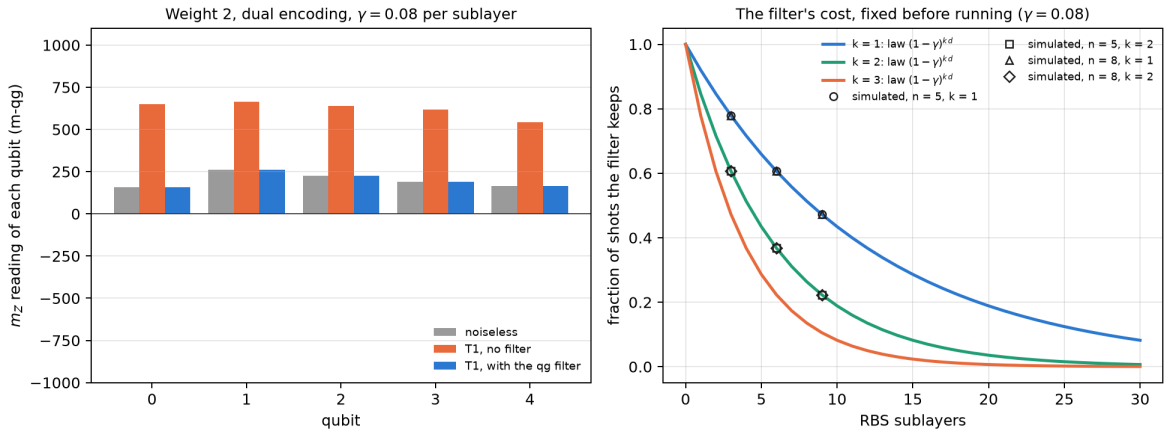


Figure 20.1: Left: m_Z reading of each qubit of a weight-2 network (dual encoding), noiseless, under T1 without the filter and under T1 with the filter; the raw readings drift towards the north pole and the filtered ones match the noiseless ones. Right: the fraction of shots the filter keeps against the number of sublayers; the lines are the law $(1-\gamma)^{kd}$ and the points, simulations on 5 and 8 qubits.

20.2 Architectures and encodings

The notes compare, on four real datasets with 4 features (iris, breast cancer, wine and digits 3 against 8), a standard QNN (A: angle encoding, Ry/Rz layers with a CZ ring) with the weight-conserving networks. The first weight-1 network (D, §75) dropped the norm of x and was the weakest (0.765). Network E, which keeps it, closes the gap (§76):

Table 20.2: Mean test accuracy (5 splits, notes §76).

dataset	RBF-SVM	MLP	quadratic twin	A	D	E
iris	0.913	0.920	0.900	0.907	0.647	0.913
cancer	0.967	0.973	0.957	0.980	0.910	0.983
wine	0.979	0.990	0.969	0.964	0.795	0.985
digits	0.967	0.960	0.943	0.957	0.813	0.960
mean	0.957	0.961	0.942	0.952	0.791	0.960

With 60 runs per model (5 seeds × 3 splits × 4 datasets, §80), the central result is this:

Table 20.3: Mean accuracy over 60 runs (notes §80).

reading	E with qang	E without qang	W with qang	W without qang
noiseless	0.952		0.922	
T1, trained noiseless	0.952	0.928	0.922	0.757
T1, trained under the noise	0.952	0.951	0.922	0.924
dephasing, trained noiseless	0.918	0.792	0.882	0.656

- **Trained on a simulator and run under T1, qang decides:** +2.5 points at weight 1 [0.8, 4.1] and +16.5 at weight 2 [11.8, 21.2] (95% CI across seeds). Without the filter, W falls to 0.76 because 78% of its shots decayed.
- **Trained under the calibrated noise, qang adds nothing measurable** (+0.1 and −0.1 points). The network without the filter learns the decay.
- **Dephasing requires training under the noise** (+3.2 and +4.5 points), although the filter still removes its T1 part (+12.6 and +22.6 points when trained noiseless).
- **Weight 2 cost 3 points**, and the cause turned out to be the encoding, not the circuit (§87). The pair product distorts the data: its information ceiling, measured with a regression on all products of amplitudes, is 2.4 points lower. With the **dual** encoding, which puts v on the ring of pairs and $u = (x^2, 1)/|(x^2, 1)|$ on the chords and covers all 10 weight-2 states, and with the qg_{ZZ} readout, weight 2 reaches weight 1: 0.947 = 0.947 (§90). Its gain with qang under T1 is 14 to 20 points, against 3 at weight 1.

20.3 Readouts: qg_Z, qg_{ZZ} and two channels

The readout is a linear function of the sector distribution, and its rank says how much of that distribution it sees:

Table 20.4: Rank of the readout on the sector distribution (with the constant term).

sector	states	rank of the qg_Z readout	rank of qg_Z plus qg_{ZZ}
weight 1, 5 qubits	5	5	5
weight 2, 5 qubits	10	5	10

At weight 1 the five m_Z readings already determine the whole distribution, and the correlations add nothing. At weight 2 the five readings see only 5 of the 10 directions, and the correlations $qg_{ZZ} = \langle Z_i Z_j \rangle$ are needed. But rank was not the cause of the gap: with the pair encoding, adding qg_{ZZ} changed nothing (§81). Both things were needed together, an encoding that does not distort and a full-rank readout (§90).

The shots that decayed. In §81 there was one case against qang: with the pair encoding, the qg_{ZZ} readout and training under T1, the network without the filter scored 1.2 points more [0.4, 2.0], because the decayed shots still say where the excitations were. The two-channel readout (§95), which gives the classifier the filtered and the raw readings together, recovered nothing, and with 60 new runs the effect dropped to +0.5 [−0.7, 1.6] (§101). No readout trained under T1 beat the filter significantly, and with the simple readout the filter was better by 0.6 points [0.1, 1.0].

20.4 Classical simulability: no quantum advantage

In the weight-1 sector the whole network is an $n \times n$ orthogonal matrix O. With $u = Ov$,

$$z = \sum_q c_q \, (1 - 2u_q^2) + b,$$

a quadratic classifier in v that is computed in $O(n^2)$ (§76 F2). The chapter's script checks it: the network's output and the quadratic twin's agree to $9 \cdot 10^{-16}$. At weight k the sector has dimension $\binom{n}{k}$, and the network stays simulable as long as that number is manageable. That is why the notes say it plainly: on real data, the best classical model ties or beats the best QNN (§75: RBF-SVM 0.963 against 0.933 on digits), and a well-regularised quadratic twin sits within 0.3–1.6 points of E (§76). **What these networks offer is robustness on noisy hardware, not quantum advantage.**

20.5 Scaling and multiclass

With 5, 6 and 8 qubits and a fixed depth of 9 sublayers (§111), the filter keeps every network at its noiseless accuracy. Without it, the weight-1 loss stays at 3.1–3.4 points, but the weight-2 loss grows from 8.7 to 26.0 points: an 8-qubit weight-2 network falls from 0.965 to 0.705 (Figure 20.2, left). Training under the noise still catches up with the filter (within 0.9 points in every cell).

In the multiclass version each qubit is a class (`qang.qml.MultiClassQNN`, §110 and §114). With 8 qubits and 3, 5 and 8 classes of digits, the filter keeps every model exactly at its noiseless accuracy; without it, the linear-head readout loses 0.7, 15.9 and 11.1 points. Unlike at 5 qubits, training under the noise no longer fully catches up: it stays 2–3 points behind with 5 and 8 classes.

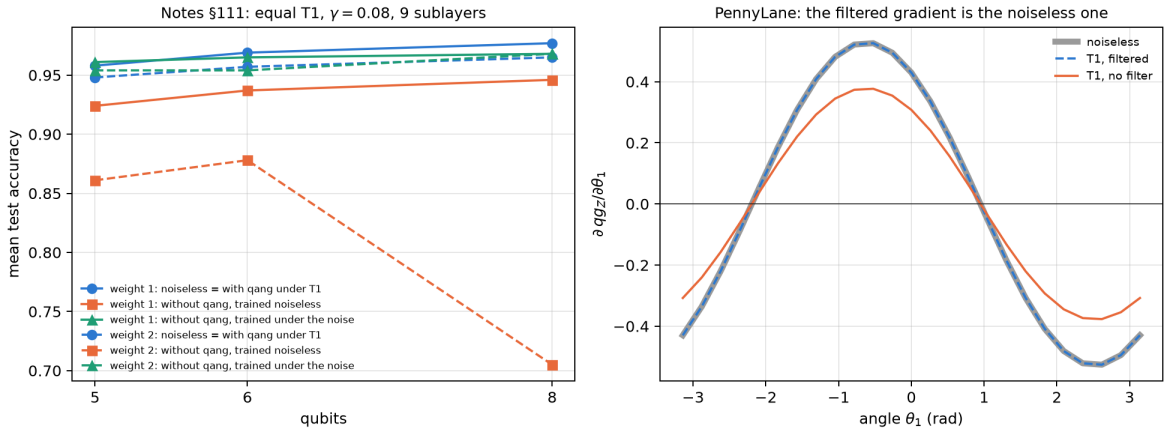


Figure 20.2: Left: mean test accuracy against the number of qubits (table of the notes, §111), with the filter under T1 (equal to the noiseless accuracy), without qang trained noiseless, and without qang trained under the noise, for weight 1 and weight 2. Right: gradient of one qubit's qg_Z with respect to the first angle in a three-qubit PennyLane circuit with Givens gates and amplitude damping ($\gamma = 0.08$); the filtered gradient matches the noiseless one, and the raw one is the noiseless one times the kept fraction.

20.6 Training under noise or with the filter

If the filtered readout under T1 is the noiseless one for every parameter, so are the loss and every gradient. That is why training under T1 with the filter gives

exactly the parameters of noiseless training (§77 F3: to $5 \cdot 10^{-9}$ after 120 Adam steps). The script checks it in PennyLane, with Givens gates (`SingleExcitation`) and amplitude damping on `default.mixed`: the gradient of the filtered readout matches that of the noiseless circuit to 10^{-16}, and the gradient of the raw readout is the noiseless one multiplied by $K = 0.716393 = (1-\gamma)^4$, component by component (Figure 20.2, right). At weight 1 the raw gradient is K times smaller; with many qubits the sector itself flattens the gradient (§112), with or without the filter. It does so according to its dimension: like a polynomial at weight 1, and like $1/\binom{n}{n/2}$ at half filling. This is the expressivity plateau, and Chapter 23 compares it with the other two (Section 23.8). The filter fixes the T1 noise plateau, not that one.

20.7 The ecosystem: Qiskit, qiskit-machine-learning and PennyLane

The network does not live in NumPy only. The notes take it to three environments:

- **Qiskit.** `qang.interop` reads the qg readings from a SamplerV2, and writes the filtered readout as a ratio of two diagonal observables, $\langle P_w Z_i \rangle / \langle P_w \rangle$, for an EstimatorV2 (§131). `qang.ibm.qnn_template` writes the weight-1 network as one parameterised circuit of native XX+YY gates and sends a whole test set in one job (§139).
- **qiskit-machine-learning.** `qang.qiskit_ml.qang_sampler_qnn` turns the network into a `SamplerQNN` (§140). The notes found that qiskit-machine-learning 0.9.1 reads the measured bits at the physical positions of the layout, so a circuit measured before transpilation reads every shot as 0; qang transpiles by itself and drops the layout.
- **PennyLane.** `qang.pennylane_gate.filtered_qg_z` acts on `qml.probs()` with `qml.math`, so the filtered readout is differentiable and can be the loss of a training loop (§131). Checking it for this chapter turned up a problem in PennyLane 0.45.1: on `default.mixed`, the `backprop` mode returns NaN for the jacobian of `qml.probs` through `SingleExcitation` (with `RY` it works, and on `default.qubit` too). The parameter-shift rule gives the right value.

`qml.math` also covers PennyLane's JAX interface, but JAX is not installed in the book's environment, and the filtered readout was not tested with it. Those two library problems and the path to JAX are proposed at the end of the book, together with the new version of qang.

20.8 The qg encoding and quantum signal processing

Encoding a feature $x \in [-1, 1]$ with $\theta = \arccos x$ is the encoding of Chapter 1: the state has $m_Z = 1000\,x$, and the encoding angle is its qang angle (§16). Loaded L times, the model is a Chebyshev polynomial of degree L ($\cos L\theta = T_L(x)$, checked to $7 \cdot 10^{-16}$), so L layers fit any bounded polynomial of degree L to machine precision and $L - 1$ layers do not. It is a polynomial inductive bias, not an advantage: with periodic targets angle encoding wins by 1.5 to 4.5 orders of magnitude.

The same rotation is the signal rotation of quantum signal processing (QSP), and `qang.qsp` builds it in qg terms (§134). As activations of a one-layer Kolmogorov–Arnold network, weighted QSP gets 179 of 189 inputs right, 182 with odd and even activations, against 185 for qang's network E (§135–§136): the limit is the parity of the activations, not the number of phases.

20.9 What remains open

The accuracies in this chapter are those of the notes, with predictions committed before running; they were not rerun here, because the book's environment has no scikit-learn. What was checked with code is structural: the filter's exactness, the law of the kept fraction, the readout ranks, the quadratic twin and the gradient in PennyLane. A run of a filtered network on real hardware is missing (Chapter 21 develops the filter and Chapter 32 the execution on devices), and so is a test of the filtered readout with JAX.

20.10 Without qang and with qang

Table 20.5: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read the network under T1	every reading; they drift towards the north pole (132–506 m-qg mean error)	the sector only: exact readout (10^{-12} m-qg)	**in favour**: exactness by construction, with equal T1

action	without qang	with qang	difference
know how many shots it costs	estimate the loss of each run	$K = (1-\gamma)^{kd}$, before running	**in favour**: the cost depends only on depth and the calibrated T1
train on a simulator and run under T1	loses 2.5 (weight 1) and 16.5 points (weight 2)	the noiseless accuracy	**in favour**, decisive at weight 2 and with more qubits (26 points at 8 qubits)
train under the calibrated noise	catches up with the filter	the same	**none** measurable (±1 point); the filter saves the noise model and the noisy runs
unequal T1 and dephasing	—	biases of 30–90 m-qg	**against** if ignored: the filter does not correct them; with dephasing, train under the noise
gradient under T1	K times the noiseless one	the noiseless one, exactly	**in favour**: the filtered readout is differentiable in PennyLane
quantum advantage	—	—	**none**: the network is a simulable quadratic classifier; classical models match it

Balance. qang's contribution in this chapter is not accuracy but robustness: the qg filter makes noiseless training remain valid under T1, exactly, and says in advance how many shots it costs. That gain belongs to the sector filter, and the base of Chapter 1 explains it: the shots that decay drift the readings towards the north pole. The networks are known (Landman et al. 2022), they are simulable, and with calibrated noise a noise-aware training reaches the same place without qang.

20.11 In code

```
import numpy as np
from qang import qml as QML

m = QML.WeightQNN(n_qubits=5, weight=2, encoding="dual", readout="zz")
rng = np.random.default_rng(0)
psi = m.encode(rng.uniform(-1, 1, size=(3, 4)))
theta = rng.uniform(-np.pi, np.pi, size=m.n_theta)
clean = m.qg_z(m.probs(theta, psi), qang=False)
noisy = m.probs(theta, psi, gamma=0.08)
np.abs(m.qg_z(noisy, qang=True) - clean).max()  # ~1e-15: the filter is exact under equa
QML.kept_fraction(0.08, m.weight, m.depth)      # 0.2229 = 0.92**18: the cost, before ru
```

20.12 Summary

- In the weight-1 sector the network is an orthogonal matrix and each qubit reads $m_Z = 1000(1-2u_q^2)$. Under equal T1 the qg filter returns the noiseless readout exactly, at a cost $K = (1-\gamma)^{kd}$ known before running; without it, the readings drift towards the north pole.
- Trained on a simulator, the network gains 2.5 points with qang at weight 1 and 16.5 at weight 2; trained under the calibrated noise, it gains nothing measurable. The filter does not correct unequal T1 or dephasing.
- Weight 2 reaches weight 1 with an encoding that does not distort and a full-rank readout (qg_{ZZ}).
- The network is a simulable quadratic classifier: robustness on noisy hardware, not quantum advantage.
- In PennyLane the filtered gradient under T1 is the noiseless one; the check turned up two library problems (PennyLane and qiskit-machine-learning) and a path to JAX, which are proposed at the end of the book.

Bibliographic notes

Orthogonal networks with RBS gates are Landman et al. (2022). Quantum kernels and feature spaces are Havlíček et al. (2019) and Schuld and Killoran (2019); the relation between absence of barren plateaus and classical simulability, Cerezo et al. (2023). Quantum signal processing is Low and Chuang

(2017) and Martyn et al. (2021). The ecosystem is Qiskit (Javadi-Abhari et al. 2024), PennyLane (Bergholm et al. 2018) and JAX (Bradbury et al. 2018). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

21 The qg symmetry filter and exactness under T1 relaxation

Chapter 20 used the qg filter as a tool: discarding the shots that left the Hamming-weight sector returns the noiseless readout under T1. This chapter treats it as an object. First the theorem and why it holds, on the scale of Chapter 1. Then what it costs, with a closed rule that predicts before running when it pays, and what it does not correct: unequal T1 across qubits, dephasing, and gate and readout errors. For each limit there is an exact property (F5 and F6 in the notes) that says how to go on. The structural facts are checked with code on random states; the classification results are those of the notes, with predictions committed before running.

21.1 From the base: decay pushes the readings to the north pole

Amplitude damping of one qubit has two Kraus operators:

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \qquad K_1 = \sqrt{\gamma}\,|0\rangle\langle 1|.$$

K_1 is the jump: the qubit falls to $|0\rangle$, and its reading goes to $+1000$ m-qg, the north pole. K_0 is the no-jump evolution: it leaves $|0\rangle$ alone and multiplies the amplitude of $|1\rangle$ by $\sqrt{1-\gamma}$. In a register, the product of the K_0 multiplies every basis state with k ones by $(1-\gamma)^{k/2}$, the same factor for every state of the weight-k sector. And any jump lowers the weight by one.

Hence the theorem (notes, §77 and §78 F4): **with the same T1 on every qubit, keeping the weight-k shots returns the noiseless state.** The no-jump factor is a scalar on the sector, and it disappears on normalising; everything that jumped is in another sector. It holds between layers of weight-conserving gates, for any strength and any depth. The fraction of shots that survives is $K = (1-\gamma)^{k\,d}$, with d the number of damping steps.

Table 21.1: The theorem in every sector of 6 qubits: 20 random states, 2 RBS layers (6 sublayers), $\gamma = 0.08$.

sector (6 qubits)	dimension	max $\lvert$filtered − noiseless$\rvert$ (m-qg)
weight 1	6	$4 \cdot 10^{-13}$
weight 2	15	$7 \cdot 10^{-13}$
weight 3	20	$6 \cdot 10^{-13}$
weight 4	15	$7 \cdot 10^{-13}$
weight 5	6	$9 \cdot 10^{-13}$

For a weight-2 state with three damping steps, the fraction left in the sector is 0.606355, exactly 0.92^6. The filter is a postselection, and what it returns is a conditional probability: the equation of Chapter 1 applied inside the sector.

21.2 What it costs: the mean-squared-error rule

The filter removes the bias and pays in shots: of S shots only KS remain. The notes (§94) wrote down before running a closed rule for the mean squared error of the qubits' readings m:

$$\mathrm{MSE}_{\mathrm{qang}} \approx \frac{1000^2 - m^2}{KS}, \qquad \mathrm{MSE}_{\mathrm{raw}} = \mathrm{bias}^2 + \frac{1000^2 - m_{\mathrm{raw}}^2}{S}.$$

The first is the shot noise of Chapter 1 with KS shots; the second adds the bias towards the north pole, which does not shrink with more shots. The chapter's script checks it with 10 random circuits of 6 qubits, weight 2, 12 sublayers and $\gamma = 0.02$ ($K = 0.616$), 200 repetitions per point:

Table 21.2: The rule of §94 against simulation (Figure 21.1, right).

shots S	30	100	300	1000	3000	10,000
MSE with qang, simulated / rule	1.05	1.01	1.02	1.00	1.01	1.00

shots S	30	100	300	1000	3000	10,000
MSE without qang / MSE with qang	1.3	3.0	7.9	25	74	249

The raw error stalls at about 31,000 m-qg² (a bias of about 175 m-qg), while the filtered one keeps falling as $1/S$. Setting the two expressions equal gives the break-even number of shots S^*: above it, the filter wins. For the circuits of this chapter S^* runs from 9 to 61 shots (Figure 21.2, left); on the notes' grid, from 3 to 110 (108 configurations up to 8 qubits, weight 4 and depth 24, with up to 1000 times less error with the filter).

With very few shots kept (§96, 405 configurations with K down to $2.6 \cdot 10^{-7}$) the generalised rule, which includes the bias left in the filtered readout, picks the winner in 342 of 342 cases with $KS \geq 5$, also with unequal T1 and dephasing. When less than one shot survives, the filter returns 0 for lack of data, and even so it beats the raw readout in 51 of 57 cases, because the raw one collapsed towards $|0 \ldots 0\rangle$, the north pole, which is a worse guess than the equator. **No regime was found where the shot cost decides against the filter.**

21.3 Training through the filter

Since the filtered readout under equal T1 is the noiseless one for every parameter, so are the loss and its gradients, and training with filtered noise is training without noise (§77; Chapter 20 checked it in PennyLane to 10^{-16}). The notes repeated it training from shots (§109), with 100 or 1000 fresh shots per evaluation and SPSA gradients:

Table 21.3: Network E under equal T1 ($\gamma = 0.08$, $K = 0.47$), trained from shots (notes §109).

training	with qang	without qang
exact noiseless probabilities	0.945	–
shots, $S = 1000$	0.939	0.929
shots, $S = 100$	0.930	0.915

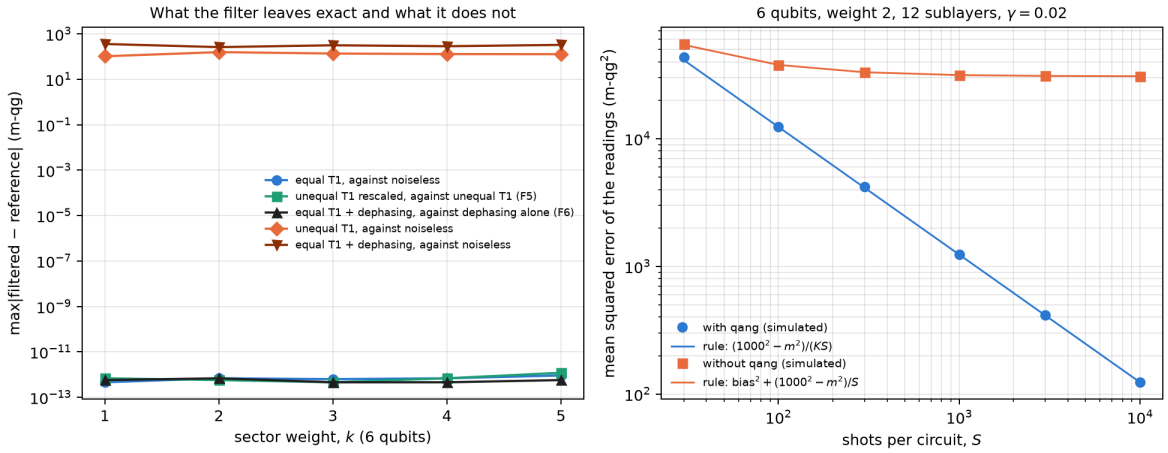


Figure 21.1: Left: maximum deviation of the filtered readout from its reference, by sector weight (6 qubits), on a log scale. With equal T1, with rescaled unequal T1 (F5) and with T1 plus dephasing against dephasing alone (F6) the deviation is rounding; with unequal T1 or with dephasing, against the noiseless readout, a bias of 100 to 360 m-qg remains. Right: mean squared error of the readings with and without the filter against shots, simulated (points) and from the rule of §94 (lines).

The prediction was that at 100 shots the filter's cost would show. It did not: the network trained with qang ended 1.5 points ahead (CI [−1.6, 4.6]). The gradient flattens with the size of the sector, with or without the filter (§112).

21.4 F5: unequal T1 depends only on the ratios

With a different T1 on each qubit, the no-jump factor is no longer a scalar on the sector: it is the product of the $\sqrt{1-\gamma_q}$ of the excited qubits, different for each state. The filter then leaves a bias: 100 to 160 m-qg with a ±50% spread of the γ_q (Figure 21.1). But the notes found an exact property (§117 F5): **a common factor in all the $(1-\gamma_q)$ cancels on normalising**, so the filtered readout depends only on the ratios $(1-\gamma_q)/(1-\gamma_r)$. The script checks it by multiplying all the $(1-\gamma_q)$ by 0.85: the filtered readout moves by no more than 10^{-12} m-qg in any sector, and the raw one moves by 640 to 850 m-qg.

The practical consequence is a recipe. Training with the filter under the calibrated rates needs only the relative T1s, tolerates calibration errors and survives a common drift of all the T1s:

Table 21.4: Unequal T1, spread 1 (γ_q from 0 to 0.16), notes §117. A: trained noiseless, with the filter; B: trained noiseless, raw; C: trained with the filter under the calibrated rates (10% error); C*: under the exact rates; D: trained raw under the calibrated rates. Drift: all T1s 1.5 times shorter.

weight	noiseless	A	B	C	C*	D	drift: A / B / C / D
1	0.950	0.931	0.873	0.951	0.945	0.952	0.913 / 0.770 / 0.942 / 0.950
2	0.945	0.918	0.797	0.940	0.941	0.941	0.903 / 0.663 / 0.932 / 0.841

With recipe C the network lands within 0.5 points of its noiseless accuracy. Under drift, at weight 2 the raw network trained under noise loses 10.0 points and the filtered one 0.7: the raw one learned the absolute rates, the filtered one only their

ratios. If one does not want to train under noise, the rule of §82 says how much is tolerated: with the filter alone the loss stays below one point as long as the qubits' $1/T_1$ differ by up to ±80% from their mean.

21.5 F6: dephasing, with the filter and extrapolation

Dephasing does not change the weight, so the filter does not see it. But the notes proved another exact property (§118 F6): **with equal T1 and dephasing, the filtered readout is the readout with dephasing alone.** The no-jump factor is a scalar on the sector and commutes with the dephasing channel, which is diagonal. The script checks it in every sector of 6 qubits, to $7 \cdot 10^{-13}$ m-qg (Figure 21.1). The filter turns two noises into one, and that one can be extrapolated to zero (ZNE):

Table 21.5: ZNE of the filtered readings over 40 random states, exact probabilities (Figure 21.2, right).

readout (5 qubits, weight 2, T1 0.08 + dephasing 0.03)	mean error (m-qg)
raw	442
filtered	152
filtered, linear extrapolation (scales 1 and 2)	69 (45%)
filtered, Richardson (scales 1, 2 and 3)	31 (20%)

In the notes' classifiers (§118, §122, at 5 and 8 qubits) linear extrapolation leaves 50–58% of the error of the filtered readings and Richardson 26–35%; without the filter, ZNE must extrapolate T1 and dephasing together and ends 3 to 10 points behind. With 1000 shots per scale, at weight 2, ZNE loses to the filter alone, because its shot cost eats the gain. What corrects dephasing best in a classifier is something else: training with the filter under the calibrated noise, which recovers the noiseless accuracy (0.951–0.963).

One subtlety (§132–§133): dephasing between RBS gates commutes with the weight and the filter does not see it, but dephasing **inside** the decomposition of an RBS into native gates turns a phase error into a change of weight. To predict the kept fraction on a device, the simulator has to place that error inside each RBS.

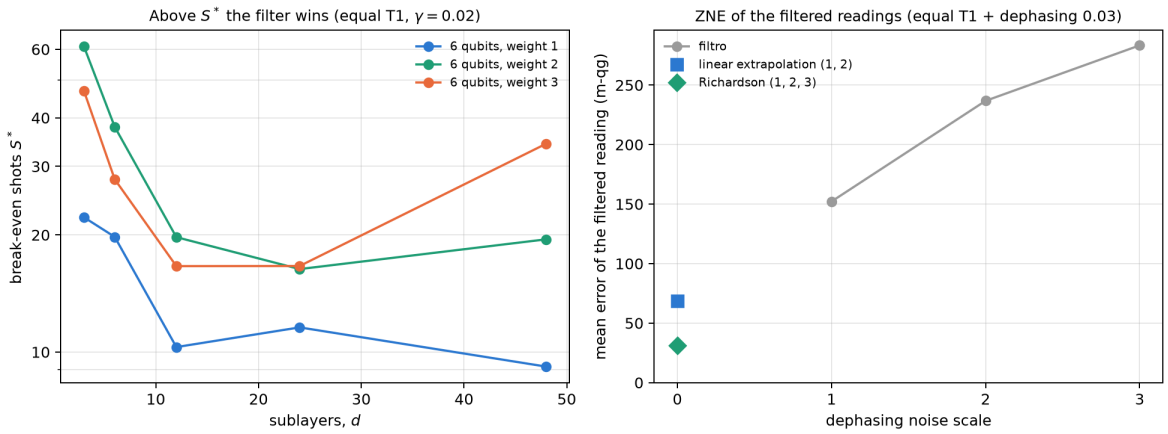


Figure 21.2: Left: break-even shots S^* from the rule of §94 against depth, for weights 1, 2 and 3 on 6 qubits ($\gamma = 0.02$); above S^* the filter has less error. Right: mean error of the filtered readout with equal T1 and dephasing 0.03 at dephasing scales 1, 2 and 3, and its linear and Richardson extrapolation to zero.

21.6 Gate and readout errors

In IBM device noise models gate errors dominate, and they are not T1. The notes added them to qang's simulator (§124: a two-qubit depolarising channel after each RBS, with the published error of the native gate) and then each qubit's asymmetric readout error (§129). Their interest here is one: predicting the kept fraction before running.

Table 21.6: Kept fraction of a weight-1 network on three IBM noise models (notes §124, §129).

simulated device	device K	predicted with T1, T2 and gates	+ readout error
fake brisbane	0.671	0.703	0.665
fake sherbrooke	0.697	0.774	0.694
fake torino	0.724	0.845	0.783

With readout error, the prediction comes within 0.006 of the device on the two devices with ECR gates. On torino (CZ gates) a gap of 0.059 remains, which §132–§133 attribute to dephasing inside the gates. In accuracy, the simplest recipe was the best on those models: train noiseless and run with the filter (182–185 of

189 correct, against 185 noiseless); training with the gate errors helped the raw readout (+1.2 points) but not the filtered one.

21.7 What remains open

The filter's properties (the theorem, the law of K, F5, F6 and the error rule) are checked with code; the classification accuracies are those of the notes. A run on real hardware is missing, where unequal T1, dephasing inside the gates and readout error act together. And a tool is missing from the libraries: weight postselection with its cost rule is proposed at the end of the book, together with the new version of qang.

21.8 Without qang and with qang

Table 21.7: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read under equal T1	bias towards the north pole, which does not shrink with shots	exact readout ($\leq 10^{-12}$ m-qg) in every sector	**in favour**: exactness by a theorem
decide whether filtering pays	try both readouts	closed rule; the winner is predicted in 99–100% of cases	**in favour**: the decision is taken before running
shot cost	—	exact $K = (1-\gamma)^{kd}$; break-even at 3–110 shots	**in favour**: the cost never decided against it on the notes' grids
unequal T1	the raw network learns the absolute rates; under drift it loses 10 points (weight 2)	depends only on the ratios (F5); training with the filter lands within 0.5 points; drift: 0.7	**in favour** with the recipe; **against** if the filter is used alone with a large spread (up to 2 points)

action	without qang	with qang	difference
dephasing	ZNE of T1 and dephasing together	the filter leaves dephasing only (F6); ZNE takes it to 20–58%	**in favour** in readout error; with few shots ZNE does not pay
gate and readout errors	—	the simulator predicts K within 0.006 (ECR)	**in favour** as a prediction; the filter does not correct them

Balance. The filter is symmetry postselection, a known mitigation technique. What the notes add are the exact properties (the theorem in every sector, F5 and F6), the closed cost rule, and the recipe for each noise the filter does not correct. The base of Chapter 1 gives the picture: the jump takes the readings to the north pole, and the filter returns the conditional probability on the sector. When the noise is not equal T1, the filter alone is no longer enough, and the answer is to train with it under the calibrated noise.

21.9 In code

```
import numpy as np
from qang import qml as QML

m = QML.WeightQNN(n_qubits=6, weight=3, layers=2)
rng = np.random.default_rng(0)
psi = np.zeros((1, m.dim))
psi[0, m.idx[3]] = rng.normal(size=20)
psi /= np.linalg.norm(psi)
theta = rng.uniform(-np.pi, np.pi, size=m.n_theta)
g = 0.08 * (1 + 0.5 * np.linspace(-1, 1, 6))          # unequal T1
f = lambda gam: m.qg_z(m.probs(theta, psi, gamma=gam), qang=True)
np.abs(f(1 - 0.85 * (1 - g)) - f(g)).max()        # ~1e-15: F5, only the ratios count
```

21.10 Summary

- The jump takes the readings to the north pole and lowers the weight; the no-jump evolution multiplies every weight-k state by $(1-\gamma)^{k/2}$. With equal

T1, the filter returns the noiseless state in every sector, with $K = (1-\gamma)^{kd}$.

- The rule MSE $\approx (1000^2 - m^2)/(KS)$ against bias$^2 + (1000^2 - m^2)/S$ predicts before running when filtering pays; the break-even is at a few tens of shots, and the cost never decided against it.
- With unequal T1 the filtered readout depends only on the ratios of the rates (F5): training with the filter under the calibrated rates recovers the accuracy and withstands a common drift.
- With T1 plus dephasing the filter leaves dephasing only (F6), which is extrapolated to zero or, in a classifier, corrected by training with the filter under the noise.
- With gate and readout errors, qang's simulator predicts a device's kept fraction; the filter does not correct them.

Bibliographic notes

The qg filter is a symmetry verification (Bonet-Monroig et al. 2018; McArdle et al. 2019); zero-noise extrapolation is Temme et al. (2017). The relaxation and dephasing channels are those of Nielsen and Chuang (2010). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

22 Echo calibration and the noise bias (pull effect)

The qg filter of Chapter 21 makes the readout exact under equal T1, but it does not correct everything. Two things remain, and this chapter deals with them. The first are errors that conserve the weight and move the excitation from one qubit to another: they pass the filter, and must be measured separately, with echo circuits. The second is what decayed shots do to a classifier without the filter: they do not add random noise, they **pull** every decision towards the same class. Both are read on the scale of Chapter 1, and for both the chapter checks the mechanism with code and quotes the notes' results on device noise models, with predictions committed before running.

22.1 From the base: two errors the filter treats differently

In the weight-1 sector, each qubit reads $m_Z(q) = 1000(1 - 2p_q)$, with p_q the probability that the excitation is on it. There are two ways to go wrong:

- **A T1 jump** takes the shot out of the sector: the excitation disappears and every qubit reads $+1000$, the north pole. The filter discards those shots. Without the filter, the raw readout is a mixture: with probability K the noiseless readout, and with probability $1 - K$ the north pole on every qubit.
- **An in-sector error** moves the excitation from one qubit to another: the shot still has weight 1, passes the filter and changes the readings of two qubits at once. On the probabilities p_q it acts as a stochastic matrix M.

The first is corrected by discarding; the second, by inverting M, as a readout confusion matrix is inverted.

22.2 In-sector errors and the echo

An echo circuit prepares a basis state of the sector, $|e_j\rangle$, applies the network's gates and then their inverses. Without errors it returns to $|e_j\rangle$; with errors, its filtered distribution over the sector states is column j of a transfer matrix M (notes, §105). Since the real circuit has half the echo's gates, the correction uses $M^{1/2}$: the filtered distribution f is unmixed by solving $M^{1/2}x = f$ with $x \geq 0$ (`qang.sectors.unmix_sector`).

The chapter's script tests it with a controlled in-sector error: after every sublayer, each pair of qubits of an RBS is swapped with probability p (it keeps the weight and moves the excitation). Five qubits, weight 1, 9 sublayers, 60 random inputs:

Table 22.1: Mean error of the m_Z readings (m-qg) with exact probabilities.

p	echo diagonal	error with the filter	+ echo with $M^{1/2}$	+ echo with M
0.005	0.93–0.94	15.0	5.7 (−62%)	13.4 (−11%)
0.01	0.87–0.89	29.4	11.3 (−61%)	25.5 (−13%)
0.02	0.76–0.80	56.4	22.4 (−60%)	47.4 (−16%)
0.04	0.58–0.65	104.2	43.7 (−58%)	87.4 (−16%)
0.08	0.38–0.45	178.8	84.5 (−53%)	164.3 (−8%)

$M^{1/2}$ removes a bit more than half of the error the filter leaves; the full M overcorrects and removes little. On the notes' device noise models, the echo removes 38–47% of the filtered error at weight 1 with 5 qubits (§105, five simulated devices), 37–43% at weight 2 with a 10×10 matrix (§107) and 27–38% at weight 2 with 8 qubits, a 28×28 matrix (§115). The larger the sector, the dirtier and harder to recover: the echo diagonal drops from 0.77–0.92 to 0.50–0.80. Filter plus echo ends 3.4 to 7 times below the raw readout.

22.3 Does it reach the classifier?

With the weight-1 network on 189 test inputs and 1000 shots (§106), on three simulated IBM devices and the IonQ simulator:

Table 22.2: Echo calibration in the classifier (notes §106; noiseless accuracy 0.979).

device	accuracy (raw / qang / qang + echo)	decisions differing from noiseless	decision-value error
fake brisbane	0.968 / 0.974 / 0.979	2 / 1 / 0	1.53 / 0.76 / 0.33
fake sherbrooke	0.963 / 0.974 / 0.979	3 / 1 / 0	1.26 / 0.61 / 0.29
fake torino	0.963 / 0.979 / 0.979	3 / 0 / 0	1.08 / 0.47 / 0.27
IonQ aria-1	0.958 / 0.963 / 0.968	4 / 3 / 2	1.67 / 0.87 / 0.44
IonQ forte-1	0.963 / 0.968 / 0.974	3 / 2 / 1	1.60 / 0.80 / 0.43

The echo cuts the decision-value error left by the filter by 43 to 56%, and filter plus echo ends 3.7 to 4.6 times below the raw readout. Decisions move only towards the noiseless network: 15 differ without qang, 7 with the filter and 3 with filter plus echo, over 945 readings.

When it does not help (§108). With narrow-margin inputs, the echo added nothing (19 flipped decisions against 17 with the filter alone). The filtered decision-value error was 0.26 to 0.35, barely above its shot noise (0.22): little bias was left, and inverting the matrix adds variance. The notes' rule is to use the echo when the filtered error is well above its shot-noise level. In this chapter's model, with 1000 shots and 4000 per echo, the echo kept helping even at $p = 0.005$ (22 → 17 m-qg); at $p = 0$ it changes nothing (15 m-qg, pure shot noise).

22.4 The pull: why the loss without qang varies between seeds

In the multiclass network of Chapter 20, the loss without the filter varied a lot from one seed to another (from −0.7 to 41.7 points). The notes explained it with an exact property (§121 F7). At weight 1 under equal T1, every decayed shot lands in $|0 \dots 0\rangle$, where each qubit reads $+1000$. Since the logits are linear in the readings, the raw logits are

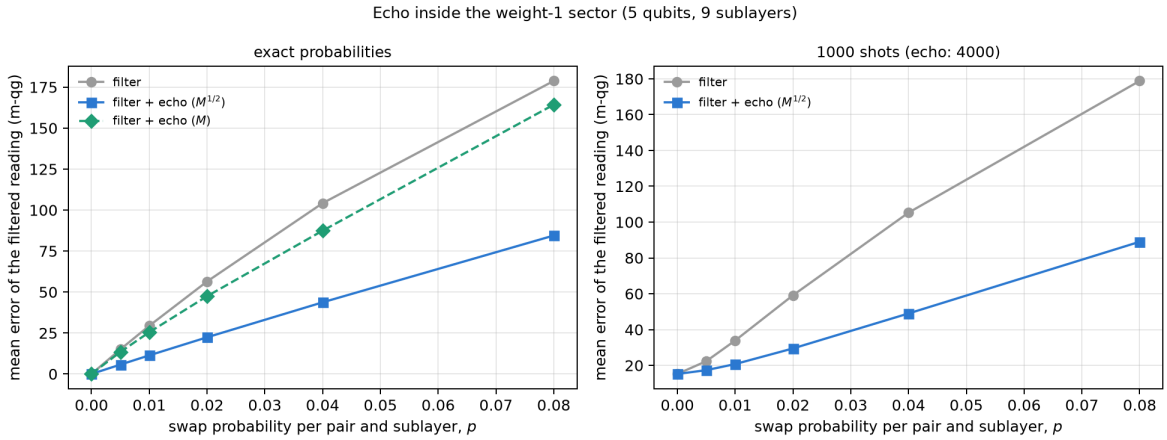


Figure 22.1: Mean error of the filtered readings in a weight-1 sector with in-sector swaps, against the swap probability. Left: exact probabilities; the echo with $M^{1/2}$ removes more than half of the error and the full M over-corrects. Right: with 1000 shots per circuit and 4000 per echo; at $p = 0$ shot noise remains.

$$\text{logits}_{\text{raw}} = K\,L + (1 - K)\,v,$$

with L the noiseless logits and v the logits of the north pole, a constant vector fixed by training (for the linear head, $v = W^{\top}\mathbf{1} + b$). Without the filter, every input is pulled towards the class with the largest v_c, and how much that weighs is measured by a pull index: $(1 - K)/K$ times the range of v, over the median training margin.

The script checks it on 40 five-class models with random parameters ($K = 0.472$, 300 inputs each):

- the formula for the raw logits holds to $7 \cdot 10^{-15}$;
- without the filter, 0 to 60% of the decisions flip depending on the model, and the flips grow with the pull index (Spearman 0.60; in the notes, 0.65 with the head and 0.98 with the qubit readout);
- with the filter, none flips.

The spread between seeds is not bad luck: it is the v that training left, and it can be read before running anything.

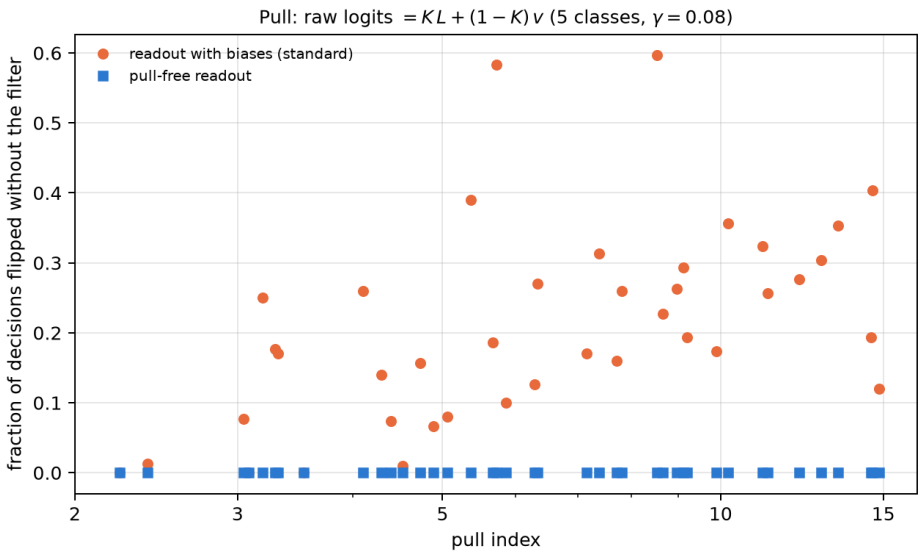


Figure 22.2: Fraction of decisions that flip without the filter against the pull index, for 40 five-class models under equal T1. With biases (standard readout), the flips grow with the pull; the pull-free readout flips none.

22.5 The pull-free readout

If the pull comes from v, it can be removed by design: logits linear in the excitation probabilities with no constant term (for the head, $W^\top(qg_Z - 1)$; for the qubit readout, $a\,p_c$). Then the north pole contributes zero, decayed shots only rescale the logits by K, and the raw decision is the noiseless one. The script checks it on the 40 models: no decision flips, and the raw logits are K times the noiseless ones to $6 \cdot 10^{-15}$. The notes (§123, 10 seeds, 5-class digits) confirm it in accuracy:

Table 22.3: The pull-free readout (notes §123).

readout	model	noiseless	T1 (qang / raw)	unequal T1 (qang / raw)	1000 shots (qang / raw)
qubit	standard	0.860	0.860 / 0.725	0.849 / 0.688	0.861 / 0.724
qubit	pull-free	0.862	0.862 / 0.862	0.851 / 0.851	0.863 / 0.863
head	standard	0.932	0.932 / 0.653	0.932 / 0.624	0.927 / 0.651

readout	model	noiseless	T1 (qang / raw)	unequal T1 (qang / raw)	1000 shots (qang / raw)
head	pull-free	0.938	0.938 / 0.938	0.928 / 0.928	0.932 / 0.932

Dropping the biases costs no accuracy, and the pull-free readout is immune to decay without filtering: it carries the filter inside it. Without the filter, the pull-free head is 28.5 points above the standard one. It is the picture of Chapter 1 carried into design: if the north pole is worth zero, falling into it does not pull.

22.6 Calibrating with the device itself

The notes tried two more ways of using calibration on IBM noise models (§119, §120): training with qang under the published T1 and T2, and measuring the device itself (with the echo matrix inside the model, or refitting the readout on 40 inputs run on the device). None beat the simple recipe: train noiseless and run with the filter (0.972 mean accuracy against 0.979 noiseless), plus the echo for the decision values. The reason is that gate errors dominate on those models, and a T1/T2 calibration does not describe them (the per-sublayer γ are only 0.002–0.017).

22.7 What remains open

The echo mechanism and property F7 are checked with code, and the results on simulated devices are those of the notes. A package is ready to run the multiclass network with the echo on an IonQ QPU (§116: 55 circuits, 60,000 shots, with predictions committed before sending anything); its rehearsal on the simulator gave 4 flipped decisions with the filter and 1 with the echo, against 8 raw. The hardware run is pending credits and an explicit approval of the cost.

22.8 Without qang and with qang

Table 22.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
in-sector errors	invisible: they pass any weight filter	echo matrix and unmixing with $M^{1/2}$: −27 to −62% of the filtered error	**in favour**, at a cost: $\binom{n}{k}$ extra echo circuits
decision-value error	—	−43 to −56% with the echo over the filter (§106)	**in favour** when the filtered error exceeds its shot noise; **none** with narrow margins (§108)
understand the spread between seeds	bad luck	raw logits $= KL + (1-K)v$, exact; the pull index predicts it	**in favour**: the pull is read before running
correct the pull	train under the noise	the filter removes it exactly	**in favour**
design an immune readout	—	no biases: the north pole is worth zero, identical decisions with or without the filter	**in favour**, at no accuracy cost
calibrate with the device	train with the calibration	equal to or worse than noiseless training + filter	**none** on those models: gate errors dominate

Balance. Echo calibration is known mitigation, of the kind that inverts readout confusion matrices, applied inside a sector. What the notes add is its measured reach (how much it removes depending on the sector size) and the rule for when it does not help. The cleanest contribution of the chapter is the pull: an exact property that explains the spread between seeds, and a bias-free readout that avoids it without filtering. Both come from reading decay as Chapter 1 reads it: every qubit at the north pole.

22.9 In code

```
import numpy as np
from qang import qml as QML

m = QML.MultiClassQNN(n_qubits=5, n_classes=5, readout="head")
rng = np.random.default_rng(0)
psi = m.encode(rng.uniform(-1, 1, size=(4, 4)))
th = rng.uniform(-np.pi, np.pi, size=m.n_theta)
head = rng.normal(size=m.n_head)
K = QML.kept_fraction(0.08, 1, m.depth)
L = m._logits(m.qg_z(m.probs(th, psi), qang=False), head)
raw = m._logits(m.qg_z(m.probs(th, psi, gamma=0.08), qang=False), head)
v = m._logits(np.ones((1, 5)), head)     # the north pole: +1000 m-qg on every qubit
np.abs(raw - (K * L + (1 - K) * v)).max()  # ~1e-15: F7
```

22.10 Summary

- A T1 jump takes every qubit to the north pole and the shot out of the sector; an in-sector error moves the excitation and passes the filter. The first is discarded; the second is measured with echoes and inverted.
- The echo with $M^{1/2}$ removes more than half of the filtered error in the chapter's model and 27–47% on the notes' simulated devices; the full M overcorrects. It helps when the filtered error exceeds its shot noise.
- Without the filter, the raw logits are exactly $K\,L + (1 - K)\,v$: decay pulls every decision towards one class, and the pull index predicts how much.
- A bias-free readout makes the north pole worth zero: the raw decisions are the noiseless ones, without filtering and without losing accuracy.
- On IBM noise models, training noiseless and running with the filter (and the echo) was the best recipe; calibrating with the device did not improve it.

Bibliographic notes

The echo is that of Hahn (1950); read-out error mitigation, Bravyi et al. (2021). The sector filter is the symmetry verification of Bonet-Monroig et al. (2018). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

23 Information geometry, the natural gradient (QNG) and kernels

Chapters 20 to 22 read in qg what comes out of a circuit. This chapter asks about what goes in: the parameters an optimizer moves. An attractive idea circulates: if qg is the natural readout, writing the circuit in qg coordinates should speed up the gradient and flatten the landscape. The chapter shows with code that it is the other way round, and why. The coordinate to optimize in is the qang angle θ, and qg is the coordinate to read out in. Four honest uses of qang follow from that distinction: regularizing the pole singularity, reading the natural-gradient metric from qg correlators, damping the step near the poles, and writing a local cost that flattens more slowly. The structural results are checked with code; the kernel results on device noise models are the notes', with predictions committed before running.

23.1 From the base: the angle and the readout

In Chapter 1, a qubit prepared with $R_y(\theta)|0\rangle$ reads $m_Z = 1000\cos\theta$ m-qg, and θ is the qang angle. Both coordinates describe the same state, but they are not equivalent for optimizing, because the change between them is not linear:

$$\frac{d\,qg_Z}{d\theta} = -\sin\theta, \qquad \frac{dE}{d\,qg_Z} = -\frac{1}{\sin\theta}\,\frac{dE}{d\theta}.$$

At the poles ($m_Z = \pm 1000$) the readout barely moves when θ changes, and so the gradient in qg diverges. The whole chapter follows from that line.

23.2 The singularity and its regularizations

The factor $-1/\sin\theta$ is -10 at $\theta = 0.1$, -100 at 0.01 and -1000 at 0.001. The notes (§1) propose two bounded versions, in `qang.gradients`:

- **clipping:** $-1/(\text{sign}(\sin\theta)\max(|\sin\theta|,\varepsilon))$, never larger than $1/\varepsilon$;
- **Tikhonov:** $-\sin\theta/(\sin^2\theta+\varepsilon^2)$, smooth everywhere and zero at the pole itself.

With $\varepsilon = 0.05$, at $\theta = 0.001$ the raw factor is 1000, the clipped one 20 and the Tikhonov one 0.40. The chapter's script repeats the notes' experiment: minimize $E = -\sin\theta$ (minimum at $\pi/2$, away from the edges) starting at the pole, $\theta_0 = 0.01$, for 150 iterations.

Table 23.1: Descent from a pole (notes §1; the script's figures match the notes').

space	rate	final θ	final E	largest step	behaviour
in θ	0.05	1.5698	−1.0000	0.05	converges to $\pi/2$
in qg, raw	0.05	3.1416	0.0000	3.14	jumps from one pole to the other and never stops
in qg, clipped	0.01	1.4184	−0.9884	0.63	converges, more slowly
in qg, Tikhonov	0.01	1.4043	−0.9862	0.27	converges, more slowly, smaller steps

With the same rate as in θ, raw descent in qg never recovers: each huge step crosses to the opposite pole, where the gradient is huge again. Both regularizations fix it, but they need a rate about five times smaller (the notes' rule of thumb is rate $\leq \varepsilon$) and still end worse than descent in θ. Regularizing removes the divergence; it does not make qg a good coordinate to optimize in.

23.3 The angle is the Fisher-flat coordinate

The underlying reason is statistical. Measuring in Z, the Fisher information per shot is (notes §15.1, §159)

$$F_\theta = 1, \qquad F_{qg} = \frac{1}{1 - qg_Z^2} = \frac{1}{\sin^2\theta}.$$

In the angle it is 1 everywhere; in qg it depends on where one is and diverges at the poles. It is the same $1/\sin\theta$ as in the previous section, now as the Fubini-Study metric written in qg coordinates. In Chapter 1 units, the standard deviation of the readout with N shots is $\sigma_m = 1000\sin\theta/\sqrt{N}$ m-qg (the sample mean attains the Cramér-Rao bound), and that of the estimated angle is $1/\sqrt{N}$ radians, the same all around the circle. The angular qga scale of Chapter 1 (Section 1.4) is also a flat coordinate, with constant $F = 4\pi^2 \cdot 10^{-6}$ qga^{-2}: measuring the parameter in qga instead of radians multiplies the gradient by a constant and, as shown below (§12), does not change the plateau.

The script checks it with 4000 repetitions of 1000 shots at 12 angles: $\sigma_{qg}\sqrt{N}$ reproduces $\sin\theta$ to the third digit (0.149 against 0.149 at $\theta = 0.15$; 0.995 against 0.992 at 1.70), and $\sigma_\theta\sqrt{N}$ stays between 0.98 and 1.05 at every point (Figure 23.1, left). The angle is the flat coordinate: in it, each optimizer step is worth the same anywhere.

A coordinate does not save evaluations. For a gate $e^{-i\theta G}$ with $G^2 = I/4$, the energy is $E(\theta) = a + b\cos\theta + c\sin\theta$. A single shifted evaluation is not enough for the derivative: the script builds two functions of that form that agree at $\theta + \pi/2$ (both -0.644218) and have different derivatives at θ (-0.644 and $+0.765$). Two evaluations suffice, so two per parameter is the minimum in any coordinate. In qg the same two numbers are used and the pole is added.

23.4 The natural gradient and the qg_YY metric

The natural gradient (QNG) (Stokes et al. 2020) corrects the step with the metric: $\Delta\theta = -\eta\, g^{-1}\nabla E$. For one qubit, the natural step in qg is exactly plain descent in θ (notes §15.1): invariance under the change of coordinates means that writing it in qg adds nothing.

With several qubits there is something qang does contribute (§159). In a real ansatz of R_y and CZ, every state along the circuit is real, so $\langle Y_k\rangle = 0$ before every R_y and **every diagonal entry of the metric is exactly 1/4**. The angles are flat one by one, and all the curvature sits off the diagonal. Within a layer those terms are $g_{ij} = \langle Y_i Y_j\rangle/4$: qg_YY correlators read from one Y-basis circuit per layer. It is the block-diagonal metric of Stokes et al. written as qg readouts. The script checks it on the 4-qubit TFIM with 3 layers: the diagonal is 1/4 to

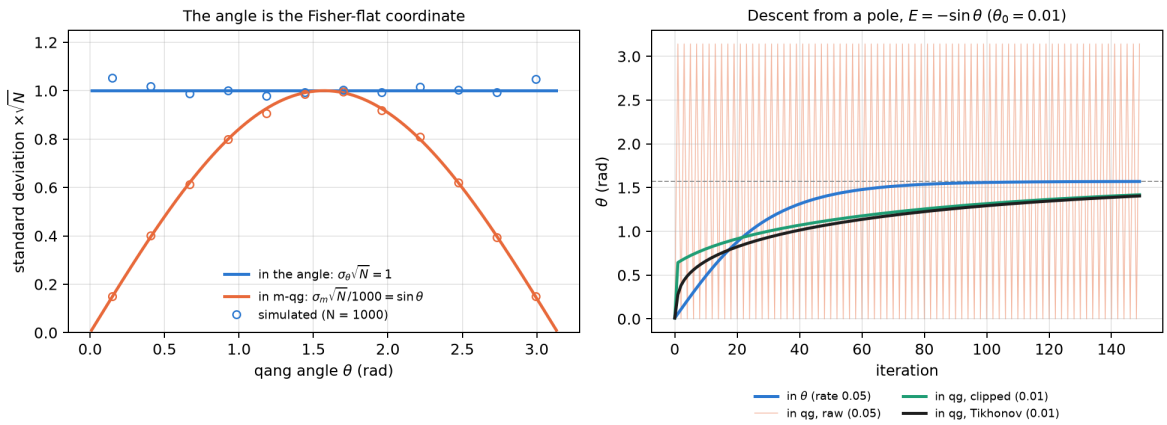


Figure 23.1: Left: standard deviation times $\sqrt{N}$ of the readout in m-qg (divided by 1000) and of the estimated angle, with 1000 shots; the readout follows $\sin\theta$ and the angle stays at 1. Right: descent from the pole on $E = -\sin\theta$; in raw qg the optimizer jumps between the poles, and the regularized versions converge more slowly than in θ.

$1.7 \cdot 10^{-16}$, the qg_YY blocks match the full metric within each layer to $1.7 \cdot 10^{-16}$, and off the diagonal there are entries up to 0.248, which is where the metric acts.

The notes' experiment, with predictions committed before running: 4-qubit TFIM ($h = 1$), 12 parameters, success if $E - E_0 \leq 0.03$ within 300 steps (the ansatz floor is 0.0245), rate chosen on 6 seeds and 20 test seeds. The script re-runs the 20 test seeds with the stored rates and gets the same figures:

Table 23.2: Natural gradient on the 4-qubit TFIM (notes §159, reproduced by the script).

method	steps (median)	successes	rate
descent in θ	103	19/20	0.2
descent in qg (clipped Jacobian)	never	0/20	0.005
QNG, qg_YY blocks	79.5	19/20	0.05
QNG, full metric	22.5	20/20	0.1

- **Descent in qg fails** at every rate of the grid. One optimizes in the angle and reads out in qg.
- **The qg_YY metric helps a little**: 1.3 times fewer steps than descent in θ. The committed prediction asked for 1.5 and failed. The full metric is 4.6 times faster.

- Counting circuits (24 per gradient, 3 more for the block metric and about 78 more for the full one on hardware), the three methods end within 15% of each other. This count was made after the run.
- With 1000 shots, each entry of the qg_YY metric comes out with an error of 0.013 to 0.022.

23.5 Pole-damped descent

If the problem is at the poles, one can brake there without leaving the angle (notes §9):

$$\theta \leftarrow \theta - \eta \, \max(|\sin\theta|, \varepsilon) \, \frac{dE}{d\theta}.$$

The factor is $|d\,qg_Z/d\theta|$, the same geometry as in the first section, and it needs no tuning. It is the natural step multiplied by $|\sin\theta|$, which for these real states is also the coherence $2|\rho_{01}|$: the optimizer slows down the less coherence the qubit has left. **It is not a new class of optimizer**: it is a position-dependent step, of the trust-region kind.

The script repeats the robustness test: 300 random landscapes $E = h_z \cos\theta + h_x \sin\theta$ per rate, starting near a pole.

Table 23.3: Success fraction starting near a pole (300 landscapes per rate, seed 0).

rate	descent in θ	damped	damped, trapped
0.3	1.000	1.000	0.000
1	0.773	1.000	0.000
2	0.193	0.440	0.013
3	0.087	0.327	0.017
5	0.043	0.203	0.030
10	0.007	0.093	0.077

At safe rates nothing changes; at rate 1 it always succeeds against 77%; from rate 2 on it wins by 2.3 to 13 times, although at large rates both fail almost always. The fraction trapped oscillating at its pole grows with the rate, up to 7.7% at rate 10. The figures differ somewhat from the notes' (§9.2): there, at rate 0.3, the damped version lost a few points (0.963 against 0.997), and the gains from rate

2 were 2.5 to 5 times. The pattern is the same; the exact figures depend on the landscapes drawn.

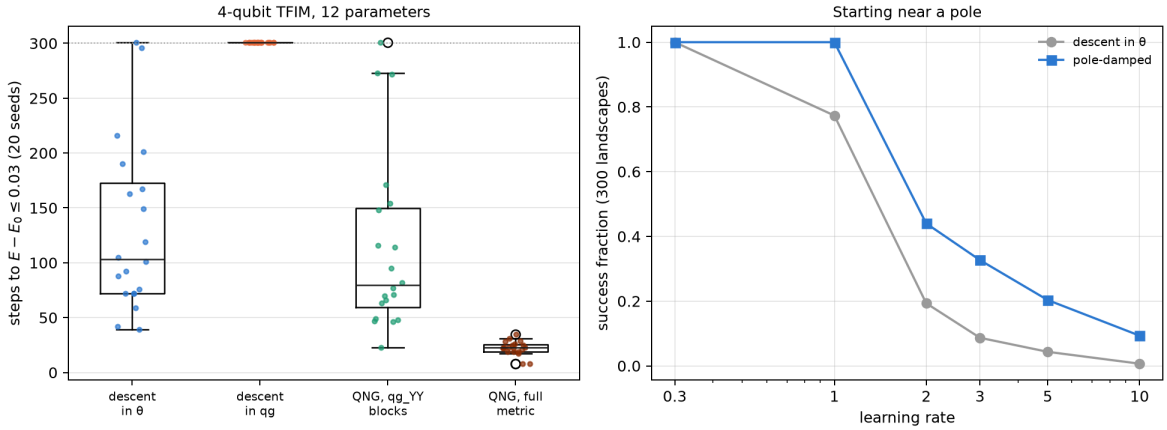


Figure 23.2: Left: steps to $E - E_0 \le 0.03$ on the 4-qubit TFIM for 20 seeds (300 = did not get there); descent in qg never gets there and the full metric is the fastest. Right: success fraction starting near a pole, with and without damping.

23.6 Quantum kernels and the qg filter

A fidelity kernel $K_{ij} = |\langle\psi(x_i)|\psi(x_j)\rangle|^2$ estimated on a noisy device comes out shrunk ($K_{ii} < 1$) and pulled towards the noise value. The usual repair, $K_{ij}/\sqrt{K_{ii}K_{jj}}$, costs the diagonal circuits. If the kernel circuit conserves the Hamming weight, the qg filter of Chapter 21 is another repair (notes §161). With the weight-1 encoding (one qubit per component of the normalized vector), the kernel is $K = (x \cdot y)^2$, and in qg it is a qang angle: $\arccos(2K - 1) = 2\varphi$, with φ the angle between the two vectors. The script checks it to $1.3 \cdot 10^{-15}$.

The notes' experiment: breast-cancer data reduced to 4 components plus a constant (5 qubits), 40 + 40 points, SVM with a precomputed kernel, 2000 shots and three IBM device noise models.

Table 23.4: The weight-1 kernel with the qg filter (notes §161).

backend	Gram-matrix error (raw / normalized / filtered / both)	correct (exact 35/40)	diagonal (raw / filtered)
FakeBrisbane	0.097 / 0.016 / 0.016 / 0.015	32 / 34 / 34 / 35	0.674 / 0.944
FakeTorino	0.101 / 0.021 / 0.017 / 0.014	34 / 35 / 34 / 34	0.717 / 0.971
FakeFez	0.039 / 0.027 / 0.011 / 0.011	34 / 34 / 35 / 35	0.953 / 1.000

- **The filter repairs the Gram matrix best**: 2.4 to 6 times below raw, and 1.0 to 2.5 times below the normalization, without the 80 diagonal circuits. Normalizing after the filter adds almost nothing.
- **Accuracy barely notices**: an SVM compares kernel values, and a common shrink cancels. The committed prediction on accuracy failed on one raw kernel, by three points.
- Kernel-target alignment does not see the shrink, because it divides by $\|K\|$.
- This kernel is a degree-2 polynomial, easy to compute classically: it is a test of repair, not of advantage.

23.7 Barren plateaus: global cost against the qg local cost

With many qubits, the gradient of a random circuit becomes exponentially small (McClean et al. 2018). The notes looked at it in qg twice. In §12, converting the gradient to qg changes nothing: the plateau stays exponential, regularization only rescales it, and at the poles the singularity is added on top. **The two are independent and compound.**

In §23 the coordinate does matter, by another route: the cost. An ansatz is trained to reach a product state, and two costs that vanish only at the target are compared: the global one, $C_G = 1 - P(0 \ldots 0)$, and the local one of Cerezo et al. (2021), which on the Chapter 1 scale is the register's mean readout, $C_L = (1 - \overline{qg_Z})/2$. The script recomputes the variance of the derivative with respect to the first parameter over 200 random initialisations:

Table 23.5: Variance of $\partial C/\partial p_0$ (200 initialisations).

	$n = 2$	$n = 4$	$n = 6$	$n = 8$	$n = 10$
global, $L = 2$	$2.7 \cdot 10^{-2}$	$2.5 \cdot 10^{-3}$	$2.5 \cdot 10^{-4}$	$2.6 \cdot 10^{-5}$	$1.2 \cdot 10^{-6}$
qg local, $L = 2$	$1.8 \cdot 10^{-2}$	$3.3 \cdot 10^{-3}$	$1.7 \cdot 10^{-3}$	$9.4 \cdot 10^{-4}$	$5.6 \cdot 10^{-4}$
global, $L = 4n$	$2.4 \cdot 10^{-2}$	$1.9 \cdot 10^{-3}$	$1.2 \cdot 10^{-4}$	$6.7 \cdot 10^{-6}$	$5.7 \cdot 10^{-7}$
qg local, $L = 4n$	$1.6 \cdot 10^{-2}$	$1.9 \cdot 10^{-3}$	$3.4 \cdot 10^{-4}$	$7.3 \cdot 10^{-5}$	$1.7 \cdot 10^{-5}$

- With short circuits ($L = 2$), the qg local cost decays like a polynomial and the global one like an exponential: at $n = 10$ the local one has 470 times more variance (in the notes, 5600 times at $n = 12$), that is, about 470 times fewer shots for the same signal-to-noise ratio.
- With deep circuits ($L = 4n$) **the local cost does not cure the plateau**: both decay exponentially.
- With finite shots (notes §23, 100 shots per circuit), the global gradient comes out exactly zero in 22% of the initialisations at $n = 10$ and 35% at $n = 12$; the local one, never. Training at $n = 10$, the local cost reaches fidelity 0.5 in a median of 11 steps against 20 for the global one.
- **The classical control.** With depth L, each $\langle Z_i \rangle$ depends only on the $2L+1$ qubits in its light cone, so the local cost and its gradient are computed classically in time $O(n\, 2^{2L+1})$. The notes train 50 and 100 qubits that way, without shots. Where the qg local cost escapes the plateau it is also simulable: the advantage is in training, not quantum.

23.8 Three plateaus and what qang does with each

A barren plateau can come from three different sources. They are worth separating, because qang answers each one differently:

- **Expressivity.** A circuit that explores the whole 2^n-dimensional space has gradients that fall like the inverse of that dimension (McClean et al. 2018). Current theory states it in terms of the Lie algebra the gates generate (Ragone et al. 2024): the variance follows the dimension of the space the circuit actually explores.
- **Global cost.** A cost that compares all qubits at once, such as the fidelity, flattens even for a short circuit (Cerezo et al. 2021). This is Section 23.7.

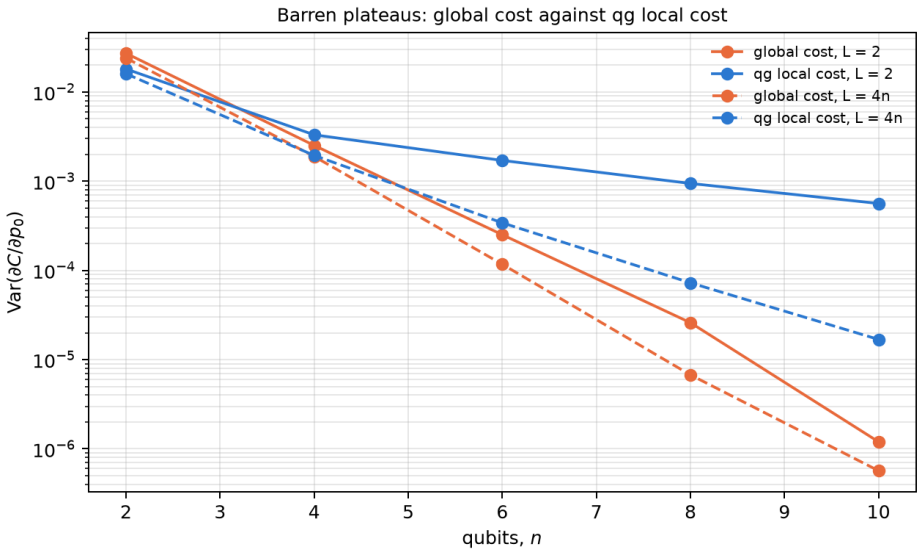


Figure 23.3: Gradient variance against the number of qubits for the global cost and the qg local cost, with short ($L = 2$) and deep ($L = 4n$) circuits. Only the local cost with short circuits escapes exponential decay.

- **Noise.** Noise accumulated with depth flattens the landscape even without either of the other two causes (Wang et al. 2021).

The unit does nothing. Measuring the parameters in qga or the readout in m-qg multiplies every derivative by a constant and changes none of the three (§12; Chapter 1).

Global cost → the qg local cost. This is Section 23.7: it escapes the plateau with short circuits, not deep ones, and where it escapes, the light cone simulates it.

Expressivity → a weight sector. The networks of Chapter 20 conserve the weight. The circuit never leaves the weight-k sector, whose dimension is $\binom{n}{k}$, not 2^n. The chapter's script checks this without noise, with `qang.qml.WeightQNN`, $2n$ layers, 100 draws per point and the derivative of the qg_Z readout of qubit 0. The calculation is exploratory: it was not registered in advance.

Table 23.6: Noiseless gradient variance of the weight network ($L = 2n$, 100 draws).

weight	$n = 4$	$n = 6$	$n = 8$	$n = 10$	$n = 12$	decay
1 (dimension n)	0.300	0.107	0.041	0.026	0.014	like $n^{-2.8}$: polynomial
$n/2$ (dimension $\binom{n}{n/2}$)	0.283	0.075	0.012	0.0043	—	variance × dimension between 0.8 and 1.7: exponential

At weight 1 the variance falls 22 times from 4 to 12 qubits, like a polynomial. At half filling it follows $1/\binom{n}{n/2}$, which grows almost like 2^n (Figure 23.4). Restricting to a sector does avoid the plateau, but only when the sector is small, and a small sector is exactly what a classical computer simulates. At weight 1 the whole network is an $n \times n$ matrix, the quadratic twin of Chapter 20 (Section 20.4). This is the same trade-off the local cost met, and the one Cerezo et al. (2023) describe: where there is no plateau, the problem tends to be simulable.

Noise → the weight filter. Under equal T1 the filter returns the noiseless gradient exactly (Chapter 20, §112 T1). Without the filter, at weight 1 the raw gradient is K times smaller. Here $K = (1-\gamma)^{k\,d}$ is the kept fraction, which falls exponentially with the depth d: this is the noise-induced plateau. The filter does not make it free. It trades a variance K^2 times smaller for a fraction K of useful shots, so the shot cost grows like $1/K$ instead of $1/K^2$. In the notes the filter resolved the sign of the gradient with fewer shots in 10 of 12 cells. At half filling with $K \leq 0.14$ it needed 1.3–2.2 times more (§112: prediction T5 failed). With dephasing or gate errors inside the sector, the filter does not return the noiseless gradient (Chapters 21 and 22).

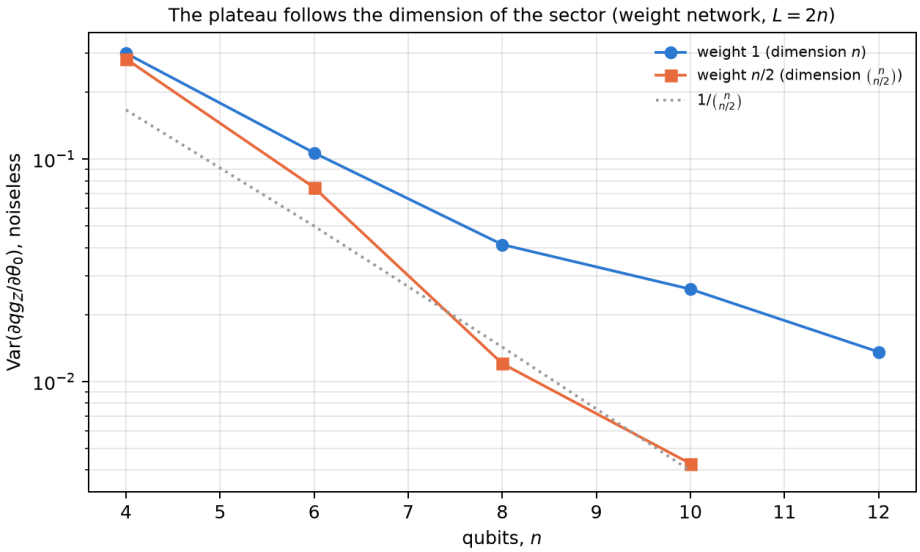


Figure 23.4: Noiseless gradient variance of the weight network against the number of qubits. At weight 1 (a sector of dimension n) it falls like a polynomial; at half filling it follows the inverse of the sector's dimension (dotted line), which grows almost like 2^n.

Table 23.7: The three plateaus and qang.

source of the plateau	without qang	with qang	reach
expressivity	ansatz over the whole space	weight-k sector, dimension $\binom{n}{k}$	no plateau at weight 1, but simulable; at half filling, a plateau
global cost	fidelity	local cost $(1 - \overline{qg_Z})/2$	only with short circuits, which the light cone simulates
noise (T1)	gradient K times smaller	filter: exact gradient with a fraction K of the shots	shots from $1/K^2$ to $1/K$; no fix for dephasing or in-sector errors
unit	radians	qga, m-qg	none

23.9 What remains open

Everything in this chapter is noiseless or on noise models. The §159 natural gradient was tested with exact gradients on one small model; measuring the qg_YY metric with shots on hardware, and seeing whether the 1.3× gain survives the metric's shot noise, is still to do. The filtered kernel was tested on one small, easy kernel, without the exponential concentration that appears with more qubits.

23.10 Without qang and with qang

Table 23.8: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
optimize the parameters	descent in the angle	descent in qg: diverges at the poles and fails (0/20)	**against**: optimize in θ, read out in qg
regularize near the poles	—	clipping or Tikhonov on $1/\sin\theta$	**none** against not using qg: it only repairs the problem qg creates
understand why	Fubini-Study metric	$F_{qg} = 1/(1-qg_Z^2)$: the qg singularity is the metric; $\sigma_m = 1000\sin\theta/\sqrt{N}$	**in favour** as a reading: it explains the failure and gives each readout's error
count gradient evaluations	2 per parameter	2 per parameter	**none**: no coordinate goes below two
natural-gradient metric	block-diagonal metric (Stokes et al.)	the same, read as qg_YY correlators; diagonal exactly 1/4	**none** in the method; **in favour** in the reading: one circuit per layer

action	without qang	with qang	difference
step near the poles	fixed rate	damped by $\lvert\sin\theta\rvert$, no tuning	**in favour** at aggressive rates (2.3 to 13 times more successes from rate 2); costs little or nothing at safe rates
repair a noisy kernel	normalize by the diagonal	qg filter on fixed-weight circuits	**in favour**: Gram 1.0 to 2.5 times better without the diagonal circuits; **none** in accuracy
training cost	global fidelity	local cost $(1-\overline{qg_Z})/2$	**in favour** with short circuits (470 times more variance at $n=10$); **none** with deep circuits
expressivity plateau	ansatz over the whole space	weight sector	**in favour** at weight 1 (polynomial variance), but simulable there; **none** at half filling
noise plateau (T1)	gradient K times smaller	weight filter: exact gradient	**in favour**: shots from $1/K^2$ to $1/K$; **against** in 2 of 12 cells of §112 with $K \le 0.14$

Balance. The chapter is largely a correction: qg is not a coordinate to optimize in, and Fisher geometry says exactly why. What remains in qang's favour is not a coordinate but a reading: the natural-gradient metric as qg_YY correlators, damping as $\lvert d\,qg_Z/d\theta\rvert$, the kernel as a qang angle repaired by the filter, and the local cost as the register's mean readout. They are known results (Stokes et al., Cerezo et al.) written on the Chapter 1 scale, plus two measured gains: the kernel repair with the filter and the robustness of pole damping.

23.11 In code

```
import math
import numpy as np
from qang import fisher as F
from qang import gradients as GR

F.fisher_theta(1.0), F.fisher_qg(math.cos(1.0))     # 1 in the angle; 1/sin^2(1) = 1.41
GR.inverse_jacobian_raw(0.001)                    # -1000: the singularity
GR.inverse_jacobian_tikhonov(0.001)                # -0.40: regularized
th = np.random.default_rng(0).uniform(0, 2 * np.pi, 12)
B = F.block_metric(th, 4, 3)                         # qg_YY blocks / 4
G = F.full_metric(th, 4, 3)                         # full Fubini-Study
np.diag(G)                                           # all 0.25
np.abs(B[:4, :4] - G[:4, :4]).max()               # ~1e-16 within the layer
```

23.12 Summary

- On the Chapter 1 scale, $dE/d\,qg_Z = -(1/\sin\theta)\,dE/d\theta$ diverges at the poles. Raw descent in qg jumps from pole to pole; clipping and Tikhonov fix it, but converge more slowly than descent in θ.
- The Fisher information is 1 in the angle and $1/(1-qg_Z^2)$ in qg: the qang angle is the flat coordinate. Optimize in θ and read out in qg, with $\sigma_m = 1000\sin\theta/\sqrt{N}$ m-qg.
- No change of coordinate goes below two evaluations per parameter.
- In a real ansatz the metric's diagonal is exactly 1/4 and the blocks are qg_YY correlators. The block metric helps a little (1.3 times); the full one, 4.6 times; counting circuits, the three methods end within 15%.
- Damping by $|\sin\theta|$ is not a new optimizer, but it withstands aggressive rates when starting near a pole.
- The qg filter repairs the Gram matrix of a fixed-weight kernel better than the normalization and without its circuits; accuracy barely changes.
- The qg local cost escapes the plateau only with short circuits, and there it is also classically simulable.
- Of the three sources of a plateau, qang answers two with tools: the weight sector against expressivity (no plateau at weight 1, which is simulable) and the filter against T1 noise (exact gradient, shots from $1/K^2$ to $1/K$). The unit changes none.

Bibliographic notes

Fisher geometry is Braunstein and Caves (1994), and the natural gradient, Amari (1998) and, for quantum circuits, Stokes et al. (2020). Damped descent is that of Levenberg (1944) and Marquardt (1963). Quantum kernels are Havlíček et al. (2019) and Schuld and Killoran (2019); barren plateaus, McClean et al. (2018) and Cerezo et al. (2021); noise-induced ones, Wang et al. (2021), and their relation to simulability, Cerezo et al. (2023). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part VIII

VIII. Many-body dynamics, open channels, matter and chemistry

24 Hamiltonian simulation, quantum walks and VarQTE

Simulating the time evolution of a quantum system is the most natural application of a quantum computer, and it is also where the qg filter of Chapter 21 reaches furthest: many Hamiltonians of physics conserve the number of excitations, and so do their Trotter circuits. Then the Hamming weight of each shot is a free witness, and the filter applies to every observable read in Z. This chapter goes through time evolution on the Chapter 1 scale: XY and XXZ chains, a quantum walk, noise channels as maps of the readouts, variational evolution at fixed depth and, more briefly, differential equations, Hubbard, a gauge theory, coherence in Ising and Boltzmann machines. The structural results are checked with code; the results on device noise models are those of the notes' committed runs, read from their stored files.

24.1 From the base: conserving N means having a filter

On the Chapter 1 scale, each qubit reads $m_Z = 1000(1-2p)$, with p the probability that it is excited. The number of excitations of a shot is its Hamming weight, and the register's mean readout counts it: $N = \sum_i (1 - q g_{Z,i})/2$. If the Hamiltonian conserves N and so does every gate of the circuit, every shot with the wrong weight had an error, and the filter discards it (Chapter 21).

With a single excitation, the state lives in n amplitudes, one per qubit: the excitation hops between neighbouring qubits. It is a **continuous-time quantum walk**, and each qubit's readout says where the walker is.

24.2 Trotter and the quantum walk on the XY chain

The XY chain with fields, $H = J\sum_k (X_k X_{k+1} + Y_k Y_{k+1})/2 + \sum_k h_k Z_k/2$, conserves N, and its first-order Trotter circuit uses only R_Z and XX+YY gates, which conserve it too (notes, §147).

How many steps. More Trotter steps lower the Trotter error and add noise, so on a device the error has a minimum at a number of steps r^*. Noiselessly, with $n = 5$, random fields and $T = 3$, the script recomputes the Trotter error: 0.381, 0.363, 0.221, 0.140, 0.078, 0.054 and 0.033 with 1, 2, 3, 4, 6, 8 and 12 steps. On the notes' 21 IBM device noise models, the filter (plus readout correction) halved the best error (median 0.147 → 0.067; 1.6 to 2.8 times depending on the device) and moved r^* from 8 to 12 steps in the median: **the filter lets the simulation use deeper circuits**. On 11 devices the optimum was 12, the deepest tested.

The walk. With $n = 7$ and the walker in the middle, the variance of its position $\sigma^2(t)$ grows like t^2 in a quantum (ballistic) walk and like t in a classical (diffusive) one. The script computes it with the exact evolution in the weight-1 sector:

Table 24.1: $\sigma^2(t)$ of the walk on 7 qubits; the exponent is the slope of $\log \sigma^2$ against $\log t$.

t	0.4	0.8	1.2	1.6	exponent
exact quantum walk	0.320	1.278	2.832	4.759	1.96
with the Trotter steps	0.327	1.291	2.832	4.737	1.94
classical walk, same rate	0.791	1.513	2.113	2.585	0.86

On the simulated devices, **raw data look like a classical walk**: median exponent 0.87, because noise spreads the walker from the start. **With the filter it is ballistic again**: median exponent 1.84, and σ^2 stays within 20% of the exact value at all four times on 15 of 21 devices (raw: none). The committed prediction asked for an exponent of at least 1.5 on 18 devices and it held on 16: it failed. The four noisiest ended between 1.26 and 1.38, and FakeKyoto at zero.

24.3 Several excitations: the XXZ chain

The XXZ chain adds an interaction $V\, Z_k Z_{k+1}/4$: the excitations interact and everything still conserves N. Here the compilation matters (notes §60). The script measures, with exact density matrices ($n = 6$, 4 steps), what share of the imbalance error the filter removes with one noise source at a time:

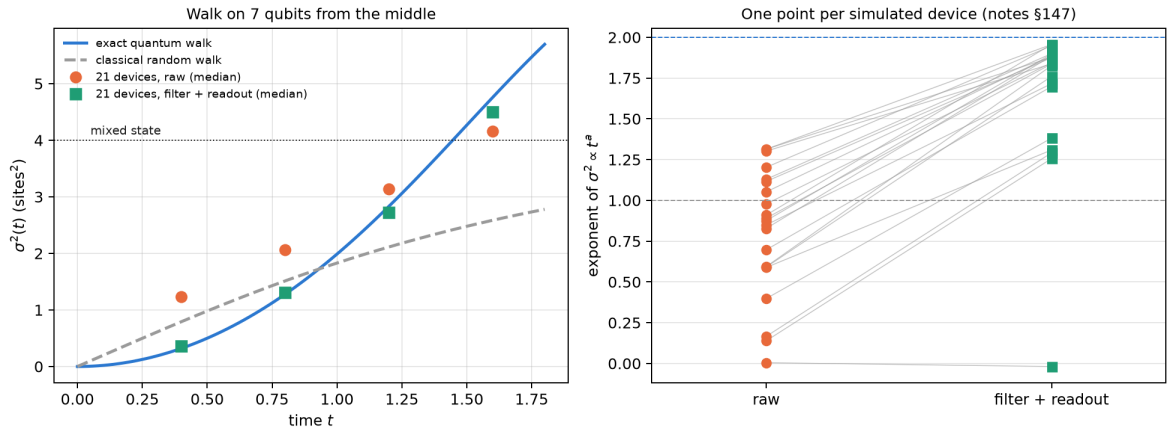


Figure 24.1: Left: $\sigma^2(t)$ of the walk on 7 qubits; the exact one is ballistic, the classical one diffusive; the medians of 21 simulated devices sit near the classical curve without the filter and near the exact one with it. Right: exponent of $\sigma^2 \propto t^a$ per device, raw and with filter plus readout (notes §147).

Table 24.2: Share of the error the filter removes on the XXZ chain (notes §60, recomputed by the script).

noise	3 CNOT per bond	N-conserving gates
2-qubit depolarizing, 0.01	41.1%	41.1%
amplitude damping (T1), 0.01	48.5%	99.5%
Z dephasing, 0.01	1.6%	0.0%
readout, 0.02	93.8%	—

- **T1 leaks through the CNOTs.** Each decay lowers N, and the filter catches it if the state is in the sector when it happens. Inside a three-CNOT bond the intermediate states are not in the sector, and a decay mid-gate can be rotated back to $N = n/2$ and pass. With gates that conserve N at every point, the filter removes 99.5%.
- Z dephasing does not change N, and the filter does not see it.
- With mixed noise, filter plus ZNE sits 2.6 to 8 times below the best single method up to 4 steps (§60).

On the 21 simulated IBM devices with $n = 6$ (§148), the filter's gain falls from 5.6 times with one excitation to 2.6 with two and with three: sectors with more excitations keep more shots, but more of the noise lands inside them. The prediction that the filter never makes things worse failed in one of 63 cases (FakeKyoto

at half filling). On IonQ's forte-1 noise model (§61), the filter removed about half of the error in both XY compilations, with a ratio of 0.82 between them, inside the committed band: on trapped ions, without dominant T1, the CNOT leak does not appear.

24.4 Channels as affine maps of the readouts

A qubit is $\rho = (I + r \cdot \sigma)/2$ with $r = (qg_X, qg_Y, qg_Z)$. Every channel with Kraus operators acts as $r' = Mr + c$ (notes §162, `qang.channels.bloch_affine`): **the readouts after any noise are affine in those before it**. The script checks it on 800 random states to $2.8 \cdot 10^{-16}$. On the Chapter 1 scale:

- amplitude damping γ: $m_Z \to (1-\gamma)m_Z + 1000\gamma$, and m_X, m_Y are multiplied by $\sqrt{1-\gamma}$ (with $\gamma = 0.3$: 0.7 and 0.837);
- dephasing: leaves m_Z alone; depolarizing: multiplies the three readouts by the same number and keeps the qang angle;
- T1/T2 relaxation for a time t: $m_Z \to e^{-t/T_1} m_Z + (1 - e^{-t/T_1})(1 - 2p_1) \cdot 1000$ and $m_{X,Y} \to e^{-t/T_2} m_{X,Y}$ (with $t = 0.5$, $T_1 = 1$, $T_2 = 0.8$: 0.6065 and 0.5353, which are $e^{-0.5}$ and $e^{-0.625}$).

It is the Chapter 1 picture carried over to noise: every channel moves the readout point inside the sphere by a linear rule.

24.5 Variational evolution at fixed depth

Trotter grows with time: more time, more steps, more gates. Variational evolution (VarQTE, McLachlan's principle) fixes the circuit and changes only its angles (notes §162). With a weight-preserving ansatz (6 layers of RBS and R_Z, 54 parameters, 48 two-qubit gates), the filter still applies. The integration is classical; the device prepares and reads the state. The script recomputes the noiseless fidelity of the angles stored by the notes' run (XXZ chain, $n = 5$, $V = 2$, two excitations): 0.9999, 0.9956, 0.967 and 0.934 at $t = 0.5$, 1, 1.5 and 2, equal to the stored ones.

Table 24.3: VarQTE against Trotter on three noise models (notes §162).

backend	filtered error at $t = 2$: variational / best Trotter (72 gates)	variational error, $t = 0.5 \to 2$
FakeBrisbane	0.036 / 0.076	0.052, 0.030, 0.030, 0.036
FakeTorino	0.030 / 0.042	0.034, 0.022, 0.012, 0.030
FakeFez	0.022 / 0.032	0.017, 0.008, 0.005, 0.022

At late times fixed depth wins, on all three devices, even though its noiseless fidelity is only 0.934 (the prediction that it would exceed 0.95 failed). The variational error stays flat and the best Trotter error grows 2.6 to 11 times. Before $t \approx 1$ a short Trotter circuit wins. Limit: at $n = 5$ everything is classically trivial, and on hardware the integration would need metric and gradient circuits at every step. The ansatz lives in the weight-2 sector of 5 qubits, of dimension 10: that is why it has no barren plateau, and for the same reason it is classically trivial (Section 23.8).

24.6 Differential equations

A one-qubit re-uploading circuit can solve a differential equation: $u(x) = u_0 + s(f(x) - f(0))$ with $f = \langle Z \rangle$, trained on the residual (notes §22). The qang question is how to encode x: in the angle, or with arccos, so that x is a qg readout. The notes' result: **qg wins clearly on exponential decay** (2.3, 17 and 6.6 times better than the best angle scale with 2, 3 and 4 layers) and **does not win in general**: on the logistic equation the angle is about 2 times better, and on the sine 25 times. qg has no scale to tune; the angle does (with scale π it fails everywhere). The classical control, Chebyshev collocation with as many coefficients as parameters, solves everything to 10^{-7} or better. The result says which encoding to use, not that the circuit beats a classical method.

24.7 Hubbard and a gauge theory: symmetry filters on every observable

1D Hubbard (notes §29): 4 sites on 8 qubits, one register per spin. The Hamiltonian conserves $N_\uparrow$ and $N_\downarrow$, so each register has zero mean readout at all times, and every observable (charge imbalance, double occupancy) is in Z. The script checks it noiselessly: mean readout 0 in both registers and kept fraction 1 at 1, 2, 4 and 8 steps. In the notes, the filters roughly halve the error on a device without routing, and ZNE plus spin filters gives the best double occupancy (6 to 9 times below raw). On the heavy-hex lattice, routing creates spin leak (8 to 47% of the right-N shots land in the wrong spin sector), and there the second filter pays. From 4 steps on everything sits on the noise floor.

$\mathbb{Z}_2$ **lattice gauge theory** (§66): 4 matter sites and 3 links. The Gauss law G_j is a product of Z's, so it is the qg readout of a parity: the same Z shots give the local Gauss witnesses and three filters (N, Gauss and both). The script checks that noiselessly each $\langle G_j \rangle$ is exactly ± 1 and no shot is lost. In the notes, with 2 steps, the Gauss filter removes more than the N filter (80 against 64% under depolarizing noise) and both together 90%. Under T1, 58–65%: the same CNOT leak. The local witnesses do not locate errors in this homogeneous model. Gauss-law post-selection is known; what qang adds is reading it as qg of parities from the same shots.

24.8 Coherence in the transverse-field Ising model

The difference between the Z-basis entropy and the entropy of the state is a known measure: the relative entropy of coherence $C(\rho)$ (Baumgratz et al. 2014). It needs the entropy of the state, which counts do not give. The notes (§34) propose a bound from two readouts per qubit:

$$C \geq \sum_i \left[qg_S(qg_{Z,i}) - qg_S(|r_i|) \right], \qquad |r_i| = \sqrt{qg_{X,i}^2 + qg_{Y,i}^2 + qg_{Z,i}^2}.$$

The script evaluates it on thermal states of the 6-qubit Ising ring at 20 points (β from 0.25 to 3, field g from 0.25 to 2): **it captures 92.3 to 99.8% of the coherence** (Figure 24.2, right). The single-qubit bound captures 15–17% (1/6 of the sum, by symmetry), and the basis-entropy bound is useless: it goes down to -5.9 bits. With 1000 shots the bound reads to ± 0.04 bits. Where it fails: on a graph state every reduced qubit is $I/2$ and the bound gives 0 with $C = 6$ bits. The tightness is a property of these thermal states, not a law.

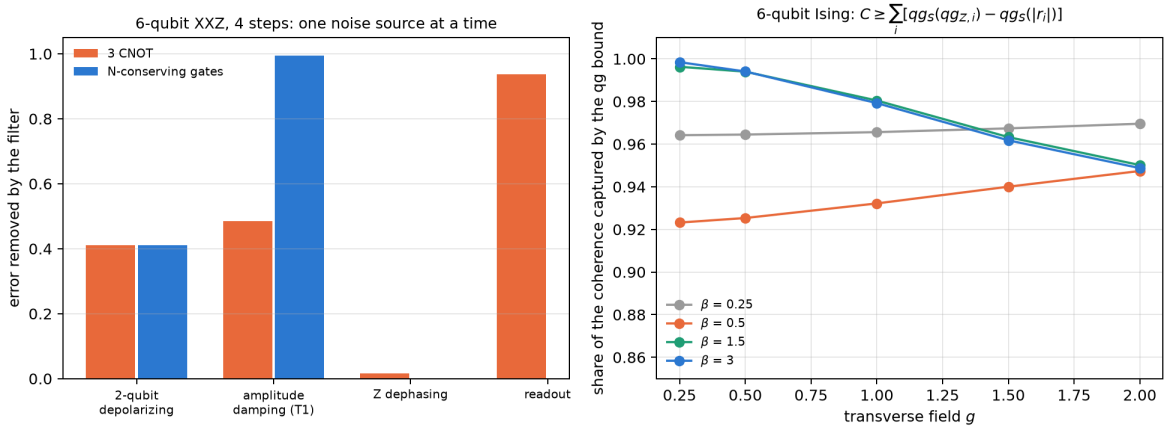


Figure 24.2: Left: share of the imbalance error the filter removes on the XXZ chain, one noise source at a time; under T1 the CNOTs let half through and N-conserving gates almost nothing. Right: share of the coherence of Ising thermal states captured by the qg bound.

24.9 Boltzmann machines: the I–η plane

A recent review (Singh et al. 2026) looks at each visible–hidden pair of an RBM through two numbers: the covariance η and the mutual information I, which lie between two bounds $LB(\eta)$ and $UB(\eta)$, and observes that trained RBMs sit on LB. In qg (notes §47), a pair of spins is fixed by three numbers, qg_v, qg_h and η, and the script checks: LB is exactly the case of zero marginals ($2.2 \cdot 10^{-16}$) and the minimum of I at fixed η; UB is the locked pair with $|qg_v| = |qg_h| = \sqrt{1-\eta}$ ($2.9 \cdot 10^{-15}$). In between, to leading order, $I - LB \approx \eta^2(qg_v^2 + qg_h^2)/(2\ln 2)$; with $\eta = 0.2$–0.4 the formula sits 10–14% above the exact value, because it is leading order in η too.

The explanation: the Ising driver has $\mathbb{Z}_2$ symmetry, a symmetric RBM has zero marginals, and so it sits on LB by symmetry, not by a learning principle. The prediction, which could have failed, held: a longitudinal field moves the same learner off LB by the amount the marginals set.

24.10 What remains open

No result in this chapter ran on hardware. The three clearest things to test on a QPU are the filtered walk (§147, the change of exponent from 0.87 to 1.84), the

T1 leak through CNOTs (§60, which does not show up on an ion simulator) and fixed-depth VarQTE against Trotter at late times (§162).

24.11 Without qang and with qang

Table 24.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
simulate N-conserving dynamics	raw readout or ZNE	weight filter on every Z observable	**in favour**: best error 1.6–2.8 times lower and deeper circuits (§147)
tell a quantum walk from a classical one	raw σ^2: exponent 0.87, looks classical	filtered: 1.84 (exact 1.96)	**in favour**: the filter recovers the physics
choose the compilation	CNOTs	N-conserving gates: the filter removes 99.5% of T1 against 48.5%	**in favour** on platforms with T1; **none** on ions (§61)
several excitations	—	gain 5.6 times with one, 2.6 with two or three	**in favour**, decreasing with N
describe a channel	Kraus matrices	$r' = Mr + c$ on the readouts	**none** in content (it is the Bloch representation); **in favour** as a reading
late-time evolution	growing Trotter	fixed-depth VarQTE with the filter	**in favour** after $t \approx 1$; **against** before
encode an ODE	angle with a scale	arccos, no scale	**in favour** on exponential decay; **against** on others; a classical method beats all
Hubbard and gauge	known post-selection	witnesses and filters as qg of parities, from the same shots	**in favour** in cost (no extra circuits); the idea is known

action	without qang	with qang	difference
bound the coherence	basis entropies: useless (−5.9 bits)	qg bound from two readouts per qubit: 92–100%	**in favour** on these states; **none** on graph states
I–η plane of an RBM	a learning principle	the bounds are the qg marginals; saturation is symmetry	**in favour**: explains and predicts

Balance. Most of what this chapter gains belongs to the weight filter, not to the scale: conserving N turns every shot into a witness, which recovers a ballistic walk, allows deeper circuits and makes the compilation matter. What belongs to qang is the reading: channels as affine maps of m_X, m_Y, m_Z, Gauss witnesses as qg of parities, a coherence bound from two readouts per qubit and the RBM bounds as marginals. The gains that depend on the scale itself (the ODE) are partial, and the chapter says so.

24.12 In code

```
import numpy as np
from qang import channels as CH

M, c = CH.bloch_affine(CH.amplitude_damping(0.3))
np.diag(M), c              # (0.837, 0.837, 0.7) and (0, 0, 0.3): m_Z -> 0.7 m_Z + 300
r = np.array([0.2, -0.1, 0.5])    # (qg_X, qg_Y, qg_Z)
CH.qg_after(CH.amplitude_damping(0.3), r)   # = M @ r + c
```

24.13 Summary

- If the Hamiltonian and the gates conserve N, each shot's weight is a witness and the filter reaches every Z observable.
- On the XY chain, the filter halves the best error and moves the Trotter optimum to deeper circuits; the raw walk looks classical (exponent 0.87) and the filtered one is ballistic (1.84; exact 1.96).

- Under T1, CNOTs let half the error through and N-conserving gates almost none (99.5% removed).
- Every channel is an affine map of the readouts; fixed-depth VarQTE wins at late times.
- Differential equations, Hubbard, gauge, coherence in Ising and RBMs: the filter and the qg readouts help where the chapter measures it, and where they do not, it says so.

Bibliographic notes

The product formula is Trotter (1959) and Suzuki (1991), and universal quantum simulation, Lloyd (1996). The quantum walk is Aharonov et al. (1993); variational evolution, McLachlan (1964) and Yuan et al. (2019). The generators of Markovian channels are Lindblad (1976) and Gorini et al. (1976), and the Hubbard model, Hubbard (1963). The coherence measure is Baumgratz et al. (2014), and the I–η plane of RBMs, Singh et al. (2026). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

25 Quantum spectroscopy, PCA and molecular chemistry

Chapter 24 used the conservation of the number of excitations to filter dynamics. This one does the same with spectra and molecules, where the conserved number is the number of electrons. The register's mean readout, on the Chapter 1 scale, counts the electrons of each shot; if the circuit conserves them, that number is both a reference-free error witness and a filter. The chapter follows that idea from a data PCA to 12-qubit molecules, and also shows where it stops reaching: when the error lives in terms that are not measured in Z. The structural results are checked with code; the noise-model results are those of the notes' committed runs.

25.1 From the base: counting electrons with the register readout

In the Jordan-Wigner encoding each qubit is a spin-orbital, and its occupation is $(1 - qg_Z)/2$, that is $(1000 - m_Z)/2000$ in m-qg. Summing over the register,

$$N_e = \sum_i \frac{1 - qg_{Z,i}}{2}, \qquad \overline{qg_Z} = 1 - \frac{2N_e}{n}.$$

For any circuit that conserves the number of electrons, the ideal mean readout is known before running (0 for H_2: 2 electrons on 4 qubits). A departure from it is a witness that the hardware added or removed electrons (T1 pushes it towards $+1$), and shot by shot it is a filter. **The filter is the known symmetry verification** (Bonet-Monroig et al. 2018; McArdle et al. 2019); what qang adds is the witness and its reading.

25.2 PCA in the weight-1 sector

Quantum PCA diagonalizes $\rho = C/\mathrm{tr}\, C$. In the unary encoding, a mixture of the loaded data vectors (with weights $|x_j|^2$) is exactly that ρ, and an RBS circuit applying the PCA rotation leaves eigenvalue λ_k on qubit k (notes §146). The script checks it with exact states: for 5 components and 16 data points, each qubit's excitation probability is the normalized eigenvalue (0.640, 0.203, 0.122, 0.023, 0.011) to $2.2 \cdot 10^{-16}$, with 10 RBS in the rotation. Every gate conserves the weight, so the filter applies. It is not the quantum-PCA speed-up (5 components use 5 qubits and the rotation is computed classically): it tests how well a device reads a known spectrum.

On the notes' 21 IBM device noise models, with all four committed predictions passing, the raw estimate is biased upwards: shots leaving the sector add excitations on every qubit (median L1 error 0.266). Filter plus readout correction brings it to 0.048 (4.6 times lower; 1.6 to 7.7 depending on the device) and orders all five eigenvalues correctly on 18 of 21 devices (raw: 13). What is left is a flattening of the spectrum that grows when few shots are kept (FakeKyoto keeps 15% and fails).

25.3 H_2: the witness and the filter

The script checks the structure on the notes' H_2 Hamiltonian (STO-3G, 0.735 Å, 4 qubits): it commutes exactly with N, $N_\uparrow$ and $N_\downarrow$; the N-conserving ansatz reaches the exact (FCI) energy to $5 \cdot 10^{-11}$ mHa, and the register's ideal mean readout is 0. Hartree-Fock sits 20.3 mHa above. The notes (§21) ran fake brisbane's noise model with 20,000 shots and an idle delay before measurement that amplifies T1:

Table 25.1: H_2 energy error against FCI, in mHa (notes §21).

method	0 µs	20 µs	50 µs
classical Hartree-Fock	20.3	20.3	20.3
quantum, raw	74	140	231
+ readout correction	18	90	188
+ readout + qg filter	**5.5**	**20**	**30**
witness: mean readout (ideal 0)	+0.011	+0.057	+0.122
shots kept	91%	83%	72%

The witness rises as electrons are lost. With the filter the energy is 3.2 times better than with readout correction alone without delay (6 times at 50 µs) and 3.7 times better than Hartree-Fock, but still 3.4 times above chemical accuracy: the rest sits in the 4 XXYY-type terms, which a Z filter cannot reach, and in errors that conserve N. FCI is exact and cheap at this size: the table measures what the correction buys a noisy quantum computation, not an advantage.

Spin filters (§28). The Hamiltonian conserves $N_\uparrow$ and $N_\downarrow$ separately, and each spin register has its own mean readout. A second free filter, from the same shots: under depolarizing noise it removes a further 34–37% of the error the total filter leaves, but on fake brisbane the gain is within the seed spread (6.5 → 5.9 mHa), because the spin leak is below 1%. The leak witness says when it pays.

25.4 Why the filter helps H_2 and not LiH

With LiH (6 qubits) the §21 filter did nothing. The notes (§50) separated the causes: reach (the filter acts only on Z shots), routing and depth.

Table 25.2: Where the error sits, in mHa against the noiseless circuit (notes §50).

case	2q gates	Z group raw → filtered	X/Y groups raw	total raw → filtered
H_2, all-to-all	3	13.0 → 4.9	0.9	13.9 → 5.9
H_2, brisbane	6	20.6 → 2.8	2.4	23.0 → 5.2
LiH, all-to-all, 1 layer	20	−2.2 → −1.9	26.2	24.0 → 24.3
LiH, brisbane, 3 layers	60	−30.5 → −18.0	165.8	135.3 → 147.9

Reach is the main reason. In H_2 the error is in the Z group, where the filter acts; in LiH it is in the 16 X/Y groups, where a Z filter does not look, and not even a perfect filter would go below that column. Routing is not the cause: the LiH ansatz is nearest-neighbour and needs the same gates on an all-to-all device. The qang attempt to rescale the X/Y groups by the kept fraction failed: it over-corrects 2–3 times, because local gate noise is not a global depolarization.

25.5 Checks that reach every group

If the problem is reach, the remedy is a check that commutes with every term (§52): the parity of the electron number on an ancilla (6 CNOTs), and N modulo 4 on a second ancilla, which also catches the errors that change N by 2. With 3-layer LiH and all-to-all noise, parity plus mod 4 take the error from 83.9 to 13.4 mHa (§50's best was 76.6), discarding 28% of the shots: for the first time in the repository, LiH ends (on that model) below Hartree-Fock. On IonQ's forte-1 model (§53) the mod-4 check survives (3.8 times less error, parity alone 1.25 times), but the result stays above Hartree-Fock.

Larger molecules (§54, §55, §69), with ideal projections:

Table 25.3: The ladder of checks, error in mHa against the noiseless circuit (notes §69).

molecule	qubits	raw	N	$N_\uparrow$, $N_\downarrow$	$+ S^2 = 0$
H_2O	8	363.9	54.6	43.3	26.7
H_4	8	287.5	79.9	70.5	44.7

- Number checks keep working up to 12 qubits, but **they reach a shrinking share** of the error: the part that changes N falls from 92% (LiH, 60 CX) to 67% (H_6, 132 CX). The rest conserves N and no number check sees it.
- Splitting by spin recovers part of it (from 75.7 to 66.8 mHa on H_4).
- **The projection onto the singlet $S^2 = 0$ removes a further 37–38%**, the largest gain since §52, but S^2 is not diagonal in Z: it needs 640 Pauli strings for the projector and 2064–3456 for the projector times H, against 105–185 for H.

Λ as a third filter (§153). A linear molecule conserves Λ, the projection of the orbital angular momentum on its axis. With the π orbitals written as π± ($m_l = \pm 1$), Λ is diagonal and is checked shot by shot with only the four π qubits, so it reaches 29 of LiH's 317 measurement groups (10 qubits), against 1 for the number filter. The committed prediction half held: **where $N_\uparrow, N_\downarrow$ already applies, Λ adds almost nothing** (13.9 → 13.3 mHa on FakeBrisbane), because changing Λ without changing N takes two errors. **The extra reach is real and does not pay**: in the 28 groups where only Λ acts, it discards 4–10% of the shots and their error barely moves, because the error that matters there is on the σ qubits measured in X/Y.

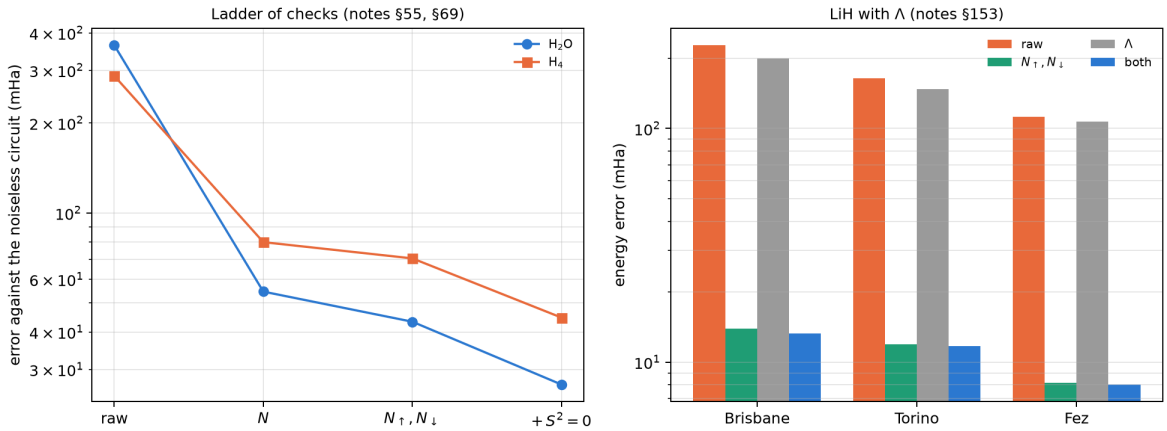


Figure 25.1: Left: ladder of checks on H_2O and H_4 (notes §55 and §69); each symmetry removes a smaller share. Right: 10-qubit LiH with the spin filter, the Λ filter and both, on three noise models (notes §153); Λ alone barely moves the error.

25.6 Overlaps: inversion test against swap test

The fidelity $F = |\langle x|y\rangle|^2$ of two unary states can be estimated two ways (§149). The **inversion test** loads x, undoes the loader of y and reads the probability of one basis state: 5 qubits, everything conserves the weight. The **swap test** uses 11 qubits and an ancilla, and reads $P(\text{ancilla} = 0) = (1 + F)/2$. The script checks both formulas with exact states ($3 \cdot 10^{-16}$ and $6 \cdot 10^{-16}$). On six IBM noise models:

- with the filter, the inversion test cuts its error 3.8 to 11.2 times, down to 0.010–0.036;
- in the swap test the answer is on the ancilla, outside the symmetry, and the filter gains only 1.4–1.8 times (error 0.089–0.239);
- a fidelity classifier (40 flowers) keeps its noiseless accuracy with both methods, because comparing two fidelities cancels the common shrink;
- but the test "are they equal?" needs the absolute value: for identical states the filtered inversion test reads 0.92–1.00 and the raw swap test 0.20–0.65.

25.7 Constrained optimization: QAOA with "exactly K"

Choosing exactly K of n vertices is the same identity as the number of electrons: $\overline{qg_Z} = 1 - 2K/n$, 0.25 for $K = 3$ of $n = 8$ (notes §30). The script checks that

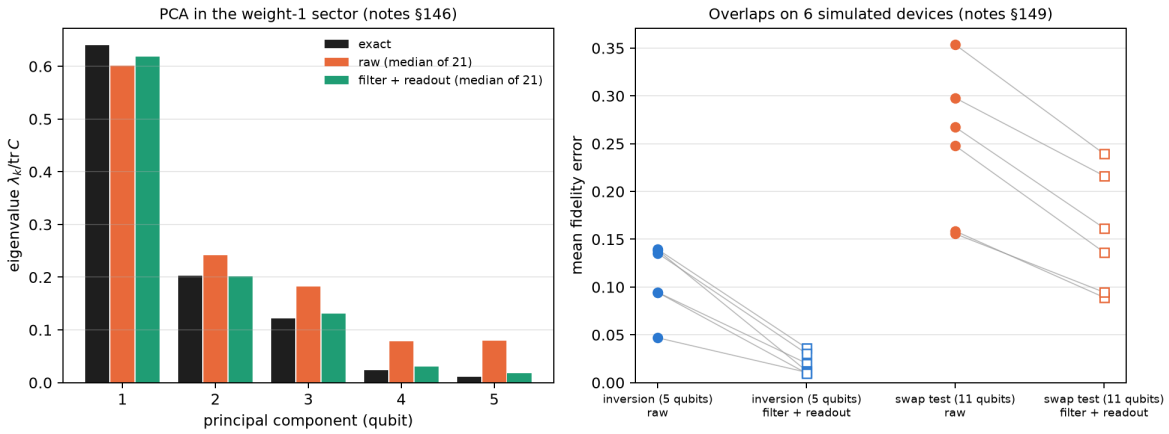


Figure 25.2: Left: PCA spectrum, exact and median of 21 simulated devices without the filter and with filter plus readout (notes §146); raw data lift the small eigenvalues. Right: fidelity error on six devices, inversion test and swap test, raw and filtered (notes §149).

with the XY mixer (which conserves the weight, from a Dicke state) the ideal distribution is 100% feasible at random angles, while penalty QAOA leaves 10 to 45% feasible mass. On an all-to-all device the filter doubles the XY mixer's probability of finding the optimum (0.150 → 0.278 at $p = 1$; 0.177 → 0.366 at $p = 2$), and on the IonQ models too (1.8–2.6 times, §38). On fake brisbane, with routing, everything falls to the level of picking a random feasible set, and the kept fraction warns of it (0.24–0.33, close to 0.22 for a mixed register). **Classical wins**: a greedy algorithm reaches 98.5% of the optimum. The qg ingredients improve QAOA against QAOA, not against classical optimization.

25.8 What remains open

No result in this chapter ran on hardware. The notes committed a criterion for LiH with the mod-4 check on an IonQ QPU (§53: it holds if parity plus mod 4 cut the error at least 2.5 times in every run and parity alone less than 1.5), at an estimated cost of about 1,100–2,200 dollars, which did not fit in the budget.

25.9 Without qang and with qang

Table 25.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
read a PCA spectrum	raw: biased upwards (L1 0.266)	weight filter: L1 0.048	**in favour**: 4.6 times less error; it belongs to the filter
detect electron loss	compare with a reference	register mean readout, known ideal	**in favour**: reference-free witness
H_2 energy	readout correction: 18 mHa	+ filter: 5.5 mHa	**in favour**: 3.2 to 6 times; symmetry verification is known
LiH energy with a Z filter	—	no gain: the error is in X/Y	**none**: the witness diagnoses, the filter does not reach
rescale X/Y by the kept fraction	—	over-corrects 2–3 times	**against**
ancilla checks	parity: a third	+ N mod 4: 6–8 times less error	**in favour**, with 1–2 ancillas; an extension of known methods
larger molecules	—	the reachable share falls from 92 to 67%	**in favour**, decreasing
Λ filter	—	reaches 29 more groups and does not pay	**none** where a spin filter already applies
absolute fidelity	raw swap test: 0.20–0.65 for equal states	filtered inversion: 0.92–1.00	**in favour** of the filtered inversion test
constrained QAOA	penalty: chance level	XY mixer + filter: doubles success	**in favour** against QAOA; **against** against classical greedy

Balance. As in Chapter 24, almost everything gained belongs to symmetry verification, which is known. What belongs to qang is the witness: the register's mean readout says how many electrons there are, and its kept fraction says how

much is left to filter. That witness is what led from the §21 failure on LiH to the §50 reach diagnosis and the §52 checks. The chapter also leaves the measured limits: the error that conserves every symmetry grows with size, and one more symmetry (Λ) can reach more groups without lowering the error.

25.10 In code

```
import numpy as np
from qiskit.quantum_info import Statevector
from qang import ibm

x = np.array([0.3, -0.5, 0.2, 0.7, 0.4])
y = np.array([0.1, 0.6, -0.3, 0.5, 0.5])
x, y = x / np.linalg.norm(x), y / np.linalg.norm(y)
sv = Statevector(ibm.unary_overlap_circuit(x, y, measure=False)).data
abs(sv[1]) ** 2, (x @ y) ** 2    # inversion test: P(only qubit 0 excited) = F
```

25.11 Summary

- In Jordan-Wigner, the register's mean readout counts the electrons: $\overline{qg_Z} = 1 - 2N_e/n$. Its ideal value is known before running, so it is a reference-free witness and, shot by shot, a filter.
- In the unary PCA, the filter cuts the error of the read spectrum 4.6 times; on H_2, the energy error 3.2 to 6 times.
- On LiH the Z filter does not reach the error, which sits in the X/Y groups; ancilla checks (parity and N mod 4) do, and cut the error 6–8 times.
- With larger molecules each symmetry reaches a smaller share of the error; S^2 removes a further 37–38% at the cost of many more measurements, and Λ adds reach without lowering the error.
- The filtered inversion test gives the absolute fidelity; the raw swap test does not. In constrained QAOA the filter doubles the XY mixer's success, and a classical algorithm still wins.

Bibliographic notes

Quantum PCA is Lloyd et al. (2014). The Jordan–Wigner transformation is Jordan and Wigner (1928); the VQE, Peruzzo et al. (2014), O'Malley et al. (2016) and Kandala et al. (2017), with integrals from PySCF (Sun et al. 2018). Symmetry verification is Bonet-Monroig et al. (2018) and McArdle et al. (2019). The swap test comes from Buhrman et al. (2001), and QAOA and its mixers, from Farhi et al. (2014) and Hadfield et al. (2019). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part IX

IX. Error correction, continuous syndromes and characterization

26 Error correction and continuous syndromes

Error correction protects a logical qubit by spreading it over several physical ones and measuring syndromes: parities that say which error happened without touching the information. On the Chapter 1 scale, each syndrome is the qg readout of a parity, a number between -1000 and 1000 m-qg that can be averaged and tracked in time. This chapter uses that readout for four things: choosing the code according to how much T1 and how much dephasing the device has, tracking noise drift with the syndromes themselves, decoding the surface code knowing that T1 leaves zeros, and deciding where to cut a circuit. The code-capacity results (perfect syndromes, one round) are checked with code; the circuit-level, qutrit, erasure-qubit and device-model results are those of the notes' committed runs.

26.1 From the base: the syndrome is a qg readout

A stabilizer such as Z_0Z_1 is $+1$ or -1 on each shot; its average is the qg of the parity, and the rate of non-trivial syndromes is $(1-qg)/2$, the same probability equation as Chapter 1. The two noises of an idle qubit are read with two experiments (notes §27): preparing $|1\rangle$ and measuring Z gives $qg_Z = -1+2\gamma$ (damping, T1), and preparing $|+\rangle$ and measuring X gives $qg_X = \sqrt{1-\gamma}\,(1-2p)$ (dephasing, T_φ). Inverting them gives γ and p, and with them one decides which code to use.

26.2 Which code to use under T1 and dephasing

At code capacity, with one round of damping γ and dephasing p on every qubit, the script recomputes the exact logical infidelity of four options: no code, the repetition code against bit flips, the phase-flip code, and the 4-qubit code of Leung et al. (1997), built for T1.

Table 26.1: Repetition codes under T1 and dephasing (notes §27, recomputed).

noise	no code	bit flip	phase flip
T1 only, $\gamma = 0.02$	**0.0067**	0.0101	0.0197
dephasing only, $p = 0.02$	0.0133	0.0384	**0.0008**
$\gamma = p = 0.02$	**0.0199**	0.0474	0.0204

- **No repetition code corrects T1**, and the bit-flip code is worse than not coding in every case. The boundary between "phase code" and "no code" is at $p = \gamma$ (ratio 1.00, 1.02 and 1.10 at $\gamma = 0.002$, 0.01 and 0.04).
- **The Leung code does correct T1**: its infidelity is $\approx 0.92\,\gamma^2$, 36 times below no code at $\gamma = 0.01$ and 364 times at $\gamma = 0.001$. But it is fragile to dephasing: it wins only if $p < \gamma/4$ (boundary 0.249, 0.247 and 0.237).

The rule in device times. With one idle round, $\gamma = t/T_1$ and $p = t/2T_\varphi$, so $p = \gamma/4$ is $T_2 = T_1$ and $p = \gamma$ is $T_2 = 0.4\,T_1$. The best option is **Leung if $T_2 > T_1$, no code if $0.4\,T_1 < T_2 < T_1$, and the phase code if $T_2 < 0.4\,T_1$**; the script confirms it at nine T_2/T_1 ratios. With both qg witnesses measured with 1000 shots each, the policy chooses right in 79% of 300 random cases, at a cost of 6.4% over the oracle, against 14.7% for the §27 policy (without Leung) and 64–220% for always taking one option. **The qg_Z readout alone does not work**: it sees T1 but not dephasing, and always recommends not coding.

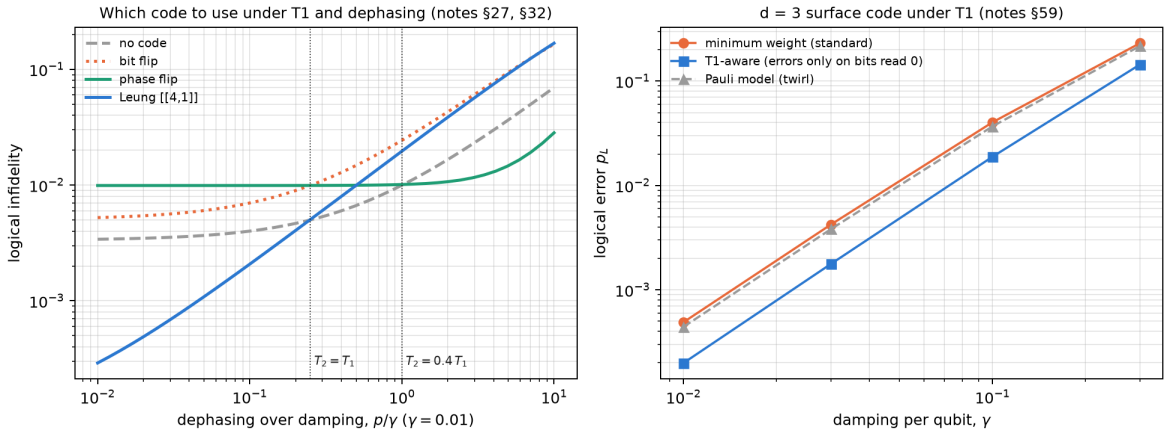


Figure 26.1: Left: logical infidelity of four options against the dephasing/damping ratio at $\gamma = 0.01$; the vertical lines are the $T_2 = T_1$ and $T_2 = 0.4\,T_1$ boundaries. Right: logical error of the d = 3 surface code under T1, with the standard decoder, the T1-aware one and the Pauli model.

26.3 The syndrome as a continuous witness

Once the code is chosen, its ancillas keep reading the noise. In the Leung code the stabilizer expectations are, in closed form (notes §33),

$$\langle Z_0 Z_1 \rangle = 1 - 2\gamma(1-\gamma), \qquad \langle XXXX \rangle = q g_X^4,$$

and the script confirms it against the Kraus computation (for instance, rates 0.0099 and 0.02161 at $\gamma = 0.01$ and $p = 0.003$). So each ancilla's qg readout inverts in one line to γ and p. In the notes, with T_2/T_1 jumping at random between 1.6, 0.6 and 0.3, tracking the noise with the syndromes alone costs 13.1% over the oracle **without any probe shots**, better than recalibrating every 10 windows (20.4%); a hybrid reaches 7.2% with 4 times fewer probe shots than probing every window (6.6%). Calibrating once at the start is worse than never coding (57.7% against 25.9%).

On device circuits. Reading the Leung syndromes on three IBM noise models (§88), the damping came out within 1–10% of the calibration on two of them and failed on the third; the causes were routing and a 9% error floor (§92). Reading them destructively, without ancillas or SWAPs, on the device's best 4-qubit path (§93), all four committed predictions held: the damping comes out within 1–4% of the calibration and the floor drops to 2–3%. The dephasing came out 1.17–1.51 times high, and §99 found why: $\langle XXXX \rangle$ reads the **product mean** over the four qubits, which is what limits the logical coherence, and against it the fit agrees to 0.99–1.04.

26.4 The surface code and the T1-aware decoder

In the rotated distance-3 surface code (9 data qubits), each averaged stabilizer is the qg readout of a parity. The script recomputes at code capacity, enumerating the 512 patterns: $p_L = 1.79 \cdot 10^{-5}$, $1.73 \cdot 10^{-3}$ and $1.44 \cdot 10^{-2}$ at $p = 0.001$, 0.01 and 0.03, and a pseudo-threshold of 0.0753 (notes §59).

The syndromes barely see T1. Under damping, a weight-w syndrome reads exactly $(1-\gamma)^w + \gamma^w$, against $(1-\gamma)^w$ for the equivalent Pauli model: at $\gamma = 0.01$ and $w = 2$, 0.98020 against 0.98010, a difference that would need about 10^8 shots. **The register's mean readout does see it**: it is γ under T1 and 0 for symmetric flips, and it comes from the same bits.

Decoding knowing that T1 leaves zeros. A decay only turns a 1 into a 0, so a decoder that places errors only on bits read 0 lowers the logical error:

Table 26.2: d = 3 surface code under T1, code capacity (notes §59, recomputed).

γ	Pauli model	minimum weight	T1-aware	maximum likelihood	ratio
0.01	$4.41 \cdot 10^{-4}$	$4.90 \cdot 10^{-4}$	$1.99 \cdot 10^{-4}$	$1.99 \cdot 10^{-4}$	2.46
0.03	$3.82 \cdot 10^{-3}$	$4.22 \cdot 10^{-3}$	$1.77 \cdot 10^{-3}$	$1.77 \cdot 10^{-3}$	2.39
0.1	$3.69 \cdot 10^{-2}$	$4.03 \cdot 10^{-2}$	$1.89 \cdot 10^{-2}$	$1.89 \cdot 10^{-2}$	2.14
0.3	0.217	0.231	0.144	0.144	1.60

The T1-aware decoder equals maximum likelihood, and the Pauli model underestimates the standard decoder's error by 6–11%. **Does it survive at circuit level?** In the notes (§64), with full syndrome extraction and d rounds, the gain is 2.2 times at $d = 3$ and 3.3 at $d = 5$ under pure T1, 1.6 and 2.0 when T1 dominates, and it grows with distance. **It is not free**: when depolarizing noise dominates, the reweighting costs 15% at $d = 3$. A qg witness (final mean readout over detector rate) chooses the decoder without labels: 1.3–2.6 where the T1-aware decoder wins and 0.34–0.54 where it loses, although the threshold of 1 was chosen after seeing the cases.

26.5 Qutrits and erasure qubits

Before extending qang beyond the qubit, the notes asked whether a three-level readout gives it an advantage (§70–§74), with criteria committed before each run:

- **If the leaked qubit keeps its value, the $|2\rangle$ flag is useless** (§70): all the gain was the T1 reweighting, and adding the flag made it worse. Recorded verdict: no qutrit advantage.
- **If the leak scrambles the qubit, it helps** (§70b): reweighting plus flag give 1.78 times over standard, more than the product of the two parts (1.55).
- **A rule for the flag** (§71): weighting it by where the leak came from (posterior $1/(1 + a)$) is never worse than standard, and the qutrit-plus-qg gain grows with distance (1.48 times at $d = 3$, 1.77 at $d = 5$). One of the four predictions failed, so it stands as a partly confirmed rule.
- **A real transmon** (§72, coherent three-level CZ): leakage comes only from $|11\rangle$, the asymmetry is zero and the flag is nearly useless; the qg reweighting is not (1.8–2.0 times).
- **Erasure qubits** (§73, §74): if the device heralds decays with probability h, naively combining the herald with the reweighting makes things worse (0.46 times at $h = 0.99$). The $(1 - h)$ rule, committed before running, fixes it: qg

still gives 1.4–1.9 times at $h = 0.25$–0.5 and fades towards 1 when heralding is nearly perfect.

26.6 Cutting circuits: where and at what cost

Cutting a two-qubit gate to run a large circuit in pieces multiplies the shots needed by γ^2. For a Pauli rotation such as RZZ(θ), in qg (notes §15.2):

$$\gamma = 1 + 2|\sin\theta| = 1 + 2\sqrt{1 - qg^2}, \qquad qg = \cos\theta.$$

Cutting is free at the poles ($qg = \pm 1$) and costs 9 times at the equator. A shortcut that circulates, $\gamma^2 = 1 + \sqrt{1 - qg^2}$, underestimates the cost by up to 4.5 times (2 instead of 9 at $qg = 0$).

Where to cut? An 8-qubit circuit with 12 RZZ couplings must run in blocks of at most 4 qubits (§26). The script compares three rules on 200 instances per regime, choosing among the 3795 partitions:

Table 26.3: Shot multiplier, geometric mean (notes §26, recomputed).

regime	fewest cuts	weakest bonds	qg rule (optimum)
Trotter (small angles)	7254	3838	3642
QAOA (large angles)	976	2174	753

With small angles, "cut the weakest bonds" nearly matches the qg rule. With large angles it is the worst rule: RZZ$(\theta \approx \pi)$ is almost a local $Z \otimes Z$ and nearly free to cut. **The reference tool already does this**: qiskit-addon-cutting's `find_cuts` reaches the optimum on all 400 of the notes' instances (it is not installed in this environment, so the script does not run it). What qg adds is the closed form: scoring the 3795 partitions takes 1–2 ms per instance. In the notes' end-to-end run, the qg choice lowers the error 2.9 times with shots proportional to the weights, and raises it with equal shots per subexperiment: the γ^2 law assumes importance sampling.

26.7 Randomized benchmarking, ZNE and qg_S as a signal

- **Randomized benchmarking** (§11.1): with per-gate depolarizing noise, the survival is exactly $qg_Z(m) = (1-p)^{m+1}$, the same for any Clifford sequence.

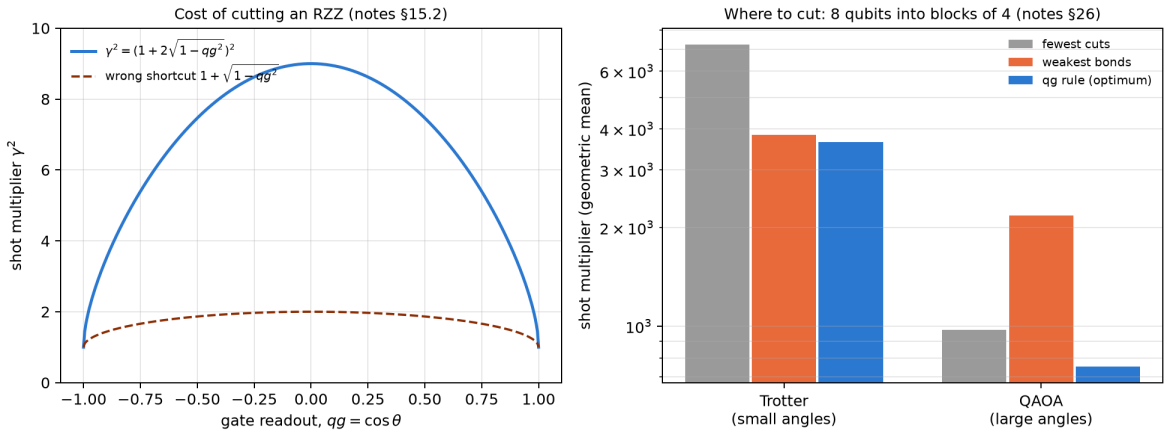


Figure 26.2: Left: shot multiplier of cutting an RZZ against its readout $qg = \cos\theta$, and the wrong shortcut. Right: cost of three rules for choosing the cuts in two angle regimes.

The script confirms it on four seeds (0.95, 0.7351 and 0.5688 at $m = 0$, 5 and 10).

- **ZNE: which space to extrapolate in** (§11.2): if the noise shrinks the readout, extrapolating in qg is exact (10^{-16}) and in the angle it is not (10^{-4}); if the noise rotates the angle, the other way round. Rule: extrapolate where the noise is linear, and the $\theta \leftrightarrow qg$ change makes trying both cheap.
- qg_S **as a signal** (§10): the output entropy follows the heavy-output probability without ideal simulation (correlation -0.926 under depolarizing noise), but under T1 it stops being monotone: on a fake brisbane model with a 200 µs idle it falls below the noiseless value, and would rank the most damaged point as the least noisy. With the register's mean readout alongside, which rises at every step, the pair is no longer fooled.

26.8 What remains open

The code-capacity results are verified; the circuit, qutrit and erasure results are from the notes' simulations with committed predictions. The destructive Leung-syndrome readout (§93, 24 four-qubit circuits) is ready to run on a real IBM device and has not been run.

26.9 Without qang and with qang

Table 26.4: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
choose a code	pick one fixed: 64–220% over the oracle	qg_Z and qg_X witnesses and the T_2/T_1 rule: 6.4%	**in favour**; the T_2/T_1 rule is equivalent, qg makes it measurable with two readouts
see T1 with one witness	—	qg_Z alone: always recommends no code	**against** if used alone: both components are needed
track drift	recalibrate every 10 windows: 20.4%	syndromes as qg readouts: 13.1% without probe shots	**in favour**
read dephasing from syndromes	arithmetic mean: 1.17–1.51 times high	product mean: 0.99–1.04	**in favour**: the right number is the one that limits the code
see T1 in the surface code	syndromes: a γ^w difference, invisible	register mean readout: equals γ	**in favour**
decode under T1	minimum weight	errors only on bits read 0: 2.1–2.5 times (capacity), 2.2–3.3 (circuit)	**in favour** when T1 dominates; **against** (15%) when depolarizing noise dominates
qutrits and erasure	flag erased blindly	flag weighted by its origin; $(1-h)$ rule	**in favour** with scrambling leaks or partial heralds; **none** with near-perfect heralds
choose cuts	count cuts or cut the weakest	$\gamma = 1+2\sqrt{1-qg^2}$ per gate	**in favour** against angle-blind rules (up to 2.9 times); **none** against qiskit-addon-cutting, which already does it

action	without qang	with qang	difference
extrapolate to zero noise	one fixed space	choose qg or θ by the noise	**in favour**, on idealized models
certify with the entropy	qg_S alone: fooled under T1	qg_S plus the mean readout	**in favour** of the pair

Balance. None of the codes or decoders is new, and the reference cutting tool already chooses the way qg does. What the qang scale brings here is measurable and concrete: the syndrome as a readout that inverts in one line to γ and p, a code-choice rule with its cost, a reading of T1 that the syndromes lack, a decoder that uses the fact that T1 leaves zeros, and the closed form of the cost of cutting. The gains depend on the regime, and the table says where they are lost.

26.10 In code

```
from qang import knitting as KN

KN.cut_sampling_overhead(0.0)   # 9: cutting an RZZ at the equator (qg = 0)
KN.cut_sampling_overhead(0.95)   # 2.64: cheap near the pole
KN.total_sampling_overhead([0.95, 0.2, -0.99])  # 38.0: product of the gamma^2
```

26.11 Summary

- Each syndrome is the qg readout of a parity; qg_Z of $|1\rangle$ and qg_X of $|+\rangle$ read T1 and dephasing.
- No repetition code corrects T1; the Leung code does. Rule: Leung if $T_2 > T_1$, no code if $0.4\,T_1 < T_2 < T_1$, phase code if $T_2 < 0.4\,T_1$.
- The Leung syndromes track drift without probe shots, and read the product mean of the dephasing.
- In the surface code the syndromes do not see T1 and the mean readout does; placing errors only on bits read 0 lowers the logical error 2–3 times when T1 dominates.
- With qutrits and erasure qubits, qg adds up if the flag is weighted by its origin and the herald by $(1 - h)$.
- Cutting an RZZ costs $(1 + 2\sqrt{1 - qg^2})^2$: free at the poles, 9 times at the equator.

Bibliographic notes

The first quantum code is Shor (1995), and the 4-qubit code for T1, Leung et al. (1997). The surface code is Kitaev (2003), Dennis et al. (2002) and Fowler et al. (2012); erasure conversion, Wu et al. (2022). Circuit cutting is Peng et al. (2020); randomized benchmarking, Magesan et al. (2011), and zero-noise extrapolation, Temme et al. (2017). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

27 Hardware characterization in qg

Before running any algorithm one has to know what the device is like: how long it takes to decay (T1), how long to lose its phase (T2), how hot it is and how often it misreads. This chapter reads those quantities on the Chapter 1 scale. Each one shows up as a qg readout, and the scale helps in three concrete cases: separating temperature from readout error, telling T1 from unital noise with the mean readout, and choosing the grid with which the electronics sets the angles. All three are recomputed with code. The results on IBM device models are those of the notes' committed runs.

27.1 From the base: an idle qubit is a readout

A real qubit is never completely cold: at rest it has a residual excited population p, and its equilibrium readout is $qg_{eq} = 1 - 2p$, in m-qg $m_{eq} = 1000(1 - 2p)$. For a qubit of frequency f in thermal equilibrium,

$$qg_{eq} = \tanh\left(\frac{hf}{2k_B T_{\text{eff}}}\right),$$

so the equilibrium readout **is** a thermometer. The instrument's readout is also affine in qg: if it confuses 0→1 with probability e_{01} and 1→0 with e_{10}, the measured readout is $a + b\,qg$, with $a = e_{10} - e_{01}$ and $b = 1 - e_{01} - e_{10}$ (notes §31). The problem is that in the Z basis only $a + b\,qg_{eq}$ is seen: temperature and readout error are mixed.

27.2 T1, T2, temperature and readout in one sweep

The standard calibration measures "$|0\rangle$" and "$|1\rangle$", which are really the thermal state and X applied to it, and counts the thermal population as readout error. The qg protocol adds a herald: measure, apply I, X or $R_y(\pi/2)$, wait and measure again. The pairs at zero delay give, in closed form, the mean m, the readout covariance $V = b^2(1 - qg_{eq}^2)$ and $C_X = a^2 - b^2$, from which a, b and qg_{eq} follow.

The delays add T1 (after X) and T2 (the decay of qg_X). The script repeats the notes' experiment with 40 repetitions of 16 circuits of 2000 shots (T1 = 100 µs, T2 = 70 µs, $e_{01} = 0.015$, $e_{10} = 0.04$, 5 GHz):

Table 27.1: Standard calibration against the heralded qg sweep (notes §31, recomputed).

true thermal population (temperature)	e_{01} standard	e_{01} qg	temperature standard	temperature qg
0 (0 mK)	0.0145	0.0149	0	11 ± 15 mK (bound)
0.01 (52 mK)	0.0240	**0.0149**	0	52.5 ± 1.0
0.03 (69 mK)	0.0418	**0.0149**	0	68.9 ± 0.9
0.08 (98 mK)	0.0915	**0.0153**	0	98.3 ± 0.8

- **The standard calibration reads temperature as readout error**: it overestimates e_{01} by exactly $p \cdot b$ and sees the qubit perfectly cold at any temperature.
- **The qg sweep is unbiased** in e_{01}, e_{10}, p and the temperature, and resolves it to 1 mK from 52 mK up (at $p = 0$ it only gives a bound).
- T1 and T2 come out unbiased in both, but the joint qg fit is more precise (T1 spread 1.5–2.1 µs against 2.8–3.1; T2, 2.2–2.5 against 3.1–4.8), partly because of the maximum-likelihood fit itself.

As IBM circuits (§86). With 32 one-qubit circuits with mid-circuit measurement, on three IBM device noise models, all three committed predictions held: T1 and T2 within 25% and readout errors within 0.01. The qg T2 sits closer to the reported value than the standard one, which comes out 10% high on two of them. What a simulated model cannot test is precisely the thermal population (Aer has none). The real run is pending and is the first one worth doing on hardware.

What a real chip already says (§168, §170). On IQM Garnet the readout was corrected with the two poles, on the same qubit: $e_0 = 0.0055$ and $e_1 = 0.029$ in §168, and 0.0030 and 0.0265 in §170. That is the standard calibration of the table above, so any thermal population the qubit had is counted inside e_0. The two numbers do not separate it, and the change of e_0 between runs does not say whether the readout or the temperature changed. The two runs also found something the affine readout model does not have: an error that depends on the angle. With the poles corrected, R_y read 16–55 m-qg low at the intermediate angles of §168 (over-rotations of $1.3°$ to $3.9°$), and the §170 eigenstate 26–71 m-qg low (Section G.6). Calibrating the readout is not enough: the angle of the gates has to be characterized too, and in m-qg that error shows without any model.

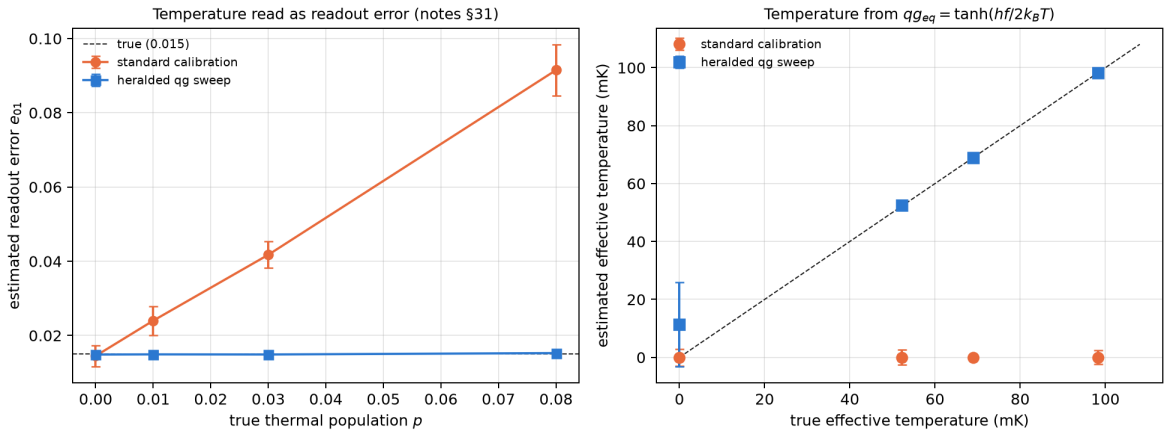


Figure 27.1: Left: estimated readout error e_{01} against the true thermal population; the standard calibration inflates it by $p \cdot b$, the qg sweep stays at the true value. Right: estimated effective temperature against the true one; the standard one is always 0.

27.3 Which noise is it? The mean readout against XEB and HOP

With N shots of a 4-qubit quantum-volume circuit, can one decide whether the noise is dissipative (T1) or unital (dephasing or depolarizing)? The notes (§18) set it up with 24 circuits, three channels and 10 strengths, and each measure is turned into a classifier with a single threshold fitted on 16 circuits and scored on 8 it has not seen. XEB and HOP need the ideal distribution (a classical simulation); the mean readout and qg_S only need counts. The script recomputes it:

Table 27.2: Accuracy on held-out circuits (0.67 = always answer 'unital').

noise per gate	shots	mean readout	qg_S	XEB	HOP
strong, 2–50%	100	**0.88**	0.72	0.70	0.71
strong, 2–50%	1000	**0.91**	0.72	0.74	0.73
strong, 2–50%	10,000	**0.91**	0.73	0.76	0.75
weak, 0.5–5%	100,000	0.67	0.67	0.67	0.67

- **With strong noise, the mean readout identifies T1** much better than the standard measures, from only 100 shots and without the ideal distribution: 0.88 with 100 shots and 0.91 with 1000 or more. Its accuracy grows with the strength (0.75 below 10% per gate, 0.98 above 25%). XEB and HOP measure how much noise there is, not which kind.

- **At realistic strengths everything is at chance**, even with 100,000 shots. The limit is not shot noise: the ideal mean readout already varies from circuit to circuit (spread 0.139, from −0.325 to +0.334), more than the T1 shift.

The practical consequence: the mean readout is a T1 detector on circuits whose ideal value is known, such as an idle wait before measurement, not a passive detector on any circuit.

27.4 Coherent errors and the always-on ZZ

A conditional-phase error (an extra ZZ) conserves the weight, so one might expect it to pass the weight filter unnoticed. The notes tested this twice with committed predictions, and both times the prediction failed in an informative way:

- **A coherent RZZ after every native gate** (§145): the kept fraction fell by 0.009–0.020 at $\varepsilon = 0.10$ rad, and the filtered error grew only 4% on the ECR devices but 62% on FakeTorino (CZ). The error acts inside the decomposition of the logical gate, where the qubits are not in the computational basis, and part of it leaves the sector.
- **Always-on ZZ between neighbours** (§150, 5 and 15 kHz): the filter removed 32–60% of the added error; the filtered error grew 14–38% at 15 kHz and stayed 4–5 times below raw. Short circuits suffer least.

On these devices coherent errors are filtered in proportion to the time the qubits spend outside the computational basis.

27.5 Control quantization: an even grid in the angle or in qg

The electronics sets each angle from a grid of 2^b values. The usual choice is a grid even in θ; a grid even in $qg = \cos\theta$ is even in the Born probability and, on the Bloch sphere, is made of equal-area bands (notes §17). The script recomputes it:

- **One qubit, random targets**: the qg grid's worst probability error is exactly $\pi/2$ times smaller (1.568, 1.570 and 1.570 with 4, 8 and 10 bits; it saves $\log_2(\pi/2) \approx 0.65$ bits). But its worst infidelity is worse, by a factor that grows like 2^b (from $2.7{\cdot}10^{-3}$ to $1.7{\cdot}10^{-2}$ with 4 bits), because it is coarse near the poles.

- **Loading a distribution** with a tree of R_y rotations (the state-preparation step of quantum Monte Carlo): in qg each rotation is set directly by a conditional probability, $qg = 2p - 1$. In the 45 cases (4, 6 and 8 qubits; 4, 6 and 8 bits; five distributions), the qg grid lowers the total variation distance in **43 of 45** (median 1.56 times, up to 2.9).
- For smooth distributions the infidelity is also lower with the qg grid (25 of 27, median 1.98 times); for sparse or sharply peaked ones the θ grid wins (15 of 18, median 3.3 times, up to 56).

Design rule: a grid even in qg when the task is to reproduce probabilities (sampling, Monte Carlo, smooth distributions); a grid even in θ when state fidelity matters and the targets sit near the poles.

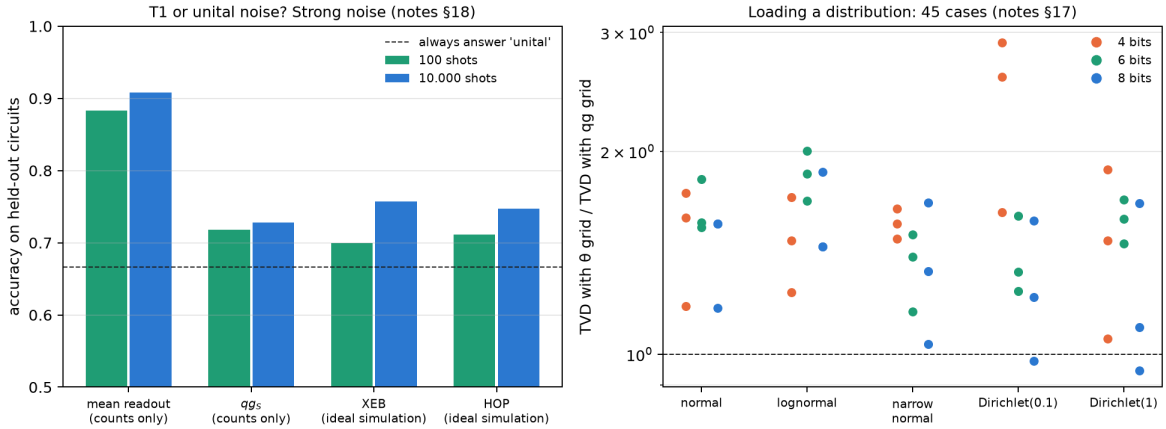


Figure 27.2: Left: accuracy in telling T1 from unital noise with four measures, at 100 and 10,000 shots and strong noise. Right: ratio between the total variation distance with a θ grid and with a qg grid in the 45 loading cases; above 1 the qg grid wins.

27.6 What remains open

The heralded sweep (§86, 32 one-qubit circuits) is ready for a real IBM device and is the first thing worth running: the thermal population, its central point, does not exist in the simulators. T1 detection at realistic strengths needs circuits with a known ideal value.

27.7 Without qang and with qang

Table 27.3: The chapter's actions without the qang scale and with it.

action	without qang	with qang	difference
calibrate the readout	two circuits, "$\|0\rangle$" and "$\|1\rangle$": e_{01} inflated by $p\,b$	heralded sweep: unbiased	**in favour**; separating by repeated measurement is known practice, qg adds the closed form
measure the temperature	the standard calibration always sees 0	T_{eff} from qg_{eq}, to 1 mK	**in favour**
measure T1 and T2	separate fits	joint qg fit: 1.5 times more precise	**in favour**, partly because of the joint fit
tell T1 from unital noise	XEB or HOP, with ideal simulation: 0.70–0.76	mean readout, counts only: 0.88–0.91	**in favour** with strong noise; **none** with realistic noise
anticipate coherent ZZ errors	"the filter does not see them"	the filter removes 32–60% of the always-on ZZ	**in favour**, although the committed prediction failed
quantize the control	grid even in θ	grid even in qg: probability $\pi/2$ times better, 43 of 45 loads better	**in favour** for probabilities; **against** for fidelity near the poles

Balance. This is the chapter where the qang scale contributes most on its own, without depending on the weight filter: the equilibrium readout is a thermometer, the instrument's readout is affine in qg and separates from the temperature in closed form, the mean readout tells dissipation from unital noise, and a grid even in qg is the right grid for probabilities. The limits are written down: with realistic noise the mean readout does not classify arbitrary circuits, and the qg grid loses in fidelity near the poles.

27.8 In code

```
import numpy as np

f = 5e9                                      # qubit frequency (Hz)
h, kB = 6.62607015e-34, 1.380649e-23
p = 0.03                                     # residual thermal population
qg_eq = 1 - 2 * p                           # equilibrium readout: 940 m-qg
T_eff = h * f / (2 * kB * np.arctanh(qg_eq))
T_eff * 1e3                                  # ~69 mK
```

27.9 Summary

- The equilibrium readout $qg_{eq} = \tanh(hf/2k_BT_{\text{eff}})$ is a thermometer, and the instrument's readout is affine in qg.
- The standard calibration counts temperature as readout error; the heralded sweep separates the two without bias and measures the temperature to 1 mK.
- The mean readout tells T1 from unital noise with 100 shots when the noise is strong (0.88 against 0.70–0.72); with realistic noise nothing manages it on arbitrary circuits.
- Coherent ZZ errors are not invisible to the filter: they are filtered according to the time spent outside the computational basis.
- A grid even in qg improves 43 of 45 distribution loads; one even in θ is better for fidelity near the poles.

Bibliographic notes

For T1, T2, read-out and ZZ coupling in superconducting qubits, see Krantz et al. (2019). XEB is Arute et al. (2019); HOP and quantum volume, Cross et al. (2019); randomized benchmarking, Magesan et al. (2011). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part X

X. Experimental evidence, real hardware and the falsification audit

28 Methodology: the pre-registered falsification protocol

The previous chapters quoted “predictions committed before running” again and again. This chapter explains what that means, how many held and how many failed, and something a book about a new scale owes its reader: **an audit of whether that “before” can be checked**. There is no new physics here. The chapter’s script simulates nothing: it inspects the repository itself, its git history and the files the runs stored.

28.1 From the base: a prediction is a readout that can be wrong

On the Chapter 1 scale, nearly every claim in the notes is a number in m-qg or a ratio between two errors. That the filter “helps” says nothing; that it lowers the readout error from 104 to 23 m-qg on three devices does, and it can fail. The discipline is to write that number and the threshold that would make it false **before** running, to pin it in a way that cannot be changed later, and to publish the outcome whatever it is.

28.2 Two ways of pinning a prediction

The notes used two:

- **By commit**: the predictions are written and pushed to the repository before the run, and the section cites that commit (for instance, “committed before the run (77700bd)”). It is the method of §71 to §165.
- **By sha256**: the script that does the run is frozen, the sha256 of its bytes is computed and stored in `examples/data/*.prereg` with a timestamp. If the script changes, the hash stops matching. It is the method of the neutrino, kernel, VarQTE, IQM and other sections.

28.3 The tally of hits and misses

What the code stored. 34 runs leave in their results file a dictionary of verdicts computed by the script itself. The chapter's script counts them: **148 predictions, 94 passed (64%) and 54 failed**. Three runs failed all their predictions (the IBM fleet, the coherent ZZ and collective neutrinos) and seven passed all of theirs, among them $E = mc^2$ on IQM Garnet (§170).

The QML line. For §75–§123, the repository's README keeps another tally: 176 predictions judged, 123 passed and 53 failed, with 4 more waiting for the IonQ hardware run. The two counts are separate bookkeeping that partly overlap, and must not be added.

What cannot be recounted automatically. The sections before §71 barely use pre-registration, and many later ones record their verdicts only in prose ("Four passed; V1 failed"), with varied wording. A full count over §1–§170 would have to be done by hand, section by section. This book does not do it: it gives the two counts that can be checked.

A 64% hit rate is not a weakness. A rate near 100% would signal thresholds that are too soft. Several failures were the most informative results in the notes: the coherent ZZ that the filter does see (Chapter 27), the erasure rule that had to be fixed with $(1 - h)$ (Chapter 26), the qg_YY metric that helps less than predicted (Chapter 23).

28.4 Audit of the "before"

That a prediction was written "before" has to be checkable. The script checks it with git:

- **By commit (54 citations, §71–§165)**: every cited commit exists in the history and, at that commit, the notes do not yet contain the section's results. **All 54 are consistent.**
- **By sha256 (16 files)**: no script matches its hash today, as expected, because after the run the findings are added to its opening comment. The question is whether any stored version matches.
  - In **8 of 16**, the hash matches the script at the commit that added the `.prereg`, and that commit is earlier than the results'. The "before" is **checked**.

- In **8 of 16**, the `.prereg` and the results came in the same commit, and no version of the script stored in git matches the hash: the script was frozen, run and edited before anything was pushed. For those eight the "before" rests **only on the file's own timestamp**, which the author could have written at any time.

The eight are: neutrino matter and oscillation, mass ordering and steriles, solar neutrinos, post-selection and CTCs, pulse area, the quantum kernel and VarQTE. Their results are not false because of this; their pre-registration cannot be checked by a third party. The fix for what comes next is simple: push the `.prereg` and the frozen script in a commit of their own before running.

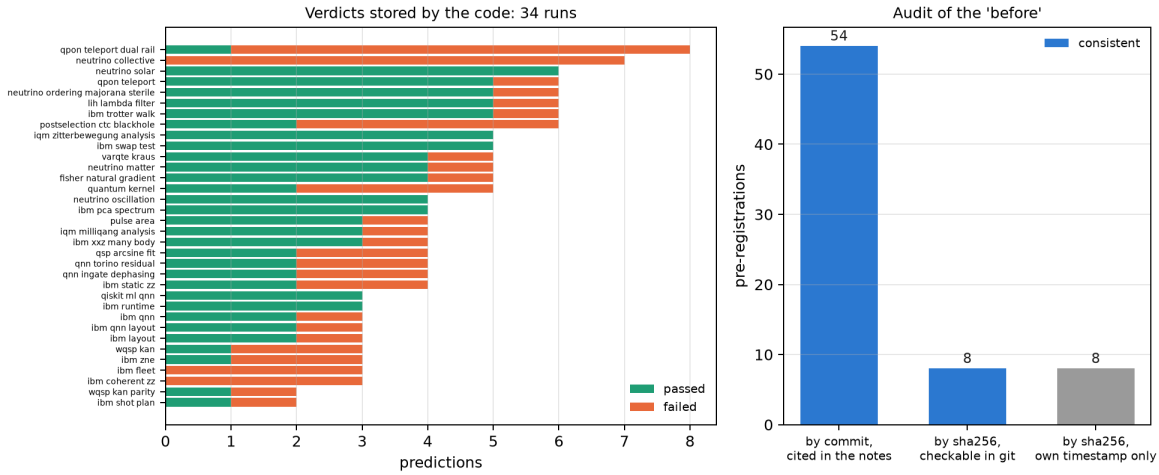


Figure 28.1: Left: predictions passed and failed in the 34 runs that store their verdicts. Right: audit of the "before": the 54 predictions pinned by commit are consistent; of the 16 sha256 pre-registrations, 8 can be checked in git and 8 only have their own timestamp.

28.5 Testbeds

Nearly all the book's results come from simulations with calibrated noise models, not from hardware:

- **IBM**: qiskit-ibm-runtime's models with real published calibrations (Brisbane, Sherbrooke, Torino, Fez and a fleet of 21), run locally in Aer or in IBM Runtime's local mode (§137, §144).

- **IonQ**: the cloud simulator with the Aria-1 and Forte-1 noise models, free (§20, §38). Without dominant T1: the mean readout does not move (+0.004 and +0.007, consistent with zero), the opposite of the IBM model.
- **IQM**: the simulated Apollo and Garnet models, used as a free rehearsal before every real run.

A calibrated noise model is not a device: it has no thermal population (Chapter 27), no drift, and none of the ways the vendor's software batches circuits. That is why the book calls "simulated" everything that did not run on hardware.

28.6 Real hardware

- **IonQ Forte-1** (§128): the predictions for H_2 with the filter were pinned by commit and the simulator rehearsal met them, but the vendor refused the first submission before accepting the job (the project did not have that QPU enabled). Nothing was charged and there is no data.
- **IQM Garnet, three-flavour neutrinos** (§165, 8 circuits of 300 shots): the script rescores the four checks from the stored run, and **all four hold**. The error falls from 0.0615 to 0.0386 (1.59 times) and 88% of the shots are kept. With the filter the survival minimum sits at 2.6 GeV (exact 2.63); without it, at 0.8 GeV.
- **IQM Garnet, the milli-qang** (§168, 26 circuits of 4000 shots on one qubit): three of four predictions held. The script recomputes the X-gate cost with a single slope: **1.64 m-qg per gate** ($\chi^2/\mathrm{dof} = 0.81$). Prediction M2 failed: the ten equator copies scatter half as much as binomial noise allows (ratio 0.50). It remains open.
- **IQM Garnet, $E = mc^2$ as a reading** (§170, 24 circuits of 2000 shots on the same qubit): the pre-registration was pushed in its own commit before the job, as the fix of Section 28.4 asked, which is why it is the eighth checkable sha256. **All five predictions hold** (Section G.6), with a 26–71 m-qg gate bias reported in the open.

The three IQM runs are the project's first real-hardware data, and Chapter 31 treats them in detail. A replication of §168 on five qubits and two days is prepared but not pushed to the repository, because at full size it would cost about 80 credits, and about 5 are left after §170 (estimated).

The $\mathrm{Var}_\theta(F_{\mathbf{echo}})$ criterion. The index announced a QPU falsification criterion based on the variance of the inversion echo over continuous trajectories. It is not defined anywhere in the repository, so the book does not use it. It remains an unwritten proposal.

28.7 Without qang and with qang

This chapter does not compare the qang scale with something else: it compares working without pre-registration with working with it.

Table 28.1: Working without pre-registration and with it.

action	without pre-registration	with pre-registration	difference
decide whether a result helps	decided after seeing the data	threshold fixed beforehand, in m-qg or in ratios	**in favour**: 54 of 148 predictions failed and were published
report failures	omitted or reinterpreted	reported with their numbers (94/54, 123/53)	**in favour**
check the “before”	impossible	by commit: 54 of 54 checked; by sha256: 8 of 16	**in favour**, with 8 cases unverified
rehearse before spending	—	free simulator with the same predictions before every QPU run	**in favour**: IonQ’s refusal cost nothing
pin the script by sha256	—	only checkable if the `.prereg` is pushed before the results	**none** in 8 cases: the procedure must be fixed

Balance. Pre-registration worked: it kept failures in view and could be checked in 61 of 69 cases. The other 8 show a defect of the procedure, not of the result, and already have a fix. The lesson holds for the whole book: a figure without its prior threshold is not evidence, and a prior threshold that cannot be checked is not entirely evidence either.

28.8 In code

```
import hashlib, subprocess
```

```
h, path = open("examples/data/iqm_neutrino_qg.prereg").read().split()[:2]
c = "a74c07c5"                              # commit that added the .prereg
frozen = subprocess.run(["git", "show", f"{c}:{path}"], capture_output=True).stdout
hashlib.sha256(frozen.replace(b"\r\n", b"\n")).hexdigest() == h  # True: this is the fro
```

28.9 Summary

- Predictions were pinned by commit (§71–§165) or by the sha256 of the frozen script.
- The 34 runs that store verdicts add up to 148 predictions: 94 passed and 54 failed; the QML line (§75–§123) records 123 and

  53. A full count of §1–§170 would have to be done by hand.

- The "before" is checked for the 54 predictions pinned by commit and for 8 of the 16 sha256 pre-registrations; for 8 there is only their own timestamp.
- Nearly everything is simulated with calibrated models. Real hardware is three runs on IQM Garnet: neutrinos (4 of 4), the milli-qang (3 of 4, with 1.64 m-qg per X gate) and $E = mc^2$ as a reading (5 of 5, §170).
- The $\mathrm{Var}_\theta(F_{\mathrm{echo}})$ criterion does not exist in the repository and is not used.

Bibliographic notes

The falsification criterion is that of Popper (1959), and pre-registration as a practice, Nosek et al. (2018). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

29 The confirmed empirical advantages

After twenty-eight chapters of "without qang and with qang" comparisons, this one gathers those that gave a measurable benefit, with their size, their limits and where each figure comes from. The bar for entry is strict: a gain measured in code or in a pre-registered run, against the same action without the scale or without the filter, with the case where it stops holding written next to it. Every factor is recomputed here, device by device, from the runs' stored files. Nothing is simulated again. What did not pass the bar goes to Chapter 30.

29.1 From the base: where the advantages come from

The book's advantages have two origins, and they should not be mixed:

- **The weight filter** (Chapters 20–25): if the circuit conserves the number of excitations, every shot with the wrong weight had an error and is discarded. It is symmetry verification, a known technique. The qang scale contributes the witness (the register's mean readout, whose ideal value is known before running) and the rule for when it pays (Chapter 21).
- **The scale itself** (Chapters 18, 23, 26, 27): readouts that separate things the raw readout mixes. Temperature from readout error, T1 from unital noise, thermal drift from an attack.

29.2 Exact restoration under T1

With the same T1 on every qubit, the filter returns exactly the noiseless readout in every sector (Chapter 21, checked to 10^{-12} m-qg). In classification, the notes (§80) measured it over 60 paired runs (5 seeds, 3 splits, 4 datasets), with predictions committed before running:

- with the weight-1 model, trained without noise and read under T1, qang recovers **+2.5 points** of accuracy (95% interval: +0.8 to +4.1), up to the exact accuracy of 0.952;

- with the weight-2 model it recovers **+16.5 points** (+11.8 to +21.2), from 0.757 to 0.922;
- with dephasing, +12.6 and +22.6 points;
- when training under the noise, the difference vanishes (+0.1 and −0.1 points): **the advantage is not having to train under the noise**, not an accuracy higher than achievable.

All five predictions held. The index's figures ("+2.5 and +16.5 over 60 runs") are in §80 as stated.

The gradient, too (§112). Under equal T1 the filter returns the noiseless gradient exactly; without the filter, at weight 1 the gradient is K times smaller. This is the noise-induced barren plateau, and the filter trades it for shots: the cost of resolving the sign of the gradient grows like $1/K$ instead of $1/K^2$ (Section 23.8). In the notes the filter needed fewer shots in 10 of 12 cells; at half filling with $K \leq 0.14$ it needed 1.3–2.2 times more (T5 failed).

29.3 Lower readout error on simulated devices

The script recomputes the error-reduction factor (raw over filter plus readout correction) on every device of every stored run:

Table 29.1: Raw error over error with filter plus readout, recomputed from the stored runs.

study	devices	min	median	max
10 weight-1 circuits (§137)	3	3.99	4.66	8.07
IBM fleet (§144)	17	2.45	3.12	6.11
PCA spectrum (§146)	21	1.57	4.63	7.67
XY Trotter, best error (§147)	21	1.59	2.04	2.82
inversion test (§149)	6	3.84	4.55	11.24
3-flavour neutrinos (§151)	3	2.97	3.43	3.83

study	devices	min	median	max
sterile neutrino (§163)	3	5.55	5.61	5.68
neutrinos on IQM Garnet (§165, real hardware)	1	—	**1.59**	—

- **The filter never made the error worse** in any of the 74 cases (the same simulated device appears in several studies). It cut it 1.6 to 11 times.
- **On real hardware it gained less** (1.59 times): the chip was better than its simulated model (raw error 0.062 against 0.126), so fewer shots left the sector. The remaining error is inside the sector, where the filter does not see.
- The IBM fleet, taking the better of filter alone and filter plus readout, gives 2.4 to 9.0 times (the notes: 2.4–8.9). The "2.7–3.6×" range the index announced is not in the notes: the figures are those of this table. For the §85 classifier the notes give 2.9–3.6 times on the readout error, with the same accuracy (0.933) with and without qang.

29.4 In-sector errors: the echo

What the filter does not see, the echo measures (Chapter 22). With the weight-1 network on 189 inputs and 1000 shots, on five simulated devices, the decisions that differ from the noiseless network are **15 without qang, 7 with the filter and 3 with the filter plus the echo**, over 945 readouts (§106). The echo removes 43 to 56% of the decision-value error the filter leaves. It does not help with narrow-margin inputs, when the filtered error is already close to its shot noise (§108).

29.5 Characterization and control

These advantages belong to the scale itself, without the filter (Chapter 27):

- **Temperature and readout**: the heralded sweep measures the readout error without bias and the effective temperature to 1 mK, where the standard calibration inflates the error by $p \cdot b$ and sees the qubit at 0 mK.

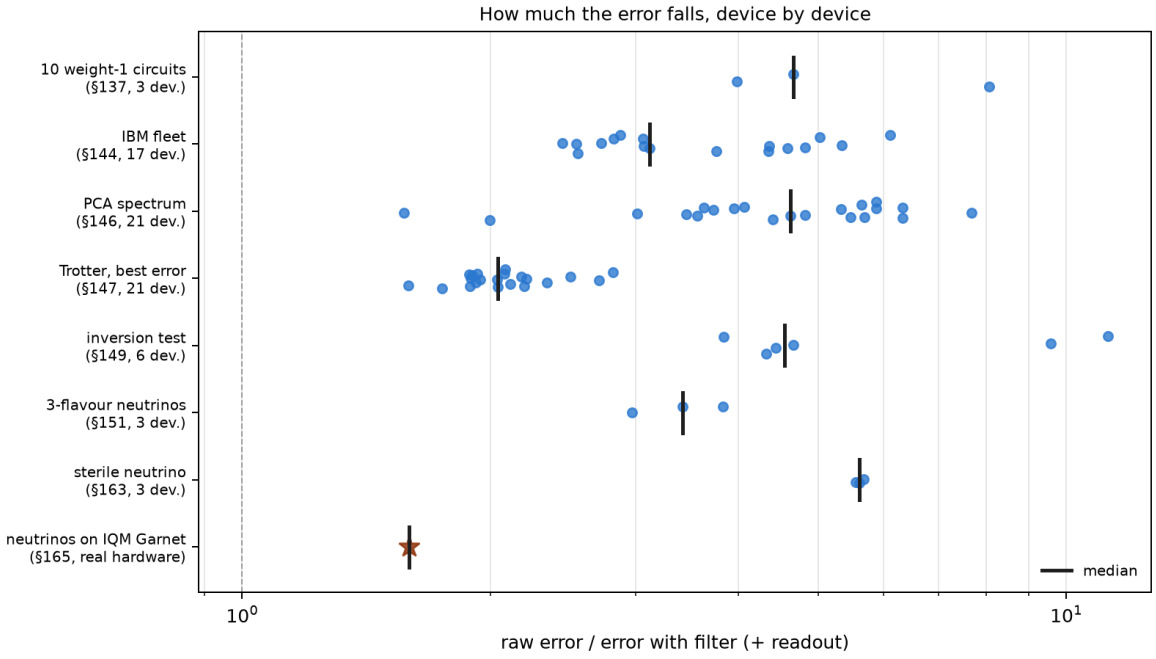


Figure 29.1: Error-reduction factor with the filter, device by device, in seven studies on simulated noise models and in the real run on IQM Garnet (star). The bars are medians.

- **T1 diagnosis without classical simulation**: the mean readout tells T1 from unital noise with 88% accuracy from 100 shots under strong noise, against 70–76% for XEB and HOP, which also need the ideal distribution. With realistic noise nothing tells the two apart on arbitrary circuits.
- **Control quantization**: a grid even in qg lowers the total variation distance when loading a distribution in 43 of 45 cases (median 1.56 times), and the worst probability error by an exact factor $\pi/2$. It loses in fidelity near the poles.

29.6 Neutrinos: mitigation without extra circuits

Neutrino oscillations in a circuit conserve the weight, so the filter applies without adding any circuit. With three flavours (§151), the error falls 3.4 times in the median of three simulated devices (the index said 3 times); with a sterile neutrino (§163), 5.6 times, and the survival dip appears at the right energy on all three. **On IQM Garnet** (§165), on real hardware, the filter cuts the error 1.59 times and puts the oscillation maximum back at 2.6 GeV (exact 2.63), where the raw readout placed it at 0.8 GeV. It is the only one of these advantages confirmed on hardware. The other two Garnet runs (§168 and §170, 8 of 9 predictions) tested the scale and the reading, not an advantage: they show that the milli-qang measures a real chip, not that qang improves it.

29.7 BB84: fewer false alarms under thermal drift

In BB84, the sum $qg(0) + qg(1)$ of the states sent in Z is a T1 witness: relaxation produces only 1→0 errors, and an intercept-resend attack produces symmetric errors (notes §40, Chapter 17). With windows of 2000 bits and 1% false alarms on the baseline:

Table 29.2: Fraction of windows with an alarm (notes §40, from the Chapter 17 cache).

scenario	QBER monitor	qg monitor
baseline	0.007	0.006
T1 drift, γ 0.02 → 0.04	0.478	**0.009**
T1 drift, γ 0.02 → 0.06	0.962	**0.006**
intercept-resend, $f = 0.05$	0.861	0.823
intercept-resend, $f = 0.10$	1.000	0.994
T1-mimicking attack	**0.962**	0.006

The qg monitor detects the attack as well as the QBER monitor and **does not fire when only the memory's T1 worsens**, where the QBER monitor fires in 48–96% of windows. The price is written down: an attacker who mimics T1 goes unnoticed by the qg monitor and not by the QBER one. The two are complementary, and there is no security gain: the key rate falls the same way in both cases.

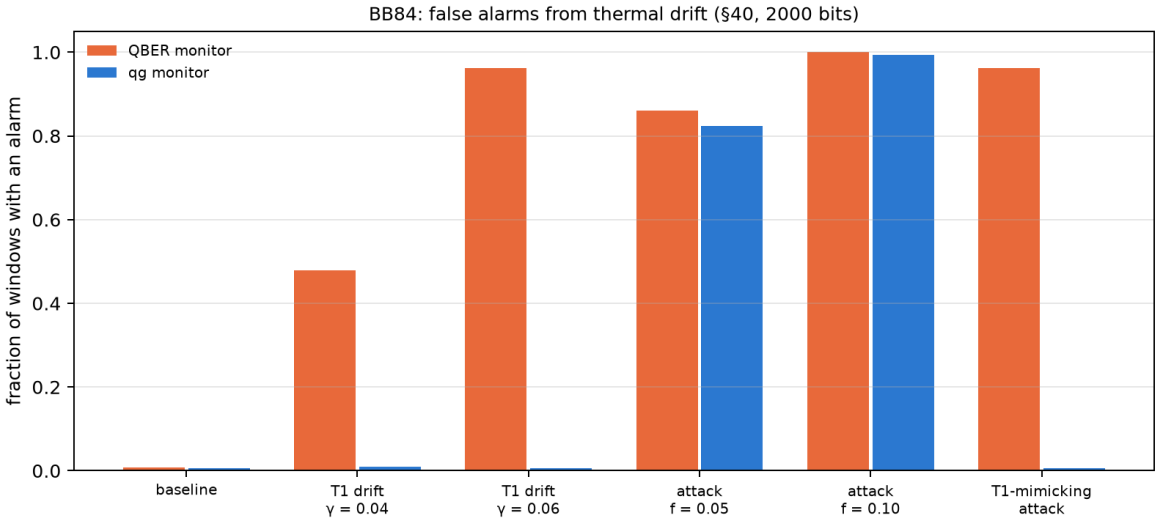


Figure 29.2: Fraction of windows with an alarm for the QBER monitor and the qg monitor in six BB84 scenarios with 2000 bits; the qg monitor does not fire under thermal drift, nor under an attack that mimics it.

29.8 What remains open

Of all these advantages, only the neutrino one is confirmed on real hardware, once, with 300 shots per circuit. The ones most worth taking to a device are the heralded characterization (thermal population does not exist in simulators) and the filter on a circuit with several excitations.

29.9 Without qang and with qang

Table 29.3: The confirmed advantages and their limits.

advantage	without qang	with qang	size	where it stops holding
accuracy under T1 without retraining	0.928 / 0.757 (weight 1 / 2)	0.952 / 0.922	+2.5 / +16.5 points	with unequal T1 it loses ≤ 1 point; training under noise, none
readout error (simulated)	raw	filter + readout	1.6 to 11 times, never worse	error inside the sector
error on real hardware	0.062	0.039	1.59 times	chips better than their model gain less
gradient under T1	K times smaller	exact with the filter	shots from $1/K^2$ to $1/K$; 10 of 12 cells	half filling with $K \leq 0.14$ (§112)
a classifier's decisions	15 of 945 differ	3 with filter and echo	5 times fewer	narrow margin (§108)
readout and temperature	e_{01} inflated, 0 mK	unbiased, to 1 mK	—	needs mid-circuit measurement
T1 diagnosis	0.70–0.76 with ideal simulation	0.88 from counts only	+12–18 points	realistic noise: chance
loading distributions	grid in θ	grid in qg	43 of 45, median 1.56 times	fidelity near the poles
BB84 alarms	48–96% false alarms from drift	0.6–0.9%	—	an attack that mimics T1

Balance. The advantages are real and measured, but of two kinds. Those of the filter are the largest (up to 11 times in simulation) and belong to symmetry verification, which is known; qang contributes the witness and the rule. Those of the scale itself are smaller and more its own: separating temperature from readout, T1 from unital noise and drift from attack. Only one advantage has been through real hardware, and there it gained less than in simulation. That is the figure to look at first.

29.10 In code

```
import json

d = json.load(open("examples/data/iqm_neutrino_qg_run.json"))
d["raw_error"] / d["filtered_error"]    # 1.59: the filter on IQM Garnet, real hardware
d["kept"]                                  # 0.88 of the shots kept
```

29.11 Summary

- Under equal T1, qang recovers the noiseless accuracy without retraining: +2.5 points with weight 1 and +16.5 with weight 2 over 60 runs.
- The filter cuts the readout error 1.6 to 11 times in 74 simulated-device cases, never making it worse; on real hardware, 1.59 times.
- Against the T1 noise plateau, the filter returns the exact gradient and lowers the shot cost from $1/K^2$ to $1/K$.
- The echo brings wrong decisions from 15 to 3 over 945 readouts.
- The scale itself separates temperature from readout error, T1 from unital noise, and thermal drift from an attack in BB84.
- Each advantage has written down where it stops holding; Chapter 30 gathers the negative results.

Bibliographic notes

The restoration under T1 is a symmetry verification (Bonet-Monroig et al. 2018; McArdle et al. 2019); the echo, Hahn (1950). The neutrino results use the parameters of Esteban et al. (2020), and the BB84 monitor, the protocol of Bennett and Brassard (2014) and the proof of Shor and Preskill (2000). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

30 Structural limits, equivalences and negative results

Chapter 29 gathered what gave a measurable benefit. This one gathers what did not: the blind spots that come from the definition of the scale itself, the equivalences that were offered as advantages, and the predictions that failed. A limit gets in if it can be shown in code or is in a stored run, with the same action without qang next to it. The new checks are in `tools/make_cap30_figures.py`; the rest is quoted from the notes and from the chapters where it was measured.

30.1 From the base: what the scale cannot see

The equation of Chapter 1 makes qg a function of a single probability, p_0, measured in the Z basis. Two kinds of limits follow:

- **What lives outside the Z populations.** Relative phases and coherences do not change p_0, so no Z-basis qg readout sees them. They only show up when measuring in another basis, and that is no longer the Z-basis qg readout. The angular qga scale of Chapter 1 (Section 1.4) gives them a unit, thousandths of a turn, but does not make them visible in Z: it changes how a phase is written, not what the Z basis measures.
- **What the filter cannot reach.** The filter discards shots outside the right weight sector. An error that leaves the shot in the sector, or a circuit that does not conserve the weight, is out of its reach.

Neither is an implementation defect: both follow from the definition.

30.2 Graph states: entanglement invisible in Z

A graph state is $H^{\otimes n}$ followed by one CZ per edge. CZ is diagonal, so it only changes phases: every Z population stays at 2^{-n}. The script checks it on the 6-qubit ring:

- its Z distribution matches that of $|+\rangle^{\otimes 6}$, which has no entanglement, with a difference of **exactly 0**;
- so $qg_Z = 0$ on every qubit, $qg_S = 1$ and the qg correlation is 0 for both states;
- the stabilizers $X_i Z_{i-1} Z_{i+1}$ are **1** on the graph state and **0** on the product. They tell the two states apart, but they are measured in the X basis.

The same happens without qang: any Z-basis statistic, raw or tomographic, is blind to this entanglement. qang inherits the limit, it does not create it (notes §8.1). Readout dephasing is the same blind spot seen from the noise side (§10.4).

30.3 qg_S is not monotonic under T1

Under depolarizing noise, qg_S (the entropy of the Z distribution divided by n) rises towards 1 as the noise grows, so it can rank devices. Under amplitude damping (T1) it cannot: it first rises, because the noise mixes, and then falls, because T1 pushes everything to $|0 \dots 0\rangle$, which has entropy 0.

The script runs 24 four-qubit quantum-volume circuits with amplitude damping on every qubit after every two-qubit block, with exact density matrices:

- qg_S **rises and falls in all 24.** In circuit 0 it starts at 0.919, reaches 0.981 at $\gamma = 0.06$ and comes back to 0.923 at $\gamma = 0.28$. On its own it cannot tell an almost noiseless device from one with a lot of T1.
- **The pair (qg_S, mean qg_Z) can.** At those two points the mean qg_Z is -0.029 and $+0.209$. T1 pushes the mean qg_Z towards +1 and depolarizing noise leaves it at 0 (it goes from -0.029 to 0.000).
- **The fix is not a theorem.** In this model the mean qg_Z rises monotonically in 21 of the 24 circuits. The notes, with their mid-circuit T1 model, found it in 23 of 24 (§10.4 D). The others have small downward steps.

30.4 The natural-gradient equivalence

Chapter 23 measured it, and here it is written down as a limit. A change of coordinate does not speed up an optimizer. The natural gradient written in qg is plain descent in θ, up to terms of order η^2 (notes §15.1). Claims that it "converges in qg where θ does not" came from clipping θ to $[0, \pi]$. That hides the minima at $\theta < 0$, and without the clip descent in θ converges.

In the pre-registered run of §159, with 20 seeds of the 4-qubit Ising model:

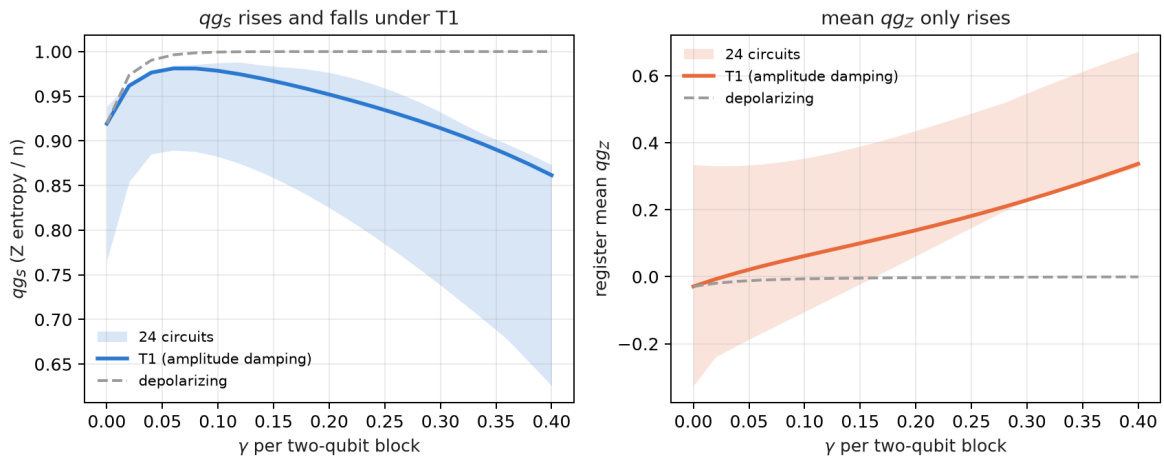


Figure 30.1: Left: under T1 (blue, band over 24 circuits) qg_S rises and falls back; under depolarizing noise (grey) it only rises. Right: the mean qg_Z rises under T1 and stays at 0 under depolarizing noise.

- descent in θ got there in 103 steps (median);
- descent in qg never got there;
- the natural gradient with the full metric took 22.5 steps.

That gain belongs to the natural gradient, not to qang: the full metric is not read in qg. The version that is read from qg correlators (qg_{YY} blocks) took 79.5 steps. It improved 1.3 times on θ, not the 1.5 that had been predicted (G2 failed).

30.5 The filter does not reach deep circuits

In Chapter 25 the electron-number filter cut the H_2 error three times and did nothing for LiH. The notes (§50) took the causes apart, and the main one turned out to be **reach**: the filter acts on the shots measured in Z, and in LiH the error is in the groups measured in X and Y.

Table 30.1: Energy error with readout correction and 200,000 shots (notes §50).

case	two-qubit gates	Z-group error, raw → filtered	X/Y groups	total raw → with filter (mHa)
H_2, all-to-all	3	13.0 → 4.9	0.9	13.9 → 5.9

case	two-qubit gates	Z-group error, raw → filtered	X/Y groups	total raw → with filter (mHa)
LiH, all-to-all, 1 layer	20	−2.2 → −1.9	26.2	24.0 → 24.3
LiH, all-to-all, 3 layers	60	−20.8 → −12.4	97.4	76.6 → 85.0
LiH, heavy-hex, 3 layers	60	−30.5 → −18.0	165.8	135.3 → 147.9

- **It is not the routing.** The LiH ansatz is nearest-neighbour and needs the same 60 gates without routing.
- **The qg remedy for the X/Y groups fails.** Rescaling them by the depolarizing strength read from the kept fraction overcorrects, by 2 to 3 times (from +97 to −124 mHa). Local gate noise is not global depolarizing noise.
- **Hubbard hits the same wall** (§29). On fake_brisbane, from 4 Trotter steps on (188 ECR or more, 18 % of the shots kept), every method sits on the noise floor: the double occupancy tends to 1/4 and the charge imbalance to 0.

For larger molecules the X/Y groups grow in number and carry most of the correlation energy. A Z filter reaches a shrinking share of the error.

30.6 The weight-1 network is classical

The filter protects a weight-conserving quantum neural network exactly against T1 (Chapter 20). The same reason makes it simulable. In the weight-1 sector the network is an orthogonal matrix O acting on the loaded data v, and each readout is $qg_Z^{(i)} = 1 - 2\,(Ov)_i^2$. It is a **quadratic form** in v, computable in $O(n^2)$ (notes §76, F2). The script checks it on 5-qubit circuits with a unary loader and a random RBS network: the largest error is $7.8 \cdot 10^{-16}$.

- §75: the best classical model ties or beats the best quantum network on all four datasets (on digits, an RBF-kernel SVM 0.963 against 0.933). There is no quantum advantage.
- §76: prediction Q3 said the network would not beat its quadratic twin by more than 1 point, and it failed, but because of the baseline. The twin used the default regularization. Rerun after the result with $C = 100$, it sits 0.3 to 1.6 points from the network.

- With weight 2 the output is quartic and no longer trivial. Even so, the ceiling of the encoding is bounded by a quadratic form on the products of amplitudes (§87).

The filter gives small networks robustness on noisy hardware. It does not give supremacy.

30.7 Where the filter does not apply: phases and classical encryption

The filter needs every gate to conserve the Hamming weight. qang's analysis pass (`qang.ibm.WeightPreservationCheck`) decides it gate by gate before running:

Table 30.2: What qang's analysis pass decides.

circuit	conserves the weight?	gates that break it
4-qubit QFT	no	4
QPE of the T gate, 3 counting qubits	no	6
4-qubit GHZ	no	4
4-qubit quantum volume	no	8
three-flavour neutrinos (Chapter 9)	yes	0
unary loader + RBS network	yes	0

Phase estimation and everything built on it (Shor, HHL) start with Hadamards, which mix every weight. There the scale gives a closed-form reading of the phase, $qg_\Phi = e^{i2\pi\varphi}$, invertible on its whole domain (notes §13). That is the reading QPE computes with a circuit, but it does not lower QPE's error. Post-quantum cryptography (Chapter 16) is classical arithmetic: there are no qubits or operators, so qang has nothing to read.

30.8 The failed predictions of §157–§164

In the eight pre-registered runs of §157–§164, **26 of 45 predictions passed**, by the verdicts the code stored. The 19 that failed:

- **Teleportation over a passive optical network** (§157, Q5): 5 of 24 cases far from the optimum, all where the ideal fidelity is 1 (measured 0.897 to 0.957). The circuit does not conserve the weight, so the only filter is the loss herald.

- **Dual rail** (§158, seven failed): the fidelity read from three correlators predicts the teleported one within 0.029, but does not decompose it (A1–A3). False heralds set a floor (B1, B2, B5).
- **Natural gradient** (§159, G2): 1.3 times against the 1.5 predicted.
- **Pulse area** (§160, P4): the period-π wrap of RZZ saves 1.14 times in error against the 1.5 predicted.
- **Kernels** (§161, K2–K4): accuracy hardly changes with the repair and target alignment does not see the shrink; K4 failed on its thresholds, not on its direction.
- **VarQTE** (§162, V1): the ansatz is not expressive enough at late times (noiseless fidelity 0.934 at $t = 2$).
- **Neutrinos** (§163, O3): the filtered values sit 0.010 to 0.033 high, because of the error inside the sector.
- **Time loops and black holes** (§164, four failed). On FakeTorino the paradox is not fully suppressed: the success at $\varphi = \pi$ is 0.056 where the prediction asked for close to 0 (C2), and the controlled rotation gives 0.600 where the exact value is 1 (C3). The decoder worked on all three devices (H2, H3); H1 missed by 0.001.

The three real-hardware runs that came afterwards, on IQM Garnet, passed 12 of 13 predictions: the neutrinos of §165 (4 of 4), the milli-qang of §168 (3 of 4) and $E = mc^2$ as a reading in §170 (5 of 5). The failed one is M2 of §168: the repeated equator copies scattered half as much as shot noise allows, and it remains unexplained (Chapter 28).

30.9 Consolidated list of limitations

The notes' list (§14), with what the book added:

1. **Jacobian singularity** at $\theta \in \{0, \pi\}$: regularizing it costs exactness near the poles (Chapter 23).
2. **arccos range trap**: qg-space updates cannot reach $\theta < 0$; on H_2 they lose all of the correlation energy (§8.2).
3. **Z-basis blindness**: graph-state entanglement and readout dephasing are invisible to qg_Z, qg_S and the qg correlation.
4. **The qg correlation is not an entanglement measure**: it is positive for classical mixtures and 0 for graph states (§7.2).
5. qg_S **is not monotonic under T1**: read it together with the mean qg_Z, a pair that is monotone in 21 to 23 of 24 circuits, not in all of them.
6. **Finite-shot bias of** qg_S: corrected with Miller–Madow (§10.2).

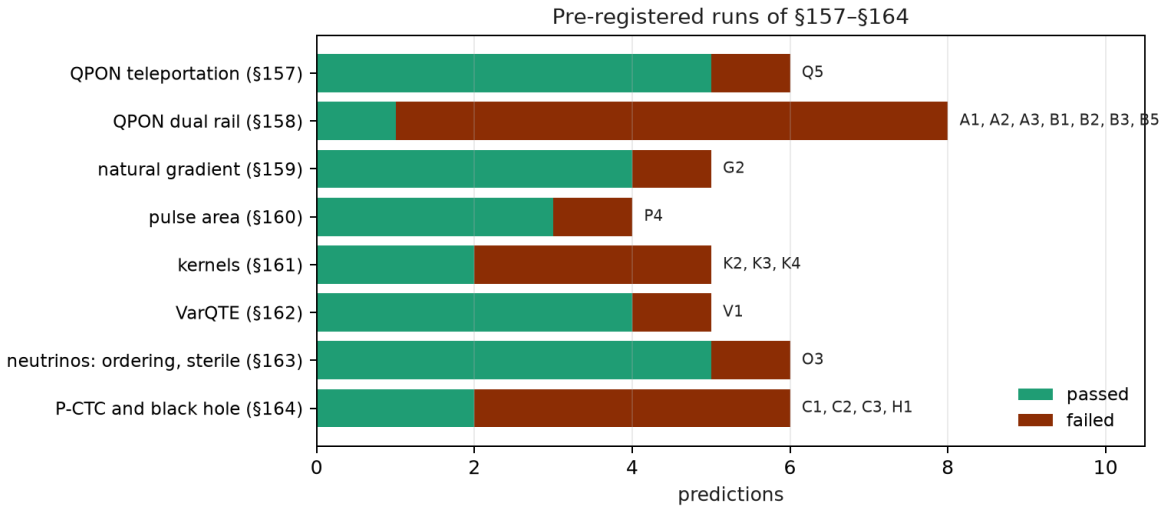


Figure 30.2: Passed and failed predictions in the eight pre-registered runs of §157–§164, read from the stored verdicts; on the right, the ones that failed.

7. **Pole damping** costs 3 to 20 times more iterations when the optimum sits on a pole (§9).
8. **No speed or energy advantage**: qg does not change how many shots a given precision needs (§14, item 9; §15.3).
9. **Reach of the filter**: it only acts on weight-conserving circuits and on what is measured in Z; it does not see the error inside the sector.
10. **The networks that admit the filter are simulable** at weight 1.
11. **On real hardware the filter gained less than in simulation** (1.59 times against 2.2 on IQM's model, Chapter 31).
12. **The correction from the poles does not see the angle-dependent error**: on Garnet, R_y read 16–71 m-qg low at intermediate angles, with the readout already corrected (§168, §170; Chapter 27).
13. **Where qang avoids a barren plateau, the problem is simulable**: the qg local cost with short circuits and the weight-1 sector. Against the T1 noise plateau, the filter returns the exact gradient but pays in shots ($1/K$) (Section 23.8).

30.10 Without qang and with qang

Table 30.3: The limits, with and without the scale.

limit	without qang	with qang	difference
graph states	Z readout blind	$qg_Z = 0$, $qg_S = 1$, same as the product	none: the limit belongs to the Z basis
T1 against unital noise	Z entropy, not monotonic	pair (qg_S, mean qg_Z)	separates the two in 21–23 of 24 circuits
optimization	descent in θ, 103 steps	natural gradient in qg = descent in θ	no speed-up of its own
LiH, 60 gates	76.6 mHa	85.0 mHa with the filter	the filter makes it worse: the error is in X/Y
weight-1 network	classical models	exact quadratic form	no quantum advantage
QPE, Shor, HHL	phase estimation	closed-form qg_Φ, filter not applicable	a new reading, no mitigation
§157–§164	—	26 of 45 predictions	19 failed, published

Balance. Almost every limit has the same root: qg reads a Z-basis probability, and what is not in that probability it does not see. The negative results are not hidden: 19 of the 45 predictions of §157–§164 failed and stay written down. Several "advantages" offered as qang's own (the natural gradient, the supremacy of the weight-1 network) turned out to be equivalences.

30.11 In code

```
import math
from qiskit.transpiler import PassManager
from qang import ibm, neutrino as nu

pm = PassManager([ibm.WeightPreservationCheck()])
pm.run(nu.three_flavour_circuit("mu", 1300.0, 2.6, -math.pi / 2))
pm.property_set["qg_weight_preserving"]      # True: the filter applies
```

30.12 Summary

- The entanglement of a graph state is invisible in Z: its distribution matches that of $|+\rangle^{\otimes n}$ exactly.
- qg_S rises and falls under T1 in all 24 circuits tested; the pair with the mean qg_Z separates them, without being a theorem.
- The natural gradient in qg is descent in θ: a change of coordinate does not speed things up.
- In LiH the error is in the X/Y groups, out of the filter's reach; Hubbard hits the same wall from 4 steps on.
- The weight-1 network the filter protects is a classical quadratic form.
- In the runs of §157–§164, 19 of 45 predictions failed; in the three real-hardware runs (§165–§170), 1 of 13.
- Where qang avoids a barren plateau, the problem is simulable; against the noise plateau, the filter pays in shots.

Bibliographic notes

Graph states are Hein et al. (2004); quantum volume, Cross et al. (2019). The natural-gradient equivalence is with Stokes et al. (2020), and the classical simulability of plateau-free models, Cerezo et al. (2023). QFT and QPE follow Nielsen and Chuang (2010). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

31 First real-hardware experiment: IQM Resonance

Up to here, almost the whole book was measured on simulators with device noise models. This chapter tells how qang was taken to a real superconducting chip, IQM's Garnet (20 qubits), on trial credits: why that machine, how it was rehearsed for free, what it cost and what came out. The figures are recomputed from the runs' stored files with `tools/make_cap31_figures.py`. Nothing is sent to IQM, and checking them does not need `iqm-client`.

31.1 From the base: what goes to the chip

The readout of a real qubit is the same as in Chapter 1: zeros and ones are counted, and $qg_Z = 2p_0 - 1$. What changes is everything around that count: the circuit is compiled to the chip's native gates, the readout has errors of its own and every shot costs credits. qang brings three things, all in post-processing:

- the **weight filter**, if the circuit conserves the weight;
- the **witness** that says how much the filter removed (the kept fraction);
- the **readout in m-qg**, where a probability difference of 0.0005 is 1 m-qg.

31.2 Why IQM, and what compiling costs

Garnet's native gates are PRX (a single-qubit rotation, `r` in Qiskit) and CZ, on a square lattice. The chosen experiment (three-flavour neutrinos) uses XX+YY gates, which conserve the weight. The script compiles on a line of three qubits to the three native sets:

Table 31.1: Two-qubit gates after compilation (`seed_transpiler` 137).

native set	RBS	XX+YY	neutrino circuit, level 1	level 3
IQM (r, cz)	2	2	18	12

native set	RBS	XX+YY	neutrino circuit, level 1	level 3
IBM Heron (cz, rz, sx, x)	2	2	18	12
IBM Eagle (ecr, rz, sx, x)	2	2	18	12

The notes (§165) counted 12 CZ on IQM against 18 two-qubit gates on IBM. **The difference is not the chip's**: the IQM rehearsal was compiled at optimization level 3 and the IBM runs of §151 at level 1. At the same level the three sets give the same: the circuit's six XX+YY gates cost two two-qubit gates each. IQM was chosen for the available credits, not for a cheaper compilation.

31.3 The `qang.iqm` module

The index announced a module `qang/iqm.py`, the analogue of `qang/ibm.py`. The two runs of this chapter were done without it: the scripts `examples/iqm_neutrino_qg.py` and `examples/iqm_milliqang_qg.py` read the counts with `qang.interop` (`canonical_counts`, `qg_z_profile`, `filtered_qg_z_profile`, `weight_filter`). The module was written afterwards, with what those runs left pending:

- **Readout correction.** §165 had none: IQM's simulated models do not publish their readout errors in the Qiskit `target`, where `qang.ibm` looks for them, but in an error profile of their own. `qang.iqm.readout_errors` reads them there. For a real device, `pole_circuits` and `readout_from_poles` measure them with two circuits, every qubit in $|0\rangle$ and every qubit in $|1\rangle$, as §168 did. On the stored counts of §168 it gives $e_0 = 0.0055$ and $e_1 = 0.029$. `iqm_qg` applies the inversion of `qang.ibm` and gives the four readings: raw, corrected, filtered, and filtered after correction.
- **Chain choice.** `best_chain` compares the default layout with chains of the square lattice and keeps the one with the highest predicted kept fraction (`predict_kept_fraction`, with IQM's gates and units).
- **The weight-preservation pass.** IQM's native single-qubit gate does not conserve the weight gate by gate, so `check_weight` checks the circuit before compiling.

The free rehearsal on the Garnet model (`examples/iqm_readout_chain_qg.py`, the eight circuits of §165, 4000 shots) gives a raw error of 0.124, filtered 0.055 and **filtered after readout correction 0.047**: the correction §165 lacked removes another 15%. The simulated kept fraction is 0.689 and the prediction brackets it between 0.640 and 0.775, less tightly than on IBM. The chain choice barely

tells chains apart (0.777 against 0.775), because the model's qubits are nearly identical. None of this has been tried on the real chip yet.

31.4 A free rehearsal before spending

IQM publishes simulated models of its chips. The neutrino circuit was run there first, with 4000 shots and 8 energies (docstring of `examples/iqm_neutrino_qg.py`):

Table 31.2: Rehearsal on IQM's simulated models (2026-10-20).

model	close to	raw error	filtered error	factor	kept
IQMFakeApollo	Garnet	0.1263	0.0584	2.16	0.688
IQMFakeAphrodite	Emerald	0.1161	0.0515	2.25	0.708

The milli-qang rehearsal (§168) also passed its four predictions on both models before running on the chip.

31.5 The budget: planned and spent

The account started with 30 credits. The plan was to spend 10 % on probes, 60 % on the main experiment and keep 30 % in reserve. The token is read only from the environment variable `IQM_TOKEN`, and no script spends credits without the `--yes-i-accept-qpu-cost` flag.

Table 31.3: IQM Resonance credits. The cost of the neutrino run is not written down: it is the difference of two balances. Neither is that of §170: it is the notes' estimate at the §168 rate.

step	shots	credits	balance
neutrino probe	100	0.50	29.50
neutrino run	2400	0.50	29.00
milli-qang probe	100	0.50	28.50
milli-qang run	104,000	16.00	12.50
$E = mc^2$ run (§170, Appendix G)	48,000	≈ 7.7 (estimated)	≈ 4.8

- **Spent up to §168:** probes 3.3 %; runs 55.0 %; 12.50 credits were left (41.7 %). The plan held with room to spare. The §170 run (Section G.6) spent about 7.7 more, outside the chapter's plan, and about 4.8 are left.

- **The cost model was corrected with data.** The first estimate, 0.50 credits per job, came from small jobs and was wrong. The cost grows with the shots (about 0.15 credits per 1000) and has a floor of about 0.50 per job.

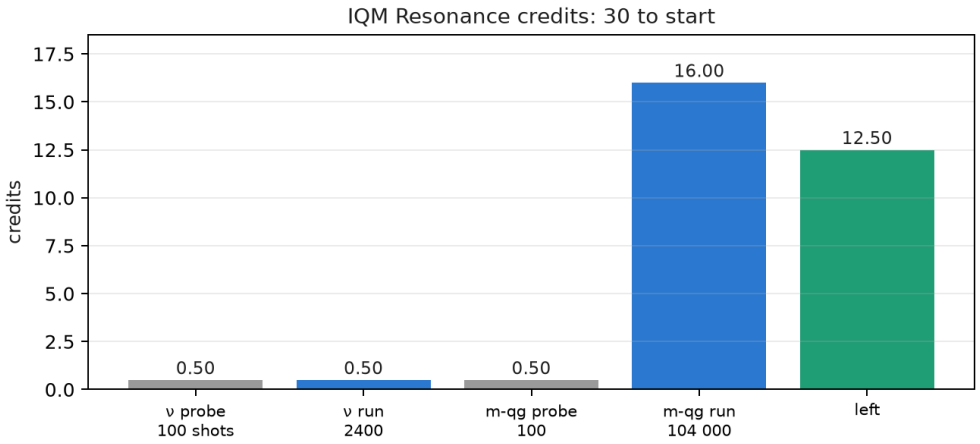


Figure 31.1: Credits spent at each step of the two Garnet runs and those left, out of 30 to start.

31.6 The main experiment: three-flavour neutrinos

It is the circuit of Chapter 9 (§151): an initial ν_μ, a 1300 km baseline, $\delta_{CP} = -\pi/2$, eight energies between 0.8 and 5 GeV, and the three flavours read at once. It has 3 qubits, 12 CZ, and every gate conserves the weight. The script was frozen (sha256 b320f51f in `examples/data/iqm_neutrino_qg.prereg`, commit a74c07c) before any job was sent, with four predictions:

Table 31.4: Job 01a15b04, Garnet, 8 circuits × 300 shots, 2026-10-20. All four passed.

prediction	result
I1: the filtered error is below the raw one	0.0386 < 0.0615
I2: the raw / filtered ratio is at least 1.5	1.59 (narrowly)
I3: the filtered survival has its minimum at 2.6 GeV	yes; **the raw one at 0.8 GeV**
I4: at least half of the shots are kept	0.880

The script recomputes everything from the stored file:

- **Per flavour**, the filter lowers the ν_e error from 0.054 to 0.023 and the ν_μ error from 0.069 to 0.041. The ν_τ error barely changes (0.062 → 0.052).
- **The oscillation maximum goes back to its place.** Without the filter the lowest survival is read at 0.8 GeV. With it, at 2.6 GeV, and the exact maximum is at 2.63.
- **What remains is inside the sector.** At 2.6 GeV the exact survival is 0.013, the raw one 0.173 and the filtered one 0.101. The filter does not see that error (Chapter 22).

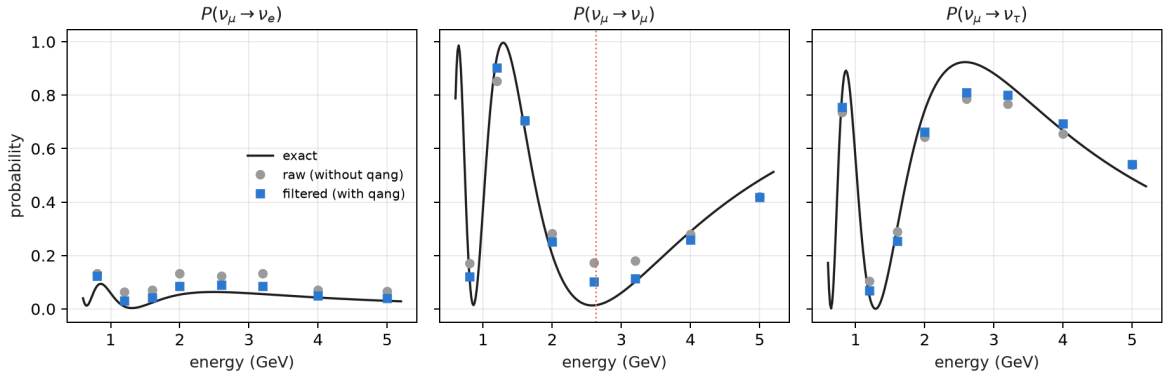


Figure 31.2: The three oscillation probabilities on Garnet: exact (line), raw (grey) and filtered (blue). The dotted line marks the oscillation maximum, 2.63 GeV.

31.7 What came next, and what remains

- **The milli-qang on Garnet** (§168, Chapter 28) used most of the credits. Three of four predictions passed: one X gate costs 1.64 m-qg and the cost grows in a straight line with the number of gates. The equator copies scattered half as much as expected (M2 failed), and that is still open.
- $E = mc^2$ **as a reading on Garnet** (§170, Appendix G): 24 circuits of 2000 shots on the same qubit, with the pre-registration pushed in its own commit before the job. All 5 predictions passed, with a 26–71 m-qg gate bias in the same direction as in §168. It cost about 7.7 credits.
- **The replication on five qubits and two days** is prepared in `examples/iqm_milliqang_replication_qg.py`, not yet pre-registered or run. At full size (130 circuits × 4000 shots per day) it would cost about 80 credits per day, far more than is left: it has to be shrunk before use.

- **The index's candidates** are still pending: the one-qubit P-CTC (§164), the Yoshida–Kitaev decoder (8 qubits and 90 two-qubit gates on the IBM models, probably too deep) and H_2 (§20, §128). The H_2 run on IonQ was refused before any job was accepted, and nothing was charged.

With about 4.8 credits after §170 at about 0.15 per 1000 shots, some 30,000 shots are left, minus the per-job floors.

31.8 Result: hits, misses and comparison with simulation

Table 31.5: The same circuit in simulation and on the chip.

	IBM models (§151)	IQM model (rehearsal)	Garnet (real)
raw error	—	0.126	**0.062**
filter factor	2.97–3.83	2.16	**1.59**
kept shots	—	0.688	0.880

The chip did better than its model: half the raw error, so fewer shots left the sector and the filter had less to remove. It is the rule of Chapter 21 seen on hardware: the filter pays in proportion to the noise it takes out of the sector. One run, one day, 300 shots per circuit: no claim about stability.

31.9 Without qang and with qang

Table 31.6: The Garnet run, with and without the scale.

action	without qang	with qang	difference
mean error of the probabilities	0.0615	0.0386	1.59 times lower
energy of the oscillation maximum	0.8 GeV	2.6 GeV (exact 2.63)	the filter corrects the physics that is read
ν_τ error	0.062	0.052	hardly any: it is in-sector error
cost	2400 shots, 0.50 credits	the same shots	none: the filter is post-processing
useful shots	2400	2113 (88 %)	12 % discarded

action	without qang	with qang	difference
compilation	12 two-qubit gates	12	none: it depends on the optimization level

Balance. The first real-hardware data point confirms what simulation said, with a smaller effect: 1.59 times against 2.2 on IQM's model and 3 to 3.8 on IBM's. The filter costs no extra shots and puts the oscillation at the right energy. From the original plan, IQM's compilation advantage was an artefact of the optimization level, and the `qang.iqm` module was written after the runs, which did not use it.

31.10 In code

```
import json
import numpy as np

d = json.load(open("examples/data/iqm_neutrino_qg_run.json"))
raw, fil, ex = (np.array(d[k]) for k in ("raw", "filtered", "exact"))
np.abs(raw - ex).mean(), np.abs(fil - ex).mean()   # 0.0615 and 0.0386
d["energies"][int(np.argmin(fil[:, 1]))]       # 2.6 GeV: the oscillation maximum
```

31.11 Summary

- On Garnet, real hardware, the filter cut the error 1.59 times and moved the oscillation maximum from 0.8 to 2.6 GeV (exact 2.63). All four pre-registered predictions passed.
- The chip did better than its simulated model, and the filter gained less (1.59 times against 2.16).
- IQM's compilation advantage (12 against 18) came from the optimization level: at the same level IBM also gives 12.
- 17.50 of 30 credits were spent on this chapter's runs, at about 0.15 per 1000 shots with a floor of 0.50 per job; with §170 (Appendix G) about 25.2 are spent and about 4.8 are left.
- `qang.iqm`, written after the runs, adds readout correction and chain choice. On the Garnet model the correction removes another 15% of the filtered error; it has not been tried on the real chip.

Bibliographic notes

Access was through IQM Resonance (IQM Quantum Computers 2026); the Garnet processor is described in Abdurakhimov et al. (2024). Compilation uses Qiskit (Javadi-Abhari et al. 2024), and the oscillation parameters are of the type of Esteban et al. (2020). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

32 Software architecture, the multi-SDK ecosystem and physical execution

A qg readout is only useful if it travels intact from the device to the number that gets reported. In between there is an SDK that orders bits its own way, primitives that return expectation values instead of counts, a compiler that picks qubits and a chip with errors no filter sees. This chapter walks that chain as the library solved it. The new checks and the rescoring of the stored IBM runs are in `tools/make_cap32_figures.py`.

32.1 From the base: one count, many paths

The equation of Chapter 1 needs one thing: how many zeros each qubit gave. Everything else is plumbing, and every pipe can break the readout without anyone noticing:

- an SDK that writes qubit 0 on the right and another that writes it on the left **swap the qubits** without an error;
- a primitive that returns $\langle Z_i \rangle$ already gives qg_Z, but **not the shots** the filter needs to discard;
- a compiler that picks bad qubits **lowers the kept fraction** before anything is measured.

32.2 One qubit order

Qiskit writes qubit 0 at the right of the bitstring; Cirq, PennyLane, Braket and pytket write it first. `qang.interop.canonical_counts` brings everything to $\{(b_0, \dots, b_{n-1}) : \text{count}\}$, qubit 0 first. The script checks it with the two SDKs installed here:

- the same outcome (only qubit 0 excited), written `001` by Qiskit and `100` by PennyLane, gives **(1, 0, 0)** in both;
- a 3-qubit W state sampled 20,000 times in each SDK gives qg_Z between 0.328 and 0.343 (ideal 1/3), every shot in weight 1 and register qg_S 0.528 (ideal $\log_2 3/3 = 0.528$), the same in Qiskit and in PennyLane.

Cirq, Braket and pytket are not installed on this machine. Their order is pinned by the library's tests (notes §130) and was not run again.

32.3 V2 primitives and filtered observables

Qiskit's V2 primitives come in two forms, and qang reads both (notes §131):

- **Sampler.** `interop.sampler_qg` takes a SamplerV2 result and gives each qubit's qg_Z, with its error, for every parameter set at once. Given the weight, it also gives the filtered profile and the kept fraction.
- **Estimator.** It returns no shots, so none can be discarded. But the filtered qg_Z is a ratio of two expectation values, $\langle P_w Z_i\rangle / \langle P_w\rangle$, where P_w projects on weight w. Both operators are diagonal, so they are sums of Z strings. `interop.filtered_observables` builds those sums.

The script checks that it is exact: on 20 random 4-qubit states, the values of `filtered_observables(4, 2)` give the kept fraction and the filtered qg_Z with a largest error of $3.3 \cdot 10^{-16}$. It has a cost: the largest observable has 2^{n-1} Z strings (8 with 4 qubits, 2048 with 12). They all commute and are measured in the same basis, so the cost is classical bookkeeping, not circuits. It still limits the method to registers of about a dozen qubits.

32.4 Sector-aware compilation

`qang.ibm` adds to Qiskit's compiler an analysis that decides, gate by gate, whether the circuit conserves the weight (Chapter 30). It also reads the published readout errors of the physical qubits and turns an IBM Runtime result into four readouts: raw, readout-corrected, filtered and filtered after readout correction. On the ten weight-1 circuits of §137, recomputed from the stored file:

Table 32.1: Mean qg_Z error, 8000 shots (§137). All three predictions passed.

device	raw	readout	filter	filter + readout
FakeBrisbane	0.104	0.074	0.023	0.022
FakeSherbrooke	0.065	0.038	0.019	0.016
FakeTorino	0.101	0.036	0.018	0.013

Choosing qubits from the calibration (§138). `predict_kept_fraction` multiplies the survival of the weight under each calibrated channel, and `best_chain` picks the chain with the best prediction:

- the prediction comes out too low: the mean error is 0.08 to 0.13 (L1 failed);
- but it ranks the layouts correctly (Spearman 0.94 to 0.98, L2 passed);
- and `best_chain` keeps 6 to 17 points more shots than Qiskit's default layout (L3 passed);
- after the result, adding the share of shots that return to the sector as for a fully mixed state brought the error down to 0.023–0.041 (`kept_mixed_floor`).

For the ring-coupled neural network (§141), `best_layout` was never worse than the default layout. It gained 0.022 on FakeBrisbane, and with filter and readout correction it got 550 of 567 inputs right. R3 asked for 551 and failed by one, which is within the variation between runs with the identical layout.

32.5 Static ZZ: simulated, not mitigated

Coupled transmons pick up a ZZ phase all the time. `ibm.add_static_zz` **adds** it to the compiled circuit, after every layer and on every edge in use, with the layer's calibrated duration. It is there to measure how much of that error the filter sees. **It does not correct it.** In §150, at 0, 5 and 15 kHz:

- **The filter sees part of the ZZ, not all of it.** Between logical gates ZZ is diagonal and passes the filter. Inside the native decomposition of the XX+YY gates it moves excitations out of the sector: the kept fraction fell by 0.019 to 0.022 on the two ECR devices, and by 0.001 on FakeTorino.
- **The filter removed 32 to 60 % of the added error.** The filtered error grew by 0.0025 to 0.0086 (14 to 38 %) at 15 kHz and stayed 3.6 to 5.1 times below raw. The notes said "4 to 5 times"; the rescoring gives 3.6–5.1.
- **Short circuits suffer least**: FakeTorino runs for 2.1 µs against 9 to 10 µs on the others.

Predictions Z2 and Z3 failed, because they assumed the filter would see none of the ZZ.

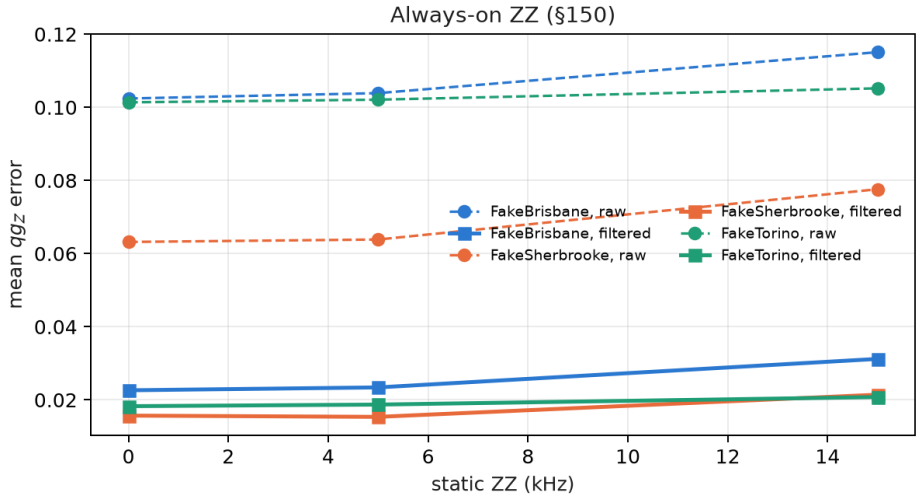


Figure 32.1: Mean raw (dashed) and filtered (solid) qg_Z error against the simulated static ZZ, on three devices (§150).

32.6 ZNE with the filter, and shot planning

In local mode, IBM Runtime's mitigation options do nothing. That is why `ibm.zne_qg` folds every two-qubit gate (scale 3), runs both scales and extrapolates linearly. Folding conserves the weight, so the extrapolation (ZNE) combines with the filter (§142):

- **The filter does more than the extrapolation**, 2 to 7 times more: filter + ZNE gives 0.012, 0.011 and 0.015 against 0.053, 0.034 and 0.085 for ZNE alone (Z2 passed).
- **ZNE does not always help.** On FakeTorino, adding it to filter + readout makes the error worse, from 0.013 to 0.014 (Z1 failed), and the best combination there is filter + readout (Z3 failed).

`ibm.plan_shots` turns a target precision of the filtered qg_Z into shots and QPU seconds, from the predicted kept fraction (§143). For an error of 0.02 it planned 3000 to 3400 shots, and over 30 repetitions the worst qubit landed at 0.021–0.023 (P1 passed). The per-qubit prediction was 1.4 to 1.6 times too small (P2 failed). The misses are the qubits near a pole, where the noise pulls qg_Z away from ±1. That is why the plan uses the worst case.

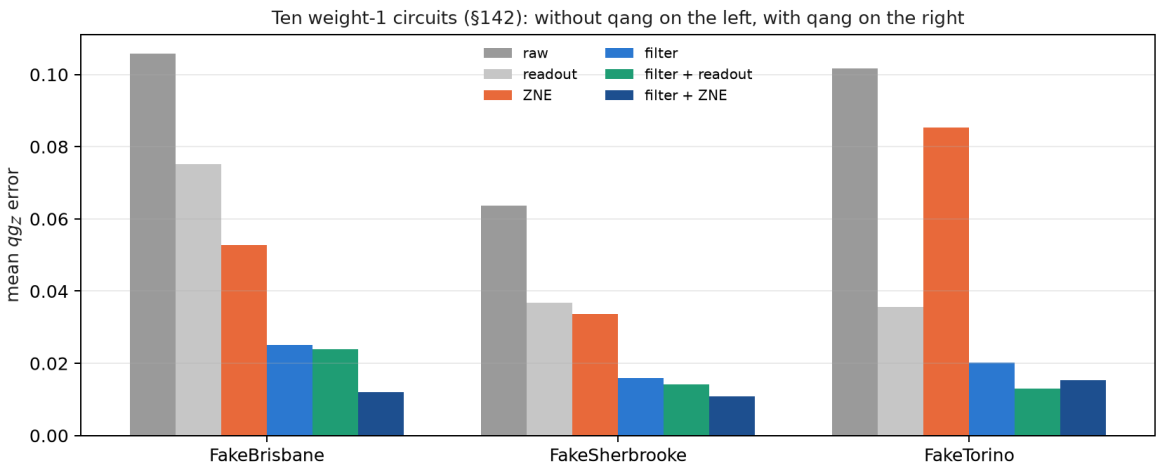


Figure 32.2: Ten weight-1 circuits on three devices (§142): raw, readout correction and ZNE without qang; filter, filter + readout and filter + ZNE with qang.

32.7 Guidelines by hardware type

Superconductors (IBM, IQM).

- Compile with weight-conserving gates (XX+YY, RBS) and check it with the analysis before running.
- Compare two-qubit gate counts at the same optimization level (Chapter 31).
- Choose the chain by the predicted kept fraction, and correct the readout from the calibration, or from the two poles where the calibration is not available (IQM).
- Prefer short circuits: ZZ and dephasing inside the gates are filtered in proportion to the time the qubits spend outside the computational frame.

Ions (IonQ, through its cloud or AWS Braket). Connectivity is all-to-all and the native gate is Mølmer–Sørensen. The H_2 rehearsal on IonQ's simulator with the Forte-1 noise model (§128) gave:

- witness +0.0003;
- 97.9 % of the shots kept;
- the error falls from 29.9 to 15.0 mHa with the filter.

The run on the real machine was refused before any job was accepted, because the project did not have Forte-1 enabled, and nothing was charged. Through Braket, `interop` handles the bit order.

On any platform.

- Rehearse for free on the simulated model.
- Measure the cost with a small probe.
- Pre-register before the main job.
- Read the token from an environment variable, and spend nothing without an explicit flag.

32.8 Without qang and with qang

Table 32.2: The software chain, with and without the scale.

action	without qang	with qang	difference
reading another SDK's counts	each SDK with its bit order	one order	no silent qubit swap
Estimator	unfiltered $\langle Z_i \rangle$	ratio of two diagonal observables	the filter without shots, exact to 10^{-16}
ten weight-1 circuits	0.065–0.104 raw	0.013–0.022 with filter and readout	4 to 8 times less error
choosing qubits	default layout	`best_chain`	6 to 17 points more shots kept
static ZZ at 15 kHz	+0.004 to +0.015 raw	+0.0025 to +0.0086 filtered	32 to 60 % of the added error removed; not mitigated
ZNE	0.034–0.085	0.011–0.015 with the filter	2 to 7 times; ZNE makes FakeTorino slightly worse
shot planning	—	3000–3400 for an error of 0.02	met on the worst qubit; the per-qubit prediction fails

Balance. The library fixes the plumbing that breaks readouts without warning: bit order, the filter inside an Estimator and qubit choice. The IBM figures come from simulated models in local mode, not from hardware. Of this chapter's predictions, L1, Z1, Z3, P2, R3 and two of the ZZ ones failed, and all stay written down. Static ZZ is the most important warning: the library adds it in order to measure it, and the filter removes only part of it.

32.9 In code

```
from qiskit.quantum_info import SparsePauliOp, random_statevector
from qang import interop as IO

obs = [SparsePauliOp.from_list(list(o.items())) for o in IO.filtered_observables(4, 2)]
sv = random_statevector(16, seed=0)
kept, qg = IO.estimator_filtered_qg([sv.expectation_value(o).real for o in obs])
# kept: share in weight 2; qg: the four filtered qg_Z, exact
```

32.10 Summary

- `qang.interop` brings any SDK's counts to a single qubit order; Qiskit and PennyLane give the same W state.
- The filtered qg_Z comes out of an Estimator as a ratio of two diagonal observables, exactly, with 2^{n-1} Z strings.
- On IBM's models, filter and readout cut the error from 0.065–0.104 to 0.013–0.022, and `best_chain` keeps 6 to 17 points more shots.
- `add_static_zz` simulates ZZ: the filter removes 32 to 60 % of what is added and does not fully mitigate it.
- The filter does 2 to 7 times more than ZNE alone; the shot plan is met by using the worst qubit.

Bibliographic notes

The kits are Qiskit (Javadi-Abhari et al. 2024), PennyLane (Bergholm et al. 2018), Cirq (Cirq Developers 2026) and Amazon Braket (Amazon Web Services 2020). For static ZZ coupling, see Krantz et al. (2019); zero-noise extrapolation is Temme et al. (2017). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part XI

XI. Nanoquantum

33 Nanoquantum: transmission, partition noise and nanoscale qubits

At the nanometre scale many measurable quantities are a probability: that an electron crosses a barrier, that a level is occupied, that a spin is up. This chapter reads them with the equation of Chapter 1. There is no notes section and no qang module for these topics, so every figure comes from a closed form or from a Monte Carlo written in `tools/make_cap33_figures.py`. Next to each reading with qang goes the usual form, and the text says when the scale adds something and when it only changes the notation.

33.1 From the base: one probability per experiment

The equation of Chapter 1 takes a probability p_0 and gives four exact readings:

$$\theta = 2\arccos\sqrt{p_0}, \qquad qg_Z = 2p_0 - 1, \qquad qg_S = H(p_0), \qquad m = 2000\,p_0 - 1000\ \text{m-qg}.$$

In this chapter p_0 is, case by case, the transmission T of a channel, the occupation of a level or the population of a spin state. Two properties from Chapter 1 do all the work:

- **the m-qg reading is linear in probability**: 1 m-qg is always $\Delta p_0 = 0.0005$;
- **the angle stabilizes binomial noise**: with N counts, $\sigma_\theta = 1/\sqrt{N}$ for any p_0, while $\sigma_{p_0} = \sqrt{p_0(1-p_0)/N}$ changes with p_0. On the angular qga scale (Section 1.4) that is $159.15/\sqrt{N}$ qga, also for any p_0.

33.2 Tunnelling transmission as a qang probability

An electron reaches a rectangular barrier and crosses it with probability T. With $p_0 = T$:

$$\theta_T = 2\arccos\sqrt{T}, \qquad \cos(\theta_T/2) = \sqrt{T} = |t|.$$

The qang half-angle is the modulus of the transmission amplitude t. It is not an analogy: it is the same definition, and the script checks it to 10^{-16}. With a 1 eV barrier, a 0.5 eV electron ($\kappa = 3.62$ nm^{-1}) and the exact rectangular barrier formula:

Table 33.1: Transmission through a 1 eV rectangular barrier for a 0.5 eV electron.

width	T	reading	θ_T
0.2 nm	0.616	+231.9 m-qg	76.6°
0.5 nm	0.101	−797.3 m-qg	142.9°
1.0 nm	$2.85\cdot10^{-3}$	−994.3 m-qg	173.9°
1.5 nm	$7.62\cdot10^{-5}$	−999.85 m-qg	179.0°

- **Without qang** one plots $\ln T$, which falls with slope -2κ.
- **With qang**, the angular distance to the pole, $\pi - \theta_T \approx 2|t|$, falls with slope $-\kappa$: the script measures -3.620 nm^{-1} against -3.623. It is the amplitude instead of the probability, the same information.
- **Thick barriers live right at the pole.** One nanometre leaves the reading 5.7 m-qg from the pole, and one and a half nanometres, 0.15 m-qg. Reading that difference with the error of an electron count needs the exact angle near the pole (Chapter 1), not the linear approximation.

33.3 The Landauer formula: one channel, one reading

The conductance of a ballistic conductor is $G = G_0 \sum_n T_n$, with $G_0 = 2e^2/h$. Written in qg:

$$\frac{G}{G_0} = \frac{N}{2} + \frac{1}{2}\sum_{n=1}^{N} q g_n.$$

The script uses the saddle-point model of a quantum point contact, $T_n = 1/(1 + e^{-2\pi(V-n-1/2)/w})$, and checks the identity to $9 \cdot 10^{-16}$:

- **Each plateau is a set of channels at the poles**: open at +1000 m-qg, closed at −1000. At $V = 1, 2, 3$ the conductance is exactly 1, 2 and 3.
- **Each step is one channel crossing the equator.** The midpoint of the step is that channel's 0 m-qg reading.
- **Conductance alone does not separate channels that open together**: it gives $\sum T_n$ and nothing else. Shot noise (next section) gives $\sum T_n(1 - T_n)$. With both, two partly open channels are recovered exactly: the script starts from $T = 0.4031$ and 0.9716 and gets them back. It is the known technique for measuring the transmissions of an atomic contact. In qg, the conductance reads the channels' mean qg and the noise reads their mean qg^2.

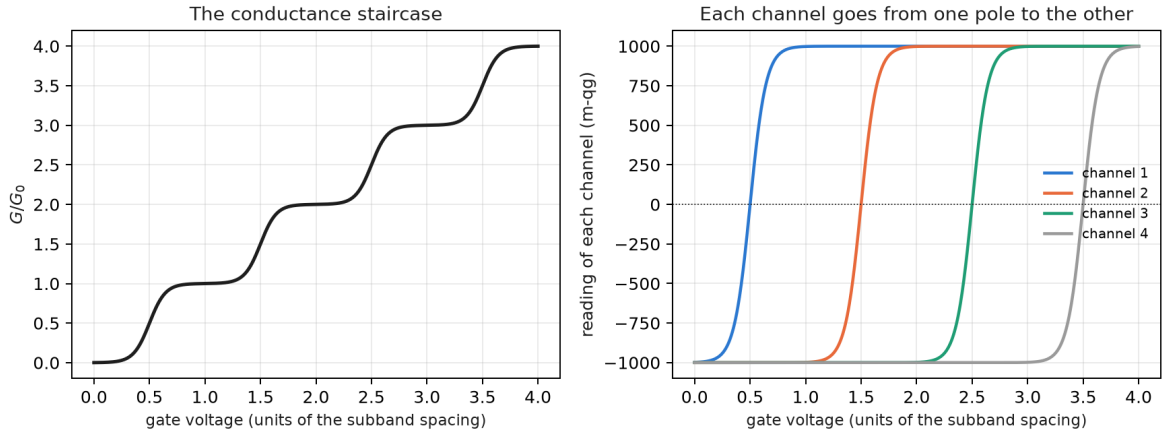


Figure 33.1: Left: the conductance staircase of a quantum point contact with four channels. Right: each channel's reading in m-qg; each step is one channel going from −1000 to +1000 m-qg and crossing the equator halfway.

33.4 Partition noise: the binomial variance, stabilized

Every electron that reaches the channel either passes or bounces back: a Bernoulli trial with probability T. That gives the partition noise, $S = 2eI(1-T)$, with Fano factor $F = 1 - T$. In qg, $T(1-T) = (1 - qg^2)/4 = \sin^2\theta/4$: it is the binomial variance of Chapter 1. The script counts 1000 electrons per window, over 20,000 windows:

- **The Fano factor**, measured as variance over mean of the transmitted electrons, matches $1 - T$ to within 0.02.

- **Without qang**, the spread of $\hat{T}$ times $\sqrt{N}$ runs from 0.14 to 0.50: it depends on what is being measured, and near the poles the error bar shrinks so much that it misleads.
- **With qang**, the spread of the angle $\hat{\theta}$ times $\sqrt{N}$ stays between 0.985 and 1.009 for T between 0.05 and 0.95, and is still 1.01–1.02 at 0.02 and 0.98. The error bar is the same over the whole range.

It is the arcsine transform, known in statistics for a long time. qang does not invent it. What it adds is that the stabilized angle is the natural coordinate of the state.

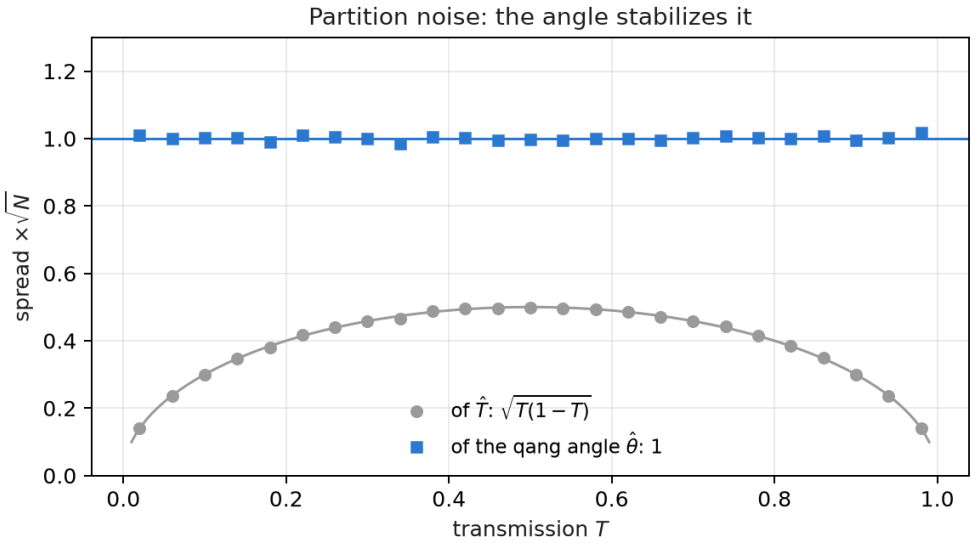


Figure 33.2: Spread of the estimated transmission and of the estimated qang angle, times $\sqrt{N}$, with 1000 electrons per window: the first follows $\sqrt{T(1-T)}$; the second stays at 1.

33.5 Coulomb blockade: the occupation of a level

In a quantum dot in Coulomb blockade, a level at energy ε from the leads' chemical potential is occupied with the Fermi probability $f = 1/(1+e^{x})$, with $x = \varepsilon/k_B T$. With $p_0 = f$:

$$qg_Z = 2f - 1 = -\tanh(x/2).$$

It is the same thermal state as $qg_Z = \tanh(\beta h)$ of Chapter 8, with the sign of the convention. The script checks it to $2 \cdot 10^{-16}$. The slope at the centre, $-1/(2k_BT)$, is a thermometer: a charge sensor reads the occupation and its width gives the electron temperature. With 25 values of ε, 200 readouts per point and 400 repetitions ($k_BT = 1$):

Table 33.2: Thermometry from the occupation of a level.

method	estimated k_BT
maximum-likelihood Fermi fit (without qang)	0.9986 ± 0.0294
straight line through artanh(qg_Z) (with qang)	1.0294 ± 0.0305

qang loses here. The straight line is closed-form and needs no non-linear fit, but it is biased by +2.9 %. In the tails, artanh amplifies the noise of readings close to ±1 (visible in the figure), and points where every readout agrees have to be dropped. The Fermi fit uses the same information without bias. The qg reading is for seeing the thermometer, not for estimating with it.

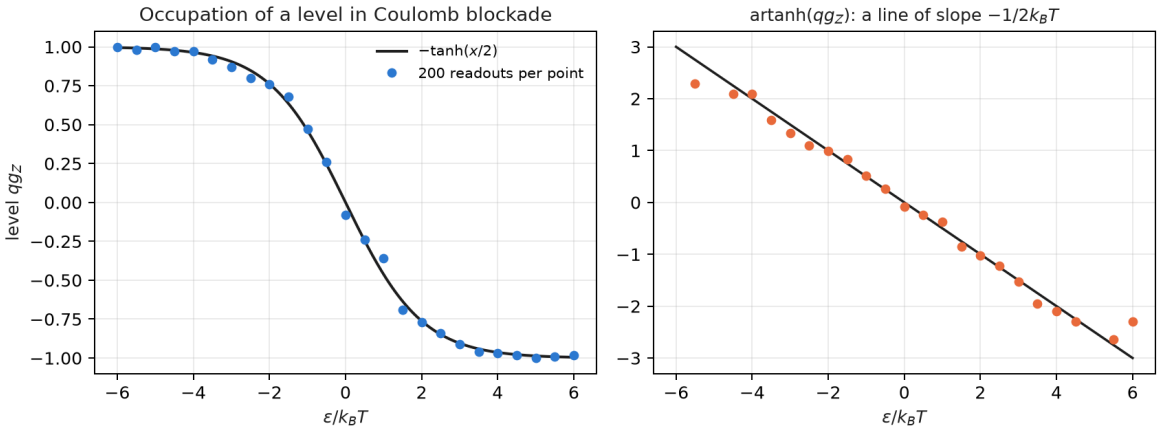


Figure 33.3: Left: a level's qg_Z against ε/k_BT, exact and with 200 readouts per point. Right: artanh(qg_Z) is a straight line of slope $-1/2k_BT$, but in the tails the noise bends it.

33.6 Charge and spin qubits in quantum dots

A spin qubit in a quantum dot is read with a charge sensor: the spin is converted into charge and the sensor decides 0 or 1, with errors $e_0 = P(1\,|\,0)$ and $e_1 = P(0\,|\,1)$. It is the same readout as a transmon's, and the rule of Chapter 27 (§31) describes it in m-qg:

$$m_{\text{read}} = (1 - e_0 - e_1)\, m + 1000\,(e_1 - e_0).$$

The script simulates a Rabi oscillation with $e_0 = 0.02$ and $e_1 = 0.08$ (illustrative values) and 2000 readouts per point:

- the affine identity holds to $2 \cdot 10^{-13}$;
- **uncorrected**, the oscillation has an amplitude of 898 m-qg instead of 1000 and is shifted by +60 m-qg. Its rms error is 95.5 m-qg;
- **corrected from the two poles** (2000 readouts on each), the error falls to 19.4 m-qg. The price is that the noise grows $1/(1 - e_0 - e_1) = 1.11$ times;
- the usual readout fidelity, $1 - (e_0 + e_1)/2 = 0.95$, is in m-qg a visibility of 900 m-qg and a shift of 60.

The correction is the standard one. Writing it in m-qg makes it a linear account: each point of readout error costs 10 m-qg of visibility.

33.7 NV centres: contrast as an m-qg budget

An NV centre in diamond is read through fluorescence: the $m_s = 0$ state is brighter than $m_s = \pm 1$, with a contrast C. If p_0 is the population of $m_s = 0$ and N_b the photons the bright state would give, the count is Poisson with mean $N_b\,[1 - C(1 - p_0)]$, and the m-qg reading of p_0 has

$$\sigma_m \approx \frac{2000}{C\sqrt{N_b}} \text{ m-qg}.$$

- With $C = 0.3$ and 444,444 photons the formula gives 10.00 m-qg. The Monte Carlo gives 9.22, the same as the formula without the approximation (the factor $\sqrt{1 - C/2}$ at $p_0 = 1/2$).
- **Resolving 10 m-qg costs about 444,000 photons.** A contrast of 0.3 makes every m-qg cost eleven times more photons than a contrast of 1 would. It is the same budget as in Chapter 1 with a factor $1/C^2$.

- **ODMR magnetometry.** A 1 MHz wide line, contrast 0.2, 61 frequencies and 1.16 million photons in total. Fitting the line gives its centre with a spread of 8.3 kHz, equal to the Cramér–Rao bound (8.3 kHz). That is 296 nT with $\gamma = 28$ GHz/T.

The line fit does not change with qang. What the scale adds is the account beforehand: how many photons a given resolution needs, in a unit that is linear in the spin population.

33.8 Without qang and with qang

Table 33.3: Nanoquantum, with and without the scale.

topic	without qang	with qang	difference
tunnelling	$\ln T$, slope -2κ	$\cos(\theta_T/2) = \lvert t\rvert$; $\pi - \theta_T$ with slope $-\kappa$	the amplitude in view; the same information
Landauer	$G/G_0 = \sum T_n$	$N/2 + \sum q g_n/2$; plateaus at the poles	conductance = mean qg, noise = mean qg^2
partition noise	error bar of $\hat{T}$: 0.14 to 0.50 ($\times 1/\sqrt{N}$)	error bar of $\hat{\theta}$: 0.985 to 1.009	constant over the whole range (arcsine transform)
thermometry	Fermi fit: 0.9986 ± 0.029	line in artanh: 1.029 ± 0.031	**qang worse**: +2.9 % bias
reading a spin	95.5 m-qg error uncorrected	19.4 corrected from the poles	the standard correction, written linearly
NV	contrast and photons	444,000 photons for 10 m-qg at $C = 0.3$	a budget in population units
ODMR	8.3 kHz from the fit	the same	none

Balance. In nanoquantum the scale gets the language right: transmissions, occupations and populations are probabilities, and qang reads them all the same way. It adds little beyond that:

- the constant error bar of the angle, which is the arcsine transform;

- the half-angle as the transmission amplitude;
- linear budgets in m-qg.

In the one case where it competes as an estimator, thermometry, it loses to the maximum-likelihood fit. Everything here is calculation and simulation: no figure in this chapter comes from an experiment. There is no qang module for these topics yet either.

33.9 In code

```
import numpy as np
from qang import probability as PB

T = 2.85e-3                                  # 1 nm barrier (table above)
PB.milliqang_from_p0(T)                       # -994.3 m-qg
np.cos(PB.theta_from_milliqang(PB.milliqang_from_p0(T)) / 2)  # 0.0534 = sqrt(T) = |t|
```

33.10 Summary

- Tunnelling transmission is a qang probability, and its half-angle is the modulus of the amplitude: $\cos(\theta_T/2) = |t|$.
- In the Landauer formula each channel is a qg reading: plateaus are channels at the poles and steps are channels crossing the equator. Conductance and noise together recover two channels opening at once.
- Partition noise is the binomial variance; the qang angle stabilizes it at $1/\sqrt{N}$ over the whole range.
- A level's occupation is $qg_Z = -\tanh(x/2)$; as a thermometer, the line in artanh has a +2.9 % bias the Fermi fit does not.
- Spin readout through a charge sensor and NV readout through fluorescence are written as linear m-qg budgets; their estimators do not change.

Bibliographic notes

Conductance as transmission is Landauer (1957) and Büttiker (1986); for mesoscopic transport in general, see Datta (1995). Shot noise is Blanter and Büttiker (2000), and the Fano factor, Fano (1947). Coulomb blockade is Beenakker (1991);

spins in quantum dots, Hanson et al. (2007), and their single-shot read-out, Elzerman et al. (2004). The NV centre is Doherty et al. (2013), and ODMR sensitivity, Dréau et al. (2011). The numbered research notes (§) cited in this chapter and the scripts that check it are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part XII

XII. Cross-domain applications

34 Angular architecture: chemistry, condensed matter, relativistic optics, thermodynamics and geometry

This chapter answers a proposal: thirteen phenomena from five fields, each with an equation, a "qg mapping" and an analytic test with its pass criterion. The proposal uses a currency the Chapter 1 base does not have: it speaks of "qg cycles" and "qg periods". Before anything is run, those criteria have to be translated into readings. Then each test is run in `tools/make_cap34_figures.py`, which takes about 25 seconds. There is no notes section and no qang module for these topics, so every figure comes from a closed form, an exact diagonalization or a Monte Carlo written there. None comes from an experiment. The verdict on each test is in Table 34.3: some hold as written, others hold only once corrected, and two fail.

34.1 From the base: readings, not cycles

The Chapter 1 equation starts from a probability p_0 and gives four exact readings:

$$\theta = 2\arccos\sqrt{p_0}, \qquad qg_Z = 2p_0 - 1 = \cos\theta, \qquad qg_S = H(p_0), \qquad m = 1000\, qg_Z \text{ m-qg}.$$

This has three consequences for the proposal:

- **qg has no period.** qg_Z lives in $[-1, 1]$ and the qang angle in $[0, \pi]$. A criterion such as "phase $= \pi \equiv 0.5$ qg cycle" cannot be checked as it stands. It has to be rewritten as a reading: for instance, $qg_X = -1$, which is -1000 m-qg. The angular qga scale of Chapter 1 (Section 1.4) is periodic, 1000 qga per turn, and on it that phase is written $\varphi = 500$ qga; but it is still an angle, not a reading: to check it one has to measure it with an interferometer,

and what is compared with the prediction is the reading that interferometer returns.

- **A phase does not change qg_Z.** A rotation about Z (Larmor precession, Aharonov–Bohm phase, the sign of an exchange) leaves the populations alone. It is read with an interferometer, in qg_X or qg_Y, like the Berry phase of Chapter 7.
- **Where there is a probability, there is a reading.** A reflectance, the weight of a determinant, the vacuum occupation or a transition probability is a p_0, so everything from Chapter 1 applies: the linear scale and the angle that stabilizes the noise.

34.2 Quantum chemistry

34.2.1 Jordan–Wigner: the exchange sign is a reading

The mapping is $a_j^\dagger = \left(\prod_{k<j}(-Z_k)\right)\sigma_j^+$. The script builds the operators of four modes and checks the anticommutation relations $\{a_i, a_j^\dagger\} = \delta_{ij}$ and $\{a_i, a_j\} = 0$ with exactly zero error. It also checks that $a_1^\dagger a_2^\dagger|0\rangle + a_2^\dagger a_1^\dagger|0\rangle = 0$. The proposal asks for the exchange to be "a half-period qg phase shift". On the Chapter 1 base that is not a reading. To make it one, the exchange has to act inside a Hadamard test: an ancilla controls the exchange of the two occupied modes and is read in X.

- **With the fermionic exchange**, the ancilla reads $qg_X = -1$, that is -1000 m-qg.
- **With a qubit SWAP**, the same circuit reads $+1000$ m-qg.

The Pauli sign is the difference between the two poles of qg_X. The Z gates of the parity string are exact, so there is no "phase loss from accumulated gates". The error of a real circuit comes from noise, not from the mapping.

34.2.2 UCC for H_2: the amplitude is half a qang angle

For H_2 in STO-3G at 0.735 Å, with the four-qubit Hamiltonian of Chapter 25, UCCSD reduces to a single double excitation. The state is $|\Psi(t)\rangle = \cos t\,|\mathrm{HF}\rangle + \sin t\,|\mathrm{double}\rangle$, and the energy has the closed form $E(t) = A + B\cos 2t + C\sin 2t$:

- the minimum is at $t^* = -0.1118$ rad, where the energy is -1.1373060 Ha, the FCI energy to within $5 \cdot 10^{-14}$ Ha; Hartree–Fock gives -1.1169990;

- **the Hartree–Fock weight is a reading**: $p_0 = \cos^2 t^* = 0.9876$, that is $+975.1$ m-qg, and its qang angle is $\theta = 2|t^*| = 0.2235$ rad. The UCC amplitude is half a qang angle, the same relation Chapter 33 finds for tunnelling;
- a sweep with step $\pi/200$ gets within $5.4 \cdot 10^{-6}$ Ha of the minimum.

The proposal asks for "sub-millihartree precision, $< 1.6 \cdot 10^{-3}$ Ha". That threshold is chemical accuracy (1 kcal/mol), not a sub-millihartree, and in an exact simulator it is beaten by eleven orders of magnitude: it discriminates nothing. The dissociation curve is not redone because the environment has no PySCF; the stored runs of Chapter 25 cover it. The gradient vanishes at the minimum in both t and m, because $dm/dt = -2000 \sin 2t \neq 0$ there. At Hartree–Fock ($t = 0$, the pole) the m coordinate is singular, as Chapter 23 shows.

34.2.3 Larmor and the Hahn echo

Free precession $\Delta\varphi = \gamma B_0 t$ is a rotation about Z. With the proton ($\gamma/2\pi = 42.577$ MHz/T) and a spin on the equator, the script gives $qg_Z = 0$ at all times and $qg_X = \cos \Delta\varphi$. The precession does not show in the Z reading.

For the echo, the script takes 20,000 spins with Gaussian detuning ($\sigma = 2\pi \cdot 50$ kHz), $\tau = 20$ µs, the sequence $\pi/2 - \tau - \pi - \tau$, and reads the magnetization in qg_X:

Table 34.1: The Hahn echo read in qg_X.

case	qg_X	theory
no π pulse (free decay)	0.004 (Monte Carlo noise)	$2.7 \cdot 10^{-9}$
echo, static detuning	**1.000000** (+1000 m-qg)	1
echo, π pulse 2 % too long	0.99901	0.99901
echo, detuning that changes by 10 % of σ between halves	0.820	0.821

The proposal's criterion, "$F = 1$ independent of $\Delta\omega$", holds only with static detuning and exact pulses. A pulse with a 2 % error costs 1 m-qg, and a drifting detuning costs 180 m-qg. Calibrating the echo with the device itself is the subject of Chapter 22.

34.3 Condensed matter

34.3.1 Transverse-field Ising: the peak moves with size

$H = -J\sum Z_i Z_{i+1} - g\sum X_i$ on an open chain. With $\tan\theta = g/J$, the transition is at $g = J$, that is at $\theta = \pi/4$, **in the thermodynamic limit**. The proposal asks to find it with 2 to 4 spins, at the peak of the derivative of the entanglement entropy. The script diagonalizes with $N = 2$ to 12 (in the even sector of the parity $\prod X$, so that the cat state below the transition is well defined):

Table 34.2: Peak position (g/J) and transverse reading at the critical point.

N	steepest drop of S	peak of $d\,qg_X/dg$	m_X at $g = J$
2	0.33	(at the edge)	894.4
4	0.65	0.43	810.1
6	0.78	0.63	771.1
8	0.85	0.70	747.9
10	0.88	0.78	732.3
12	0.90	0.83	720.9
∞	1	1	636.6 ($2000/\pi$)

The test fails with 2 to 4 spins. The peak is at $g/J = 0.33$ and 0.65, that is at $\theta = 18°$ and $33°$, not at $45°$. With two spins there is no transition: the curve is smooth. The peak approaches 1 as N grows, and the entropy converges faster than the qg_X reading (Figure 34.1). The transverse reading at $g = J$, m_X, also falls towards the exact value $1000 \cdot 2/\pi = 636.6$ m-qg (Pfeuty 1970), but with 12 spins it is still 84 m-qg above it. What gives $\theta = \pi/4$ is the limit, not the scale.

34.3.2 Kitaev chain: the protection is in the Hamiltonian, not in the angle

The parity $P = \prod_j (-i\gamma_{2j-1}\gamma_{2j})$ takes the values ± 1, that is ± 1000 m-qg. The script separates two questions:

- **Does the chain protect?** With $t = \Delta = 1$ and $\mu = 0.5$, the energy difference between the two parity sectors halves with each site: 0.207 with 2 sites and 0.0029 with 8. With $\mu = 0$ and local disorder ± 0.3 it falls from $1.5\cdot 10^{-2}$ to $3\cdot 10^{-7}$, though not monotonically. That is topological protection: exponential in the length.

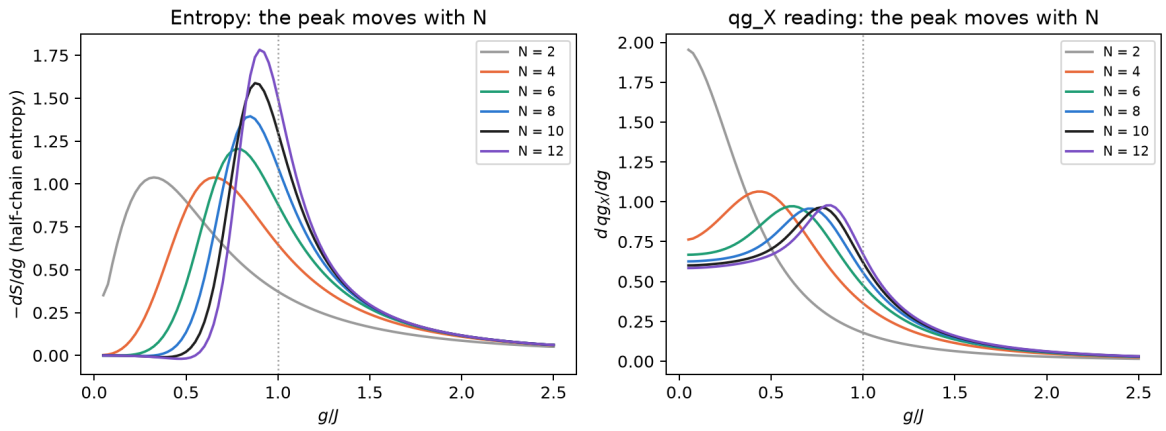


Figure 34.1: Left: $-dS/dg$ of the half-chain entropy for $N = 2$ to 12. Right: the susceptibility of the transverse reading, $d\,qg_X/dg$. Both peaks approach $g/J = 1$ as the chain grows; with 2 to 4 spins they are far from it.

- **Does the gate protect?** The braid $B = e^{\pi\gamma_2\gamma_3/4} = (1 + \gamma_2\gamma_3)/\sqrt{2}$ holds to 10^{-16}, with exact eigenphases $\pm\pi/4$ and $B^2 = \gamma_2\gamma_3$. But applied as a gate with a relative angle error ϵ, the infidelity is $6.2 \cdot 10^{-5}$ at 1 %, $1.5 \cdot 10^{-3}$ at 5 % and $6.2 \cdot 10^{-3}$ at 10 %, quadratic in ϵ, like any rotation.

Writing the braid in "quarter qg cycles" does not protect it. Invariance under local perturbations comes from the Hamiltonian with separated Majorana modes (Kitaev 2001; Ivanov 2001). A gate that imitates the braid has the error of any gate.

34.3.3 Aharonov–Bohm: one fringe per flux quantum

In a two-path ring, the output has $p_0 = \cos^2(\pi\Phi/\Phi_0)$, so

$$qg_Z = \cos\left(2\pi\,\frac{\Phi}{\Phi_0}\right), \qquad \theta = 2\pi\,\frac{\Phi}{\Phi_0} \ \text{(folded into } [0,\pi]).$$

The qang angle is the Aharonov–Bohm phase. The reading returns to its value every $\Phi_0 = h/e$ (the script checks it to 10^{-15}). What is periodic is the flux, not qg. With 2000 electrons per measurement and 4000 repetitions, the flux estimated from the angle has a spread times $\sqrt{N}$ of between 0.156 and 0.163 for Φ/Φ_0 from 0.05 to 0.45. The bound is $1/2\pi = 0.159$, the same over the whole

range (Figure 34.2). It is the Chapter 1 arcsine transformation applied to an interferometer. **What the reading does not give** is the sign: Φ and $-\Phi$ (and $\Phi_0 - \Phi$) read the same. Telling them apart needs the second quadrature, qg_Y, as in Chapter 7.

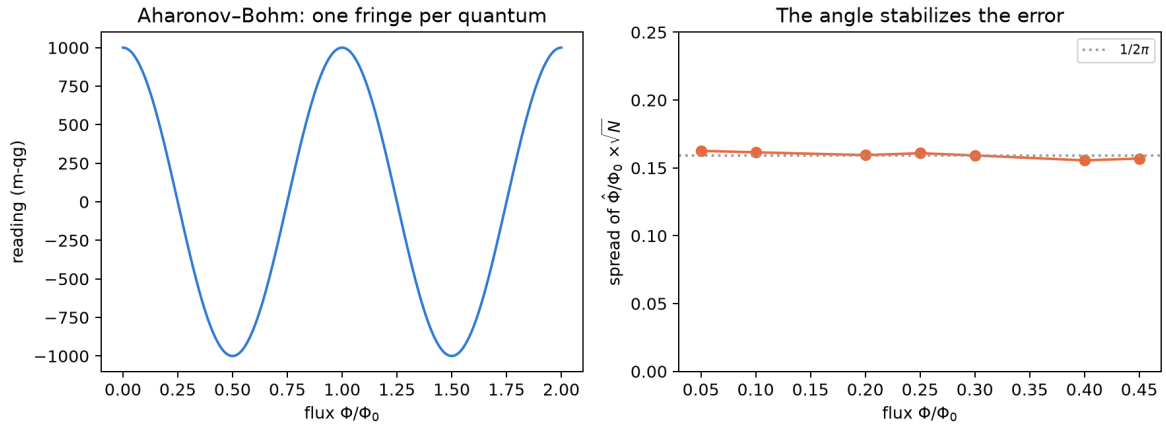


Figure 34.2: Left: the ring's reading in m-qg against flux; one fringe per quantum. Right: spread of the flux estimated from the qang angle, times $\sqrt{N}$, with 2000 electrons per measurement: it stays at $1/2\pi$ over the whole range.

34.4 Relativity and optics

34.4.1 Wigner rotation

Two non-collinear boosts, $v_1 = 0.6c\,\hat{x}$ and $v_2 = 0.8c\,\hat{y}$, compose into a boost followed by a rotation. The script multiplies them as 4×4 matrices, splits off the resulting boost and reads the angle: $\Omega_W = 0.3303$ rad ($18.92°$). The proposal's formula gives the same to $1.5 \cdot 10^{-15}$. With spinors, the same product in $SL(2,\mathbb{C})$ splits into a Hermitian part (the boost) and a unitary part (the rotation), and gives the same angle to $2 \cdot 10^{-16}$. The proposal's checks pass: $\Lambda^T \eta \Lambda = \eta$ to $9 \cdot 10^{-16}$, the rotation is orthogonal and the spinor part is unitary. A spin that pointed along x afterwards reads $qg_X = \cos\Omega_W$, that is $+945.9$ m-qg. With collinear boosts, $\Omega_W = 0$. All of this is structure of the Lorentz group: the scale only translates the angle into a spin reading. With Dirac spinors the result is the same, because a Dirac spinor is two $SL(2,\mathbb{C})$ spinors. On the angular qga scale of Chapter 1, $\Omega_W = 52.6$ qga. That chapter's energy and antimatter proposal (Section 1.4.1)

uses another angle of the same group, the θ_r with $\cos\theta_r = 1/\gamma$; it is not used here, and there it remains unverified.

34.4.2 Fresnel: Brewster is a zero of probability, total reflection is a phase

From air into glass with $n = 1.5$:

- **Brewster**, at $56.31°$: $r_p = 0$ exactly. s light reflects $R_s = 0.148$, which as a "reflected" reading is -704 m-qg. At normal incidence, $R = 0.04$, that is -920 m-qg.
- **Total reflection**, from inside the glass and above $41.81°$: $|r_s| = |r_p| = 1$ to $4\cdot10^{-16}$. The reflectance always reads $+1000$ m-qg and says nothing. All the information is in the phases. Their difference $\delta_p - \delta_s$ reaches a maximum of $45.24°$ at $51.67°$ and passes through $45°$ at $50.23°$ and $53.26°$: the Fresnel rhomb. Linear light at $45°$ incident at $45°$ leaves with $qg_Y = S_3/S_0 = \sin(\delta_p - \delta_s) = 0.600$, that is $+600$ m-qg, a Stokes reading like those of Chapter 6.
- **Goos–Hänchen** is not a "phase shift". It is a lateral displacement that comes from the derivative of the phase: with Artmann's formula, $D = -(\lambda/2\pi n)\, d\delta/d\theta_i$, it is 0.64 λ for s at $45°$ (Goos and Hänchen 1947; Artmann 1948).

The total-reflection phase as a "single-qubit phase gate" works: it acts on the polarization and is read in qg_Y.

34.4.3 Bogoliubov and Unruh: the half angle is the Gudermannian

The Minkowski vacuum seen by an accelerated observer is a two-mode squeezed state. Each reduced mode is thermal, with $\bar{n} = \sinh^2 r$ and Boltzmann factor $\tanh^2 r = e^{-2\pi\omega c/a}$ (Unruh 1976; Crispino et al. 2008). The script checks the closed-form entropy, $S = \cosh^2 r \ln\cosh^2 r - \sinh^2 r \ln\sinh^2 r$, against the sum over occupation numbers: they agree to 10^{-16} up to $r = 1$, and to $2\cdot10^{-6}$ at $r = 1.5$ because of the 80-level cutoff. The hyperbolic matrix preserves the form $\mathrm{diag}(1,-1)$ to $2\cdot10^{-16}$. The proposal calls it a "pseudo-unitary rotation parametrized in qg". An $SU(1,1)$ transformation is not a rotation, but the vacuum weight is a probability, $p_0 = 1/\cosh^2 r$, and it gives an exact identity:

$$\frac{\theta}{2} = \mathrm{gd}(r) = \arctan(\sinh r).$$

The qang half-angle of the vacuum weight is the Gudermannian function of the squeezing parameter. It maps the hyperbolic line onto a circle angle, and the script checks it to 10^{-16}. The vacuum readings are $+880.0$, $+572.9$, -160.1 and -638.6 m-qg for $r = 0.25$, 0.5, 1 and 1.5. As $r \to \infty$ (huge acceleration or zero frequency), $\theta \to \pi$.

34.5 Thermodynamics and metrology

34.5.1 Jarzynski with two measurements

A qubit with $H = -\frac{h}{2}(\cos\varphi\, Z + \sin\varphi\, X)$ is driven from $h = 1$, $\varphi = 0$ to $h = 2$, $\varphi = 1$ rad in a time 1.5 (400 steps), with $\beta = 1$. Work is measured with two projective measurements (Talkner et al. 2007). The transition probabilities $|\langle n_f|U|m_i\rangle|^2$ are readings: the two "stay" probabilities are $qg = 0.707$, that is 707.4 m-qg.

- **Exact**: $\langle e^{-\beta W}\rangle = 1.3684330 = Z_f/Z_i$, with error $3.8{\cdot}10^{-15}$. $\Delta F = -0.3137$ and $\langle W\rangle = -0.0958 \geq \Delta F$.
- **With sampled trajectories** (2000 repetitions), the estimator $-\ln\overline{e^{-\beta W}}/\beta$ has a bias of $+0.021$ with 10 trajectories, $+0.0016$ with 100 and $+1.5\cdot10^{-4}$ with 1000. With 10,000 trajectories the bias is within the noise.

The equality holds (Jarzynski 1997), and the scale changes nothing in the thermodynamics. What it adds is that each transition probability is a reading with the Chapter 1 error bar.

34.5.2 Fisher information of a GHZ state

For $N = 1$ to 8 qubits and the phase generated by $\sum Z/2$, the script gives a quantum Fisher information of N^2 for the GHZ state and N for the product state, with 10^{-15} error. The step that connects to Chapter 1 is the parity reading. Its mean is $qg = \cos N\varphi$, and the single-shot Fisher information of a ±1 reading is $(dqg/d\varphi)^2/(1-qg^2)$. That gives exactly N^2 at any φ: **the parity reading reaches the quantum bound** (Bollinger et al. 1996; Giovannetti et al. 2006). The proposal asks to use "qg as a continuous differentiable parameter". If the parameter were the reading itself, the information would be divided by $(dqg/d\varphi)^2$ and diverge at the poles. The scale reads, and the parameter is still the phase. The cost of read-out noise, $(1-2e)^N$, is in Chapter 8.

34.6 Geometry and chaos

34.6.1 Fubini–Study: pole to pole, half the qang angle

The script sums, in 20,000 steps, the distances $\arccos|\langle\psi_k|\psi_{k+1}\rangle|$ along a meridian and gets $1.5707963 = \pi/2$. With $\sqrt{1-|\langle\cdot|\cdot\rangle|^2}$ it gets the same, and also with the qg coordinate, $ds = dqg/(2\sqrt{1-qg^2})$, although that integrand is singular at the poles. The Gaussian curvature of $ds^2 = \frac{1}{4}(d\theta^2 + \sin^2\theta\, d\phi^2)$ comes out as 4.000: a sphere of radius 1/2 (Provost and Vallee 1980). The proposal speaks of "half a qg metric period". Translated to the base, **the Fubini–Study distance is half the qang angle**, $d_{FS} = \theta/2$, the Wootters distance of Chapter 1.

34.6.2 OTOC: integrable against chaotic

On an 8-spin chain, with $W = Z_1$, $V = Z_8$ and infinite temperature, the script compares the transverse Ising model ($g = 1.05$, integrable) with the same model plus a longitudinal field of 0.5 (chaotic):

- **The front arrives at the same time**: $F(t) = \mathrm{Re}\langle W(t)VW(t)V\rangle$ drops below 0.5 at $t = 4.0$ and 4.2.
- **What separates them is what comes after**: the chaotic one stays at 0.036 on average ($+36$ m-qg) for t between 10 and 30, and the integrable one revives, with mean 0.839 (Figure 34.3).
- **The usual diagnostic agrees**: the mean level-spacing ratio (10 spins, reflection-even sector) is 0.532 for the chaotic chain (GOE: 0.531) and 0.387 for the integrable one (Poisson: 0.386) (Oganesyan and Huse 2007; Atas et al. 2013).

The proposal's test cannot be done here. With 8 spins there is no exponential window separate from the front, and the bound $\lambda_L \le 2\pi k_B T/\hbar$ (Maldacena et al. 2016) is empty at infinite temperature. That bound has content in large-N models at low temperature, such as SYK (Sachdev and Ye 1993; Maldacena and Stanford 2016), not in an 8-spin chain. For Paulis, $F(t)$ is the qg_X reading of an ancilla in an interferometric protocol (Swingle et al. 2016), the same kind of reading as the decoder of Chapter 15.

34.7 The test matrix, revised

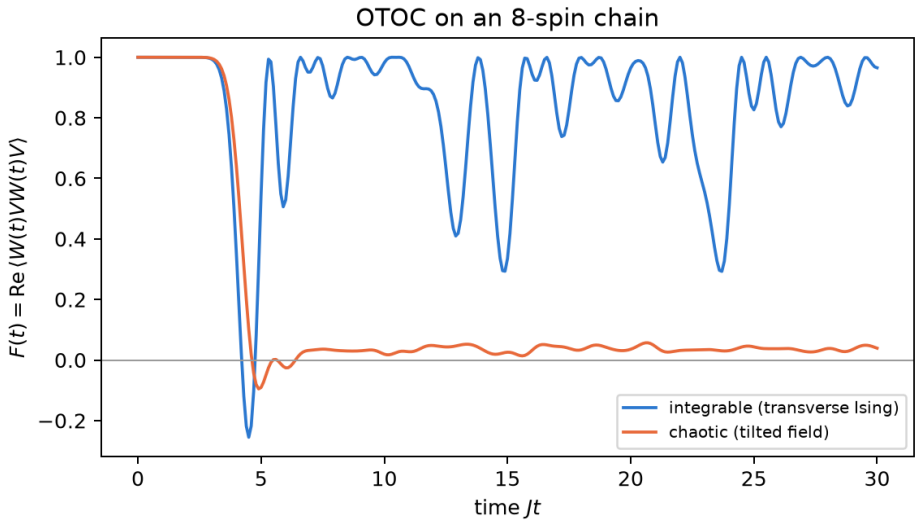


Figure 34.3: OTOC $F(t)$ of Z_1 and Z_8 on an 8-spin chain at infinite temperature. The front arrives at the same time in both; the chaotic chain stays near zero and the integrable one revives.

Table 34.3: The proposal's thirteen criteria, run.

phenomenon	proposal's criterion	result
Jordan–Wigner	accumulated phase $= \pi \equiv 0.5$ qg cycle	**corrected**: Hadamard test, -1000 m-qg (qubit SWAP: $+1000$)
UCCSD (H_2)	$E \approx -1.137$ Ha, zero gradient in qg	**holds**: FCI to $5 \cdot 10^{-14}$; HF at $+975.1$ m-qg; the curve is not redone
Hahn echo	$F = 1$ independent of $\Delta\omega$	**holds only static**: 2 % pulse error, -1 m-qg; drift, -180
transverse Ising	transition at $\theta = \pi/4$ with 2–4 spins	**fails**: the peak is at $g/J = 0.33$–0.65; it approaches 1 with N
Kitaev	braid invariant under $\epsilon(\theta)$	**fails as a gate** ($6 \cdot 10^{-5}$ at 1 %); the protection is the Hamiltonian's
Aharonov–Bohm	periodicity every quantum	**holds** to 10^{-15}; the flux is periodic, qg is not; the sign needs qg_Y

phenomenon	proposal's criterion	result
Wigner	orthogonality and Lorentz metric	**holds** to 10^{-15}; $\Omega_W = 18.92°$
Brewster	exact zero at θ_B	**holds**: $r_p = 0$ at $56.31°$
Goos–Hänchen	phase shift	**corrected**: it is a displacement, 0.64 λ
Bogoliubov	entropy of the thermal state	**holds**, and $\theta/2 = \mathrm{gd}(r)$ appears
Jarzynski	identity with theoretical ΔF	**holds** to $4 \cdot 10^{-15}$; sampling is biased
GHZ QFI	N^2 against N	**holds**; parity reaches N^2
Fubini–Study	curvature $K = 4$	**holds**: $K = 4.000$; pole to pole $\pi/2 = \theta/2$
OTOC	polynomial against exponential, MSS bound	**not applicable** with 8 spins; the late-time mean does separate

34.8 Without qang and with qang

Table 34.4: The thirteen tests, with and without the scale.

topic	without qang	with qang	difference
UCC	amplitude t	HF weight $+975.1$ m-qg, $\theta = 2\lvert t\rvert$	the amplitude is half an angle; same energy
Aharonov–Bohm	fringes in P	$\theta = 2\pi\Phi/\Phi_0$, spread $1/2\pi\sqrt{N}$ constant	even error bar (arcsine); the sign needs another reading
Bogoliubov	$\tanh^2 r$, $\bar{n}$, S	$\theta/2 = \mathrm{gd}(r)$	a new, exact identity
GHZ	QFI $= N^2$	parity Fisher $(dqg)^2/(1-qg^2) = N^2$	Chapter 1's single-shot formula reaches the bound
Fubini–Study	$d_{FS} = \arccos\lvert\langle\psi\rvert\phi\rangle\rvert$	$d_{FS} = \theta/2$	the same distance (Wootters)
echo, Jarzynski	$\langle X\rangle$, probabilities	readings in m-qg	unit only

topic	without qang	with qang	difference
Ising	entropy, $\langle X \rangle$	$m_X \to 636.6$	none: the scale does not bring the limit closer
Wigner, Kitaev, OTOC	Lorentz group, Hamiltonian, spectrum	angle translated into a reading	**none**; "qg cycles" neither protect nor measure chaos

Verdict. The proposal is right in its intuition: almost everything it looks at has an angle or a probability inside. It goes wrong in the currency, because "qg cycle" and "qg period" do not exist in the base. Translated into readings, the scale contributes four concrete things: the UCC amplitude as half an angle, the constant error bar of the flux, the Gudermannian of the squeezed vacuum, and the GHZ parity at the quantum bound from the single-shot formula. For the echo and Jarzynski it changes only the unit. For Wigner, Kitaev and the OTOC it contributes nothing: the physics is in the group, the Hamiltonian or the spectrum. Two criteria fail as written: Ising with 2 to 4 spins and the braid as a protected gate. Everything is exact computation or simulation, and there is no qang module for these topics yet.

34.9 In code

```
import math
from qang import probability as PB

# UCC for H2: Hartree-Fock weight at the minimum
t = -0.11176849913374033
PB.milliqang_from_p0(math.cos(t) ** 2)       # +975.1 m-qg; qang angle 2|t| = 0.2235 rad

# Bogoliubov: the qang half-angle of the vacuum is the Gudermannian
r = 1.0
theta = PB.theta_from_milliqang(PB.milliqang_from_p0(1 / math.cosh(r) ** 2))
theta / 2, math.atan(math.sinh(r))              # 0.8658 and 0.8658
```

34.10 Summary

- The base has no "qg cycles": qg is a reading in $[-1, 1]$. Phases are read in qg_X or qg_Y with an interferometer, and probabilities directly.
- The sign of a fermionic exchange is -1000 m-qg in a Hadamard test; the UCC amplitude of H_2 is half a qang angle ($+975.1$ m-qg at the FCI minimum).
- The Hahn echo recovers $+1000$ m-qg only with static detuning; Ising with 2 to 4 spins does not show the transition at $\theta = \pi/4$, and the Majorana braid written as a gate is not protected.
- The qang angle of an Aharonov–Bohm ring is the phase, with a constant $1/2\pi\sqrt{N}$ error bar; the half-angle of the squeezed vacuum is $\mathrm{gd}(r)$.
- GHZ parity reaches the quantum information N^2 with Chapter 1's single-shot formula, and the Fubini–Study distance is half the qang angle.
- Wigner, Jarzynski and the OTOC hold or are diagnosed with the usual tools; the scale only translates them.

Bibliographic notes

The Jordan–Wigner mapping is Jordan and Wigner (1928), and UCC for VQE, Romero et al. (2018); the H_2 energy follows O'Malley et al. (2016). The echo is Hahn (1950), and precession, Bloch (1946). The solved transverse Ising model is Pfeuty (1970), and entanglement entropy at criticality, Calabrese and Cardy (2004). The Kitaev chain is Kitaev (2001), and Majorana braiding, Ivanov (2001). The Aharonov–Bohm effect is Aharonov and Bohm (1959). Thomas precession is Thomas (1926); Wigner's representations, Wigner (1939), and the composition of boosts, O'Donnell and Visser (2011). Fresnel and Brewster follow Born et al. (1999); the lateral shift is Goos and Hänchen (1947) and Artmann (1948). The Unruh effect is Unruh (1976) (review: Crispino et al. (2008)). The Jarzynski equality is Jarzynski (1997), and work from two measurements, Talkner et al. (2007). Spectroscopy with GHZ states is Bollinger et al. (1996), and quantum metrology, Giovannetti et al. (2006) and Braunstein and Caves (1994). The Fubini–Study metric as a Riemannian structure is Provost and Vallee (1980). The chaos bound is Maldacena et al. (2016); SYK, Sachdev and Ye (1993) and Maldacena and Stanford (2016); measuring scrambling, Swingle et al. (2016), and spacing ratios, Oganesyan and Huse (2007) and Atas et al. (2013). The scripts that check this chapter are in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

Part XIII

XIII. Ecosystem: qang inside other libraries

35 qang as a module: changes to quantum and machine-learning libraries

Every chapter of this book noted what it would ask of an existing library. This chapter delivers that list. qang 0.13.0 ships a new subpackage, `qang.integrations`, with one module per library: mitiq, qiskit-experiments, Qiskit Dynamics, QuTiP, Stim, scikit-learn, JAX, PyTorch, Qibo, pytket, Amazon Braket and NVIDIA's CUDA-Q. Each module brings qang's three things (the readings, the filter and the closed-form rules) into that library's own terms and has its test in `tests/test_integrations.py`. The figures in this chapter come from `tools/make_cap35_numbers.py`, which runs in three environments because not every library fits in one. None comes from real hardware. The chapter separates what is done and tested inside qang from what would still have to be proposed to each project.

The next version, 0.14.0, adds the module `qang.qga` with the angular scale of Chapter 1 (Section 1.4): 1000 qga per turn and the conversions to radians, p_0 and m-qg. The integrations of this chapter do not change, and their figures were measured with 0.13.0.

35.1 From the base: what each library gets

The Chapter 1 base is a probability p_0 and its readings, $qg_Z = 2p_0 - 1$, $m = 1000\, qg_Z$ and $\theta = \arccos qg_Z$. What qang adds to another library is always one of three things:

- **a reading**: turn what the library already measures (counts, populations, detectors, probabilities) into qg_Z and m-qg, with a single qubit order (Chapter 32);
- **the filter**: keep the shots in the circuit's Hamming-weight sector, with the kept fraction K (Chapters 21 and 22);
- **a closed-form rule**: the filter's cost rule (Chapter 21), the closed form of heralded calibration (§31) or the minimal pulse area (§160).

35.2 How it is built

- **Nothing extra gets imported.** `import qang` still needs only NumPy. Each submodule imports its library when asked, and if the library is missing, the error says what to install: `pip install "qang[mitiq]"`, `qang[stim]`, and so on.
- **One extra per library** in `pyproject.toml`, plus one, `integrations`, with every library that can share an environment.
- **Each test skips itself** if its library is not installed. That is why the repository's CI, which installs `qang[all]`, runs the scikit-learn tests and skips the rest.

Table 35.1: The integrations of qang 0.13.0.

library	module	version tested	environment	test
mitiq	`integrations.mitiq`	1.1.0	`.venv-int`	passes
qiskit-experiments	`integrations.qiskit_experiments`	0.11.2 (Qiskit 2.6.1)	`.venv-int`	passes
Qiskit Dynamics	`integrations.qiskit_dynamics`	0.6.0 (Qiskit 1.3.0)	`.venv-dyn`	passes
QuTiP	`integrations.qutip`	5.3.1	`.venv-int`	passes
Stim + PyMatching	`integrations.stim`	1.16.0 + 2.4.0	`.venv-int`	passes
scikit-learn	`integrations.sklearn`	1.9.1	`.venv-int` and CI	passes
JAX (+ PennyLane)	`integrations.jax`	0.8.3 (PennyLane 0.45.1)	`.venv-int`	passes
PyTorch (+ PennyLane)	`integrations.torch`	2.11.1 (CPU)	`.venv-int`	passes
Qibo	`integrations.qibo`	0.3.6	`.venv-int`	passes
pytket	`integrations.pytket`	2.14.5	`.venv-int`	passes
Amazon Braket	`integrations.braket`	1.127.5	`.venv-int`	passes
CUDA-Q	`integrations.cudaq`	0.16.0	WSL (Ubuntu)	passes (CPU)

35.3 Mitigation: mitiq

mitiq can already post-select on Hamming weight: its `post_select` ships with that very example. There was no need to write it again. What qang adds on top is:

- `qg_z` and `filtered_qg_z`, to read a `MeasurementResult` in qg;
- `qg_executor`, which turns a sampler into the real-valued executor that mitiq's techniques use;
- `execute_with_filter_zne`, which combines the filter with zero-noise extrapolation;
- `filter_rule`, the Chapter 21 rule.

The script runs a three-qubit weight-1 circuit (Givens rotations in Cirq), with amplitude damping of 0.05 and dephasing of 0.03 per moment, and reads qubit 1. The exact value is -257.3 m-qg:

Table 35.2: Error in m-qg of the qubit-1 reading, Richardson ZNE with factors 1, 3 and 5 and global folding.

noise	raw	filtered	raw + ZNE	filtered + ZNE
T1	233.2	**0.0**	49.4	**0.0**
T1 + dephasing	257.9	30.3	55.9	**1.1**

With T1 alone the filter is exact (Chapter 21, F1) and ZNE has nothing to correct. With T1 and dephasing, the filter leaves only the dephasing and ZNE removes it (F6): the error drops from 30.3 to 1.1 m-qg. ZNE on the raw reading stays at 55.9 m-qg, because T1 does not scale well under folding. **The proposal for mitiq** is the combination: post-selection, then ZNE, with the cost rule deciding whether to filter. The separate pieces already existed.

35.4 Calibration: qiskit-experiments

`HeraldedReadout` is an experiment in the framework with two circuits (herald, I or X, measure) and an analysis with the §31 closed form. It runs on the Aer simulator with readout error $e_{01} = 0.015$, $e_{10} = 0.04$, and with a thermal population p prepared by purification:

Table 35.3: Standard calibration against heralded calibration, 200,000 shots per circuit.

p	e_{01} standard	e_{01} qg	e_{10} standard	e_{10} qg	p qg	T_{eff} (5 GHz)
0	0.0148	0.0151	0.0401	0.0396	0.0000	0 mK
0.01	0.0244	0.0151	0.0494	0.0397	0.0099	52.1 mK
0.03	0.0431	0.0151	0.0683	0.0396	0.0297	68.8 mK

p	e_{01} standard	e_{01} qg	e_{10} standard	e_{10} qg	p qg	T_{eff} (5 GHz)
0.08	0.0900	0.0150	0.1152	0.0401	0.0794	97.9 mK

The standard calibration adds the temperature to the readout error, $e_{01} + p\,b$, and sees the qubit as cold. The closed form separates the two. This is the result of §31 and Chapter 27, now inside a `BaseExperiment` run with `exp.run(backend)`. **It is the most concrete proposal in this chapter** for qiskit-experiments, because Chapter 18 showed that the standard calibration can leave a random-number certifier unsafe.

35.5 Open systems: QuTiP

`qg_readings` reads (qg_X, qg_Y, qg_Z) of each qubit of a `Qobj`, `mesolve_qg` reads a `mesolve` evolution, and `t1_collapse` builds the collapse operators for equal T1. With a three-qubit W state and $t/T_1 = 0.6$, the raw reading of each qubit is $+634.1$ m-qg and the filtered one is $+333.333$, the noiseless value. The kept fraction is $0.5488 = e^{-0.6}$, as Chapter 21 predicts. It is the same check Chapters 19 and 21 made by integrating Lindblad by hand, now with QuTiP's integrator.

35.6 Error correction: Stim

A detector that fires with probability p reads $qg = 1 - 2p$. With the distance-5 repetition code, 5 rounds and 200,000 samples:

Table 35.4: qg readings of a repetition code in Stim; decoding with PyMatching.

noise p	detectors (m-qg)	raw logical	decoded logical
0.01	972.4 to 1000.0	933.7	1000.0
0.03	919.0 to 1000.0	815.2	999.1
0.06	841.8 to 1000.0	660.1	993.7

Decoding is bringing the logical qubit to the pole. The last-round detectors sit at $+1000$ because the noise enters before each round. `measurement_qg` also applies the filter to the measurement records. **The proposal for Stim/PyMatching** is still the one from Chapter 26, the T1-aware reweighting (2.1 to 2.5 times), which this module does not implement yet.

35.7 Learning: scikit-learn, JAX and PyTorch

- **scikit-learn.** `WeightQNNClassifier` and `MultiClassQNNClassifier` have `fit`, `predict`, `decision_function` and `score`, accept any labels and work in `Pipeline`, `cross_val_score` and `GridSearchCV`. On the two-moons set (120 points), with T1 of 0.15 and 5-fold cross-validation, the filtered model scores 0.833 ± 0.083 and the raw one 0.842 ± 0.081. **There is no difference**, as in Chapter 20: trained under the calibrated noise, the model without the filter catches up. `GridSearchCV` can choose `qang=True/False` like any other hyperparameter.
- **JAX.** The readings and the filter written with `jax.numpy` go through `jit`, `grad` and `vmap`. The gradient of a filtered loss through a PennyLane QNode with `interface="jax"` matches finite differences to $2 \cdot 10^{-10}$. That was what Chapter 20 left untested. Getting there turned up a conflict: PennyLane 0.45 uses `jax.core.is_concrete`, which JAX 0.12 removed, and it works only with JAX < 0.9.
- **PyTorch.** `WeightQNNTorch` is an `nn.Module`: a PennyLane `TorchLayer` with `SingleExcitation` rotations, which conserve the weight, the filtered reading and a linear layer. With T1 of 0.1 per layer, the filtered reading differs from the clean one by $1.2 \cdot 10^{-16}$ and the raw one by 262.5 m-qg. The gradient flows through `backward()`.

35.8 Pulses: Qiskit Dynamics

`rabi_qg` solves a resonantly driven qubit and returns qg_Z. Against $\cos A$, with A the pulse area, the error is $1.1 \cdot 10^{-7}$ to $5 \cdot 10^{-7}$ with a square pulse and $5 \cdot 10^{-7}$ to $1.9 \cdot 10^{-6}$ with a Gaussian pulse, for A equal to $\pi/2$, π and 2.2. It is the basis of the §160 minimal-area rule. **The conflict**: Qiskit Dynamics 0.6 requires `qiskit<=1.3`, so it does not coexist with Qiskit 2.x and needs its own environment.

35.9 SDKs: Qibo, pytket, Braket and CUDA-Q

Each gets the qang rotation ($R_Y(\arccos qg)$) and the canonical reading with the filter:

- **Qibo**: the reading of $qg = 0.6$ with 100,000 shots gives 0.5978;

- **pytket**: gives the exact value, -0.25, from the statevector. pytket measures angles in half-turns, not radians, and `rqang` does the conversion. It is the kind of unit error this scale is meant to prevent. On the qga scale of Chapter 1 a half-turn is 500 qga, so going from pytket to qga is multiplying by 500;
- **Braket**: runs on the `LocalSimulator`, so it is no longer just a result format (Chapter 32); the reading of 0.2 gives 0.201;
- **CUDA-Q**: runs in WSL with CUDA-Q 0.16.0. On the CPU simulator (`qpp-cpu`), the reading of -0.5 gives -0.5018 and the filter keeps 0.7509, against 0.75 expected. The GPU simulator (`nvidia`, cuStateVec) did not start, because the CUDA driver WSL sees is older than CUDA-Q's runtime. The Windows NVIDIA driver needs updating; the code does not change, only `cudaq.set_target("nvidia")`.

35.10 What to propose to each project

What is in `qang.integrations` already works. What follows is what would be worth taking to each project, ordered by the backing it has in the book:

Table 35.5: Proposals for each project, with their backing.

project	proposal	backing	strength
qiskit-experiments	heralded calibration: separates temperature and readout	§31, Chapters 18 and 27, Table 35.3	**strong**
mitiq	post-selection + ZNE with the cost rule	Chapter 21, Table 35.2	**strong**
scikit-learn (via qang)	weight QNN as an estimator	Chapter 20	useful, no accuracy gain
PennyLane	fix compatibility with JAX ≥ 0.9 (`is_concrete`)	this chapter	a real ecosystem bug
Qiskit Dynamics	Qiskit 2.x support	this chapter	a real ecosystem bug
PyMatching / Stim	T1-aware reweighting	Chapter 26	moderate (loses 15 % under depolarizing noise)

project	proposal	backing	strength
Qiskit transpiler	keep N-conserving gates	Chapter 24	moderate
qiskit-machine-learning	filter as a Gram-matrix repair	Chapter 23	weak (a single kernel)
PennyLane	NaN in the jacobian of `qml.probs` through `SingleExcitation` with `default.mixed` and `backprop` (0.45.1)	Chapter 20	a real ecosystem bug
qiskit-machine-learning	`SamplerQNN` reads the bits at the physical *layout* positions (0.9.1)	Chapter 20, §140	a real bug, quoted from the notes

35.11 Environment conflicts

Putting everything in one environment was not possible, and the conflicts are information:

- Qiskit Dynamics 0.6 requires `qiskit<=1.3`;
- PennyLane 0.45 does not work with JAX ≥ 0.9;
- amazon-braket-sdk pins `cloudpickle==2.2.1`, while recent joblib (used by scikit-learn) needs ≥ 3.0. Both worked with cloudpickle 3.1.2, but pip warns about the conflict;
- mitiq 1.1 pulled NumPy down to 2.2.6;
- CUDA-Q has no wheels for Windows and was tested in WSL; its GPU simulator needs a newer NVIDIA driver.

That is why there are three environments: `.venv-int` with almost everything, `.venv-dyn` for Qiskit Dynamics and WSL for CUDA-Q. The book's `.venv` was left untouched.

35.12 Without qang and with qang

Table 35.6: The integrations, with and without the scale.

library	without qang	with qang	difference
mitiq	`post_select` and ZNE separately	ZNE of the filtered reading + rule	55.9 → 1.1 m-qg under T1 and dephasing
qiskit-experiments	e_{01} inflated by $p\,b$, "cold" qubit	unbiased e_{01}, p and T_{eff}	separates two errors the standard mixes up
QuTiP	operator expectations	readings in m-qg and the filter	filter exact under T1 to 10^{-6}
Stim	detection rates	detectors and logicals in m-qg	unit only
scikit-learn	no estimator for the QNN	`fit`/`predict`, CV, search	API; same accuracy (0.833 against 0.842)
JAX	filtered reading untested	`jit`/`grad`, $2 \cdot 10^{-10}$	tested, with JAX < 0.9
PyTorch	no weight layer	`WeightQNNTorch`	filter exact under T1 inside torch
Dynamics	populations	$qg_Z = \cos A$ to 10^{-6}	unit only
SDKs	different bit orders and units	one canonical reading	prevents order and half-turn errors

Verdict. Two integrations add something the library does not have: heralded calibration, which separates temperature and readout, and ZNE of the filtered reading with the cost rule. The learning ones open up use without improving accuracy, which is what Chapter 20 had already measured. The rest are translation: readings in m-qg and a single qubit order, which is little but prevents real errors, such as pytket's half-turns. No figure comes from hardware. Three proposals were posted on 25 October 2026: the heralded calibration to qiskit-experiments (issue #1662), post-selection with ZNE and its cost rule to mitiq (issue #3173), and T1-aware reweighting in a PyMatching discussion (#190). The drafts and the pre-send check are in `docs/upstream/`; the others have not been sent yet.

35.13 In code

```
# pip install "qang[mitiq]"  (or qang[experiments], qang[sklearn], qang[integrations],
from qang.integrations import mitiq as qm

executor = qm.qg_executor(sampler, weight=1)         # circuit -> mean filtered qg_Z
value = qm.execute_with_filter_zne(circuit, sampler, weight=1, scale_factors=(1, 3, 5))

from qang.integrations.qiskit_experiments import HeraldedReadout
exp = HeraldedReadout([0], backend=backend, frequency=5e9)
results = exp.run().block_for_results()            # e01, e10, p_excited, m_eq, T_eff

from qang.integrations.sklearn import WeightQNNClassifier
from sklearn.model_selection import GridSearchCV
GridSearchCV(WeightQNNClassifier(gamma=0.15), {"qang": [True, False]}, cv=3).fit(X, y)
```

35.14 Summary

- qang 0.13.0 has `qang.integrations`, with twelve libraries, one extra for each and tests that skip themselves; 0.14.0 adds `qang.qga`, the angular scale.
- mitiq: the filter with ZNE lowers the error from 55.9 to 1.1 m-qg under T1 and dephasing; post-selection already existed.
- qiskit-experiments: heralded calibration gives unbiased e_{01}, p and T_{eff}, where the standard one sees the qubit as cold.
- QuTiP, Stim, Qiskit Dynamics and the SDKs get readings in m-qg and the filter, and the four SDKs share the same qubit order.
- scikit-learn, JAX and PyTorch let weight networks train like any other model, without an accuracy gain.
- Three environment conflicts are documented: Qiskit Dynamics with Qiskit 2, PennyLane with JAX ≥ 0.9 and Braket with cloudpickle.

Bibliographic notes

mitiq is LaRose et al. (2022); QuTiP, Johansson et al. (2012); Stim, Gidney (2021), and PyMatching, Higgott (2022). qiskit-experiments is Kanazawa et al. (2023), and Qiskit Dynamics, Puzzuoli et al. (2023), on top of Qiskit (Javadi-Abhari et al. 2024). scikit-learn is Pedregosa et al. (2011); JAX, Bradbury et al. (2018);

PyTorch, Paszke et al. (2019), and PennyLane, Bergholm et al. (2018). Qibo is Efthymiou et al. (2021); pytket, Sivarajah et al. (2020); Amazon Braket, Amazon Web Services (2020); Cirq, Cirq Developers (2026), and CUDA-Q, Kim et al. (2023). The symmetry verification mitiq already does is that of Bonet-Monroig et al. (2018) and McArdle et al. (2019), and zero-noise extrapolation, Temme et al. (2017). The code is in the repository (Monteverde 2026b); the definition of the qang unit is in Monteverde (2026c).

References

Works cited in Chapters 1 to 35. Each chapter closes with its bibliographic notes; the records with a DOI were generated from doi.org metadata.

Abdurakhimov, Leonid, Janos Adam, Hasnain Ahmad, et al. 2024. *Technology and Performance Benchmarks of IQM's 20-Qubit Quantum Computer*. arXiv. https://doi.org/10.48550/ARXIV.2408.12433.

Aguilar, A., others, and LSND Collaboration. 2001. "Evidence for Neutrino Oscillations from the Observation of $\bar{\nu}_e$ Appearance in a $\bar{\nu}_\mu$ Beam." *Physical Review D* 64 (11). https://doi.org/10.1103/physrevd.64.112007.

Aharonov, Y., and J. Anandan. 1987. "Phase Change During a Cyclic Quantum Evolution." *Physical Review Letters* 58 (16): 1593–96. https://doi.org/10.1103/physrevlett.58.1593.

Aharonov, Y., and D. Bohm. 1959. "Significance of Electromagnetic Potentials in the Quantum Theory." *Physical Review* 115 (3): 485–91. https://doi.org/10.1103/physrev.115.485.

Aharonov, Y., L. Davidovich, and N. Zagury. 1993. "Quantum Random Walks." *Physical Review A* 48 (2): 1687–90. https://doi.org/10.1103/physreva.48.1687.

Alagic, Gorjan, Maxime Bros, Pierre Ciadoux, et al. 2025. *Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process*. National Institute of Standards; Technology (U.S.). https://doi.org/10.6028/nist.ir.8545.

Alicki, Robert, and Mark Fannes. 2013. "Entanglement Boost for Extractable Work from Ensembles of Quantum Batteries." *Physical Review E* 87 (4). https://doi.org/10.1103/physreve.87.042123.

Allahverdyan, A. E, R Balian, and Th. M Nieuwenhuizen. 2004. "Maximal Work Extraction from Finite Quantum Systems." *Europhysics Letters (EPL)* 67 (4): 565–71. https://doi.org/10.1209/epl/i2004-10101-2.

Amari, Shun-ichi. 1998. “Natural Gradient Works Efficiently in Learning.” *Neural Computation* 10 (2): 251–76. https://doi.org/10.1162/089976698300017746.

Amazon Web Services. 2020. *Amazon Braket Python SDK*. Https://github.com/amazon-braket/amazon-braket-sdk-python.

Anscombe, F. J. 1948. “The Transformation of Poisson, Binomial and Negative-Binomial Data.” *Biometrika* 35 (3-4): 246–54. https://doi.org/10.1093/biomet/35.3-4.246.

Artmann, Kurt. 1948. “Berechnung Der Seitenversetzung Des Totalreflektierten Strahles.” *Annalen Der Physik* 437 (1-2): 87–102. https://doi.org/10.1002/andp.19484370108.

Arute, Frank, Kunal Arya, Ryan Babbush, et al. 2019. “Quantum Supremacy Using a Programmable Superconducting Processor.” *Nature* 574 (7779): 505–10. https://doi.org/10.1038/s41586-019-1666-5.

Atas, Y. Y., E. Bogomolny, O. Giraud, and G. Roux. 2013. “Distribution of the Ratio of Consecutive Level Spacings in Random Matrix Ensembles.” *Physical Review Letters* 110 (8). https://doi.org/10.1103/physrevlett.110.084101.

Barenco, Adriano, Charles H. Bennett, Richard Cleve, et al. 1995. “Elementary Gates for Quantum Computation.” *Physical Review A* 52 (5): 3457–67. https://doi.org/10.1103/physreva.52.3457.

Bartlett, M. S. 1947. “The Use of Transformations.” *Biometrics* 3 (1): 39. https://doi.org/10.2307/3001536.

Baumgratz, T., M. Cramer, and M. B. Plenio. 2014. “Quantifying Coherence.” *Physical Review Letters* 113 (14). https://doi.org/10.1103/physrevlett.113.140401.

Beenakker, C. W. J. 1991. “Theory of Coulomb-Blockade Oscillations in the Conductance of a Quantum Dot.” *Physical Review B* 44 (4): 1646–56. https://doi.org/10.1103/physrevb.44.1646.

Bell, J. S. 1964. “On the Einstein Podolsky Rosen Paradox.” *Physics Physique Физика* 1 (3): 195–200. https://doi.org/10.1103/physicsphysiquefizika.1.195.

Bennett, Charles H., and Gilles Brassard. 2014. “Quantum Cryptography: Public Key Distribution and Coin Tossing.” *Theoretical Computer Science* 560: 7–11.

https://doi.org/10.1016/j.tcs.2014.05.025.

Bennett, Charles H., Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. 1993. "Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels." *Physical Review Letters* 70 (13): 1895–99. https://doi.org/10.1103/physrevlett.70.1895.

Bergholm, Ville, Josh Izaac, Maria Schuld, et al. 2018. *PennyLane: Automatic Differentiation of Hybrid Quantum-Classical Computations*. arXiv. https://doi.org/10.48550/ARXIV.1811.04968.

Bernstein, Ethan, and Umesh Vazirani. 1997. "Quantum Complexity Theory." *SIAM Journal on Computing* 26 (5): 1411–73. https://doi.org/10.1137/s0097539796300921.

Berry, Michael Victor. 1984. "Quantal Phase Factors Accompanying Adiabatic Changes." *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 392 (1802): 45–57. https://doi.org/10.1098/rspa.1984.0023.

Berta, Mario, Matthias Christandl, Roger Colbeck, Joseph M. Renes, and Renato Renner. 2010. "The Uncertainty Principle in the Presence of Quantum Memory." *Nature Physics* 6 (9): 659–62. https://doi.org/10.1038/nphys1734.

Bhattacharyya, Anil Kumar. 1943. "On a Measure of Divergence Between Two Statistical Populations Defined by Their Probability Distributions." *Bulletin of the Calcutta Mathematical Society* 35: 99–109.

Bilenky, S. M., J. Hošek, and S. T. Petcov. 1980. "On the Oscillations of Neutrinos with Dirac and Majorana Masses." *Physics Letters B* 94 (4): 495–98. https://doi.org/10.1016/0370-2693(80)90927-2.

Blanter, Ya.M., and M. Büttiker. 2000. "Shot Noise in Mesoscopic Conductors." *Physics Reports* 336 (1-2): 1–166. https://doi.org/10.1016/s0370-1573(99)00123-4.

Bloch, F. 1946. "Nuclear Induction." *Physical Review* 70 (7-8): 460–74. https://doi.org/10.1103/physrev.70.460.

Bollinger, J. J ., Wayne M. Itano, D. J. Wineland, and D. J. Heinzen. 1996. "Optimal Frequency Measurements with Maximally Correlated States." *Physical Review A* 54 (6): R4649–52. https://doi.org/10.1103/physreva.54.r4649.

Bonet-Monroig, X., R. Sagastizabal, M. Singh, and T. E. O'Brien. 2018. "Low-

Cost Error Mitigation by Symmetry Verification." *Physical Review A* 98 (6). https://doi.org/10.1103/physreva.98.062339.

Born, Max. 1926. "Zur Quantenmechanik Der Stoßvorgänge." *Zeitschrift für Physik* 37 (12): 863–67. https://doi.org/10.1007/bf01397477.

Born, Max, Emil Wolf, A. B. Bhatia, et al. 1999. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press. https://doi.org/10.1017/cbo9781139644181.

Bradbury, James, Roy Frostig, Peter Hawkins, et al. 2018. *JAX: Composable Transformations of Python+NumPy Programs*. Https://github.com/jax-ml/jax.

Brassard, Gilles, Peter Høyer, Michele Mosca, and Alain Tapp. 2002. "Quantum Amplitude Amplification and Estimation." In *Quantum Computation and Information*. American Mathematical Society. https://doi.org/10.1090/conm/305/05215.

Braunstein, Samuel L., and Carlton M. Caves. 1994. "Statistical Distance and the Geometry of Quantum States." *Physical Review Letters* 72 (22): 3439–43. https://doi.org/10.1103/physrevlett.72.3439.

Bravyi, Sergey, Sarah Sheldon, Abhinav Kandala, David C. Mckay, and Jay M. Gambetta. 2021. "Mitigating Measurement Errors in Multiqubit Experiments." *Physical Review A* 103 (4). https://doi.org/10.1103/physreva.103.042605.

Briegel, H.-J., W. Dür, J. I. Cirac, and P. Zoller. 1998. "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication." *Physical Review Letters* 81 (26): 5932–35. https://doi.org/10.1103/physrevlett.81.5932.

Buhrman, Harry, Richard Cleve, John Watrous, and Ronald de Wolf. 2001. "Quantum Fingerprinting." *Physical Review Letters* 87 (16). https://doi.org/10.1103/physrevlett.87.167902.

Büttiker, M. 1986. "Four-Terminal Phase-Coherent Conductance." *Physical Review Letters* 57 (14): 1761–64. https://doi.org/10.1103/physrevlett.57.1761.

Calabrese, Pasquale, and John Cardy. 2004. "Entanglement Entropy and Quantum Field Theory." *Journal of Statistical Mechanics: Theory and Experiment* 2004 (06): P06002. https://doi.org/10.1088/1742-5468/2004/06/p06002.

Campaioli, Francesco, Stefano Gherardini, James Q. Quach, Marco Polini, and

Gian Marcello Andolina. 2024. “Colloquium: Quantum Batteries.” *Reviews of Modern Physics* 96 (3). https://doi.org/10.1103/revmodphys.96.031001.

Cerezo, M., Martin Larocca, Diego García-Martín, et al. 2023. *Does Provable Absence of Barren Plateaus Imply Classical Simulability?* arXiv. https://doi.org/10.48550/ARXIV.2312.09121.

Cerezo, M., Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. 2021. “Cost Function Dependent Barren Plateaus in Shallow Parametrized Quantum Circuits.” *Nature Communications* 12 (1). https://doi.org/10.1038/s41467-021-21728-w.

Chin, Alex W., Susana F. Huelga, and Martin B. Plenio. 2012. “Quantum Metrology in Non-Markovian Environments.” *Physical Review Letters* 109 (23). https://doi.org/10.1103/physrevlett.109.233601.

Chou, C. W., D. B. Hume, T. Rosenband, and D. J. Wineland. 2010. “Optical Clocks and Relativity.” *Science* 329 (5999): 1630–33. https://doi.org/10.1126/science.1192720.

Cirq Developers. 2026. *Cirq*. Zenodo. https://doi.org/10.5281/ZENODO.4062499.

Coffman, Valerie, Joydip Kundu, and William K. Wootters. 2000. “Distributed Entanglement.” *Physical Review A* 61 (5). https://doi.org/10.1103/physreva.61.052306.

Correa, Luis A., Mohammad Mehboudi, Gerardo Adesso, and Anna Sanpera. 2015. “Individual Quantum Probes for Optimal Thermometry.” *Physical Review Letters* 114 (22). https://doi.org/10.1103/physrevlett.114.220405.

Cramér, Harald. 1946. *Mathematical Methods of Statistics*. Princeton University Press.

Crispino, Luís C. B., Atsushi Higuchi, and George E. A. Matsas. 2008. “The Unruh Effect and Its Applications.” *Reviews of Modern Physics* 80 (3): 787–838. https://doi.org/10.1103/revmodphys.80.787.

Cross, Andrew W., Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. 2019. “Validating Quantum Computers Using Randomized Model Circuits.” *Physical Review A* 100 (3). https://doi.org/10.1103/physreva.100.032328.

Datta, Supriyo. 1995. *Electronic Transport in Mesoscopic Systems*. Cambridge University Press. https://doi.org/10.1017/cbo9780511805776.

Dennis, Eric, Alexei Kitaev, Andrew Landahl, and John Preskill. 2002. "Topological Quantum Memory." *Journal of Mathematical Physics* 43 (9): 4452–505. https://doi.org/10.1063/1.1499754.

Deutsch, David. 1991. "Quantum Mechanics Near Closed Timelike Lines." *Physical Review D* 44 (10): 3197–217. https://doi.org/10.1103/physrevd.44.3197.

Deutsch, David, and Richard Jozsa. 1992. "Rapid Solution of Problems by Quantum Computation." *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439 (1907): 553–58. https://doi.org/10.1098/rspa.1992.0167.

Dirac, P. A. M. 1928. "The Quantum Theory of the Electron." *Proceedings of the Royal Society of London. Series A* 117 (778): 610–24. https://doi.org/10.1098/rspa.1928.0023.

Doherty, Marcus W., Neil B. Manson, Paul Delaney, Fedor Jelezko, Jörg Wrachtrup, and Lloyd C. L. Hollenberg. 2013. "The Nitrogen-Vacancy Colour Centre in Diamond." *Physics Reports* 528 (1): 1–45. https://doi.org/10.1016/j.physrep.2013.02.001.

Dolinski, Michelle J., Alan W. P. Poon, and Werner Rodejohann. 2019. "Neutrinoless Double-Beta Decay: Status and Prospects." *Annual Review of Nuclear and Particle Science* 69 (1): 219–51. https://doi.org/10.1146/annurev-nucl-101918-023407.

Dréau, A., M. Lesik, L. Rondin, et al. 2011. "Avoiding Power Broadening in Optically Detected Magnetic Resonance of Single NV Defects for Enhanced Dc Magnetic Field Sensitivity." *Physical Review B* 84 (19). https://doi.org/10.1103/physrevb.84.195204.

Duan, Huaiyu, George M. Fuller, and Yong-Zhong Qian. 2010. "Collective Neutrino Oscillations." *Annual Review of Nuclear and Particle Science* 60 (1): 569–94. https://doi.org/10.1146/annurev.nucl.012809.104524.

Efthymiou, Stavros, Sergi Ramos-Calderer, Carlos Bravo-Prieto, et al. 2021. "Qibo : A Framework for Quantum Simulation with Hardware Acceleration." *Quantum Science and Technology* 7 (1): 015018. https://doi.org/10.1088/2058-9565/ac39f5.

Elzerman, J. M., R. Hanson, L. H. Willems van Beveren, B. Witkamp, L. M. K. Vandersypen, and L. P. Kouwenhoven. 2004. “Single-Shot Read-Out of an Individual Electron Spin in a Quantum Dot.” *Nature* 430 (6998): 431–35. https://doi.org/10.1038/nature02693.

Englert, Berthold-Georg. 1996. “Fringe Visibility and Which-Way Information: An Inequality.” *Physical Review Letters* 77 (11): 2154–57. https://doi.org/10.1103/physrevlett.77.2154.

Esteban, Ivan, M. C. Gonzalez-Garcia, Michele Maltoni, Thomas Schwetz, and Albert Zhou. 2020. “The Fate of Hints: Updated Global Analysis of Three-Flavor Neutrino Oscillations.” *Journal of High Energy Physics* 2020 (9). https://doi.org/10.1007/jhep09(2020)178.

Fano, U. 1947. “Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions.” *Physical Review* 72 (1): 26–29. https://doi.org/10.1103/physrev.72.26.

Farhi, Edward, Jeffrey Goldstone, and Sam Gutmann. 2014. *A Quantum Approximate Optimization Algorithm*. arXiv. https://doi.org/10.48550/ARXIV.1411.4028.

Feynman, R. P. 1939. “Forces in Molecules.” *Physical Review* 56 (4): 340–43. https://doi.org/10.1103/PhysRev.56.340.

Feynman, R. P. 1949. “The Theory of Positrons.” *Physical Review* 76 (6): 749–59. https://doi.org/10.1103/PhysRev.76.749.

Fowler, Austin G., Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. 2012. “Surface Codes: Towards Practical Large-Scale Quantum Computation.” *Physical Review A* 86 (3). https://doi.org/10.1103/physreva.86.032324.

Fréchet, Maurice. 1907. “Sur Les Ensembles de Fonctions Et Les Opérations Linéaires.” *Comptes Rendus de l’Académie Des Sciences* 144: 1414–16.

Freeman, Murray F., and John W. Tukey. 1950. “Transformations Related to the Angular and the Square Root.” *The Annals of Mathematical Statistics* 21 (4): 607–11. https://doi.org/10.1214/aoms/1177729756.

Gao, Ping, Daniel Louis Jafferis, and Aron C. Wall. 2017. “Traversable Wormholes via a Double Trace Deformation.” *Journal of High Energy Physics* 2017 (12). https://doi.org/10.1007/jhep12(2017)151.

Gell-Mann, Murray. 1962. “Symmetries of Baryons and Mesons.” *Physical Review* 125 (3): 1067–84. https://doi.org/10.1103/physrev.125.1067.

Gerritsma, R., G. Kirchmair, F. Zähringer, E. Solano, R. Blatt, and C. F. Roos. 2010. “Quantum Simulation of the Dirac Equation.” *Nature* 463: 68–71. https://doi.org/10.1038/nature08688.

Gidney, Craig. 2021. “Stim: A Fast Stabilizer Circuit Simulator.” *Quantum* 5: 497. https://doi.org/10.22331/q-2021-07-06-497.

Giovannetti, Vittorio, Seth Lloyd, and Lorenzo Maccone. 2004. “Quantum-Enhanced Measurements: Beating the Standard Quantum Limit.” *Science* 306 (5700): 1330–36. https://doi.org/10.1126/science.1104149.

Giovannetti, Vittorio, Seth Lloyd, and Lorenzo Maccone. 2006. “Quantum Metrology.” *Physical Review Letters* 96 (1). https://doi.org/10.1103/physrevlett.96.010401.

Goos, F., and H. Hänchen. 1947. “Ein Neuer Und Fundamentaler Versuch Zur Totalreflexion.” *Annalen Der Physik* 436 (7-8): 333–46. https://doi.org/10.1002/andp.19474360704.

Gorini, Vittorio, Andrzej Kossakowski, and E. C. G. Sudarshan. 1976. “Completely Positive Dynamical Semigroups of N-Level Systems.” *Journal of Mathematical Physics* 17 (5): 821–25. https://doi.org/10.1063/1.522979.

Greenberger, Daniel M., and Allaine Yasin. 1988. “Simultaneous Wave and Particle Knowledge in a Neutron Interferometer.” *Physics Letters A* 128 (8): 391–94. https://doi.org/10.1016/0375-9601(88)90114-4.

Grover, Lov K. 1996. “A Fast Quantum Mechanical Algorithm for Database Search.” *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing - STOC ’96*, 212–19. https://doi.org/10.1145/237814.237866.

Hadfield, Stuart, Zhihui Wang, Bryan O’Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. 2019. “From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz.” *Algorithms* 12 (2): 34. https://doi.org/10.3390/a12020034.

Hahn, E. L. 1950. “Spin Echoes.” *Physical Review* 80 (4): 580–94. https://doi.org/10.1103/physrev.80.580.

Hanson, R., L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. K. Vandersypen. 2007. “Spins in Few-Electron Quantum Dots.” *Reviews of Modern Physics* 79 (4): 1217–65. https://doi.org/10.1103/revmodphys.79.1217.

Havlíček, Vojtěch, Antonio D. Córcoles, Kristan Temme, et al. 2019. “Supervised Learning with Quantum-Enhanced Feature Spaces.” *Nature* 567 (7747): 209–12. https://doi.org/10.1038/s41586-019-0980-2.

Hawking, S. W. 1976. “Breakdown of Predictability in Gravitational Collapse.” *Physical Review D* 14 (10): 2460–73. https://doi.org/10.1103/physrevd.14.2460.

Hayden, Patrick, and John Preskill. 2007. “Black Holes as Mirrors: Quantum Information in Random Subsystems.” *Journal of High Energy Physics* 2007 (09): 120–20. https://doi.org/10.1088/1126-6708/2007/09/120.

Hein, M., J. Eisert, and H. J. Briegel. 2004. “Multiparty Entanglement in Graph States.” *Physical Review A* 69 (6). https://doi.org/10.1103/physreva.69.062311.

Helstrom, Carl W. 1976. *Quantum Detection and Estimation Theory*. Academic Press.

Herrero-Collantes, Miguel, and Juan Carlos Garcia-Escartin. 2017. “Quantum Random Number Generators.” *Reviews of Modern Physics* 89 (1). https://doi.org/10.1103/revmodphys.89.015004.

Higgott, Oscar. 2022. “PyMatching: A Python Package for Decoding Quantum Codes with Minimum-Weight Perfect Matching.” *ACM Transactions on Quantum Computing* 3 (3): 1–16. https://doi.org/10.1145/3505637.

Householder, Alston S. 1958. “Unitary Triangularization of a Nonsymmetric Matrix.” *Journal of the ACM* 5 (4): 339–42. https://doi.org/10.1145/320941.320947.

Huang, Hsin-Yuan, Richard Kueng, and John Preskill. 2020. “Predicting Many Properties of a Quantum System from Very Few Measurements.” *Nature Physics* 16 (10): 1050–57. https://doi.org/10.1038/s41567-020-0932-7.

Hubbard, J. 1963. “Electron Correlations in Narrow Energy Bands.” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 276 (1365): 238–57. https://doi.org/10.1098/rspa.1963.0204.

Huelga, S. F., C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac. 1997. “Improvement of Frequency Standards with Quantum Entanglement.” *Physical Review Letters* 79 (20): 3865–68. https://doi.org/10.1103/physrevlett.79.3865.

IQM Quantum Computers. 2026. *IQM Resonance: Cloud Access to IQM Quantum Computers*. Https://resonance.meetiqm.com.

Ivanov, D. A. 2001. “Non-Abelian Statistics of Half-Quantum Vortices in p -Wave Superconductors.” *Physical Review Letters* 86 (2): 268–71. https://doi.org/10.1103/physrevlett.86.268.

Jafferis, Daniel, Alexander Zlokapa, Joseph D. Lykken, et al. 2022. “Traversable Wormhole Dynamics on a Quantum Processor.” *Nature* 612 (7938): 51–55. https://doi.org/10.1038/s41586-022-05424-3.

Jarzynski, C. 1997. “Nonequilibrium Equality for Free Energy Differences.” *Physical Review Letters* 78 (14): 2690–93. https://doi.org/10.1103/physrevlett.78.2690.

Javadi-Abhari, Ali, Matthew Treinish, Kevin Krsulich, et al. 2024. *Quantum Computing with Qiskit*. arXiv. https://doi.org/10.48550/ARXIV.2405.08810.

Jeffreys, Harold. 1946. “An Invariant Form for the Prior Probability in Estimation Problems.” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 186 (1007): 453–61. https://doi.org/10.1098/rspa.1946.0056.

Johansson, J. R., P. D. Nation, and Franco Nori. 2012. “QuTiP: An Open-Source Python Framework for the Dynamics of Open Quantum Systems.” *Computer Physics Communications* 183 (8): 1760–72. https://doi.org/10.1016/j.cpc.2012.02.021.

Jordan, P., and E. Wigner. 1928. “Über Das Paulische Äquivalenzverbot.” *Zeitschrift für Physik* 47 (9-10): 631–51. https://doi.org/10.1007/bf01331938.

Kanazawa, Naoki, Daniel J. Egger, Yael Ben-Haim, et al. 2023. “Qiskit Experiments: A Python Package to Characterize and Calibrate Quantum Computers.” *Journal of Open Source Software* 8 (84): 5329. https://doi.org/10.21105/joss.05329.

Kandala, Abhinav, Antonio Mezzacapo, Kristan Temme, et al. 2017. “Hardware-Efficient Variational Quantum Eigensolver for Small Molecules and Quantum

Magnets." *Nature* 549 (7671): 242–46. https://doi.org/10.1038/nature23879.

Khaneja, Navin, and Steffen J. Glaser. 2001. "Cartan Decomposition of SU(2n) and Control of Spin Systems." *Chemical Physics* 267 (1-3): 11–23. https://doi.org/10.1016/s0301-0104(01)00318-4.

Kim, Jin-Sung, Alex McCaskey, Bettina Heim, Manish Modani, Sam Stanwyck, and Timothy Costa. 2023. "CUDA Quantum: The Platform for Integrated Quantum-Classical Computing." *2023 60th ACM/IEEE Design Automation Conference (DAC)*, 1–4. https://doi.org/10.1109/dac56929.2023.10247886.

Kitaev, A Yu. 2001. "Unpaired Majorana Fermions in Quantum Wires." *Physics-Uspekhi* 44 (10S): 131–36. https://doi.org/10.1070/1063-7869/44/10s/s29.

Kitaev, A. Yu. 1995. *Quantum Measurements and the Abelian Stabilizer Problem*. arXiv. https://doi.org/10.48550/ARXIV.QUANT-PH/9511026.

Kitaev, A.Yu. 2003. "Fault-Tolerant Quantum Computation by Anyons." *Annals of Physics* 303 (1): 2–30. https://doi.org/10.1016/s0003-4916(02)00018-0.

Knill, E., R. Laflamme, and G. J. Milburn. 2001. "A Scheme for Efficient Quantum Computation with Linear Optics." *Nature* 409 (6816): 46–52. https://doi.org/10.1038/35051009.

Kobrin, Bryce, Thomas Schuster, and Norman Y. Yao. 2023. *Comment on "Traversable Wormhole Dynamics on a Quantum Processor"*. arXiv. https://doi.org/10.48550/ARXIV.2302.07897.

Konig, Robert, Renato Renner, and Christian Schaffner. 2009. "The Operational Meaning of Min- and Max-Entropy." *IEEE Transactions on Information Theory* 55 (9): 4337–47. https://doi.org/10.1109/tit.2009.2025545.

Krantz, P., M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. 2019. "A Quantum Engineer's Guide to Superconducting Qubits." *Applied Physics Reviews* 6 (2). https://doi.org/10.1063/1.5089550.

Landauer, R. 1957. "Spatial Variation of Currents and Fields Due to Localized Scatterers in Metallic Conduction." *IBM Journal of Research and Development* 1 (3): 223–31. https://doi.org/10.1147/rd.13.0223.

Landman, Jonas, Natansh Mathur, Yun Yvonna Li, et al. 2022. "Quantum Methods for Neural Networks and Application to Medical Image Classification." *Quantum* 6: 881. https://doi.org/10.22331/q-2022-12-22-881.

LaRose, Ryan, Andrea Mari, Sarah Kaiser, et al. 2022. "Mitiq: A Software Package for Error Mitigation on Noisy Quantum Computers." *Quantum* 6: 774. https://doi.org/10.22331/q-2022-08-11-774.

Leung, Debbie W., M. A. Nielsen, Isaac L. Chuang, and Yoshihisa Yamamoto. 1997. "Approximate Quantum Error Correction Can Lead to Better Codes." *Physical Review A* 56 (4): 2567–73. https://doi.org/10.1103/physreva.56.2567.

Levenberg, Kenneth. 1944. "A Method for the Solution of Certain Non-Linear Problems in Least Squares." *Quarterly of Applied Mathematics* 2 (2): 164–68. https://doi.org/10.1090/qam/10666.

Lindblad, G. 1976. "On the Generators of Quantum Dynamical Semigroups." *Communications in Mathematical Physics* 48 (2): 119–30. https://doi.org/10.1007/bf01608499.

Lloyd, Seth. 1996. "Universal Quantum Simulators." *Science* 273 (5278): 1073–78. https://doi.org/10.1126/science.273.5278.1073.

Lloyd, Seth, Lorenzo Maccone, Raul Garcia-Patron, et al. 2011. "Closed Timelike Curves via Postselection: Theory and Experimental Test of Consistency." *Physical Review Letters* 106 (4). https://doi.org/10.1103/physrevlett.106.040403.

Lloyd, Seth, Masoud Mohseni, and Patrick Rebentrost. 2014. "Quantum Principal Component Analysis." *Nature Physics* 10 (9): 631–33. https://doi.org/10.1038/nphys3029.

Low, Guang Hao, and Isaac L. Chuang. 2017. "Optimal Hamiltonian Simulation by Quantum Signal Processing." *Physical Review Letters* 118 (1). https://doi.org/10.1103/physrevlett.118.010501.

Ma, Xiongfeng, Xiao Yuan, Zhu Cao, Bing Qi, and Zhen Zhang. 2016. "Quantum Random Number Generation." *Npj Quantum Information* 2 (1). https://doi.org/10.1038/npjqi.2016.21.

Maassen, Hans, and J. B. M. Uffink. 1988. "Generalized Entropic Uncertainty Relations." *Physical Review Letters* 60 (12): 1103–6. https://doi.org/10.1103/physrevlett.60.1103.

Magesan, Easwar, J. M. Gambetta, and Joseph Emerson. 2011. "Scalable and Robust Randomized Benchmarking of Quantum Processes." *Physical Review Letters* 106 (18). https://doi.org/10.1103/physrevlett.106.180504.

Majorana, Ettore. 1937. “Teoria Simmetrica Dell’elettrone e Del Positrone.” *Il Nuovo Cimento* 14 (4): 171–84. https://doi.org/10.1007/bf02961314.

Maki, Ziro, Masami Nakagawa, and Shoichi Sakata. 1962. “Remarks on the Unified Model of Elementary Particles.” *Progress of Theoretical Physics* 28 (5): 870–80. https://doi.org/10.1143/ptp.28.870.

Maldacena, J., and L. Susskind. 2013. “Cool Horizons for Entangled Black Holes.” *Fortschritte Der Physik* 61 (9): 781–811. https://doi.org/10.1002/prop.201300020.

Maldacena, Juan, Stephen H. Shenker, and Douglas Stanford. 2016. “A Bound on Chaos.” *Journal of High Energy Physics* 2016 (8). https://doi.org/10.1007/jhep08(2016)106.

Maldacena, Juan, and Douglas Stanford. 2016. “Remarks on the Sachdev-Ye-Kitaev Model.” *Physical Review D* 94 (10). https://doi.org/10.1103/physrevd.94.106002.

Maldacena, Juan, Douglas Stanford, and Zhenbin Yang. 2017. “Diving into Traversable Wormholes.” *Fortschritte Der Physik* 65 (5). https://doi.org/10.1002/prop.201700034.

Marquardt, Donald W. 1963. “An Algorithm for Least-Squares Estimation of Nonlinear Parameters.” *Journal of the Society for Industrial and Applied Mathematics* 11 (2): 431–41. https://doi.org/10.1137/0111030.

Martyn, John M., Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. 2021. “Grand Unification of Quantum Algorithms.” *PRX Quantum* 2 (4). https://doi.org/10.1103/prxquantum.2.040203.

Matsuzaki, Yuichiro, Simon C. Benjamin, and Joseph Fitzsimons. 2011. “Magnetic Field Sensing Beyond the Standard Quantum Limit Under the Effect of Decoherence.” *Physical Review A* 84 (1). https://doi.org/10.1103/physreva.84.012103.

McArdle, Sam, Xiao Yuan, and Simon Benjamin. 2019. “Error-Mitigated Digital Quantum Simulation.” *Physical Review Letters* 122 (18). https://doi.org/10.1103/physrevlett.122.180501.

McClean, Jarrod R., Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. “Barren Plateaus in Quantum Neural Network Training Landscapes.” *Nature Communications* 9 (1). https://doi.org/10.1038/

s41467-018-07090-4.

McLachlan, A. D. 1964. "A Variational Solution of the Time-Dependent Schrodinger Equation." *Molecular Physics* 8 (1): 39–44. https://doi.org/10.1080/00268976400100041.

Meyer, David A., and Nolan R. Wallach. 2002. "Global Entanglement in Multiparticle Systems." *Journal of Mathematical Physics* 43 (9): 4273–78. https://doi.org/10.1063/1.1497700.

Mezzadri, Francesco. 2006. *How to Generate Random Matrices from the Classical Compact Groups*. arXiv. https://doi.org/10.48550/ARXIV.MATH-PH/0609050.

Mikheyev, Stanislav P., and Alexei Yu. Smirnov. 1985. "Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos." *Soviet Journal of Nuclear Physics* 42: 913–17.

Monteverde, Vicente Humberto. 2026a. "Fundamentación Geométrica y Dimensional Del Sistema Qang." Unpublished manuscript.

Monteverde, Vicente Humberto. 2026b. *Qang: The Qang (Qg) Unit, the Symmetry Filter and Their Applications. Source Code and Research Notes*. Https://github.com/Viny2030/qang.

Monteverde, Vicente Humberto. 2026c. *The Qang (Qg): A Unified Angular-Probability Unit and Metric for Parametric Quantum Circuit Design*. https://doi.org/10.5281/ZENODO.22832150.

National Institute of Standards and Technology. 2024a. *Module-Lattice-Based Digital Signature Standard*. Federal Information Processing Standards Publication FIPS 204. National Institute of Standards; Technology (U.S.). https://doi.org/10.6028/nist.fips.204.

National Institute of Standards and Technology. 2024b. *Module-Lattice-Based Key-Encapsulation Mechanism Standard*. Federal Information Processing Standards Publication FIPS 203. National Institute of Standards; Technology (U.S.). https://doi.org/10.6028/nist.fips.203.

National Institute of Standards and Technology. 2024c. *Stateless Hash-Based Digital Signature Standard*. Federal Information Processing Standards Publication FIPS 205. National Institute of Standards; Technology (U.S.). https://doi.org/10.6028/nist.fips.205.

Navas, S., C. Amsler, T. Gutsche, et al. 2024. “Review of Particle Physics.” *Physical Review D* 110 (3). https://doi.org/10.1103/physrevd.110.030001.

Nielsen, Michael A., and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information*. 10th anniversary. Cambridge University Press. https://doi.org/10.1017/cbo9780511976667.

Nosek, Brian A., Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. 2018. “The Preregistration Revolution.” *Proceedings of the National Academy of Sciences* 115 (11): 2600–2606. https://doi.org/10.1073/pnas.1708274114.

O’Donnell, Kane, and Matt Visser. 2011. “Elementary Analysis of the Special Relativistic Combination of Velocities, Wigner Rotation and Thomas Precession.” *European Journal of Physics* 32 (4): 1033–47. https://doi.org/10.1088/0143-0807/32/4/016.

O’Malley, P. J. J., R. Babbush, I. D. Kivlichan, et al. 2016. “Scalable Quantum Simulation of Molecular Energies.” *Physical Review X* 6 (3). https://doi.org/10.1103/physrevx.6.031007.

Oganesyan, Vadim, and David A. Huse. 2007. “Localization of Interacting Fermions at High Temperature.” *Physical Review B* 75 (15). https://doi.org/10.1103/physrevb.75.155111.

Page, Don N. 1993. “Information in Black Hole Radiation.” *Physical Review Letters* 71 (23): 3743–46. https://doi.org/10.1103/physrevlett.71.3743.

Pancharatnam, S. 1956. “Generalized Theory of Interference, and Its Applications: Part i. Coherent Pencils.” *Proceedings of the Indian Academy of Sciences - Section A* 44 (5): 247–62. https://doi.org/10.1007/bf03046050.

Paszke, Adam, Sam Gross, Francisco Massa, et al. 2019. *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. arXiv. https://doi.org/10.48550/ARXIV.1912.01703.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, et al. 2011. “Scikit-Learn: Machine Learning in Python.” *Journal of Machine Learning Research* 12: 2825–30.

Peng, Tianyi, Aram W. Harrow, Maris Ozols, and Xiaodi Wu. 2020. “Simulating Large Quantum Circuits on a Small Quantum Computer.” *Physical Review Letters* 125 (15). https://doi.org/10.1103/physrevlett.125.150504.

Peruzzo, Alberto, Jarrod McClean, Peter Shadbolt, et al. 2014. "A Variational Eigenvalue Solver on a Photonic Quantum Processor." *Nature Communications* 5 (1). https://doi.org/10.1038/ncomms5213.

Pfeuty, Pierre. 1970. "The One-Dimensional Ising Model with a Transverse Field." *Annals of Physics* 57 (1): 79–90. https://doi.org/10.1016/0003-4916(70)90270-8.

Pontecorvo, Bruno. 1968. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge." *Soviet Physics JETP* 26: 984–88.

Popper, Karl R. 1959. *The Logic of Scientific Discovery*. Hutchinson.

Provost, J. P., and G. Vallee. 1980. "Riemannian Structure on Manifolds of Quantum States." *Communications in Mathematical Physics* 76 (3): 289–301. https://doi.org/10.1007/bf02193559.

Pusz, W., and S. L. Woronowicz. 1978. "Passive States and KMS States for General Quantum Systems." *Communications in Mathematical Physics* 58 (3): 273–90. https://doi.org/10.1007/bf01614224.

Puzzuoli, Daniel, Christopher J. Wood, Daniel J. Egger, Benjamin Rosand, and Kento Ueda. 2023. "Qiskit Dynamics: A Python Package for Simulating the Time Dynamics of Quantum Systems." *Journal of Open Source Software* 8 (90): 5853. https://doi.org/10.21105/joss.05853.

Ramsey, Norman F. 1950. "A Molecular Beam Resonance Method with Separated Oscillating Fields." *Physical Review* 78 (6): 695–99. https://doi.org/10.1103/physrev.78.695.

Rao, C. Radhakrishna. 1945. "Information and the Accuracy Attainable in the Estimation of Statistical Parameters." *Bulletin of the Calcutta Mathematical Society* 37: 81–91.

Riesz, Frigyes. 1907. "Sur Une Espèce de géométrie Analytique Des Systèmes de Fonctions Sommables." *Comptes Rendus de l'Académie Des Sciences* 144: 1409–11.

Romero, Jonathan, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. 2018. "Strategies for Quantum Computing Molecular Energies Using the Unitary Coupled Cluster Ansatz." *Quantum Science and Technology* 4 (1): 014008. https://doi.org/10.1088/2058-9565/aad3e4.

Sachdev, Subir, and Jinwu Ye. 1993. “Gapless Spin-Fluid Ground State in a Random Quantum Heisenberg Magnet.” *Physical Review Letters* 70 (21): 3339–42. https://doi.org/10.1103/physrevlett.70.3339.

Samuel, Joseph, and Rajendra Bhandari. 1988. “General Setting for Berry’s Phase.” *Physical Review Letters* 60 (23): 2339–42. https://doi.org/10.1103/physrevlett.60.2339.

Schechter, J., and J. W. F. Valle. 1980. “Neutrino Masses in SU(2) ⊗ u(1) Theories.” *Physical Review D* 22 (9): 2227–35. https://doi.org/10.1103/physrevd.22.2227.

Schuld, Maria, and Nathan Killoran. 2019. “Quantum Machine Learning in Feature Hilbert Spaces.” *Physical Review Letters* 122 (4). https://doi.org/10.1103/physrevlett.122.040504.

Shannon, C. E. 1948. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27 (3): 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

Shor, Peter W. 1995. “Scheme for Reducing Decoherence in Quantum Computer Memory.” *Physical Review A* 52 (4): R2493–96. https://doi.org/10.1103/physreva.52.r2493.

Shor, Peter W. 1997. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer.” *SIAM Journal on Computing* 26 (5): 1484–509. https://doi.org/10.1137/s0097539795293172.

Shor, Peter W., and John Preskill. 2000. “Simple Proof of Security of the BB84 Quantum Key Distribution Protocol.” *Physical Review Letters* 85 (2): 441–44. https://doi.org/10.1103/physrevlett.85.441.

Simon, Daniel R. 1997. “On the Power of Quantum Computation.” *SIAM Journal on Computing* 26 (5): 1474–83. https://doi.org/10.1137/s0097539796298637.

Singh, Vinit, Amandeep Singh Bhatia, Mandeep Kaur Saggi, Manas Sajjan, and Sabre Kais. 2026. “Quantum Machine Learning for Complex Systems: Paradigms, Applications, and Challenges.” *Academia Quantum* 3 (2). https://doi.org/10.20935/acadquant8243.

Sivarajah, Seyon, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. 2020. “T|ket⟩: A Retargetable Compiler for NISQ Devices.” *Quantum Science and Technology* 6 (1): 014003. https://doi.org/10.

1088/2058-9565/ab8e92.

Stokes, George Gabriel. 1852. “On the Composition and Resolution of Streams of Polarized Light from Different Sources.” *Transactions of the Cambridge Philosophical Society* 9: 399–416.

Stokes, James, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. 2020. “Quantum Natural Gradient.” *Quantum* 4: 269. https://doi.org/10.22331/q-2020-05-25-269.

Sun, Qiming, Timothy C. Berkelbach, Nick S. Blunt, et al. 2018. “PySCF: The Python-Based Simulations of Chemistry Framework.” *WIREs Computational Molecular Science* 8 (1). https://doi.org/10.1002/wcms.1340.

Susskind, Leonard, Lárus Thorlacius, and John Uglum. 1993. “The Stretched Horizon and Black Hole Complementarity.” *Physical Review D* 48 (8): 3743–61. https://doi.org/10.1103/physrevd.48.3743.

Suzuki, Masuo. 1991. “General Theory of Fractal Path Integrals with Applications to Many-Body Theories and Statistical Physics.” *Journal of Mathematical Physics* 32 (2): 400–407. https://doi.org/10.1063/1.529425.

Suzuki, Yohichi, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. 2020. “Amplitude Estimation Without Phase Estimation.” *Quantum Information Processing* 19 (2). https://doi.org/10.1007/s11128-019-2565-2.

Svetlichny, George. 2011. “Time Travel: Deutsch Vs. Teleportation.” *International Journal of Theoretical Physics* 50 (12): 3903–14. https://doi.org/10.1007/s10773-011-0973-x.

Swingle, Brian, Gregory Bentsen, Monika Schleier-Smith, and Patrick Hayden. 2016. “Measuring the Scrambling of Quantum Information.” *Physical Review A* 94 (4). https://doi.org/10.1103/physreva.94.040302.

Talkner, Peter, Eric Lutz, and Peter Hänggi. 2007. “Fluctuation Theorems: Work Is Not an Observable.” *Physical Review E* 75 (5). https://doi.org/10.1103/physreve.75.050102.

Temme, Kristan, Sergey Bravyi, and Jay M. Gambetta. 2017. “Error Mitigation for Short-Depth Quantum Circuits.” *Physical Review Letters* 119 (18). https://doi.org/10.1103/physrevlett.119.180509.

Thaller, Bernd. 1992. *The Dirac Equation*. Texts and Monographs in Physics. Springer. https://doi.org/10.1007/978-3-662-02753-0.

Thomas, L. H. 1926. “The Motion of the Spinning Electron.” *Nature* 117 (2945): 514–14. https://doi.org/10.1038/117514a0.

Tomamichel, Marco, Charles Ci Wen Lim, Nicolas Gisin, and Renato Renner. 2012. “Tight Finite-Key Analysis for Quantum Cryptography.” *Nature Communications* 3 (1). https://doi.org/10.1038/ncomms1631.

Trotter, H. F. 1959. “On the Product of Semi-Groups of Operators.” *Proceedings of the American Mathematical Society* 10 (4): 545–51. https://doi.org/10.1090/s0002-9939-1959-0108732-6.

Unruh, W. G. 1976. “Notes on Black-Hole Evaporation.” *Physical Review D* 14 (4): 870–92. https://doi.org/10.1103/physrevd.14.870.

Wang, Samson, Enrico Fontana, M. Cerezo, et al. 2021. “Noise-Induced Barren Plateaus in Variational Quantum Algorithms.” *Nature Communications* 12 (1). https://doi.org/10.1038/s41467-021-27045-6.

Wigner, E. 1939. “On Unitary Representations of the Inhomogeneous Lorentz Group.” *The Annals of Mathematics* 40 (1): 149. https://doi.org/10.2307/1968551.

Wilson, Edwin B. 1927. “Probable Inference, the Law of Succession, and Statistical Inference.” *Journal of the American Statistical Association* 22 (158): 209–12. https://doi.org/10.1080/01621459.1927.10502953.

Wolfenstein, L. 1978. “Neutrino Oscillations in Matter.” *Physical Review D* 17 (9): 2369–74. https://doi.org/10.1103/physrevd.17.2369.

Wootters, W. K. 1981. “Statistical Distance and Hilbert Space.” *Physical Review D* 23 (2): 357–62. https://doi.org/10.1103/physrevd.23.357.

Wu, Yue, Shimon Kolkowitz, Shruti Puri, and Jeff D. Thompson. 2022. “Erasure Conversion for Fault-Tolerant Quantum Computing in Alkaline Earth Rydberg Atom Arrays.” *Nature Communications* 13 (1). https://doi.org/10.1038/s41467-022-32094-6.

Yaghoobianzadeh, Amir Mohammad, and Jawad A. Salehi. 2026. “Continuous-Variable Teleportation and Entanglement Distribution in Quantum Passive Optical Networks.” *Academia Quantum* 3 (3). https://doi.org/10.20935/

acadquant8460.

Yoshida, Beni, and Alexei Kitaev. 2017. *Efficient Decoding for the Hayden-Preskill Protocol*. arXiv. https://doi.org/10.48550/ARXIV.1710.03363.

Yuan, Xiao, Suguru Endo, Qi Zhao, Ying Li, and Simon C. Benjamin. 2019. "Theory of Variational Quantum Simulation." *Quantum* 3: 191. https://doi.org/10.22331/q-2019-10-07-191.

Żukowski, M., A. Zeilinger, M. A. Horne, and A. K. Ekert. 1993. "'Event-Ready-Detectors' Bell Experiment via Entanglement Swapping." *Physical Review Letters* 71 (26): 4287–90. https://doi.org/10.1103/physrevlett.71.4287.

A Catalogue of the 15 fundamental gates in qg representation

This appendix collects the action of the gates of `qang.formulation` on the qg readings. A gate U transforms each Pauli reading by conjugation,

$$qg'_P = qg_{U^\dagger P U},$$

so it is enough to say what each local reading becomes. The tables are not written by hand: `tools/make_appendices.py` generates them from `formulation.gate_table`, `is_clifford`, `conserves_weight` and `gate_radius_class`. If the code changes, they are regenerated. In Pauli strings the leftmost letter is qubit 0, which is the control of CX and CZ, both controls of Toffoli and the control of Fredkin.

Why fifteen. The module defines 16 matrices: six fixed single-qubit gates, four parametric ones and six multi-qubit gates. $R_Z(\varphi)$ and the phase gate $P(\varphi)$ differ only by a global phase, $P(\varphi) = e^{i\varphi/2} R_Z(\varphi)$, so they act identically on every reading and count as one in qg.

A.1 Fixed single-qubit gates

Table A.1: Fixed single-qubit gates.

gate	qg'_X	qg'_Y	qg'_Z	from $\lvert 0\rangle$ (m-qg)	Clifford	conserves weight
X	qg_{X}	$-qg_{\mathrm{Y}}$	$-qg_{\mathrm{Z}}$	-1000.0	yes	no
Y	$-qg_{\mathrm{X}}$	qg_{Y}	$-qg_{\mathrm{Z}}$	-1000.0	yes	no
Z	$-qg_{\mathrm{X}}$	$-qg_{\mathrm{Y}}$	qg_{Z}	+1000.0	yes	yes
H	qg_{Z}	$-qg_{\mathrm{Y}}$	qg_{X}	+0.0	yes	no
S	$-qg_{\mathrm{Y}}$	qg_{X}	qg_{Z}	+1000.0	yes	yes

gate	qg'_X	qg'_Y	qg'_Z	from $\lvert 0\rangle$ (m-qg)	Clifford	conserves weight
T	$\frac{1}{\sqrt{2}}qg_{\mathrm{X}} - \frac{1}{\sqrt{2}}qg_{\mathrm{Y}}$	$\frac{1}{\sqrt{2}}qg_{\mathrm{X}} + \frac{1}{\sqrt{2}}qg_{\mathrm{Y}}$	qg_{Z}	+1000.0	no	yes

Clifford gates permute the readings with signs, and T mixes them. The diagonal gates (Z, S, T) leave qg_Z alone: they do not move populations, only the phase, which is read in qg_X and qg_Y. Conserving the weight is what allows the filter of Chapters 21 and 22.

A.2 Parametric gates

Table A.2: Parametric gates; the script checks each coefficient at $\theta = 0.7$.

gate	qg'_X	qg'_Y	qg'_Z	from $\lvert 0\rangle$
$R_X(\theta)$	qg_X	$\cos\theta\, qg_Y - \sin\theta\, qg_Z$	$\sin\theta\, qg_Y + \cos\theta\, qg_Z$	$m = 1000\cos\theta$
$R_Y(\theta)$	$\cos\theta\, qg_X + \sin\theta\, qg_Z$	qg_Y	$-\sin\theta\, qg_X + \cos\theta\, qg_Z$	$m = 1000\cos\theta$
$R_Z(\varphi)$, $P(\varphi)$	$\cos\varphi\, qg_X - \sin\varphi\, qg_Y$	$\sin\varphi\, qg_X + \cos\varphi\, qg_Y$	qg_Z	$m = +1000$

From $\lvert 0\rangle$, $R_X(\theta)$ and $R_Y(\theta)$ are the Chapter 1 probability equation: the gate angle is the qang angle and the reading is $m = 1000\cos\theta$. R_Z and P do not change the Z reading. No rotation about X or Y conserves the weight: it moves population between $\lvert 0\rangle$ and $\lvert 1\rangle$. That is why the layers used with the filter are made of Givens (RBS) rotations between pairs, not of single-qubit rotations.

On the angular qga scale of Chapter 1 (Section 1.4), with 1000 qga per turn, a gate angle is written without factors of π: $\theta = 2\pi\,\theta_{qga}/1000$ and, from $\lvert 0\rangle$, $m = 1000\cos(2\pi\,\theta_{qga}/1000)$. So $R_X(500\ \mathrm{qga})$ is X up to a global phase and reads -1000 m-qg, $R_Y(250\ \mathrm{qga})$ leaves the qubit on the equator ($m = 0$) and $R_Z(\varphi)$ turns the phase by φ_{qga} thousandths of a turn. The script checks the first two.

A.3 Multi-qubit gates

Table A.3: Multi-qubit gates: how the local Z readings transform (those of X and Y are in `gate_table`).

gate	qubits	what happens to each local qg_Z	Clifford	conserves weight	radius	maximal deficit
CX	2	$qg_{\mathrm{ZI}} \to qg_{\mathrm{ZI}}$; $qg_{\mathrm{IZ}} \to qg_{\mathrm{ZZ}}$	yes	no	entangles	1
CZ	2	$qg_{\mathrm{ZI}} \to qg_{\mathrm{ZI}}$; $qg_{\mathrm{IZ}} \to qg_{\mathrm{IZ}}$	yes	yes	entangles	1
SWAP	2	$qg_{\mathrm{ZI}} \to qg_{\mathrm{IZ}}$; $qg_{\mathrm{IZ}} \to qg_{\mathrm{ZI}}$	yes	yes	permutes	0
iSWAP	2	$qg_{\mathrm{ZI}} \to qg_{\mathrm{IZ}}$; $qg_{\mathrm{IZ}} \to qg_{\mathrm{ZI}}$	yes	yes	entangles	1
Toffoli	3	$qg_{\mathrm{ZII}} \to qg_{\mathrm{ZII}}$; $qg_{\mathrm{IZI}} \to qg_{\mathrm{IZI}}$; $qg_{\mathrm{IIZ}} \to \frac{1}{2}qg_{\mathrm{IIZ}} + \frac{1}{2}qg_{\mathrm{IZZ}} + \frac{1}{2}qg_{\mathrm{ZIZ}} - \frac{1}{2}qg_{\mathrm{ZZZ}}$	no	no	entangles	1

gate	qubits	what happens to each local qg_Z	Clifford	conserves weight	radius	maximal deficit
Fredkin	3	$qg_{\mathrm{ZII}} \to qg_{\mathrm{ZII}}$; $qg_{\mathrm{IZI}} \to \frac{1}{2}qg_{\mathrm{IIZ}} + \frac{1}{2}qg_{\mathrm{IZI}} - \frac{1}{2}qg_{\mathrm{ZIZ}} + \frac{1}{2}qg_{\mathrm{ZZI}}$; $qg_{\mathrm{IIZ}} \to \frac{1}{2}qg_{\mathrm{IIZ}} + \frac{1}{2}qg_{\mathrm{IZI}} + \frac{1}{2}qg_{\mathrm{ZIZ}} - \frac{1}{2}qg_{\mathrm{ZZI}}$	no	yes	entangles	1

- **Radius.** "Preserves": every local radius stays 1 for pure product inputs. "Permutes": the radii are exchanged (SWAP). "Entangles": some product input ends with deficit $1 - r_q^2 > 0$. The maximal deficit, 1, takes a qubit to the centre of the sphere (Chapter 5).
- **On basis states**, none of these gates creates a deficit: they act as classical permutations.
- **CZ, SWAP, iSWAP and Fredkin conserve the weight.** CX and Toffoli do not, because they flip the target bit. A circuit made only of the former has a fixed weight sector and admits the qg filter.
- Toffoli and Fredkin are not Clifford: a local reading becomes a mixture of four strings with coefficients $\pm\frac{1}{2}$.

B Transition formulas, Gell-Mann matrices and Chebyshev polynomials

This appendix gathers in one place the closed forms used in Chapters 2 and 3: the universal gate and the reflections written in $q = qg_Z$, the transition probability, the generalization of qg_Z to qudits with the diagonal Gell-Mann matrices, and the Grover trajectory as a Chebyshev polynomial (notes, §37 and §166). There are no new results. What it adds is the check: `tools/make_appendices.py` evaluates each formula of `qang.algebra` against direct linear algebra on 2000 random states and writes the maximal error here. If a formula were wrong, the script would stop.

B.1 Single-qubit formulas

With $q = qg_Z = \cos\theta$, $\cos\frac{\theta}{2} = \sqrt{(1+q)/2}$ and $\sin\frac{\theta}{2} = \sqrt{(1-q)/2}$. The pure state with polar bias q and azimuth φ is

$$|\psi(q,\varphi)\rangle = \sqrt{\tfrac{1+q}{2}}\,|0\rangle + e^{i\varphi}\sqrt{\tfrac{1-q}{2}}\,|1\rangle, \qquad \vec{r} = (\sqrt{1-q^2}\cos\varphi,\ \sqrt{1-q^2}\sin\varphi,\ q).$$

On the angular qga scale of Chapter 1 (Section 1.4), with $\theta = 2\pi\,\theta_{qga}/1000$, the bias is $q = \cos(2\pi\,\theta_{qga}/1000)$ and the amplitudes are $\cos(\pi\,\theta_{qga}/1000)$ and $\sin(\pi\,\theta_{qga}/1000)$: the formulas below do not change, only the unit of the angle.

The formulas are:

$$U_3^{qg}(q,\varphi,\lambda) = \begin{pmatrix} \sqrt{\frac{1+q}{2}} & -e^{i\lambda}\sqrt{\frac{1-q}{2}} \\ e^{i\varphi}\sqrt{\frac{1-q}{2}} & e^{i(\varphi+\lambda)}\sqrt{\frac{1+q}{2}} \end{pmatrix}, \qquad |U_{ij}|^2 = \begin{pmatrix} \frac{1+q}{2} & \frac{1-q}{2} \\ \frac{1-q}{2} & \frac{1+q}{2} \end{pmatrix} \tag{B.1}$$

$$I - 2|\psi\rangle\langle\psi| = -q\,Z - \sqrt{1-q^2}\,X, \qquad 2|\psi\rangle\langle\psi| - I = q\,Z + \sqrt{1-q^2}\,X \tag{B.2}$$

$$|\langle\psi_1|\psi_2\rangle|^2 = \frac{1+\vec{r}_1\cdot\vec{r}_2}{2} \xrightarrow{\varphi_1=\varphi_2} \frac{1+q_1q_2+\sqrt{(1-q_1^2)(1-q_2^2)}}{2} \tag{B.3}$$

Table B.1: Check of the formulas on 2000 random states.

formula	`qang.algebra` function	compared with	maximal error
universal gate	`u3_qang`	Qiskit's `UGate(arccos q, φ, λ)`	$1.7\cdot10^{-16}$
transition matrix	`transition_matrix`	$\lvert U_{ij}\rvert^2$	$5.6\cdot10^{-16}$
Householder	`householder_qg`	$I-2\lvert\psi\rangle\langle\psi\rvert$	$1.2\cdot10^{-15}$
Grover diffuser	`grover_diffuser_qg`	$2\lvert\psi\rangle\langle\psi\rvert-I$	$1.2\cdot10^{-15}$
transition probability	`transition_probability`	$\lvert\langle\psi_1\rvert\psi_2\rangle\rvert^2$	$5.6\cdot10^{-16}$
same, on one meridian	(closed form)	$\lvert\langle\psi_1\rvert\psi_2\rangle\rvert^2$	$6.7\cdot10^{-16}$

The transition matrix is doubly stochastic, rational in q and independent of the phases: the probability a gate leaves is read without trigonometry. With $(q,\varphi,\lambda)=(1,0,0)$, $(-1,0,\pi)$, $(1,0,\pi)$ and $(0,0,\pi)$, Equation B.1 gives exactly I, X, Z and H; the script checks it. The reflector and the diffuser are the same operator with opposite sign, and both are linear combinations of Z and X with coefficients q and $\sqrt{1-q^2}$ (Chapter 3).

B.2 Diagonal Gell-Mann matrices

A d-level qudit has d^2-1 generators of $\mathfrak{su}(d)$, the generalized Gell-Mann matrices. Only $d-1$ of them are diagonal: they form the Cartan subalgebra and are the only ones read by measuring in the computational basis. They are

$$\lambda_k = \sqrt{\tfrac{2}{k(k+1)}}\,\mathrm{diag}(\underbrace{1,\dots,1}_{k},-k,0,\dots,0), \qquad k=1,\dots,d-1, \qquad \mathrm{Tr}(\lambda_j\lambda_k)=2\delta_{jk}, \tag{B.4}$$

and their expectation values are linear in the populations p and determine them:

$$\langle \lambda_k \rangle = \sum_i p_i \, (\lambda_k)_{ii}, \qquad p = \frac{1}{d} + \frac{1}{2} \sum_{k=1}^{d-1} \langle \lambda_k \rangle \, \mathrm{diag}(\lambda_k). \tag{B.5}$$

Table B.2: The diagonal matrices up to $d = 4$, generated by `cartan_generators`.

d	k	λ_k
2	1	$\mathrm{diag}(1, -1)$
3	1	$\mathrm{diag}(1, -1, 0)$
3	2	$\frac{1}{\sqrt{3}} \mathrm{diag}(1, 1, -2)$
4	1	$\mathrm{diag}(1, -1, 0, 0)$
4	2	$\frac{1}{\sqrt{3}} \mathrm{diag}(1, 1, -2, 0)$
4	3	$\frac{1}{\sqrt{6}} \mathrm{diag}(1, 1, 1, -3)$

For $d = 2$ the only one is $\lambda_1 = Z$ and its value is qg_Z. For $d = 3$ they are Gell-Mann's λ_3 and λ_8. The script recovers the populations from the $d-1$ values for $d = 2, \dots, 6$, 200 random distributions each, with maximal error $2.2 \cdot 10^{-16}$. An example: the qutrit with $p = (0.5, 0.3, 0.2)$ reads $\langle \lambda_3 \rangle = 0.2000$ and $\langle \lambda_8 \rangle = 0.2309$.

This is a definition and an identity. Whether reading more levels gives qang any advantage was asked in §70 for leakage, and none was found (Chapter 2).

B.3 Chebyshev polynomials and the Grover trajectory

With amplitude $a = \sin^2 t$ and the good/bad flag read as a qubit, $qg_0 = 1 - 2a = \cos 2t$. Each Grover iteration adds $2t$ to the angle, so after k iterations

$$qg_k = T_{2k+1}(qg_0), \qquad T_m(\cos x) = \cos mx, \tag{B.6}$$

with T_m the Chebyshev polynomial of the first kind. Since $T_m' = m \, U_{m-1}$ and $1 - T_m^2 = (1 - q^2) \, U_{m-1}^2$, the one-shot Fisher information at depth m is

$$F_m(qg_0) = \frac{m^2}{1 - qg_0^2}, \tag{B.7}$$

the factor m^2 over the depth-1 information, the same along the whole fringe. It is the qg form of the known $4m^2$ for the angle (Suzuki et al. 2020).

Table B.3: The first polynomials of the trajectory. At $a = 0.1$, $qg_0 = 0.8$; success is -1000 m-qg.

k	$m = 2k + 1$	$T_m(q)$	F_m/F_1	qg_k at $a = 0.1$ (m-qg)
0	1	q	1	+800.0
1	3	$4q^3 - 3q$	9	-352.0
2	5	$16q^5 - 20q^3 + 5q$	25	-997.1
3	7	$64q^7 - 112q^5 + 56q^3 - 7q$	49	-206.4
4	9	$256q^9 - 576q^7 + 432q^5 - 120q^3 + 9q$	81	+881.5

The script applies the diffuser of Equation B.2 and the oracle Z to 200 random initial states, up to $k = 5$, and finds Equation B.6 with maximal error $1.7 \cdot 10^{-15}$. It checks Equation B.7 by finite differences for $m = 1, \dots, 9$, with maximal relative error $1.7 \cdot 10^{-10}$. The practical side of §37 is in Chapter 3: the noise has to enter the likelihood, and the advantage of amplitude estimation ends near depth $1/p$.

Bibliographic notes

The matrices are those of Gell-Mann (1962). The reflector is that of Householder (1958); Grover search, Grover (1996); amplitude estimation, Brassard et al. (2002), and its maximum-likelihood version, Suzuki et al. (2020). The transition probability is the distance of Wootters (1981). The numbered notes (§) and the script of this appendix are in the repository (Monteverde 2026b).

C The echo-calibration protocol and inversion in Hamming sectors (§105–§108)

Chapter 22 explains why the echo is needed: the qg filter discards the shots that left the sector, but it does not see the errors that move an excitation from one qubit to another without changing the weight. This appendix gives the protocol step by step, what it costs, an exact example generated by `tools/make_appendices.py` with the functions of `qang.sectors`, and the record of the four pre-registered runs of the notes.

C.1 The protocol

1. **Sector.** n qubits and weight w: the sector has $D = \binom{n}{w}$ basis states, $|e_1\rangle, \ldots, |e_D\rangle$ (`sector_states`).
2. **Echoes.** For each j: prepare $|e_j\rangle$ with X gates, apply the network W and then $W^\dagger$, and measure in Z. Without errors the circuit returns to $|e_j\rangle$. The notes use 4000 shots per echo.
3. **Matrix.** Filter each echo to the sector (`filter_distribution`). Its distribution over the D states is column j of M, with $M_{ij} = P(e_i \text{ measured} \mid e_j \text{ prepared})$ (`echo_transfer_matrix`).
4. **Root.** The circuit to be read has half the gates of the echo, so it is corrected with $A = M^{1/2}$: the principal, real, non-negative, column-normalized root (`matrix_power_stochastic`).
5. **Unmixing.** Filter the circuit's distribution, f, and solve $Ax = f$ with $x \geq 0$ by non-negative least squares; normalize x (`unmix_sector`).
6. **Reading.** Each qubit reads $m_Z(q) = 1000\,(1 - 2p_q)$ with the populations of x.
7. **Before using it.** Simulate without noise at the same shots to know the sampling error. The echo helps when the filtered error is well above that level, and not when it is already almost all shot noise (§108).

C.2 Cost

Table C.1: One echo per sector state; the matrix is measured once per network, not per input.

qubits	weight	sector states D	echo circuits	echo shots (4000 each)	used in
5	1	5	5	20,000	§105, §106, §108
5	2	10	10	40,000	§107
8	1	8	8	32,000	—
8	2	28	28	112,000	§115

The cost grows as $\binom{n}{w}$, not as 2^n: that is what makes the inversion practical. At weight 2 with 8 qubits it is 28 circuits; a full readout confusion matrix on 8 qubits would need 256.

C.3 An exact example

The script builds an in-sector channel K on 5 qubits and weight 1, with probability 0.90 of staying on the right qubit, and a T1 loss of 15% of the shots to $|00000\rangle$, which reads $+1000$ m-qg on every qubit. The echo applies K twice, so $M = K^2$ (mean diagonal 0.81), and the principal root recovers K exactly. That is the assumption behind $M^{1/2}$: that the way back has the same error as the way out. In §105 the full M did better on Forte-1, so it does not always hold.

Table C.2: No shot noise. The filter removes the loss; $M^{1/2}$ undoes K; the full M over-corrects and ends near the filter alone.

readout	mean qg_Z error (m-qg)
without qang	71.9
qg filter	37.5
filter and echo, $M^{1/2}$	$< 10^{-12}$
filter and echo, M	36.7

Why non-negative least squares and not the inverse. With 1000 shots per run, in 37 of 300 runs the direct inverse $A^{-1}f$ gives some negative population,

which is not a probability. The non-negative solution never does. The mean qg_Z error is almost the same: 21.6 m-qg with the inverse and 21.5 m-qg with the non-negative solution. What is gained is that the result is always a distribution, not accuracy; that is why the usual readout mitigation also restricts the solution to the simplex.

C.4 The four pre-registered runs

Table C.3: The runs of the notes, on the noise models of three IBM and two IonQ devices.

§	predictions (commit)	passed	result
§105	C1–C4 (`052752d`)	4 of 4	weight 1, 5 qubits, five simulated devices: the echo removes 38–47% of the error the filter leaves; filter plus echo, 6–7× more accurate than the raw readout
§106	L1–L4 (`fa8eda6`)	4 of 4	the weight-1 classifier, 189 inputs: the decision-value error falls 43–56%; decisions differing from the noiseless model, 15 → 7 → 3
§107	W1–W4 (`4d83029`)	4 of 4	weight 2, a 10 × 10 matrix: the echo removes 37–43% of the qg_Z error and 39–48% of the qg_{ZZ} error

§	predictions (commit)	passed	result
§108	M1–M4 (`4ae2674`)	0 of 4	narrow-margin inputs: the filter cuts flipped decisions from 47 to 17 and the echo adds nothing (19)

In total, 12 of 16 predictions passed. The script checks that each cited commit exists and that in it the notes did not yet have the section: the four ’before’s are consistent (Appendix E). The numbers in the table come from the notes; the `examples/` scripts each section names reproduce them, but they take hours and this appendix does not run them.

The failure of §108 is the rule of step 7. In that model the decision value scatters by about 0.22 from shots alone, and with the filter the error was already 0.26–0.35: little in-sector bias was left, and inverting the matrix only added variance. In §106 the filtered error was 0.47–0.87, mostly bias, and the echo halved it.

Bibliographic notes

The echo is that of Hahn (1950); readout-error mitigation by inverting a confusion matrix, Bravyi et al. (2021). The sector filter is the symmetry verification of Bonet-Monroig et al. (2018). The numbered notes (§) and the scripts are in the repository (Monteverde 2026b).

D Quick reference of the qang 0.14.0 API

This is the list of every module of the `qang` package at version 0.14.0, with its public names: those in `__all__` or, if the module does not define it, the public functions and classes it declares. `tools/make_appendices.py` generates it by importing each module, so the list is the code's own and not a copy. The signature and documentation of each name are in `help(qang.<module>.<name>)`.

- `import qang` needs only NumPy. The SDK modules (`qiskit_gate`, `cirq_gate`, `pennylane_gate`, `ibm`, `qiskit_ml`, `qec`, `circuits`, `algorithms`) and those of `qang.integrations` import their library when used.
- There is no `crypto` module: the book's cryptography is in the examples (Chapters 16–18). `qang.iqm` needs `iqm-client`, which pins its own Qiskit: it goes in a separate environment (`qang[iqm]`).
- The "names" column gives the number of public names, and the full list follows the table.

Table D.1: The modules of qang 0.14.0.

module	what it does	names
`algebra`	Closed-form qg algebra: the universal gate, reflections and transition probabilities; Cartan subalgebra (§166).	9
`algorithms`	Textbook algorithms in Qiskit: teleportation, superdense coding, Grover.	9
`angular`	The four quantum numbers of an electron in qg (§155).	17
`ansatze`	Reusable variational ansätze, including qg layers.	4

module	what it does	names
`channels`	Single-qubit channels as affine maps of the qg vector (§162).	7
`circuits`	Circuit constructors (Bell, GHZ, W, graph, Dicke) and their readings.	14
`cirq_gate`	qg as native Cirq gates.	5
`clock_comparison`	Differential phase between two clocks and the gravitational shift.	5
`complementarity`	Wave–particle duality and entropic uncertainty in qg.	8
`core`	The qang (qg) unit and the `Qang` class.	2
`fisher`	Fisher information, shift rules and the natural gradient in qg (§159).	14
`formulation`	The 15 gates and 14 algorithms in qg language; the local radius (Appendix A, Chapter 5).	48
`geometric`	Geometric phase (Berry, Pancharatnam) in qg units.	8
`gradients`	The coordinate singularity at the poles and its regularizations.	17
`ibm`	qg readings on IBM Runtime: weight-preservation pass, readout correction, chains, ZNE (§137).	30
`integrations`	qang inside other libraries: one submodule per library (Chapter 35).	12
`interop`	qg readings from any SDK’s results with one qubit order (§130).	10

module	what it does	names
`iqm`	qg readings on IQM: readout from the poles or the error profile, filter, chain choice (Chapter 31).	7
`kernels`	Quantum kernels with the qg filter (§161).	6
`knitting`	Sampling cost of circuit cutting.	4
`mixed`	The qang unit for mixed states and POVMs.	11
`multiqubit`	The qang unit for multi-qubit registers: marginals, entropies, correlations.	14
`neutrino`	Neutrino flavour oscillations in qg (§151; Chapters 9–11).	38
`pennylane_gate`	qg as native, differentiable PennyLane operations; filtered readout.	6
`phase`	The deferred qg_Φ phase unit.	2
`polarization`	Polarized light in qg units: Stokes, Jones, Mueller (§83; Chapter 6).	17
`postselection`	Teleportation with postselection: closed timelike curves and black-hole decoding (Chapters 14–15).	8
`probability`	The qang probability equation and the milli-qang (§167; Chapter 1).	13
`proportions`	The qang angle as the arcsine transformation of a proportion.	9

module	what it does	names
`pulse_area`	Pulse area of fractional gates: the minimal-area rule (§160).	6
`qec`	The 3-qubit bit-flip code with syndromes read in qg.	6
`qga`	The angular qga scale: 1000 qga per turn; conversions to radians, p_0 and m-qg, Fisher and Berry phase in qga (Chapter 1).	15
`qiskit_gate`	qg as a native Qiskit gate.	5
`qiskit_ml`	The weight QNN inside qiskit-machine-learning (§140).	4
`qml`	Weight-conserving quantum neural networks, with and without the filter (§75–§80; Chapter 20).	4
`qpon`	Teleportation through a passive optical network, discrete variables (§157; Chapters 12–13).	17
`qsp`	Quantum signal processing in qg (§134).	8
`sectors`	Conserved-sector tools for the qg filter and the echo (§20, §60, §66, §68; Appendix C).	9
`statistics`	Error propagation between probability and angle, estimators and intervals.	25
`transformations`	Closed-form transformations between qang-parameterized states.	1

module	what it does	names
`varqte`	Variational real-time evolution (McLachlan) with a weight-preserving ansatz (§162).	6

D.1 Public names by module

- **algebra**: u3_qang, transition_matrix, householder_qg, grover_diffuser_qg, qg_vector, transition_probability, cartan_generators, cartan_values, populations_from_cartan
- **algorithms**: teleportation_circuit, teleportation_output_fidelity, superdense_coding_circuit, decode_superdense_coding, grover_oracle, grover_diffusion, grover_optimal_iterations, grover_circuit, grover_success_probabil
- **angular**: m_from_weight, weight_from_m, m_distribution, qg_z_spin, qg_s_spin, coherent_m_distribution, qang_angle_of_eigenstate, vector_model_angle, transverse_spread, dicke_state, p_orbital, spin_operators, shell_quantum_numbers, shell_register_qubits, shell_l_spectrum, j_squared_from_swaps, l_from_j_squared
- **ansatze**: single_excitation_ansatz, hardware_efficient_ansatz, qg_ry_layer, hardware_efficient_ansatz_qg
- **channels**: apply_kraus, bloch_affine, qg_after, amplitude_damping, phase_damping, depolarizing, thermal_relaxation
- **circuits**: qg_z_profile_of_circuit, joint_qg_s_of_circuit, bell_circuit, ghz_circuit, w_circuit, graph_state_circuit, linear_cluster_state_circuit, dicke_state_circuit, ghz_qg_z_profile, ghz_joint_qg_s, w_qg_z_profile, graph_state_qg_z_profile, graph_state_joint_qg_s, dicke_qg_z_profile
- **cirq_gate**: rqang_gate, full_rqang_gate, signed_rqang_gate, rqang_symbol, append_qang
- **clock_comparison**: differential_phase_from_c, ellipse_estimate, moment_estimate, redshift_phase, simulate_pairs
- **complementarity**: predictability, visibility, duality_sum, von_neumann_entropy_bits, entropic_sum, entropic_bound, duality_from_counts, entropic_sum_from_counts
- **core**: qg_s_from_qg_z, Qang
- **fisher**: fisher_theta, fisher_qg, cramer_rao_qg, shift_gradient, indistinguishable_pair, ansatz_states, ansatz_state, tfim_hamiltonian, energy, gradient, full_metric, block_metric, y_basis_circuit, block_metric_from_counts
- **formulation**: kron, pauli, rx, ry, rz, phase, clean, conjugate, conserves_weight, z_parity, single_qg_z, hadamard_n, phase_oracle_output,

simon_output, grover_state, grover_qg_z_prediction, qft_matrix, local_qg, kickback_control, counting_qg_x, order_finding_output, purity_from_qg, walk_register_mean, qg_values, state_from_qg, apply_gate, gate_table, is_clifford, born_p0, deutsch_jozsa_is_constant, bernstein_vazirani_secret, simon_nonzero_parities, grover_marked_from_signs, count_from_qg_x, hhl_readout, energy_from_qg, maxcut_from_qg, product_kernel, radius2, radius_profile, radius_deficit, sphere_area, meyer_wallach, gate_radius_class, weight_sector_radius, hadamard_test_radius, algorithm_radii, direction_from_qg

- **geometric**: state_from_bloch, bloch_from_state, pancharatnam_phase, solid_angle, berry_phase_from_solid_angle, cone_phase_turns, curvature_flux, geometric_phase_qg
- **gradients**: jacobian, inverse_jacobian_raw, inverse_jacobian_clipped, inverse_jacobian_tikhonov, pole_damping_factor, qfi_qg, natural_gradient_step_qg, toy_vqe_energy, toy_vqe_grad_theta, GDHistory, run_gradient_descent, benchmark, MultiGDHistory, multi_param_gradient_descent, signed_theta, signed_qg_step, signed_natural_qg_step
- **ibm**: WeightPreservationCheck, measured_physical_qubits, readout_errors, readout_inverse_distribution, ibm_qg, estimator_qg, predict_kept_fraction, linear_chains, best_chain, best_layout, plan_shots, fold_two_qubit_gates, zne_qg, qnn_template, qnn_parameter_values, run_qnn, unary_loader, unary_orthogonal, pca_rotation, unary_spectrum, xy_trotter, xy_exact_occupations, occupations, xxz_trotter, xxz_exact_occupations, unary_overlap_circuit, swap_test_circuit, overlap_fidelity, swap_test_fidelity, add_static_zz
- **interop**: canonical_counts, qg_z_profile, weight_filter, filtered_qg_z_profile, joint_qg_s, joint_qg_s_se, sampler_qg, qg_z_observables, filtered_observables, estimator_filtered_qg
- **iqm**: check_weight, readout_errors, pole_circuits, readout_from_poles, iqm_qg, predict_kept_fraction, best_chain
- **kernels**: unary_embedding, exact_kernel, kernel_from_counts, normalize_gram, alignment, kernel_qang_angle
- **knitting**: pauli_rotation_cut_gamma, controlled_rotation_cut_gamma, cut_sampling_overhead, total_sampling_overhead
- **mixed**: density_from_statevector, mix, is_valid_density_matrix, qg_z_density, purity, von_neumann_entropy, standard_z_povm, is_valid_povm, povm_outcome_probabilities, qg_s_povm, trine_povm
- **multiqubit**: product_state, bell_state, ghz_state, partial_trace, per_qubit_qg_z, marginal_qg_s, marginal_von_neumann_entropy, joint_qg_s, shannon_entropy_plugin_bits, shannon_entropy_miller_madow_bits, joint_qg_s_from_counts, qg_correlation, mean_qg_z, mean_qg_z_from_counts
- **neutrino**: phase, two_flavour_probability, two_flavour_qg_z, two_flavour_qang_angle, mixing_angle_from_maximum, mixing_angle_std, two_flavour_circuit, pmns, three_flavour_probabilities, three_flavour_circuit, cp_asymmetry,

matter_potential, matter_probabilities, matter_trotter_circuit, first_maximum_energy, occupation_profile, collective_mass_hamiltonian, collective_exact, mass_sector_probabilities, collective_trotter_circuit, solar_core_potential, solar_matter_angle, solar_pee_from_p2, solar_pee_adiabatic, solar_slab_edges, solar_slab_unitaries, solar_p2, solar_p2_exact, solar_circuit, pmns_majorana, majorana_invisible, masses, m_bb, majorana_circuit, sterile_angles, sterile_mixing, sterile_probabilities, sterile_circuit

- **pennylane_gate**: rqang, full_rqang, append_qang, qg_z_from_probs, kept_fraction, filtered_qg_z
- **phase**: QangPhi, anchor_value
- **polarization**: stokes_to_qg, qg_to_stokes, degree_of_polarization, purity_from_stokes, jones_to_stokes, jones_to_qg, mueller_from_jones, ptm_from_mueller, mueller_from_ptm, rotator, linear_polarizer, wave_plate, half_wave_plate, quarter_wave_plate, depolarizer, malus_intensity, qg_from_counts
- **postselection**: pctc_success, pctc_controlled_qg_z, pctc_circuit, pctc_from_counts, is_scrambler, find_scrambler, yk_decoder_circuit, yk_from_counts
- **probability**: readings_from_p0, milliqang_from_p0, p0_from_milliqang, theta_from_milliqang, milliqang_from_theta, angle_per_milliqang, pole_angle, shot_noise_milliqang, shots_for_resolution, pulse_error_milliqang, statistical_angle, Reading, analyze
- **proportions**: theta_from_proportion, proportion_from_theta, arcsine, anscombe_theta, theta_variance, proportion_interval, cohens_h, compare_proportions, z_value
- **pulse_area**: minimal_area_angle, rx_qg_z_random_amplitude, rx_qg_z_systematic_amplitu
  qg_sensitivity, minimal_area_circuit, area
- **qec**: encode_bit_flip, extract_syndrome, syndrome_from_qg_z, decode_syndrome, apply_correction, decode_bit_flip
- **qga**: QGA_PER_TURN, MILLI_QGA_PER_QGA, FISHER_PER_QGA2, radians_from_qga, qga_from_radians, degrees_from_qga, qga_from_degrees, p0_from_qga, qga_from_p0, milliqang_from_qga, qga_from_milliqang, azimuth_qga, angle_sigma_qga, fubini_study_metric_qga, cone_phase_qga
- **qiskit_gate**: RQangGate, FullRQangGate, SignedRQangGate, append_qang, append_rqang_parameter
- **qiskit_ml**: loader_inputs, qang_sampler_qnn, qg_features, QangQNNClassifier
- **qml**: WeightQNN, MultiClassQNN, kept_fraction, compare_qang
- **qpon**: fibre_transmission, link, heralding_probability, singlet_fraction, average_fidelity, optimal_angle, optimal_fidelity, optimal_qg_z, resource_angle_from_qg_z, teleport_circuit, fidelity_from_counts, fidelity_from_correlators, resource_circuit, correlator_from_counts,

dual_rail_teleport_circuit, dual_rail_fidelity_from_counts, dual_rail_fidelities
- **qsp**: signal_from_qg, qg_of_signal, wqsp_unitary, wqsp_poly, hadamard_test_p0, sample_hadamard_test, fit_phases, wqsp_ops
- **sectors**: hamming_weights, filter_distribution, reweight_sectors, sector_exposure, SectorExposurePass, sector_states, echo_transfer_matrix, matrix_power_stochastic, unmix_sector
- **statistics**: p0_from_theta, var_qg_z_shot_noise, std_qg_z_shot_noise, propagated_theta_variance, propagated_theta_std, confidence_interval_theta, BootstrapResult, empirical_theta_std, beta_cdf, beta_ppf, QgEstimate, bayes_qg_estimate, wilson_qg_estimate, delta_qg_estimate, qg_estimate, propagated_theta_variance_mixed, theta_qcrb_variance, qg_covariance, delta_method_variance, register_witness_variance, qg_s_from_qg_z_interval, qg_s_estimate, qg2_unbiased, radius2_estimate, direction_estimate
- **transformations**: transition_probability_qang
- **varqte**: n_params, ansatz, mclachlan_rhs, evolve, xxz_matrix, occupations_from_counts

D.2 qang.integrations

Each submodule imports only its library (Chapter 35); the pip extra is in parentheses.

- **integrations.mitiq** (qang[mitiq]): qg_z, filtered_qg_z, qg_executor, execute_with_filter_zne, filter_rule
- **integrations.qiskit_experiments** (qang[experiments]): HeraldedReadout, HeraldedReadoutAnalysis, heralded_closed_form, t_eff_from_qg
- **integrations.qiskit_dynamics** (qang[dynamics]): rabi_qg, area_error_qg
- **integrations.qutip** (qang[qutip]): qg_readings, filtered_qg_z, mesolve_qg, t1_collapse
- **integrations.stim** (qang[stim]): detector_qg, observable_qg, measurement_qg, decoded_logical_qg
- **integrations.sklearn** (qang[sklearn]): WeightQNNClassifier, MultiClassQNNClassifier
- **integrations.jax** (qang[jax]): qg_from_p0, theta_from_qg, milliqang, qg_z_from_probs, filtered_qg_z, pole_damped_update
- **integrations.torch** (qang[torch]): WeightQNNTorch, unary_state
- **integrations.qibo** (qang[qibo]): rqang, canonical_counts, qg_readout
- **integrations.pytket** (qang[pytket]): rqang, qg_readout, statevector_qg
- **integrations.braket** (qang[braket]): rqang, run_qg
- **integrations.cudaq** (qang[cudaq]): rqang, canonical_counts, qg_readout

E Falsification register: the predictions the repository allows one to check

The book's plan announced a full count of §1–§164 here. Chapter 28 explains why it was not done: the sections before §71 hardly use pre-registration, and many later ones leave their verdicts only in prose, worded differently. This appendix lists what can be checked with code. `tools/make_appendices.py` generates it with the same functions as the Chapter 28 figure (`tools/make_cap28_figures.py`), so its totals are the chapter's.

E.1 Falsification statement

- **Definitions and identities are not falsified by experiments.** The qg unit, the milli-qang, the probability equation and the formulas of Appendices A and B are proven or checked to machine precision in `tests/`. If one fails, it is a code bug, not a result.
- **Empirical claims are.** Each says that a reading with qang is more accurate, cheaper or more stable than without qang, and it is fixed with a threshold before running. A failed prediction is published with its number and not rewritten.
- **What would falsify the programme as a whole.** That the advantages of Chapter 29 did not repeat on real hardware with thresholds fixed in advance. Today there are three runs on IQM Garnet, with 12 of 13 predictions passed, but only one tests an advantage: §165, the weight filter on neutrinos. §168 and §170 test the scale and the reading, not an advantage. That is too little to decide.

E.2 Verdicts stored by the code

Each file of `examples/data/` with a `verdict` dictionary is a run whose predictions the script itself scored. The § column is the notes section that names the script.

Table E.1: 34 runs, 148 predictions: 94 passed (64%) and 54 failed.

run	§	passed	failed	which failed
fisher_natural_gradient_qg	§158	4	1	G2
ibm_coherent_zz_qg	§145	0	3	X1, X2, X3
ibm_fleet_qg	§144	0	3	F1, F2, F3
ibm_layout_qg	§138	2	1	L1
ibm_pca_spectrum_qg	§146	4	0	—
ibm_qnn_layout_qg	§141	2	1	R3
ibm_qnn_qg	§139	2	1	Q2
ibm_runtime_qg	§137	3	0	—
ibm_shot_plan_qg	§143	1	1	P2
ibm_static_zz_qg	§150	2	2	Z2, Z3
ibm_swap_test_qg	§149	5	0	—
ibm_trotter_walk_qg	§147	5	1	W1
ibm_xxz_many_body_qg	§148	3	1	K1
ibm_zne_qg	§142	1	2	Z1, Z3
iqm_milliqang_qg_analysis	§168	3	1	M2
iqm_zitterbewegung_qg_analysis	§170	5	0	—
lih_lambda_filter_qg	§153	5	1	L4
neutrino_collective_qg	§154	0	7	C1, C2, C3, C4, C5, C6, C7
neutrino_matter_qg	§152	4	1	M1
neutrino_ordering_majorana_sterile_qg	§163	5	1	O3
neutrino_oscillation_qg	§151	4	0	—
neutrino_solar_qg	§156	6	0	—
postselection_cdc_blackhole_qg	§164	2	4	C1, C2, C3, H1
pulse_area_qg	§160	3	1	P4
qiskit_ml_qnn_qg	§140	3	0	—
qnn_ingate_dephasing	§133	2	2	D2, D3
qnn_torino_residual	§132	2	2	T2, T3
qpon_teleport_quasi_rail_qg	§158	1	7	A1, A2, A3, B1, B2, B3, B5
qpon_teleport_qg	§157	5	1	Q5
qsp_arcsine_fit	§134	2	2	Q2, Q4
quantum_kernel_qg	§161	2	3	K2, K3, K4
varqte_kraus_qg	§162	4	1	V1
wqsp_kan	§135	1	2	K1, K2
wqsp_kan_parity	§136	1	1	P2

run	§	passed	failed	which failed
total		**94**	**54**	

E.3 Pre-registrations by sha256

Each `examples/data/*.prereg` holds the sha256 of a frozen script. The ‘before’ is checked if the hash matches the script at the commit that added the `.prereg` and that commit is earlier than the one of the results.

Table E.2: 16 pre-registrations: 8 checked in git, 8 with the file’s own timestamp only.

pre-registration	§	matches today	matches at its commit	before the results	‘before’
`fisher_natural_gradient_qg`	§159	no	yes	yes	checked
`iqm_milliqan_qg`	§168	no	yes	yes	checked
`iqm_neutrino_qg`	§165	no	yes	yes	checked
`iqm_zitterbewegung_qg`	§170	yes	yes	yes	checked
`lih_lambda_flavor_qg`	§153	no	yes	yes	checked
`neutrino_collective_qg`	§154	no	yes	yes	checked
`neutrino_mass_qg`	§152	no	no	no	own timestamp only
`neutrino_ordering_majorana_sterile_qg`	§163	no	no	no	own timestamp only
`neutrino_oscillation_qg`	§151	no	no	no	own timestamp only
`neutrino_solar_qg`	§156	no	no	no	own timestamp only
`postselection_bhc_blackhole_qg`	§164	no	no	no	own timestamp only
`pulse_area_qg`	§160	no	no	no	own timestamp only
`qpon_teleportation_rail_qg`	§158	no	yes	yes	checked

pre-registration	§	matches today	matches at its commit	before the results	‘before’
qpon_teleport_qg	§157	no	yes	yes	checked
quantum_kernel_qg	§161	no	no	no	own timestamp only
varqte_kraus_qg	§162	no	no	no	own timestamp only

That none matches today is expected: after the run the findings are added to the script’s opening comment. In those with their own timestamp only, the `.prereg` and the results entered in the same commit.

E.4 Pre-registrations by commit

Each “committed … (abcdef1)” in the notes: the commit exists and in it the notes do not yet have the section.

Table E.3: 54 citations in 54 sections, §71–§165: 54 consistent.

§	commits	consistent
§71	daf0b6f	1 of 1
§72	092580f	1 of 1
§73	e877a72	1 of 1
§74	68d8d00	1 of 1
§75	1a516b4	1 of 1
§76	8d32f92	1 of 1
§77	1450374	1 of 1
§78	e999d7f	1 of 1
§79	00520dc	1 of 1
§80	77700bd	1 of 1
§81	1f2da5c	1 of 1
§82	1f4f8ac	1 of 1
§85	d0191cb	1 of 1
§86	5239ad3	1 of 1
§87	1f33029	1 of 1
§88	0b266fe	1 of 1
§90	14fab61	1 of 1

§	commits	consistent
§91	`d292c7e`	1 of 1
§92	`7a3bdad`	1 of 1
§93	`52d5cb1`	1 of 1
§94	`ebd795b`	1 of 1
§95	`505943d`	1 of 1
§96	`5648765`	1 of 1
§98	`4c11ff3`	1 of 1
§100	`517902b`	1 of 1
§101	`d4fcbaa`	1 of 1
§102	`6c0a1d7`	1 of 1
§103	`1b408dd`	1 of 1
§104	`22554eb`	1 of 1
§105	`052752d`	1 of 1
§106	`fa8eda6`	1 of 1
§107	`4d83029`	1 of 1
§108	`4ae2674`	1 of 1
§109	`9c06b36`	1 of 1
§110	`59f855a`	1 of 1
§111	`34760c3`	1 of 1
§112	`ea3fd3a`	1 of 1
§113	`d2dcd56`	1 of 1
§114	`d6e4f50`	1 of 1
§115	`12d12f0`	1 of 1
§117	`3c0fca9`	1 of 1
§118	`637e800`	1 of 1
§119	`d5d31e7`	1 of 1
§120	`fcfdf3b`	1 of 1
§121	`efba91a`	1 of 1
§122	`1bca045`	1 of 1
§123	`fc1e0ea`	1 of 1
§124	`b37b344`	1 of 1
§125	`b42cde2`	1 of 1
§126	`b2e92c4`	1 of 1
§127	`b2e92c4`	1 of 1
§128	`8a16974`	1 of 1
§129	`1a5507f`	1 of 1
§165	`a74c07c`	1 of 1

E.5 What this register leaves out

- **The QML line (§75–§123)** keeps its own count in the README: 176 predictions judged, 123 passed and 53 failed, with 4 more waiting for the IonQ hardware run (§116). It partly overlaps the tables above and must not be added to them.
- **The verdicts in prose** ("Four passed; V1 failed") of the sections that store no dictionary. They are in the notes, section by section, but counting them requires reading them one by one.
- **The sections before §71**, which hardly use pre-registration: their numbers are reproducible, but they have no threshold fixed in advance.

Bibliographic notes

The falsification criterion is that of Popper (1959), and pre-registration as a practice, Nosek et al. (2018). The numbered notes (§), the files of `examples/data/` and the scripts are in the repository (Monteverde 2026b).

F Functional-analysis foundation: Riesz–Fréchet

This appendix is conceptual and adds no results. It is Appendix A of the research notes, which Chapter 4 cites: it explains why qg_Z is the canonical linear readout of a state and why qg_S cannot be one. The three numerical checks are done by `tools/make_appendices.py`.

F.1 The setting

For n qubits, the operators on $\mathbb{C}^{2^n}$ form a finite-dimensional Hilbert space with the Hilbert–Schmidt inner product

$$\langle A, B\rangle = \mathrm{Tr}(A^\dagger B),$$

and density matrices live in it.

F.2 The theorem

Riesz–Fréchet representation (Riesz 1907; Fréchet 1907). Every continuous linear functional f on a Hilbert space H has the form $f(x) = \langle y_f, x\rangle$ for a unique $y_f \in H$. In finite dimension continuity is automatic and the theorem reduces to linear algebra: with an orthonormal basis $\{E_i\}$, $y_f = \sum_i \overline{f(E_i)}\, E_i$. That is why it is a foundation and not a result.

F.3 Application: qg_Z has a representative

The readout $\rho \mapsto qg_Z(\rho) = \mathrm{Tr}(\rho\,\sigma_z)$ is linear in ρ, so it has a unique Hilbert–Schmidt representative: σ_z on one qubit, and $\sigma_z^{(i)} \otimes I$ on qubit i of a register. Every linear readout of a state, be it an expectation value, a population or a per-shot average such as linear XEB, has this form. qg_Z is the one whose representative is the Z observable.

The script checks it on two qubits with the normalized Pauli basis, $E_P = P/2$. It rebuilds the representative of $\rho \mapsto \mathrm{Tr}(\rho\, Z \otimes I)$ and gets $Z \otimes I$ with error 0. For a population readout with random weights, $\rho \mapsto \sum_x w_x \langle x|\rho|x\rangle$, it gets $\mathrm{diag}(w)$ with error $8.7 \cdot 10^{-18}$.

F.4 What has no representative: qg_S

$qg_S(\rho) = H(\langle 0|\rho|0\rangle)$ is strictly concave in ρ, not linear, so no operator represents it. The shortest example: $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ both have $qg_S = 0$, and their equal mixture has $qg_S = 1$ bit; a linear readout would give the average, 0. With qg_Z the average does hold: $(+1 - 1)/2 = 0$.

The same structural fact is behind two results of the notes:

- **§10.3 B:** linear readouts have unbiased shot-average estimators, and entropy does not (Jensen's inequality). The table gives the exact bias, by binomial sum, of the two direct estimators.
- **§7.1:** the identity $qg_S = H((1 + qg_Z)/2)$ writes the non-linear metric as a fixed scalar function of the linear one. qg_S carries no information about the state beyond what the representative of qg_Z already extracts.

Table F.1: The qg_Z average is unbiased to machine precision; the entropy of the frequency always underestimates.

p_0	shots N	bias of the qg_Z average	bias of $H(k/N)$ (bits)
0.5	10	0	-0.07650
0.5	100	0	-0.00725
0.5	1000	0	-0.00072
0.9	10	0	-0.08440
0.9	100	0	-0.00735
0.9	1000	0	-0.00072

F.5 On Riesz’s lemma (1918)

The almost-orthogonal-vector lemma is used to show that the closed unit ball of an infinite-dimensional normed space is not compact. In the finite dimensions of this book it is trivial, because an exactly orthogonal vector always exists. It plays no role in qang and should not be cited as if it did.

Bibliographic notes

The representation theorem is Riesz (1907) and Fréchet (1907). The binary entropy is that of Shannon (1948). The numbered notes (§) and the script are in the repository (Monteverde 2026b).

G Relativistic energy and antimatter as qang readings

Chapter 1 left the extension of the qang scale to $E = mc^2$ and to antimatter as an **unverified proposal** (Section 1.4.1). This appendix develops it with the probability equations of Chapter 1, and separates carefully what is exact from what remains a proposed reading. The central result is exact, although known: for a free Dirac particle, **the expectation value of the matrix β is mc^2/E.** Since β has eigenvalues ± 1, like σ_z, that expectation value is a qg_Z, and all of Chapter 1 applies to it: it has a probability, a qang angle, an m-qg reading and an error bar. Every figure comes from `tools/verify_apendice_g.py` (25 claims, all pass), and nothing below is new physics.

G.1 The relativistic angle

Dividing $E^2 = (mc^2)^2 + (pc)^2$ by E^2 gives a Pythagorean identity, which defines an angle θ_r:

$$\cos\theta_r = \frac{mc^2}{E} = \frac{1}{\gamma}, \qquad \sin\theta_r = \frac{pc}{E} = \beta, \qquad \tan\theta_r = \beta\gamma = \frac{p}{mc}, \qquad E = \frac{mc^2}{\cos\theta_r}.$$

At rest $\theta_r = 0$ and the formula gives $E = mc^2$. When $\beta \to 1$, $\theta_r \to 90^\circ = 250$ qga and $E = pc$: the limit of light, with no rest mass. On the angular scale of Chapter 1 (Section 1.4), $E = mc^2/\cos(2\pi\,\theta_{r,qga}/1000)$. Table G.1 gives some values.

Table G.1: The relativistic angle and the reading $1000/\gamma$.

β	γ	θ_r	θ_r (qga)	$1000/\gamma$ (m-qg)	$T/mc^2 = \gamma - 1$
0	1	0°	0	+1000	0
0.5	1.155	30°	83.3	+866.0	0.155

β	γ	θ_r	θ_r (qga)	$1000/\gamma$ (m-qg)	$T/mc^2 = \gamma - 1$
0.6	1.25	36.87°	102.4	$+800.0$	0.25
0.8	1.667	53.13°	147.6	$+600.0$	0.667
0.99	7.089	81.89°	227.5	$+141.1$	6.089
0.999	22.37	87.44°	242.9	$+44.7$	21.37

The angle θ_r is not the rapidity η ($\beta = \tanh\eta$): they are related by $\cos\theta_r = 1/\cosh\eta$, the Gudermann function. Rapidities add for collinear boosts; angles do not. Two boosts of $0.6c$ give $\beta = 0.8824$ and $\gamma = 2.125$, that is, $\theta_r = 61.93^\circ = 172.0$ qga, not the sum $2 \times 102.4 = 204.8$ qga. The qga scale measures the angle, but does not make it additive.

G.2 $E = mc^2$ as a reading: the expectation value of β

In the free Dirac equation, $H = c\,\vec{\alpha}\cdot\vec{p} + \beta\,mc^2$. By the Hellmann–Feynman theorem (Feynman 1939), differentiating an eigenvalue with respect to a parameter gives the expectation value of the derivative of the Hamiltonian: $\partial E/\partial(mc^2) = \langle\beta\rangle$. Since $E = \pm\sqrt{(pc)^2 + (mc^2)^2}$,

$$\boxed{\langle\beta\rangle = \frac{mc^2}{E} = \cos\theta_r}$$

in every free eigenstate, with the sign of E. The script checks it on 1200 random eigenstates of both branches, to $2 \cdot 10^{-15}$. The matrix β has eigenvalues ± 1, so $\langle\beta\rangle$ is a qg_Z. In the usual encoding of the bispinor in two qubits it is literally the Z of the first one, $\beta = \sigma_z \otimes I$, and Qiskit confirms it: the qg_Z of that qubit is mc^2/E. Call that qubit the β **qubit**. The probability equation of Chapter 1 gives it its four readings:

$$p_0 = \frac{E + mc^2}{2E} = \cos^2\frac{\theta_r}{2}, \qquad p_1 = \frac{E - mc^2}{2E} = \frac{T}{2E}, \qquad m = 1000\,\frac{mc^2}{E} = \frac{1000}{\gamma}\ \text{m-qg}, \qquad \theta = \theta$$

Read backwards, it is a way of writing the energy: $E = mc^2/qg_Z = 1000\,mc^2/m$, and the kinetic energy is $T = 2E\,p_1$. $E = mc^2$ **is the pole**: at rest, $m = +1000$

m-qg, $p_1 = 0$, no kinetic energy. Light is on the equator. A particle at $0.6c$ reads $+800$ m-qg, with $p_0 = 0.9$; at $0.99c$, $+141.1$ m-qg.

Two caveats. First, the reading does not depend on the representation: $\langle\beta\rangle = \bar{\psi}\psi/\psi^\dagger\psi$ is a Lorentz scalar divided by a density, and in the chiral representation, where β is not diagonal, it gives the same (checked). What does depend on the representation is the picture of p_0 and p_1 as the weights of the "large" and "small" components of the spinor, which only holds in the standard Dirac representation. Second, the β qubit is a reading of a **simulation** of the Dirac equation, like those with trapped ions (Gerritsma et al. 2010): in a particle physics laboratory nobody measures β shot by shot.

G.3 Antimatter: the negative branch and charge conjugation

On the negative-energy branch the same identity gives $\langle\beta\rangle = -mc^2/|E|$: the reading is $m = -1000/\gamma$, the mirror of the positive branch's reading through the equator. Dirac (1928) found those solutions, and Feynman's reinterpretation (Feynman 1949) reads them as positive-energy antiparticles travelling forward in time. In qang the branch of the solution shows in the **sign of the reading** of the β qubit: northern hemisphere, the positive branch; southern hemisphere, the negative one.

The symmetry that goes from one branch to the other is **charge conjugation** C, $\psi \mapsto i\gamma^2\psi^*$. The script checks that C takes an energy $+E$ eigenstate of $H(\vec{p})$ to an energy $-E$ eigenstate of $H(-\vec{p})$, that it **inverts the reading**, $m \mapsto -m$, and that parity P and time reversal T preserve it. Hence CPT inverts it too. In the angle of the β qubit, C is the reflection $\theta_{qga} \mapsto 500 - \theta_{qga}$; for the reading it is the same as half a turn, $\theta_{qga} \mapsto \theta_{qga} + 500$.

This recovers, corrected, what the source document claimed: that CPT "induces an angular translation of 500 qg". There is a correct core: charge conjugation, and with it CPT, acts on the β reading as half a turn in qga. But it is not a translation of the state: it is a reflection of its reading. And it is true, as the document says, that the probability density does not change: $i\gamma^2$ is unitary, so $\psi^{c\dagger}\psi^c = \psi^\dagger\psi$ point by point (checked).

G.4 The probabilities of the source document

The source document uses another pair of probabilities, $P_{\text{rest}} = (mc^2/E)^2$ and $P_{\text{kinetic}} = \beta^2$. They also add to 1, by the Pythagorean identity, and they are also a qubit, but with **twice the angle**: $R_y(2\theta_r)|0\rangle$ gives $p_0 = 1/\gamma^2$ and $p_1 = \beta^2$ (QuTiP). Call it the **Pythagoras qubit**. That factor 2 is the one Chapter 1 flagged as an error. Neither formula is wrong; they are two different qubits for the same particle:

Table G.2: Two qang qubits for the same particle.

	angle	p_0	p_1	at $0.6c$
β qubit (expectation value of β)	θ_r	$(E+mc^2)/2E$	$T/2E$	$p_0 = 0.9$; 102.4 qga
Pythagoras qubit (split of E^2)	$2\theta_r$	$(mc^2/E)^2$	β^2	$p_0 = 0.64$; 204.8 qga

Only the first is the expectation value of an observable of the Dirac equation. The second splits E^2 into two squares: a well-defined probability, but with no observable behind it.

G.5 Measuring the relativistic angle with shots

If the angle θ_r is encoded in a qubit, with $R_y(\theta_r)$ and $qg_Z = 1/\gamma$, Chapter 1 says what reading it costs. Since the qang angle is the arcsine transformation, $\sigma(\theta_r) = 1/\sqrt{N}$ for any velocity, that is, $159.15/\sqrt{N}$ qga. The velocity, $\beta = \sin\theta_r$, then has $\sigma(\beta) = 1/(\gamma\sqrt{N})$: it is resolved better the faster the particle. With 1000 shots the angle has 5.03 qga of uncertainty over the whole range, and at $0.99c$ the velocity is resolved to 0.0045. The script checks it by sampling in Qiskit at $0.3c$, $0.6c$ and $0.9c$.

G.6 The reading of β on a real chip (§170)

In one dimension and at fixed momentum the Dirac Hamiltonian is a single qubit, $H = p\,X + m\,Z$ with $\beta = Z$, so the β qubit can be prepared and read on one

physical qubit. With $m = 1$ and $p = 0.75$, $E = 1.25$, $v = 0.6c$, $\gamma = 1.25$ and $\theta_r = 102.4$ qga. The exact readings: the positive-energy eigenstate, $R_y(\theta_r)|0\rangle$, reads $+800$ m-qg at every time; the negative-energy one reads -800. The state $|0\rangle$ mixes both branches instead, and its reading oscillates at frequency $2E$ between $+1000$ and $+280$ m-qg around $1000/\gamma^2 = +640$: the *Zitterbewegung*, Dirac's trembling motion. That mean, $1000\,(mc^2/E)^2$, is the source document's P_{rest} (Section G.4): the document's probability is the time average of the β reading.

It ran on IQM Garnet, qubit 0: 24 circuits of 2000 shots in one 32 s job, with the readout corrected from the poles ($e_0 = 0.0030$, $e_1 = 0.0265$). The five predictions were frozen before the run (the sha256 of the file was committed to the repository before the job; Chapter 28). They carry a 60 m-qg systematic allowance taken from §168, where R_y read low at intermediate angles.

Table G.3: The reading of β on IQM Garnet (notes §170).

prediction	measured (m-qg, corrected)	predicted	verdict
Z1: E_+ at $t = 0$	$+759.9 \pm 14.9$	$+800$	pass ($z = -2.7$; also without the allowance)
Z2: E_- at $t = 0$; C reflection	-809.4 ± 14.9; -49.5	-800; 0	pass (within 3σ without the allowance)
Z3: E_+ does not tremble (span over a period)	45.3	≤ 108	pass
Z4: fitted frequency	2.511	$2E = 2.5$ ($\pm 5\%$)	pass (0.44% off)
Z5: mean; amplitude	625.6; 356.1	640; 360 (± 60)	pass (14 and 4 m-qg off)

All five pass. On a real chip the Dirac eigenstate reads $1000/\gamma$, the negative branch mirrors it, and the *Zitterbewegung* shows at frequency $2E$ with the predicted mean $1000/\gamma^2$ and amplitude (Figure G.1). Z1 and Z5 did discriminate. The Pythagoras qubit would have read $+640$ at Z1, and a mean of $1000/\gamma$ would have given $+800$ at Z5. Both lie outside the tolerance.

Not clean, and said so. The four E_+ readings sit 26–71 m-qg below $+800$. The worst, at $t = T/2$, has $z = -4.5$ and is not part of any prediction. This is the same direction as in §168. The 16 *Zitterbewegung* points scatter 30 m-qg rms around the exact curve, more than shot noise (χ^2 per degree of freedom 2.4). So

gate errors are visible, but they move neither the frequency nor the mean. As a speedometer, E_+ at $t = 0$ gives $\gamma = 1.316$ (true value 1.25: +5%) and E_- gives 1.236 (−1%). This is one run on one qubit. A replication on another qubit and another day is pending, as for §168.



Figure G.1: The reading of β on IQM Garnet (§170), with $v = 0.6c$ and 2000 shots per point. The curve is the exact *Zitterbewegung* from $|0\rangle$ and the circles are measured. The squares are the E_+ eigenstate, which does not tremble and reads near $+800$; the diamonds are the E_- eigenstate, its mirror at -800.

G.7 What is claimed and what is not

- **Exact and checked:** the relativistic angle and its identities; $\langle\beta\rangle = mc^2/E$ with the sign of the branch (Hellmann–Feynman), independent of the representation; its four qang readings ($p_1 = T/2E$, $m = 1000/\gamma$); C and CPT invert the reading, and P and T preserve it; the Pythagoras qubit; the angle's error bar.
- **Checked on hardware (one run, §170):** on IQM Garnet, the reading $1000/\gamma$ of an eigenstate, its mirror on the negative branch, and the *Zitterbewegung* at $2E$ with mean $1000/\gamma^2$. All five predictions pass, and a gate bias of 26–71 m-qg is reported in the open.

- **Known, not new:** all of the above is textbook relativistic quantum mechanics (Thaller 1992) written on the scale of Chapter 1. What the scale adds is the reading: the energy in m-qg, the angle in qga and its shot noise.
- **Proposed reading, unverified as physics:** calling the hemispheres of the β qubit "matter" and "antimatter". It is the branch of the Dirac solution, not a probability that a particle is antimatter; a free positive-energy particle always reads in the northern hemisphere.
- **Not adopted:** that CPT is a translation of the state by 500 qga, or that $P_{\text{rest}} = p_0$ of the same angle. The correct version of both is above.

H Catalogue of qang sub-measures

The qang unit does not reformulate any physical law. Each sub-measure in this catalogue is a reading of one probability, or of a few, and equals a quantity that physics or statistics already has. What a sub-measure can bring is a **decision**: fewer shots, a shorter pulse, a compilation chosen before running, a reading that noise does not bias. For each one the catalogue records its definition, what it equals, which decision it improves and with what evidence, and classifies it with one of the statuses of Table H.1.

The tables are generated from the registry `qang.measures` with `tools/make_apendice_h.py`, the same source for both languages of the book. A new sub-measure enters the catalogue as **unverified** and changes status only with a test, preferably pre-registered (Chapter 28).

Table H.1: Statuses of the catalogue.

status	meaning	sub-measures
exact	exact identity: it names or re-scales something; no gain is claimed	18
measured gain	tested in simulation or analysis, with a measured gain over the usual practice	14
tested on hardware	tested on a real device	3
mixed	a gain in some regimes, none or a loss in others	7
no gain	tested: equivalent to the usual practice	4
against	tested: worse than the usual practice	2
unverified	untested proposal: a candidate for a pre-registered experiment	3

status	meaning	sub-measures
not adopted	proposal rejected as wrong or unsupported	1

H.1 Base readings

Table H.2: Base readings.

sub-measure	definition	equals	decision it improves	status	evidence	where
qg_Z	$2p_0 - 1 = \langle\sigma_z\rangle = \cos\theta$	expectation value of σ_z	working coordinate measured without inversion	**exact**	identity; holds for mixed states	ch. 1, §7
m-qg	$m = 1000\, qg_Z$	$\langle\sigma_z\rangle$ times a thousand, linear in p_0	budget gate error and shot noise on one scale	**tested on hardware**	IQM Garnet: 1.64 m-qg per X, over-rotations of 16–55 m-qg hidden from the readout model; M2 failed	ch. 1, §168
θ	$\arccos qg_Z$	arcsine transformation of a proportion	error bar $1/\sqrt{N}$ equal over the whole circle	**exact**	exact binomial sums; asymptotic (fails at $Np \approx 2.5$)	ch. 1, §125

sub-measure	definition	equals	decision it improves	status	evidence	where
qg_S	$H(p_0) = H\big((1 + qg_Z)/2\big)$	Shannon entropy of the outcome	certify randomness together with the mean reading	**mixed**	alone it is fooled under T1; paired with the mean reading it is not	chs. 1 and 26, §7, §51
qga	1 turn $=$ 1000 qga	an angle in another unit	write gate angles without factors of π	**exact**	convenience; changes neither gradients nor plateaus	ch. 1
(qg_Z, s)	$s = \mathrm{sign}(qg_X)$	the sign of the angle's sine	cross a pole during an optimization	**measured gain**	H_2: with qg_Z alone it stops at Hartree–Fock; with the sign it reaches the exact energy	ch. 1, §65

H.2 Derived quantities

Table H.3: Derived quantities.

sub-measure	definition	equals	decision it improves	status	evidence	where
qg_κ	$1/[\ln 2\,(1-qg_Z^2)]$	Fisher information about qg_Z, in bits	damp an optimizer's step near the poles	**no gain**	removes the divergence of the gradient in q, but it is descent on θ: 103 steps both (ch. 30); the JAX bench repeats it	chs. 1, 23 and 30
qg_{SNR}	$\lvert qg_Z\rvert/\sqrt{1-qg_Z^2}$	signal-to-noise of one shot, $\lvert\cot\theta\rvert$	shots to read a sign: $(z/qg_{SNR})^2$	**exact**	27 shots at 60° for 3σ; exact binomial: error < 1 %	chs. 1 and 3
qg_L	$\lvert\theta-\theta'\rvert/2$	Bhattacharyya angle; Fubini–Study on a common meridian	distance between states from two Z histograms	**exact**	with different azimuths it is only a lower bound	chs. 1 and 3
qg_τ	$\sqrt{1-r^2}=2\sqrt{\det\rho}$	lack of purity	see a pulse's area where the m-qg reading does not shift	**exact**	0.277 for $R_X(324^\circ)$ against 0.031 for $R_X(-36^\circ)$; needs three bases	chs. 1–3

sub-measure	definition	equals	decision it improves	status	evidence	where
κ_{qg}	$1000\,\lvert\cos\theta\rvert$	curvature of the reading with respect to the angle	predict the shift from amplitude error, $\frac{1}{2}\theta^2\sigma^2\kappa_{qg}$	**exact**	QuTiP: agrees to second order	chs. 1 and 2, §160
qFi	$4\pi^2\cdot 10^{-6}$ qga^{-2}	Fisher information of the qang angle	fixed angular resolution: $159.15/\sqrt{N}$ qga	**exact**	not a flat manifold: the sphere has curvature 4	ch. 1
qHz	qga$/$s	angular frequency, $2\pi/1000$ rad/s	none so far	**exact**	unit with no use in the book	ch. 1

H.3 States and registers

Table H.4: States and registers.

sub-measure	definition	equals	decision it improves	status	evidence	where
$1-r^2$	$4p_0(1-p_0)$ at definite weight	local purity	measure purity with one basis instead of three	**measured gain**	holds only with a definite Hamming weight	ch. 5

sub-measure	definition	equals	decision it improves	status	evidence	where
$\arccos(qg_Z/r)$	qg_Z/r	Bloch polar angle of a mixed state	direction of a noisy state	**mixed**	in favour with depolarizing noise; against with pure dephasing	ch. 5, §45
qg_{XX}, qg_{YY}	$\langle qg \otimes \sigma_i \rangle$	Pauli correlators	diagnose a Bell pair and predict teleportation without teleporting	**measured gain**	3 settings instead of 9; 4–5 two-qubit gates instead of 10–11, maximum error 0.029	chs. 12 and 13
m_E	$\sqrt{1000^2 - m_\perp^2}$	Eve's information (Helstrom with purification)	how many test rounds certifying bits needs	**exact**	the computation is standard; the reading makes it geometric	ch. 18
$(1000 - m_r)/2000$	$m_r = 1000\, r$	energy of the passive state	know how much energy nobody can extract	**exact**	modest: the split is read at a glance	ch. 19

sub-measure	definition	equals	decision it improves	status	evidence	where
syndrome qg_Z, qg_X	ancilla readings	the T_2/T_1 rule	choose and re-calibrate an error-correcting code without test shots	**measured gain**	6.4 % over the oracle against 64–220 % for a fixed code; drift 13.1 % against 20.4 %	ch. 26
γ	$1 + \frac{}{2\sqrt{1-qg^2}}$	overhead of cutting a gate	choose which gates to cut	**mixed**	up to 2.9 times against angle-blind rules; none against qiskit-addon-cutting	ch. 26
qg_Z on graph states	$qg_Z = 0,\ qg_S = 1$	the local Z reading	detect entanglement with local Z readings	**no gain**	cannot tell a graph state from a product: the limit belongs to the Z basis	ch. 30

sub-measure	definition	equals	decision it improves	status	evidence	where
parity qg	$(dqg)^2/(1-qg^2) = N^2$	quantum Fisher information of a GHZ state	reach the Heisenberg bound with a parity reading	**exact**	the one-shot formula of Chapter 1 reaches the bound	ch. 34

H.4 Gates and compilation

Table H.5: Gates and compilation.

sub-measure	definition	equals	decision it improves	status	evidence	where
minimum area	θ mod 2π; R_{ZZ} with period π	shortest representative of the same gate	compile with less pulse area: less error	**measured gain**	4-qubit QAOA: area 4.4× and error 4.1× lower; the 1.5× prediction for R_{ZZ} failed (1.14×)	ch. 2, §160

sub-measure	definition	equals	decision it improves	status	evidence	where
exposure	population outside the sector after each two-qubit gate	noiseless computation on the circuit	choose the compilation before spending shots	**measured gain**	rank correlation 0.80 with T1 leakage; zero exposure removes ≥ 98 % of T1	ch. 2, §68
N-conserving gates	$[U, \sum_i Z_i] = 0$	compilation within the sector	compile so that the filter removes T1	**mixed**	the filter removes 99.5 % of T1 against 48.5 % with CNOT; none on ions	ch. 24, §61
qg grid	equal steps in qg_Z	control quantization	quantize the control angles	**mixed**	probabilities $\pi/2$ times better, 43 of 45 loads; against for fidelity near the poles	ch. 27

sub-measure	definition	equals	decision it improves	status	evidence	where
best_chain	qubit chain of highest predicted K	layout choice from calibration data	choose the physical qubits before running	**measured gain**	6 to 17 points more shots kept than the default layout (simulated)	ch. 32

H.5 Noise and readout

Table H.6: Noise and readout.

sub-measure	definition	equals	decision it improves	status	evidence	where
weight filter, K	$K = (1-\gamma)^{kd}$	post-selection on the Hamming sector	read, train and differentiate under T1 without bias; shot cost known before running	**tested on hardware**	IQM Resonance: error 0.0615 → 0.0386 (1.59×) with 88 % of the shots; simulated: exact to 10^{-12} m-qg; against with $K \le 0.14$ (2 of 12) and on deep LiH (76.6 → 85.0 mHa)	chs. 20, 21, 23, 30 and 31, §112
echo, $M^{1/2}$	echo transfer matrix	correction of errors within the sector	remove what the filter cannot see	**measured gain**	−27 to −62 % of the filtered error, with $\binom{n}{k}$ more echo circuits	ch. 22, §106

sub-measure	definition	equals	decision it improves	status	evidence	where
qg_{eq}, T_{eff}	heralded affine readout	readout calibration with repeated measurements	calibrate readout without mistaking temperature for error	**measured gain**	effective temperature to 1 mK; avoids an unsafe certifier and recovers 10–50 % of bits	chs. 18 and 27, §31
step $\times\lvert\sin\theta\rvert$	$\theta - \eta\max(\lvert\sin\theta\rvert, c_{\mathrm{reg}})\,g$	trust region	train with aggressive rates near the poles	**measured gain**	2.3 to 13 times more successes from rate 2; little or nothing with safe rates	ch. 23
local cost	$(1 - \overline{qg_Z})/2$	local cost of Cerezo et al.	avoid the barren plateau of a global cost	**mixed**	470 times more gradient variance with 10 qubits and short circuits; none with deep circuits	ch. 23

sub-measure	definition	equals	decision it improves	status	evidence	where
ZNE of the filtered reading	zero-noise extrapolation on the filtered qg_Z	post-selection and ZNE separately (mitiq)	mitigate T1 and dephasing together	**measured gain**	55.9 → 1.1 m-qg with T1 and dephasing; with few shots ZNE does not pay	chs. 21 and 35
pair $(qg_S, \overline{qg_Z})$	entropy and mean reading of the register	XEB or HOP with an ideal simulation	tell T1 from unital noise from counts alone	**mixed**	0.88–0.91 against 0.70–0.76 with strong noise; chance with realistic noise	chs. 27, 29 and 30

H.6 Readings of known physical laws

Table H.7: Readings of known physical laws.

sub-measure	definition	equals	decision it improves	status	evidence	where
$\langle \beta \rangle$	mc^2/E	expectation value of the Dirac matrix β	read $E = mc^2$ and charge conjugation as a qg_Z	**tested on hardware**	exact and known; charge conjugation flips the reading; run §170 on IQM Garnet	app. G, §170
$\widehat{qg^2}$	unbiased squares	standard bias correction	compare two clocks with few shots	**measured gain**	bias from +11.8 % to +0.9 % with 100 shots	ch. 8
unbiased $\sqrt{S_1^2 + S_2^2 + S_3^2}$	unbiased Stokes squares	degree of polarization	estimate polarization with few photons	**measured gain**	the bias drops about three times, without vanishing	ch. 6
$\arccos(m/l)$	$\langle J_z \rangle / l$	direction of $\langle \mathbf{J} \rangle$	read angular momentum as a qg_Z	**exact**	the vector model $\arccos(m/\sqrt{l(l+1)})$ draws something else	ch. 11
flavour qg_Z	$1 - 2P_{\alpha\to\beta}$	oscillation probability	read the mixing angle with a fixed error bar	**exact**	identity on the two-flavour formula	ch. 9

sub-measure	definition	equals	decision it improves	status	evidence	where
m_Z, m_X	readings of the chronology qubit	comparing density matrices	tell P-CTC from D-CTC with one Z reading	**measured gain**	P-CTC moves m_Z, D-CTC does not	ch. 14
$(1000 + m_P)(1000 + m_F)$	$= 10^6$	Yoshida–Kitaev law	judge a wormhole decoder	**exact**	the law is known; the scale shows that $m_P = -500$ proves nothing	ch. 15
$N/2 + \sum_n q g_n/2$	$G/G_0 = \sum_n T_n$	Landauer formula	read conductance and partition noise on one scale	**exact**	conductance = mean qg, noise = mean qg^2; the bar of $\widehat{\theta}$ is constant (0.985–1.009)	ch. 33
line in artanh	$\operatorname{artanh} qg$	Fermi fit	measure electron temperature	**against**	bias of +2.9 % (Fermi fit: −0.14 %)	ch. 33
$\theta/2 = \mathrm{gd}(r)$	Gudermannian of the squeezing r	$\tanh^2 r$, $\bar{n}$ and the entropy	read squeezing as a qang angle	**exact**	a new identity, exact	ch. 34

sub-measure	definition	equals	decision it improves	status	evidence	where
$\theta = 2\pi\Phi/\Phi_0$	interference fringes	Aharonov–Bohm phase	estimate the flux with an even error bar	**exact**	deviation $1/2\pi\sqrt{N}$ constant; the sign needs another reading	ch. 34

H.7 Open and rejected proposals

Table H.8: Open and rejected proposals.

sub-measure	definition	equals	decision it improves	status	evidence	where
κ_{qg} compilation	minimum arrival curvature	—	mitigate barren plateaus	**unverified**	untested	chs. 1 and 2
E_{total}	$\langle H\rangle + \lambda_\theta L(\theta) + \lambda_\varphi L(\varphi)$	—	regularize a VQE	**unverified**	L and λ undefined	ch. 1
$R(\theta,\varphi)$	$N[p_1 + \alpha\cos\varphi]$	—	model emission from the excited population	**unverified**	N and α not fixed; not compared with an emission model	ch. 1

sub-measure	definition	equals	decision it improves	status	evidence	where
500 qga CPT	$E_{_} \leftrightarrow$ a 500 qga shift	—	place anti-matter on the scale	**not adopted**	no support; the correct reading is charge conjugation on $\langle\beta\rangle$	ch. 1, app. G
signs of $qg_{Z,i}$	sign $qg_{Z,i}$	bit-string histogram	read the marked item in a noisy search	**against**	66 % against 91 % for the histogram mode (5 qubits, 100 shots)	ch. 3, §63
stopping qg_{SNR}, qg_{κ}	threshold on $\lvert qg_k \rvert$	the formula for k^*	stop Grover iterations	**no gain**	they cannot tell whether the flag passed the pole; they cost one run per k	ch. 3
qg cycles	angle-to-reading translation in Wigner, Kitaev and OTOC	—	protect states or measure chaos	**no gain**	they neither protect nor measure chaos	ch. 34

H.8 Balance

Of 52 sub-measures, 18 are exact identities, 14 have a measured gain in simulation or analysis, 3 were tested on a real device (with a mixed result for the milli-qang), 7 are mixed, 4 turned out equivalent to the usual practice, 2 lost to it, 3 are unverified and 1 was not adopted. The gains do not come from the unit but from the decisions the reading allows before spending resources: choosing the minimum-area representative, the compilation of least exposure, the code from two readings, or reading only the weight sector that T1 does not contaminate. The unverified sub-measures are the list of pending experiments; the catalogue keeps them in sight so that they are not cited as results.

H.9 In code

```
from qang import measures

measures.summary()                        # sub-measures per status
measures.by_status("sin_verificar")       # the pending experiments
m = measures.find("min_area")
m.decision[1], m.evidence[1]          # the decision it improves and the evidence, in Engl
measures.resolve(m)                # qang.pulse_area.minimal_area_circuit
```