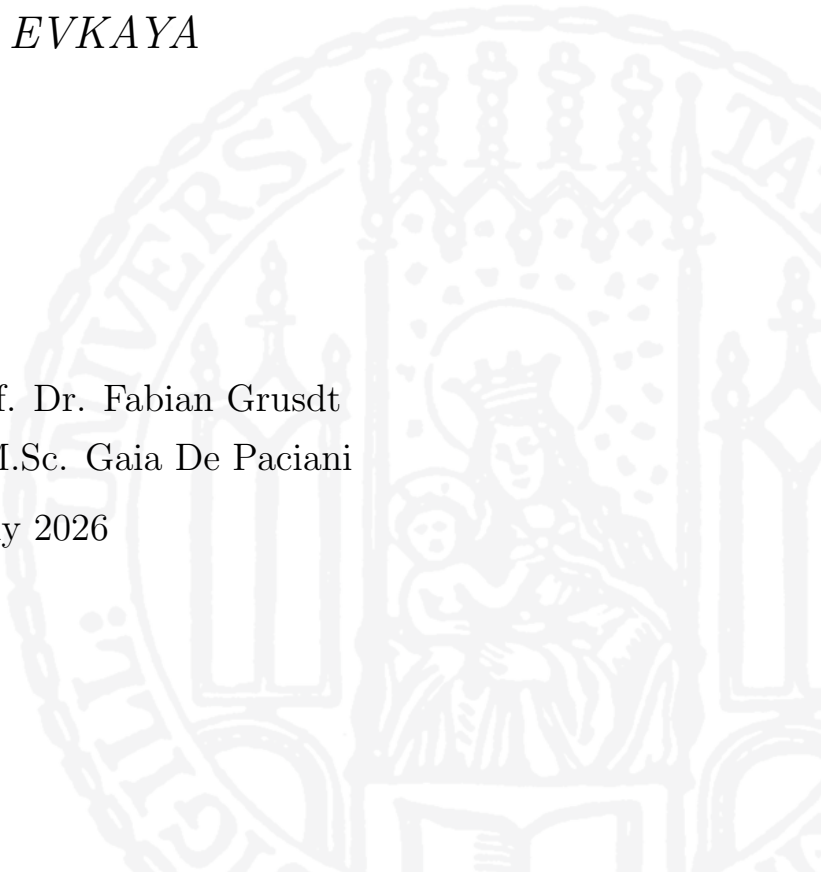
Ludwig-Maximilians-Universität München

Faculty of Physics

Master's Thesis

# QGP-Scope:
## Non-Abelian Gauge Dynamics in a (2+1)D SU(2) Quantum Link Model

*Yaman EVKAYA*



Supervision: Prof. Dr. Fabian Grusdt
Co-supervision: M.Sc. Gaia De Paciani

July 2026

Ludwig-Maximilians-Universität München

Fakultät für Physik

Masterarbeit

# QGP-Rahmen:
## Nichtabelsche Eichdynamik in einem (2+1)D-SU(2)-Quantenlinkmodell

*Yaman EVKAYA*

Betreuung: Prof. Dr. Fabian Grusdt
Mitbetreuung: M.Sc. Gaia De Paciani

Juli 2026

# Contents


**List of Notations** vi

**List of Figures** viii

**List of Tables** viii

**Summary** x

**I Introduction** 1

I.1 Non-Abelian Gauge Theories and Confinement . . . 1
I.2 Quark–Gluon-Plasma Motivation and Scope . . . 1
I.3 Quantum Link Models and Rishon Encoding . . . 1
I.4 Gauge-Aware Relaxation and Finite Temperature . . . 2
I.5 Cold-Atom Quantum Simulation . . . 2
I.6 Positioning and Original Contributions . . . 2

**II Model and Methods** 3

II.1 Rishon Representation and Gauge Constraints . . . 3
II.1.1 $(2+1)$D Honeycomb Geometry . . . 3
II.2 Hamiltonian . . . 5
II.2.1 $(2+1)$D Honeycomb Hamiltonian . . . 5
II.3 Observables and Diagnostics . . . 7
II.3.1 $(2+1)$D Honeycomb Diagnostics . . . 7
II.4 Tensor-Network Representation . . . 9
II.4.1 $(2+1)$D Honeycomb Ordering . . . 9
II.5 Numerical Methods and Validation . . . 10
II.5.1 Honeycomb Operator and State Validation . . . 10
II.5.2 Exact Sector Diagonalization of the C2 Patch . . . 10
II.5.3 Physical-Sector Thermal Representation . . . 14
II.5.4 Reproducibility . . . 15

**III Static Response and SU(2) Structure** 16

III.1 $(2+1)$D Source Sectors . . . 16
III.1.1 Mass Dependence at Fixed Separation . . . 16
III.1.2 Static Potential and Finite-Distance Force . . . 24
III.1.3 Coupling Dependence on C2 . . . 29
III.1.4 Wider-Patch Comparison at Fixed Separation . . . 33
III.2 $(2+1)$D Local Response . . . 37
III.3 Local Fusion Information . . . 38
III.3.1 Vertex, Plaquette, and Loop Information . . . 38

**IV Dynamics** 39

IV.1 $(2+1)$D Honeycomb Dynamics . . . 39

IV.1.1 Protocols and Observables . . . 39
IV.1.2 Breaking Quench: String-Breaking-Like Rearrangement . . . 40
IV.1.3 Formation Ramp: Flux-Tube-Formation-Like Growth . . . 42
IV.1.4 Exact Propagation and Tensor-Network Validation . . . 42
IV.1.5 Thermal Initial States and Non-Abelian Scope . . . 45
IV.1.6 Physical Interpretation and Experimental Implications . . . 46

**V Gauge-Aware Relaxation and Thermalization Criteria 47**
V.1 Closed-System Honeycomb Reference . . . 47
V.2 Gauge-Aware Honeycomb Thermal Dynamics . . . 47
V.2.1 Physical Sectors and Canonical Equilibrium . . . 47
V.2.2 Sector-Preserving Contacts and the Davies Reference . . . 48
V.2.3 Algorithmic KMS Thermostats . . . 49
V.2.4 Accessibility and the Conditional Uniqueness Result . . . 50
V.2.5 Independent Full-Sector Gibbs Estimates . . . 51
V.2.6 Finite-Temperature Trajectory Diagnostics . . . 52
V.2.7 Heat, Work and Ensemble Entropy . . . 53
V.2.8 Structural Response and Its Interpretation . . . 54
V.2.9 Established Results and Physical Scope . . . 55
V.3 Thermal-Response Protocols and Limits . . . 55
V.3.1 Stationary, Post-Preparation and Driven Contact . . . 55
V.3.2 Controls, Convergence and Statistical Decisions . . . 56

**VI Cold-Atom Implementation 58**
VI.1 Honeycomb Encoding and Implementation Tests . . . 58
VI.1.1 Exact Matter and Link Encoding . . . 58
VI.1.2 Microscopic Control Families and Effective Reduction . . . 59
VI.1.3 Star, Loop, and Source-Pair Validation . . . 61
VI.1.4 Preparation, Readout, and Error Separation . . . 62
VI.1.5 Numerical References and Nominal Statistical Planning . . . 63
VI.2 QGP-Motivated Correlation Readouts . . . 64
VI.3 Experimental Realization and Scope . . . 64
VI.3.1 Operator Validation and Physical Interpretation . . . 65
VI.3.2 Equilibrium Preparation and Driven Protocols . . . 66
VI.3.3 Thermal Coupling and Calibrated Readout . . . 67
VI.3.4 Atomic Platforms and Experimental Outlook . . . 67

**VII Conclusion 69**

**A Supplementary C2 Results and Numerical Accuracy 71**
A.1 Operator Conventions and Physical Coordinates . . . 71
A.2 Detailed Mass Ladders and Spatial Responses . . . 73
A.3 Gauss-Exact Static Certification . . . 76
A.4 Numerical Reliability and Reproducibility . . . 78

**B Sector-Resolved Thermal Methods and Numerical Controls 79**
B.1 Physical Representation and State Transfer . . . 79
B.2 Filtered Actions and Numerical Conditioning . . . 80
B.3 Constructive Accessibility and Uniqueness . . . 81
B.4 Equilibrium Estimators and Uncertainty . . . 82
B.5 Distinct Small-System Reference Calculations . . . 83

B.6 Stochastic Evolution, Replay and Error Separation . . . . . 84
B.7 Memory and Measured Action Costs . . . . . 85
B.8 Spectral and Tensor-Network Alternatives . . . . . 86
B.9 Ensemble Entropy and Response Design . . . . . 86

**C Wider-Patch Numerical Status** **88**

**Declaration** **95**

# List of Notations

| Symbol | Definition and use |
|---|---|
| **Geometry and local gauge structure** | |
| $v, \ell, p$ | Honeycomb vertex, oriented link and plaquette indices. |
| $A, B$ | Honeycomb sublattices; full-link ends follow this orientation. |
| $L_x, L_y$ | Open-patch dimensions. |
| $N_v, N_\ell$ | Numbers of vertices and links. |
| $C = N_\ell - N_v + 1$ | Independent cycle count of a connected patch. |
| $R, R_x$ | Shortest-path graph separation and cell-coordinate separation. |
| $\psi_v, n_v, N$ | Fermionic two-colour matter spinor, local occupation and total matter number. |
| $Q_v^a$ | SU(2) matter colour charge. |
| $c_{\ell,L/R,a}$, $n_{\ell,L/R}$ | Link-end rishon modes and their occupations. |
| $U_\ell^{ab}, U_p$ | Link transporter and ordered plaquette product. |
| $L_\ell^a, R_\ell^a, E_\ell^2$ | Link-end generators and electric Casimir, with $E_\ell^2 = 3n_{\ell,L}n_{\ell,R}/4$. |
| $G_v^a, \boldsymbol{G}_v^2$ | Local non-Abelian Gauss generator and Casimir. |
| $q_v^3$ | Classical source target in the selected colour direction. |
| $D_v, d_v$ | Conserved vertex imbalance and its selected eigenvalue. |
| **Hamiltonian, sectors and numerical controls** | |
| $H_{\mathrm{hop}}, H_m, H_E$ | Gauge–matter hopping, staggered mass and electric terms. |
| $H_p, H_G, H_D$ | Plaquette, Gauss-protection and imbalance-penalty terms. |
| $t_{\mathrm{h}}, m, g^2$ | Hopping amplitude, matter mass and gauge coupling. |
| $\kappa, \lambda, \mu$ | Plaquette, Gauss-protection and imbalance strengths. |
| $H_{\mathrm{opt}}, H_0$ | Optimization Hamiltonian and penalty-free physical Hamiltonian. |
| $P_s, \mathcal{U}_s, H_s$ | Physical-sector projector, embedding and sector Hamiltonian. |
| $\chi, \chi_{\max}, \varepsilon_{\mathrm{SVD}}$ | MPS bond dimension, imposed cap and truncation threshold. |
| $\mathrm{Var}(H)$ | Variational energy variance. |
| **Static and dynamical observables** | |
| $F_\ell$ | Active-link projector $4E_\ell^2/3$. |
| $P_{\mathrm{any}}, P_{\mathrm{no}}$ | Connection through any active path and its complement. |
| $P_{\mathrm{sel}}, P_{\mathrm{alt}}$ | Selected-corridor and alternative-path weights. |
| $P_{\mathrm{ff}}(v)$ | Flux-free single-particle source projector. |
| $P_{1\mathrm{ff}}, P_{2\mathrm{ff}}, P_{0\mathrm{ff}}$ | Exclusive source-disconnection classes. |

| Symbol | Definition and use (continued) |
|---|---|
| $\rho_\ell$ | Source-minus-vacuum electric response. |
| $V_D(R;m), \sigma_{\text{eff}}^D$ | Registered-sector potential and finite-distance secant force. |
| $\tau_E^{\text{resp}}, \tau_E^{\text{geo}}$ | Uniform and geodesic source-minus-vacuum electric responses per graph link. |
| $P_{ij}^{(0)}, W_{\text{plaq}}$ | Singlet projector and plaquette recoupling witness. |
| $t, \Delta t, \tau, F_{\text{GS}}$ | Time, integration step, ramp duration and ground-state population. |
| **Thermal structure and microscopic mapping** | |
| $\beta, \rho_{\beta,s}, Z_s$ | Inverse temperature, sector Gibbs state and partition function. |
| $S_a, A_a(\omega), \Gamma_a(\omega)$ | Contact, Davies jump and spectral transition rate. |
| $L_a, \mathcal{D}[L_a], G_{\text{KMS}}$ | KMS jump, Lindblad dissipator and coherent correction. |
| $N_{\text{traj}}, S_{\text{vN}}$ | Trajectory count and ensemble von Neumann entropy. |
| $g_A, g_B, \Delta$ | Molecular conversion amplitudes and detuning. |
| $P_A, P_B, b$ | Singlet pair-conversion operators and mediator mode. |
| $\Sigma_A, \Sigma_B$ | Independent self-energy counterterms. |
| $\varepsilon_{\text{op}}, \mathcal{L}$ | Projected operator error and leakage measure. |

# List of Figures


II.1 Geometry and source separation of the C2 patch . . . . . 4

III.1 Ground-state connectivity and screening on C2 . . . . . 18
III.2 Electric and matter response profiles of C2 . . . . . 19
III.3 Spatial maps of C2 electric and matter responses . . . . . 20
III.4 C2 energy budget and classical competition . . . . . 22
III.5 Source-relocation control on C2 . . . . . 24
III.6 Static separation dependence on C2 . . . . . 25
III.7 Energy budget of the C2 finite-distance force . . . . . 28
III.8 Electric-coupling dependence on C2 . . . . . 31
III.9 Plaquette-coupling dependence on C2 . . . . . 33
III.10 C2, C4 and C8 patches with their source corridors . . . . . 34
III.11 Connectivity and source potential on C2, C4 and C8 . . . . . 36

IV.1 String-breaking-like and flux-tube-formation-like dynamics on C2 . . . 41
IV.2 Gauss-exact TDVP accuracy on C2 . . . . . 44

V.1 Physical sectors, equilibrium references and thermal-response scope . . 50
V.2 Energy and all-path connectivity in the C2 canonical ensembles . . . . 53

VI.1 Encoding, microscopic reduction and readout of the honeycomb model 59

A.1 Static C2 separation cost: exact versus Gauss-exact DMRG . . . . . 77


# List of Tables


II.1 C2 sector dimensions and gaps . . . . . 12
II.2 Gauss-exact DMRG against exact C2 ground states . . . . . 14

III.1 Exact C2 mass ladder . . . . . 17
III.2 Classical C2 energy branches . . . . . 22
III.3 Controls on the exact C2 ladder . . . . . 23
III.4 C2 static data along the separation scan . . . . . 26

III.5 C2 finite-distance force and its energy budget . . . 27
III.6 Exact and Gauss-exact DMRG secants on C2 . . . 29
III.7 Coupling dependence of the exact C2 response . . . 30
III.8 Plaquette-coupling dependence of the exact C2 response . . . 32
III.9 Static response on C2, C4 and C8 at fixed separation . . . 35

IV.1 Characteristic C2 real-time values . . . 40
IV.2 Gauss-exact TDVP against exact C2 propagation . . . 43

V.1 C2 full-sector Gibbs energy and connectivity estimates . . . 52
V.2 C2 modular-KMS first-step diagnostics . . . 54
V.3 Scope of the C2 thermal results . . . 56

VI.1 Conditional microscopic validation of the two-rishon honeycomb model 62
VI.2 Nominal C2 endpoint shot-count oracle . . . 64
VI.3 Physical scope of the cold-atom construction . . . 66

A.1 C2 sector energies and gaps . . . 73
A.2 C2 configuration-projector ladder . . . 73
A.3 C2 source–vacuum response ladder . . . 73
A.4 C2 endpoint energy components . . . 74
A.5 C2 physical-energy decomposition . . . 74
A.6 All C2 electric-response profiles . . . 75
A.7 All C2 matter-response profiles . . . 75
A.8 C2 connectivity and screening controls . . . 75
A.9 Nonminimal Gauss penalty in the colour-diagonal control . . . 76
A.10 C2 relocation on the shared mass grid . . . 76
A.11 C2 Gauss-penalty comparison . . . 76

B.1 Constructive scalar-commutant witnesses for C2 . . . 81
B.2 Davies and KMS small-system numerical references . . . 83
B.3 Measured cost of the bounded C2 thermal calculations . . . 85

C.1 Numerical status of the C4 and C8 states . . . 89
C.2 Bond entropies of the stored C4 and C8 states . . . 90

# Summary

This thesis investigates source-supported electric flux, matter screening and real-time response in a finite two-rishon SU(2) quantum-link model on an open honeycomb lattice. Its representation, local conserved sectors and numerical accuracy are specified explicitly. The results are motivated by QCD colour dynamics, but concern finite patches rather than continuum or thermodynamic phases.

The two-cycle C2 patch is solved exactly in pre-specified source and vacuum sectors with gauge-invariant dimensions up to 46,829. At graph separation four, a thirteen-mass ladder resolves a crossover from screening-like to confinement-like source response. A separation scan at graph distances $R = 2, 4, 6$ gives a heavy-mass finite-distance force $1.03 \pm 0.01\, t_{\mathrm{h}}/a$ per graph link, of which about 69% is electric. At light mass the force softens by 47–53%, connected flux collapses and flux-free endpoint matter grows. The remaining positive cost is conditional on the conserved U(1) corridor labels and is not interpreted as a string tension. Increasing the electric coupling from $g^2 = 1$ to 8 reduces the scaled offset of the half-connectivity mass from the classical string/one-pair crossing from 0.65 to below 0.01, but at $g^2 = 6$ and 8 it lies slightly below the crossing, against the registered expectation; the heavy-mass force approaches the strong-coupling estimate (within 0.03% at $g^2 = 8$). By the registered criterion the heavy-mass force over $R = 2$–4 and the $R = 4$ connectivity are insensitive to the plaquette coupling up to $\kappa = 2$ (force within 0.05%), whereas the $R = 4$–6 forces fall by 5–11% at $\kappa = 2$; plaquette-dominated sector ground states appear at light mass from $\kappa = 1$ and at heavy mass at $\kappa = 4$. On wider four- and eight-cycle patches, whose corridor coordination differs from C2, heavy-mass variational states differ from the exact connectivity and source potential of the relocated C2 sources by +0.5% and −1.0%, with bounded numerical support on the eight-cycle patch; light-mass source states meeting the sector criteria have not been obtained there.

In real time, a mass quench from $m = 4$ to 0.5 strongly suppresses the source–vacuum electric response and connection probability while increasing flux-free endpoint support. Connectivity remains suppressed through the exact window $t \leq 20$. A ramp from $m = 0.5$ to 4 increases electric response and connectivity, but the subsequent hold does not establish a persistent formed tube. Independent first-principles reconstruction and Gauss-exact matrix-product calculations reproduce the static and real-time C2 references.

The thermal analysis separates equilibrium structure from reservoir-induced dynamics. Complete physical-sector Gibbs estimates, small-system Davies and KMS tests, a conditional uniqueness result and four finite-temperature first-step calculations establish bounded numerical and algebraic results. They do not yet measure converged ensemble thermalization, structural lifetimes or a QCD transition.

Finally, an exact local fermionic encoding is used to test microscopic cold-atom controls.

A restricted tunnelling family is insufficient, whereas a detuned molecular mediator generates the target hopping conditionally at second order. Star and hexagon calculations validate this finite-patch reduction under the stated controls; they are not an integrated apparatus or an experimental feasibility demonstration.

# Chapter I

# Introduction

## I.1 Non-Abelian Gauge Theories and Confinement

A local gauge symmetry constrains how matter exchanges charge with a field. For SU(2), noncommuting colour generators also permit distinct ways of coupling local spins to the same total representation. Finite quantum-link models make these structures accessible to exact diagonalization and tensor-network methods [12, 6].

This thesis asks how an imposed source pair organizes electric flux and how dynamical matter screens it on a loop-bearing honeycomb graph. The analysis separates source energies, connected electric response, matter screening, local recoupling and driven preparation. The terms *confinement-like* and *screening-like* refer to the specified finite-source observables, not to thermodynamic phases.

## I.2 Quark–Gluon-Plasma Motivation and Scope

QCD supplies the motivation: colour fields, screening and strongly interacting matter are central to its low-energy and plasma regimes. The present SU(2) model isolates selected mechanisms while omitting SU(3), a continuum limit, expanding matter and a thermodynamic extrapolation. String breaking into screened configurations therefore does not establish a deconfinement transition or the formation of a plasma droplet.

The source-free state is a matched reference, while the source sector contains imposed classical colour targets. Connected differences isolate source-associated response only when geometry, couplings and conserved labels are specified. The two-cycle C2 patch tests string-breaking-like and flux-tube-formation-like responses; persistent flux-tube formation is not established.

## I.3 Quantum Link Models and Rishon Encoding

Quantum-link models replace the infinite-dimensional link space of a Wilson Hamiltonian by a finite representation of the gauge algebra. A rishon construction realizes its generators through auxiliary fermionic modes [25, 31, 12]. The current honeycomb model uses canonical two-colour fermionic matter and a fixed two-rishon full link. Link orien-

tation, Jordan–Wigner registration and the conserved vertex imbalances are part of the Hamiltonian definition.

## I.4 Gauge-Aware Relaxation and Finite Temperature

Relaxation, thermal equilibrium and coupling to a reservoir are separate physical problems. The C2 analysis constructs complete physical-sector Gibbs references, checks accessibility assumptions and implements bounded numerical tests of specified thermostats. Stationary Gibbs structure and algebraic uniqueness do not themselves demonstrate finite-temperature mixing or a flux-tube lifetime.

## I.5 Cold-Atom Quantum Simulation

A cold-atom realization must encode the same matter and link spaces, derive the required couplings and quantify unwanted processes. The exact fermionic encoding below supports a comparison of microscopic control families. Ordinary tunnelling has insufficient independent controls in the tested family, whereas a detuned molecular mediator can generate the target interaction conditionally. Preparation, readout and laboratory calibration remain additional requirements.

## I.6 Positioning and Original Contributions

Real-time string dynamics and string breaking have been studied directly in $(2+1)$D non-Abelian calculations and quantum-simulator settings [10, 21], while honeycomb quantum-link constructions provide nearby theoretical references [7, 30]. The contribution here is a controlled finite-patch evidence hierarchy rather than novelty from the appearance of a string alone.

The C2 contribution combines exact fixed-sector ground states, a mass ladder, a registered source-separation scan, energy-resolved finite-distance forces, exclusive endpoint classes and controlled real-time quench and ramp protocols. An independent first-principles implementation and Gauss-exact tensor networks test the same operators and states. Exact scans of the electric and plaquette couplings test the mass-ladder results against the classical branch structure and a magnetic term. Variational states on the wider C4 and C8 patches compare the heavy- and intermediate-mass response with C2 under explicit numerical status labels. Complete-sector thermal references and a constructive accessibility witness are kept distinct from open-system response. Finally, the local encoding and finite-patch microscopic reductions identify which atomic controls would be required for a honeycomb realization.

Chapter II defines the model and methods. Chapters III and IV present static and driven response. Chapter V develops the thermal framework, Chapter VI the microscopic implementation tests, and Chapter VII states the conclusions and remaining limits.

# Chapter II

# Model and Methods

This chapter specifies the honeycomb geometry, local spaces, Hamiltonian, source sectors, observables and numerical methods. Every result is attached to its geometry, source definition and conserved sector.

## II.1 Rishon Representation and Gauge Constraints

The Hamiltonian formulation follows the Wilson–Kogut–Susskind framework [43, 26], while finite-dimensional links follow the gauge-magnet and quantum-link construction [25, 31, 12]. Gauge invariance follows from the local algebra; finite-dimensionality is a representation choice whose consequences are tested explicitly.

### II.1.1 $(2+1)$D Honeycomb Geometry

Honeycomb SU(2) gauge theories also admit Gauss-integrated spin-Hamiltonian descriptions [30]. That construction is a nearby theoretical reference rather than an identity: the present model retains explicit matter and rishon degrees of freedom with their conserved sectors specified separately.

The honeycomb model uses genuine two-colour fermionic matter and a fixed *two-rishon full link*. Its local spaces are

$$\mathcal{H}_v = \mathrm{span}\{|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}, \tag{II.1}$$

$$\mathcal{H}_\ell = \mathrm{span}\{|0,2\rangle, |\uparrow,\uparrow\rangle, |\uparrow,\downarrow\rangle, |\downarrow,\uparrow\rangle, |\downarrow,\downarrow\rangle, |2,0\rangle\}. \tag{II.2}$$

Here the two entries of a link label refer to its left and right ends; $n_{\ell,L} + n_{\ell,R} = 2$ is enforced by the basis itself. Empty and doubly occupied ends are colour singlets; a singly occupied end carries spin 1/2. Matter has the analogous singlet–doublet–singlet structure. Jordan–Wigner parity strings enforce anticommutation between matter sites, including graph neighbours separated by other matter sites in the tensor ordering. This registration is part of the Hamiltonian, not a measurement option.

**Matter statistics.** The honeycomb model uses the four-state local matter basis and registers the sites with full Jordan–Wigner parity strings. Matter operators on distinct

sites therefore obey canonical fermionic anticommutation. This registration is part of the $(2+1)$D Hamiltonian definition and is not interchangeable with an alternative matter-statistics convention.

Every graph edge is oriented from sublattice $A$ to $B$, identified with its $L$ and $R$ ends. The open patches have

$$N_v = 2L_xL_y, \qquad N_\ell = 3L_xL_y - L_x - L_y, \qquad C = N_\ell - N_v + 1 = (L_x-1)(L_y-1). \quad \text{(II.3)}$$

$C$ counts independent cycles; it is neither spatial dimension nor the number of paths between sources. A strip with $L_y = 1$ is a tree. The $2 \times 2$ patch ($C = 1$, hereafter C1) has eight vertices, eight links and one cycle; it is the smallest loop-bearing pilot used below. Source separation must specify graph distance and positions: for same-row $A$ sites a column displacement $R_x$ corresponds to graph distance $2R_x$. Throughout the honeycomb chapters $R$ denotes the graph distance, the number of links on a shortest source path; the principal C2 sources have $R_x = 2$ and $R = 4$. Two wider patches keep $L_x = 3$: C4 ($L_y = 3$; 18 vertices, 21 links, four cycles) and C8 ($L_y = 5$; 30 vertices, 37 links, eight cycles). Their source placements and corridors are shown in Fig. III.10.

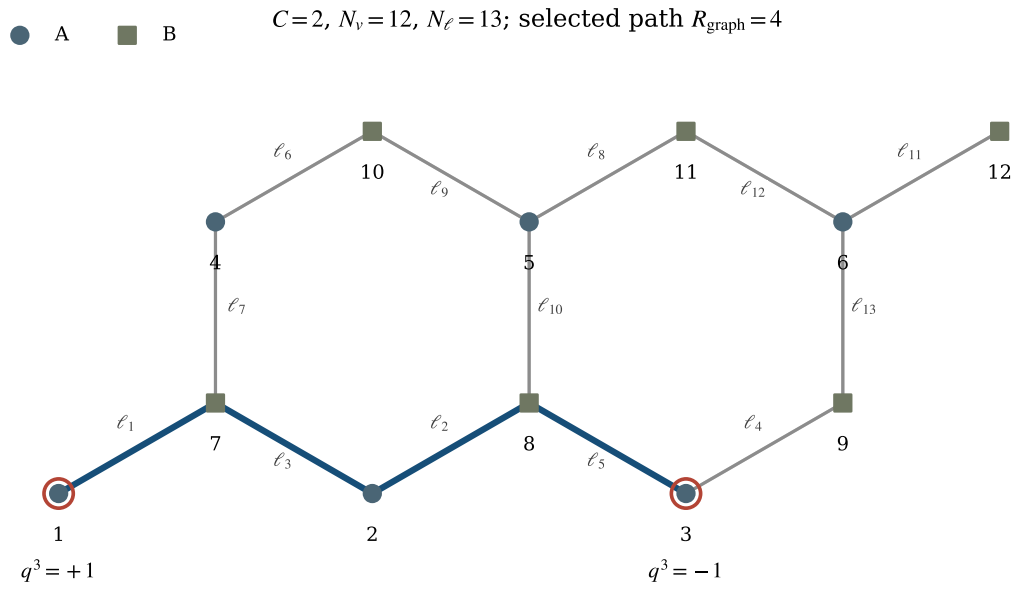



**Figure II.1:** The actual unit-bond embedding of the open $3 \times 2$ honeycomb patch. Circles and squares distinguish the $A$ and $B$ matter sublattices; vertex IDs appear below the sites and link IDs beside the bonds. The blue four-link corridor connects the circled classical sources at vertices 1 and 3. Its graph distance is four, while the same-row cell displacement is two. The tensor ordering does not replace this spatial geometry.

With $\boldsymbol{Q}_v = \psi_v^\dagger \boldsymbol{\sigma} \psi_v/2$, the source-free local generator and additional conserved charge are

$$\boldsymbol{G}_v = \boldsymbol{Q}_v + \sum_{\ell \ni v} \boldsymbol{E}_{\ell,v}, \qquad D_v = n_v - \sum_{\ell \ni v} n_{\ell,v}. \quad \text{(II.4)}$$

The end generators use the all-plus convention in this expression. Number the link basis states above from 1 to 6. The nonzero matrix entries of the implemented transporter are fixed explicitly by

$$U_{11} = |4\rangle\langle 1| - |6\rangle\langle 3|, \qquad U_{12} = -|5\rangle\langle 1| - |6\rangle\langle 2|, \tag{II.5}$$

$$U_{21} = |2\rangle\langle 1| + |6\rangle\langle 5|, \qquad U_{22} = -|3\rangle\langle 1| + |6\rangle\langle 4|. \tag{II.6}$$

They obey $[E^c_L, U_{ab}] = -(\sigma^c/2)_{aa'}U_{a'b}$ and $[E^c_R, U_{ab}] = U_{ab'}(\sigma^c/2)_{b'b}$, with summed colour indices. This $\epsilon$-conjugated convention moves a rishon from $R$ to $L$ when matter hops from $B$ to $A$. Thus each hopping event preserves both endpoint imbalances. The determinant-free Hamiltonian below obeys

$$[H, D_v] = [H, N] = 0, \qquad \sum_v D_v = N - 2N_\ell, \qquad N = \sum_v n_v. \tag{II.7}$$

A fixed $D_v$ profile is consequently a dynamical restriction. Equal total matter number does not imply that two states can be connected by this Hamiltonian.

## II.2 Hamiltonian

### II.2.1 $(2+1)$D Honeycomb Hamiltonian

The Hamiltonian used by the current honeycomb calculations is

$$\begin{aligned} H = & -t_{\mathrm{h}} \sum_{\ell=(A,B)} (\psi^\dagger_A U_\ell \psi_B + \mathrm{h.c.}) + m \sum_v \eta_v n_v + \frac{g^2}{2} \sum_\ell E^2_\ell \\ & - \kappa \sum_p (\mathrm{Tr}\, U_p + \mathrm{h.c.}) + \lambda \sum_v (\boldsymbol{G}_v - \boldsymbol{q}_v)^2 + \mu \sum_v (D_v - d_v)^2, \end{aligned} \tag{II.8}$$

where $\eta_A = +1$, $\eta_B = -1$. The ordered hexagonal product $U_p$ uses $U_\ell$ or $U^\dagger_\ell$ according to the traversal direction. $t_{\mathrm{h}}$ distinguishes the hopping amplitude from evolution time; $\hbar = 1$ throughout. The electric operator is

$$E^2_\ell = \frac{3}{4} n_{\ell,L} n_{\ell,R} = \mathrm{diag}(0, \tfrac{3}{4}, \tfrac{3}{4}, \tfrac{3}{4}, \tfrac{3}{4}, 0). \tag{II.9}$$

This is state dependent. The expectation of the occupation product must be contracted as an operator; replacing it by the product of occupations is not valid. The local quantum-link construction retains a finite SU(2) algebra, without asserting a continuum or large-representation limit [25, 12, 6].

**Sources and physical energies.** The impurities are classical $q^3 = +1, -1$, with $q^1 = q^2 = 0$. There is no additional quantum external-doublet site. At a source the energetically selected block has $j = 1/2$, $G^3 = q^3/2$, and

$$\langle (\boldsymbol{G} - \boldsymbol{q})^2 \rangle = j(j+1) - 2q^3 G^3 + (q^3)^2 = \frac{3}{4}. \tag{II.10}$$

Neutral vertices target $j = 0$. Two sources at $\lambda = 20$ therefore contribute a floor of 30, not 40. Source energy offsets must be removed consistently before comparing physical energies. The source-free gauge part commutes with all local generators; the classical source term selects a colour axis, retaining $G^3$ and the local Casimir but not full transverse commutation at the source. Gauge quality is assessed against the appropriate source or neutral block, not against zero at every vertex.

The C2 exact ground states explicitly realize this minimal source block: each source contributes 3/4, the pair contributes 30 at $\lambda = 20$, and all neutral vertices are singlets. The vacuum contribution is zero. At the examined full-SU(2) parameter points this makes the ground state unchanged for every $\lambda' \geq 20$, with only the penalty constant shifting. The argument and the direct $\lambda = 40$ check are given in Sec. II.5.2; the statement does not extend to the non-minimal Cartan control. That section also distinguishes the product-sector restriction from the energetically selected Casimir block. For the C2 comparison, $H_{\mathrm{opt}} \equiv H$ and $H_0 \equiv H_{\mathrm{phys}}$; this notation applies to the static and thermal sector calculations, not the unperturbed microscopic cold-atom Hamiltonian.

The pinning term selects $D_v = d_v$. In that exact sector it vanishes; at finite numerical precision its mean and fluctuations are measured. Define $H_{\mathrm{phys}} = H - H_G - H_D$, where $H_G$ and $H_D$ denote the two penalties. Its expectation and variance are reported together with those of the optimization Hamiltonian. Subtracting two approximate variational energies does not produce a variational upper bound on their difference.

The determinant-free model preserves the additional local $D_v$ charges, so the thermal ensemble fixes their declared pattern as well as the non-Abelian Gauss representation and $N$. In a retained physical sector, $H_{\mathrm{opt}} = H_0 + c_s I_s$, with $c_{\mathrm{src}} = 30$ and $c_{\mathrm{vac}} = 0$ at the stated penalty. This constant cancels from normalized Gibbs states and transition frequencies but must be retained when converting partition functions or free energies. It is not constant on arbitrary states outside the physical block.

The loop-bearing pilot uses $t_{\mathrm{h}} = 1$, $g^2 = 2$, $\lambda = 20$, $\kappa = 0$, $N = 8$, sources at vertices 1 and 2 (both on $A$), and $d = (-1, -1, 0, 0, -2, -2, -2, 0)$. Its sector matrix has dimension 17431. The zero plaquette coupling is an explicit restriction: matter-mediated gauge rearrangement is active, but an independently driven magnetic plaquette mechanism is not tested by these calculations. On C2 the plaquette coupling is varied up to $\kappa = 4$ in Sec. III.1.3.

**What pair production means.** Relative to the staggered reference with empty $A$ and doubly occupied $B$ vertices, define

$$N_{\mathrm{p}} = \sum_{v \in A} n_v, \qquad N_{\bar{\mathrm{p}}} = \sum_{v \in B} (2 - n_v), \tag{II.11}$$

$$N_{\mathrm{p}} - N_{\bar{\mathrm{p}}} = N - 2M_B, \qquad H_m = m(2N_{\mathrm{pair}} - 2M_B), \tag{II.12}$$

where $M_B$ is the number of $B$ vertices and $N_{\mathrm{pair}} = (N_{\mathrm{p}} + N_{\bar{\mathrm{p}}})/2$. At $N = 2M_B$, $N_{\mathrm{pair}} = N_{\mathrm{p}} = N_A$. Matter hopping can create a particle–hole pair relative to that reference while conserving $N$. The rishons encode the gauge link; they are not the produced dynamical matter. The interacting ground state is already dressed, so its pair expectation need not vanish. A change in this expectation is a net occupation change, not a count of individual

creation events. Opposite impurity colour signs on two $A$ vertices do not make those vertices a particle and an antiparticle site.

# II.3 Observables and Diagnostics

Source and vacuum states are evaluated in separately registered sectors with the same couplings and geometry. Source-minus-vacuum subtraction is therefore a comparison of declared sectors, not a source-switching operation inside one Hilbert-space block.

## II.3.1 $(2+1)$D Honeycomb Diagnostics

The electric Casimir profile $\langle E_\ell^2\rangle$, matter occupations and pair number resolve distinct aspects of a source state. A connected profile $\delta E_\ell^2 = \langle E_\ell^2\rangle_q - \langle E_\ell^2\rangle_0$ requires a control with the same couplings, geometry, number and declared imbalance constraints. Different source and source-free sectors must be labelled; no subtraction silently changes them. Likewise a potential $V(R) = E_{q,\mathrm{phys}}(R) - E_{0,\mathrm{phys}}$ is meaningful only with a specified sector-selection rule. Comparing two constrained branches is not a minimization over all superselection sectors.

Let $F_\ell = 4E_\ell^2/3$. It is a projector onto an active link. For each basis configuration the active links define a graph. The diagonal projector $P_{\mathrm{any}}$ tests whether *any* active path connects the sources; $P_{\mathrm{no}} = 1 - P_{\mathrm{any}}$. A selected-corridor projector instead requires all links of that corridor to be active. On a loop, a loss on one corridor can represent rerouting, so its complement alone is not an all-path breaking probability. These graph projectors use occupations and commute with the local colour generators.

A second diagonal witness, $P_{\mathrm{ff,m}}$, requires both source vertices to carry matter while every link incident on them is inactive. It measures joint local flux-free matter support. With classical impurities it is not an external-source–matter singlet projector; its expectation is a probability, not proof that every configuration is screened. Corroboration comes from matter redistribution, connected electric profiles and the energy budget.

**C2 local projectors and structural restrictions.** For C2, let $\Pi_v^{(n=1)}$ project onto exactly one matter particle and define

$$P_{\mathrm{ff}}(v) = \Pi_v^{(n=1)} \prod_{\ell \ni v} (1 - F_\ell), \qquad P_{\mathrm{both}} = P_{\mathrm{ff}}(1) P_{\mathrm{ff}}(3). \tag{II.13}$$

These projectors measure flux-free single-particle support, not bare flux freedom on an unrestricted Hilbert space. In the registered source sector, vanishing incident flux and $G_v^3 = q_v^3/2 = \pm 1/2$ already force a single matter particle. Source vertex 1 has degree one; calling its sole incident link $\ell_1$, the exact sector identity is

$$\langle P_{\mathrm{ff}}(1)\rangle = 1 - \frac{4}{3}\langle E_{\ell_1}^2\rangle. \tag{II.14}$$

Thus this probability and the electric profile are not independent evidence. Both sources can be flux free only with at least two particle–hole excitations; in the fixed-$D_v$ sector at

$R = 4$, a connection avoiding the corridor requires at least four pairs. The local flux-free single-particle probabilities vanish identically in the vacuum sector by construction; this does not imply vanishing vacuum all-path connectivity. For the source-minus-vacuum electric response, define

$$\rho_\ell = \langle E_\ell^2 \rangle_{\rm src} - \langle E_\ell^2 \rangle_{\rm vac}, \qquad \epsilon_{\rm off} = 1 - \frac{\sum_{\ell \in \mathcal{C}} \rho_\ell}{\sum_\ell \rho_\ell}. \tag{II.15}$$

The ratio requires a nonzero denominator. Because the connected response is signed, it is a spatial-response ratio rather than a probability constrained to $[0, 1]$.

**Exclusive source classes.** For the two source vertices $s, s'$ of a registered source sector (one specified before the corresponding data were generated), write $B_v = \Pi_v^{(n=1)} \prod_{\ell \ni v}(1 - F_\ell)$, so that $B_v = P_{\rm ff}(v)$ as an operator, and define

$$\Pi_{\rm 2ff} = B_s B_{s'}, \qquad \Pi_{\rm 1ff} = B_s(1 - B_{s'}) + (1 - B_s)B_{s'}, \qquad \Pi_{\rm 0ff} = (1 - \Pi_{\rm any})(1 - B_s)(1 - B_{s'}). \tag{II.16}$$

A connected configuration has active links at both sources, so $\Pi_{\rm any}$ and the three projectors above are mutually exclusive and sum to the identity; no probability is discarded. Their expectations $P_{\rm any}$, $P_{\rm 1ff}$, $P_{\rm 2ff} = P_{\rm both}$ and $P_{\rm 0ff}$ resolve one- and two-endpoint flux-free matter support and disconnection without either endpoint event; $P_{\geq 1{\rm ff}} = P_{\rm 1ff} + P_{\rm 2ff}$. The C2 separation scan uses $(s, s') = (1, 2)$, $(1, 3)$ and $(1, 6)$ at $R = 2, 4, 6$. Several C2 figures label $\Pi_{\rm any}$ and $P_{\rm any}$ as $\Pi_{\rm conn}$ and $P_{\rm conn}$. Growth of $P_{\rm 1ff}$ or $P_{\rm 2ff}$ with declining $P_{\rm any}$ resolves one- or two-endpoint flux-free rearrangement; growth of $P_{\rm 0ff}$ alone does not establish matter screening, and a high $P_{\rm any}$ is not a measurement of tube width, persistence or a pure connected-string state.

**Static potential and finite-distance force.** For a registered source-sector family $D_R$ at graph separation $R$ (Sec. II.5.2), the physical static potential and the finite-distance force per graph link are

$$V_D(R; m) = \langle H_0 \rangle_{{\rm src}, D_R, m} - \langle H_0 \rangle_{{\rm vac}, m}, \qquad \sigma_{\rm eff}^D(R_i, R_j; m) = \frac{V_D(R_j; m) - V_D(R_i; m)}{a(R_j - R_i)}, \tag{II.17}$$

with $a = 1$ and $R$ the number of links on a shortest source path. In the registered exact sectors the $D_v$ penalty vanishes and the source Gauss-protection offset is 30, so that $V_D = E_{\rm opt,src} - 30 - E_{\rm opt,vac}$. The secant is defined in analogy with the potential-difference force of Ref. [36, Sec. 2.2, Eq. (2.5)], which refers to Euclidean displacements along a fixed lattice direction; on the honeycomb graph it is a force per graph link rather than a derivative along one axis. For a linearly rising potential $V_D = \sigma R + c$ the secant equals $\sigma$, and screening makes it decrease with $R$. Where $V_D$ rises approximately linearly in $R$, as at heavy mass below, the secant is referred to as the *effective string tension* of the patch; otherwise it is quoted only as a finite-distance force. In neither case is it an asymptotic string tension. Splitting $H_0 = H_E + H_m + H_{\rm hop}$ gives the budget $\sigma_{\rm eff}^D = \sigma_E^D + \sigma_m^D + \sigma_{\rm hop}^D$.

**Electric response.** With $H_E = (g^2/2)\sum_\ell E_\ell^2$, define for $R > 0$

$$\widehat{\mathcal{T}}_E^{(R)} = \frac{H_E}{aR}, \qquad \tau_E^{\mathrm{resp}}(t) = \langle \widehat{\mathcal{T}}_E^{(R)} \rangle_{\mathrm{src}}(t) - \langle \widehat{\mathcal{T}}_E^{(R)} \rangle_{\mathrm{vac}}(t), \tag{II.18}$$

where the sum runs over all links and source and vacuum states follow the same protocol in their own sectors. Its ground-state counterpart is $\tau_{E,\mathrm{GS}}^{\mathrm{resp}}(R;m) = V_E^D(R;m)/(aR)$ with $V_E^D = \langle H_E \rangle_{\mathrm{src}} - \langle H_E \rangle_{\mathrm{vac}}$. These normalizations are operational definitions of this work. $\tau_E^{\mathrm{resp}}$ is an energy per graph link, not an evolution time; equal units do not make it a force. Even for a linear electric contribution $V_E = \sigma_E aR + c_E$, one has $V_E/(aR) = \sigma_E + c_E/(aR)$ whereas $dV_E/d(aR) = \sigma_E$, and matter, hopping and endpoint contributions prevent identifying $\tau_E^{\mathrm{resp}}$ with $\sigma_{\mathrm{eff}}^D$. A geodesic control replaces the uniform link sum by path weights, $\tau_E^{\mathrm{geo}} = (g^2/2a)\sum_\ell w_\ell \Delta\langle E_\ell^2 \rangle$, where $w_\ell$ is the fraction of shortest source paths containing link $\ell$, divided by $R$, so that $\sum_\ell w_\ell = 1$. It measures the electric response carried along the source paths; a difference from $\tau_E^{\mathrm{resp}}$ can reveal off-path or loop contributions or path multiplicity.

For two spin-half constituents at a neutral vertex, local colour fusion can be measured by the scalar

$$P_{ij}^{(0)} = \frac{1}{4} P_i^{(1/2)} P_j^{(1/2)} - \boldsymbol{S}_i \cdot \boldsymbol{S}_j. \tag{II.19}$$

The active-spin projectors are essential if empty or doubly occupied states are allowed. Conditional expectations within a four-spin sector must be reported with that sector's probability. A Casimir profile, a binary path measurement or a fusion population is not, alone, a uniquely non-Abelian signature. A matched control must test what information is lost when colour recoupling is removed.

## II.4 Tensor-Network Representation

### II.4.1 $(2+1)$D Honeycomb Ordering

The MPS contains one dimension-four matter tensor per vertex and one dimension-six tensor per full link. The linear tensor ordering determines operator range and entanglement cost, not the physical graph adjacency. An ordering change must transform fermionic parity strings consistently. Small patches are also assembled as sector-resolved sparse matrices, providing an independent many-body contraction route using the registered local algebra. Larger bond dimensions address variational error; increasing patch width addresses a different, physical finite-size question. Neither can replace the other.

The Gauss-exact C2 calculations use a different site arrangement: one site per vertex combines the matter state with the rishon ends of its links, restricted to the local Gauss and $D_v$ targets (Sec. II.5.2). In this vertex ordering several hoppings connect sites that are not nearest neighbours in the linear tensor sequence. A two-site tangent space can then miss directions generated by the Hamiltonian, which can limit projected time integrators (Sec. IV.1.4).

## II.5 Numerical Methods and Validation

Small patches are solved in complete registered physical sectors using sparse matrix actions and Lanczos iteration. The independent implementation rebuilds the fermionic local operators and physical intertwiners before evaluating the same observables. Gauss-exact matrix-product states impose the local Casimir, Cartan and imbalance targets in their basis. Ground-state and real-time claims are accompanied by residual, observable, fidelity and step-size comparisons. The wider C4 and C8 patches are treated with penalty-protected DMRG and enter only with the numerical status labels of Sec. III.1.4.

### II.5.1 Honeycomb Operator and State Validation

The current model is checked at three levels. Local identities test the representation, hopping convention, Hermiticity and conserved charges. Full hexagonal link-space checks use all $6^6 = 46656$ link configurations. The loop-bearing matter sector is assembled under two matter orderings; after fermionic basis transport the matrices coincide, while deliberately omitted exchange signs produce a nonzero mismatch. These checks establish the tested operator registration, not the absence of all possible implementation errors.

For the two pilot ground states, the residual norms are below $2 \times 10^{-13}$. Independent finite differences test the response identities

$$\frac{\partial E_0}{\partial m} = \left\langle \sum_v \eta_v n_v \right\rangle , \qquad \frac{\partial E_0}{\partial g^2} = \frac{1}{2} \left\langle \sum_\ell E_\ell^2 \right\rangle . \tag{II.20}$$

Halving the parameter step from $10^{-3}$ to $5 \times 10^{-4}$ gives agreement within $1.2 \times 10^{-7}$ for the mass derivative and $2.1 \times 10^{-10}$ for the electric derivative over the two masses. A charged-star ED–DMRG comparison at the same two masses agrees in energy within $4.1 \times 10^{-14}$ and in measured local observables within $2.7 \times 10^{-15}$. This is a tree benchmark; it does not certify large loop-bearing MPS states. Shared local matrices also limit how independent the two solvers are.

An MPS ground-state approximation is characterized by its energy and variance, Hamiltonian parameters, operator ordering, source geometry, $N, D_v$, and source-adjusted Gauss diagnostics. Reproducibility requires the same state and numerical conventions. Ground-state confidence further requires independent initializations and observable convergence under bond dimension and cutoff. A small variance establishes proximity to an eigenstate, not by itself to the ground state. These requirements apply before promoting an energy difference or a profile to a physical result.

### II.5.2 Exact Sector Diagonalization of the C2 Patch

The C2 reference resolves the ground states of the open $3 \times 2$ honeycomb patch in explicitly specified sectors. Its twelve four-state matter vertices and thirteen six-state quantum links give 25 sites in the matter–link ordering; the Gauss-exact representation below uses twelve vertex sites. The full product space has dimension $4^{12}6^{13} = 219\,122\,084\,616\,339\,456$. Sector factorization makes this problem accessible without an energy cutoff. The parameters are

$t_{\mathrm{h}} = 1$, $g^2 = 2$, $\lambda = \mu = 20$, with $\kappa = 0$ in the main mass ladder and in the separation scan. All energies are in units of $t_{\mathrm{h}}$; several C2 figures write the hopping amplitude as $J \equiv t_{\mathrm{h}}$.

**Sources, symmetries and sector selection.** The sources are classical colour charges $q_1^3 = +1$, $q_3^3 = -1$; no external quantum-doublet site is added. Both vertices belong to sublattice $A$. The retained source blocks have $j = 1/2$, $\boldsymbol{G}_v^2 = 3/4$, and $G_v^3 = q_v^3/2$; neutral vertices are singlets. With all incident link-end generators added to the matter charge, $\boldsymbol{G}_v = \boldsymbol{Q}_v + \sum_{\ell \ni v} \boldsymbol{E}_{\ell,v}$, the source contribution is

$$\lambda(\boldsymbol{G}_v - \boldsymbol{q}_v)^2 = \lambda\big(3/4 - 2q_v^3 G_v^3 + (q_v^3)^2\big) = 3\lambda/4.$$

Thus the two-source Gauss minimum is 30, not zero or 40. The penalty-free Hamiltonian $H_0$ commutes with every local colour generator. The classical source penalty selects a colour axis: at a source, the full Hamiltonian preserves $G^3$ and $\boldsymbol{G}^2$, but not generally $G^1, G^2$. This is a fixed-source representation, not a demonstrated equivalence to a quantum-source global singlet for all observables.

The determinant-free Hamiltonian also preserves $D_v = n_v - \sum_{\ell \ni v} n_{\ell,v}$. In vertex order $1, \ldots, 12$, the fixed-separation sectors with sources at vertices 1 and 3 are

$$\begin{aligned} d_{\mathrm{src}} &= (-1, -2, -1, 0, 0, 0, -2, -2, -2, -2, -2, 0), \\ d_{\mathrm{vac}} &= (0, 0, 0, 0, 0, 0, -4, -4, -2, -2, -2, 0). \end{aligned} \tag{II.21}$$

Both profiles sum to $-14$, fixing $N = 2N_\ell + \sum_v D_v = 12$. Each mass ladder remains in its own fixed $D_v, G_v^3, N$ sector. Source–vacuum subtraction compares these two declared sectors; it is neither a unitary source-switching protocol within one sector nor a minimization over all possible imbalance profiles.

The separation scan uses a registered sector family $D_R$, that is, one specified before the corresponding data were generated: the sources occupy vertices $(1, 2)$, $(1, 3)$ and $(1, 6)$ at graph separations $R = 2, 4, 6$, with Euclidean separations $\sqrt{3}$, $2\sqrt{3}$ and $\sqrt{21}$ bond lengths. Each member fixes the labels $D_v$ along the source corridor by the same geometric balance rule; the $R = 4$ member is the sector of Eq. (II.21). At $R = 6$ three shortest paths join the sources and three sectors are registered: a predeclared primary sector and two alternative routes, B and C. Their spread is a systematic sector sensitivity, not a statistical error. All source records share the vacuum reference of the same mass.

This additional local conservation law is physically significant. For the two-rishon matrices, $\det U = 2|2, 0\rangle\langle 0, 2|$ and $[n_L, \det U] = 2 \det U$. A determinant interaction would be an SU(2) scalar but would change the local imbalance labels. Its absence leaves the source-free local algebra $\mathfrak{su}(2) \oplus \mathfrak{u}(1)$, a U(2)-type quantum-link structure. The results therefore concern the SU(2) colour algebra in specified Abelian sectors of the determinant-free model, not an unconstrained realization of every pure-SU(2) sector.

**Kinematic and physical Hilbert spaces.** Enumerating $3^{13} = 1\,594\,323$ link-end occupation patterns, rather than all $6^{13}$ colour-resolved link configurations, yields the complete fixed-$G^3, D, N$ product bases, for example of dimensions 6 713 288 (source, $R = 4$) and 4 498 399 (vacuum). These bases do not impose the minimal Casimir on every product

vector. Local intertwiners further reduce them to the gauge-invariant physical blocks of Table II.1. For $k$ spin halves, the multiplicity of spin $j$ is

$$M(k,j) = \binom{k}{k/2-j} - \binom{k}{k/2-j-1},$$

with inadmissible binomial entries set to zero. Four spin halves admit two singlets; three admit two fundamental intertwiners. Occupation labels alone therefore need not resolve the colour-fusion multiplicity.

| Sector | Product $(G^3, D, N)$ | Physical | Gap, $m = 4$ | Gap, $m = 0.5$ |
|---|---|---|---|---|
| Vacuum | 4 498 399 | 16 502 | 9.230566 | 2.682117 |
| $R = 2$ | 4 363 471 | 21 972 | 7.389744 | 1.462568 |
| $R = 4$ | 6 713 288 | 27 580 | 5.548551 | 0.756744 |
| $R = 6$ primary | 8 726 960 | 33 215 | 3.687790 | 0.501157 |
| $R = 6$ route B | 16 149 597 | 46 829 | 3.684068 | 0.515613 |
| $R = 6$ route C | 8 662 564 | 32 637 | 3.678736 | 0.587944 |

**Table II.1:** Fixed-$G^3, D, N$ product-sector and gauge-invariant physical dimensions of the six registered C2 sectors, with the first excitation gap in each sector at the two endpoint masses ($t_\mathrm{h} = 1$, $g^2 = 2$, $\kappa = 0$). The values were reproduced by the independent first-principles implementation described below.

**Sparse solution.** A grouped sparse matrix–vector application and Lanczos iteration determine the four lowest algebraic eigenvalues at a solver tolerance of $10^{-12}$. The fixed-separation set contains thirteen masses in two sectors (26 records), ten colour-diagonal controls and ten $\kappa = 0.5$ controls: 46 records in total. The six relocated-source records are additional. The separation scan comprises ten source records, $R = 2$, 4 and the three $R = 6$ sectors at both masses, of which the two $R = 4$ records are the ladder endpoints; all share the two vacuum references. The ground levels are nondegenerate on the sampled grid.

Every stored fixed-separation ground vector was re-evaluated on a regenerated basis for norm, energy, residual, sector invariance, local profiles and mechanism projectors. The largest ground-state residual is $3.10 \times 10^{-13}$; energy, profile and projector differences are at most $4.20 \times 10^{-12}$, $2.7 \times 10^{-13}$, and $6.7 \times 10^{-16}$. The largest out-of-sector component of $H\psi$ is approximately $2.0 \times 10^{-16}$. Independent fermionic Fock constructions, 144 small-system term comparisons at 24 parameter points and a production-MPO comparison of 35,428 sampled matrix elements (maximum difference $7.39 \times 10^{-13}$) test the operator assembly. Removing Jordan–Wigner signs produces 2,925 predicted sign errors: agreement of Gauss commutators alone would not establish the correct fermionic statistics.

**Independent first-principles re-derivation.** Every exact datum used for the separation scan was re-derived by a separate implementation that does not use the production Hamiltonian assembly. Its local operators are constructed from four rishon modes per link and two colour modes per vertex. They equal the production matrices up to a

pure basis permutation, with no sign changes, and satisfy the $\mathfrak{su}(2)$, anticommutation and covariance identities exactly. The Jordan–Wigner Hamiltonian action was written independently and checked against explicit fermionic Kronecker products on a test graph, with difference zero; a negative control without Jordan–Wigner strings fails that check. For each sector the gauge-invariant physical block is spanned by the embedding $\mathcal{U}_s$ of local Casimir intertwiners, built by the routine of the thermal calculation (Sec. V.2.1) from the production local operators; the Hamiltonian is applied matrix-free, and ground states follow from Lanczos iteration with full reorthogonalization. For all twelve ground states (six sectors, two masses) the re-derived values agree with the production results within $1.9 \times 10^{-12}$ for $\langle H_0 \rangle$, $2.8 \times 10^{-13}$ for every $\langle E_\ell^2 \rangle$, $9.9 \times 10^{-13}$ for every $\langle n_v \rangle$, $8.2 \times 10^{-14}$ for every class probability of Eq. (II.16) and $3.6 \times 10^{-13}$ for the gap. Eigenvector residuals are at most $2.4 \times 10^{-11}$, and every neutral-vertex Casimir is at most $1.9 \times 10^{-13}$. For the $R = 4$ source and vacuum sectors used in all real-time protocols, the hopping blocks of the independently derived physical-block Hamiltonian equal the production ones element by element within $5.5 \times 10^{-13}$ and $6.6 \times 10^{-14}$, and the electric diagonals agree within $2.3 \times 10^{-12}$ and $8.9 \times 10^{-16}$; the mass diagonal was not compared. Hopping energies are also evaluated directly as named expectation values. The exact re-propagation of the real-time protocols (Sec. IV.1.4) has its own diagonal, integrators and observables but shares the production operator layer, hopping term list and Cartan assembly of the hopping part, projected with the same embedding; its hopping blocks are those compared above.

**Gauss-exact tensor networks.** Approximate states are constructed in a matrix-product representation in which the non-Abelian Gauss law holds by construction. One MPS site per vertex carries the matter state and the rishon ends of its links, restricted to the local Gauss target (a singlet at neutral vertices and the registered doublet component at the sources) and to the local $D_v$ target. Every local basis state meets its Casimir, $G^3$ and $D_v$ targets within $4.4 \times 10^{-16}$. The thirteen link constraints $n_L + n_R = 2$ and $N = 12$ are additive quantum numbers. Hopping becomes a product of two vertex-local bilinears without a Jordan–Wigner string, and the Hamiltonian is a matrix-product operator of bond dimension 10. The constrained sector dimensions equal the physical dimensions of Table II.1 exactly. The operator equals the certified physical-block Hamiltonian on 400 sampled product-state columns per mass (of 27,580 and 16,502) within $4.4 \times 10^{-16}$ (source) and zero (vacuum), and on three random matrix-product states per mass to a relative error of at most $3.7 \times 10^{-16}$; lifting a state to the physical block and back reproduces it within $1.2 \times 10^{-14}$.

DMRG [41] starts in each sector from a random matrix-product state of bond dimension 16 (a random circuit applied to the dominant basis configuration) and runs 40 sweeps, raising the bond dimension stepwise to $\chi_{\max} = 768$ (1024 for route B). Noise decreases from $10^{-5}$ and is zero from the ninth sweep on; the cutoff is $10^{-15}$ and the eigensolver tolerance $10^{-14}$. Four acceptance criteria, fixed before any production data, compare each state with the first-principles exact diagonalization (ED): SG1 $|E - E_0| \leq 10^{-10}$; SG2 the named energy parts $H_E, H_m, H_{\text{hop}}$, every $\langle E_\ell^2 \rangle$, $\langle n_v \rangle$ and class probability within $10^{-8}$; SG3 $1 - F \leq 10^{-9}$; SG4 every derived quantity within $10^{-8}$. All twelve states pass (Table II.2), with energy variance at most $1.9 \times 10^{-9}$. The derived quantities of Sec. III.1.2 agree within $7.9 \times 10^{-11}$ for every $V_D(R)$, $4.1 \times 10^{-11}$ for every secant $\sigma_{\text{eff}}^D$, $7.9 \times 10^{-10}$ for every electric,

mass and hopping budget entry, $5.7 \times 10^{-11}$ for every $\tau^{\mathrm{resp}}_{E,\mathrm{GS}}$ and $1.3 \times 10^{-10}$ for every $P_{\mathrm{any}}$; the signed differences are shown in Fig. A.1 of Appendix A.

| | | $m = 4$ | | | | $m = 0.5$ | | | |
|---|---|---|---|---|---|---|---|---|---|
| Sector | Dim. | $\lvert E - E_0\rvert$ | $1 - F$ | SG2 max | $\chi$ | $\lvert E - E_0\rvert$ | $1 - F$ | SG2 max | $\chi$ |
| Vacuum | 16 502 | $4.5\times10^{-12}$ | $8.7\times10^{-14}$ | $6.1\times10^{-11}$ | 151 | $1.7\times10^{-13}$ | $3.8\times10^{-15}$ | $6.1\times10^{-13}$ | 311 |
| $R = 2$ | 21 972 | $3.1\times10^{-12}$ | $5.8\times10^{-14}$ | $4.0\times10^{-11}$ | 154 | $4.7\times10^{-14}$ | $1.8\times10^{-15}$ | $7.6\times10^{-13}$ | 364 |
| $R = 4$ | 27 580 | $2.3\times10^{-12}$ | $4.5\times10^{-14}$ | $3.3\times10^{-11}$ | 158 | $7.2\times10^{-14}$ | $2.6\times10^{-15}$ | $3.9\times10^{-13}$ | 409 |
| $R = 6$ primary | 33 215 | $4.6\times10^{-13}$ | $1.0\times10^{-14}$ | $1.1\times10^{-11}$ | 183 | $5.0\times10^{-14}$ | $2.7\times10^{-15}$ | $5.4\times10^{-12}$ | 527 |
| $R = 6$ route B | 46 829 | $1.4\times10^{-12}$ | $3.3\times10^{-14}$ | $9.5\times10^{-11}$ | 270 | $1.3\times10^{-13}$ | $2.7\times10^{-15}$ | $8.4\times10^{-13}$ | 675 |
| $R = 6$ route C | 32 637 | $8.3\times10^{-11}$ | $3.8\times10^{-12}$ | $1.7\times10^{-9}$ | 286 | $6.8\times10^{-14}$ | $2.7\times10^{-15}$ | $1.6\times10^{-12}$ | 514 |

**Table II.2:** Gauss-exact DMRG from random initial states compared with the first-principles exact ground states of all six registered sectors at both masses. Energies are in units of $t_{\mathrm{h}}$; $F$ is the squared ground-state overlap and SG2 max the largest deviation among the named energy parts, local observables and class probabilities. $\chi$ is the bond dimension reached. All twelve states meet the acceptance criteria SG1–SG4.

An earlier engine without symmetry quantum numbers, which imposed the Gauss law and the $D_v$ labels only through the penalties $H_G$ and $H_D$, reproduced the high-mass states ($1 - F = 5.8 \times 10^{-10}$ and $1.5 \times 10^{-10}$ for source and vacuum at $R = 4$) but not the low-mass source states; a refined low-mass vacuum passed individually, but no low-mass pair was certified, and its two-site time evolution violated the Gauss constraints (Sec. IV.1.4). For C2 it is superseded by the Gauss-exact construction and enters none of the C2 results. It is the only engine applied to the wider C4 and C8 patches (Sec. III.1.4), where its C2 record limits what is claimed. With the gauge constraint imposed in the basis, and with larger bond dimensions ($\chi_{\max} = 768$–$1024$ against the earlier cap 400), DMRG reaches the low-mass ground states that the penalty-based engine missed; the two changes are not separated. Constraint leakage is removed by construction.

**Gauss-penalty independence at the examined points.** Writing $H(\lambda) = H_0 + \lambda K_G$, the verified minimal block has $K_{\min} = 3/2$ in the source sector and zero in vacuum. If the ground state at $\lambda = 20$ lies in this block, then

$$H(\lambda') - E_0(\lambda') = H(20) - E_0(20) + (\lambda' - 20)(K_G - K_{\min}) \succeq 0 \quad (\lambda' \geq 20). \tag{II.22}$$

The same vector remains a ground state; only the known constant changes. This statement applies to the examined parameter points and sectors. Separate Lanczos solutions at $\lambda = 40$ for both endpoint masses and both sectors reproduce the expected source shift $+30$ and vacuum shift zero within $2.8 \times 10^{-13}$; profile changes are below $9 \times 10^{-14}$. Across the 36 full-colour main and plaquette-control states, $\|(H_G - h_{\min})\psi\| \lesssim 2.734 \times 10^{-13}$. The imbalance penalty vanishes identically in the fixed-$D$ basis. The colour-diagonal control does not share the minimal Gauss block and is therefore excluded from this argument.

### II.5.3 Physical-Sector Thermal Representation

The complete C2 source and vacuum spaces have dimensions 27,580 and 16,502, with fixed $N, D_v, G_v^3$, full local Gauss representations and no energy cutoff. Canonical traces use

the penalty-free $H_0$. Sparse pure-state thermal steps are treated separately from exact GKSL identities and from ensemble convergence; Chapter V and Appendix B give the corresponding controls.

### II.5.4 Reproducibility

The physical meaning of a stored state depends on its basis order, geometry, local matrices, fermionic convention, sector labels, parameters and measurement Hamiltonian. The static and dynamic C2 figures are generated from certified records and checked against independently reconstructed operators. Hashes bind figure inputs to the archived campaign data; they identify files but do not replace a physical validation argument.

# Chapter III

# Static Response and SU(2) Structure

The static analysis first resolves source response on the C2 patch, compares it with variational states on the wider C4 and C8 patches, and then separates connectivity, matter dressing, energy cost and local colour recoupling. All conclusions are finite-sector statements.

## III.1 $(2+1)$D Source Sectors

The exact honeycomb static analysis is carried out on the two-cycle C2 patch of Fig. II.1. Two questions are separated. At fixed graph separation $R = 4$, a thirteen-mass ladder resolves how the character of the ground state changes within one conserved sector (Sec. III.1.1). A separation scan at $R = 2, 4, 6$ and the two endpoint masses then measures the static potential, the finite-distance force and its energy budget (Sec. III.1.2). All C2 values are exact sector results. The separation-scan data are independently re-derived from first principles, and Gauss-exact DMRG reproduces all twelve ground states of the scan, six sectors at two masses (Sec. II.5.2). A third subsection repeats the exact calculation with the electric and plaquette couplings varied one at a time (Sec. III.1.3). A fourth compares the response at $m \geq 2$ with variational states on the wider C4 and C8 patches (Sec. III.1.4); these carry their own numerical status and do not enter the C2 statements.

The interpretive labels follow fixed criteria. *Confinement-like* character denotes flux-supported source connectivity together with a positive, approximately constant separation cost over the sampled window. *Deconfinement-like* (screening-like) character denotes weakening of the connected flux support and of its energetic cost. Both labels characterize responses of this finite patch, not thermodynamic phases. Static string breaking in QCD [5] provides the reference phenomenology, although its gauge theory and source prescription differ from the present model.

### III.1.1 Mass Dependence at Fixed Separation

Exact sector diagonalization resolves a mass-dependent reorganization of the C2 ground state at graph separation four, from string-like electric support to a screening-like (deconfinement-like) configuration mixture. The conclusion follows jointly from all-path connectivity, source-local matter support, spatial responses and energy decomposition. It

is distinct from the smaller loop-bearing C1 pilot of Sec. III.2.

**Geometry and fixed-sector comparison.** The open $3 \times 2$ patch contains twelve matter vertices, thirteen links and two independent cycles. Vertices 1–6 belong to $A$, and 7–12 to $B$. The classical sources at vertices 1 and 3 have degrees 1 and 2; their cell-coordinate separation is $R_x = 2$, whereas the graph distance is four links. The selected path is $1 \to 7 \to 2 \to 8 \to 3$, with ordered links $[1, 3, 2, 5]$. All thirteen masses are evaluated at the same source or vacuum labels specified in Eq. (II.21), with $N = 12$, $t_\mathrm{h} = 1$, $g^2 = 2$, $\lambda = \mu = 20$, and $\kappa = 0$. The ground-state change is therefore not produced by choosing a different imbalance sector at each mass. Neither the source nor the vacuum ladder establishes a global minimum over all superselection sectors.

**Connectivity and local screening support.** The active-link projector is $a_\ell = 4E_\ell^2/3$, with eigenvalues zero and one. In each configuration, the active links form a subgraph. The projector expectation $P_\mathrm{any}$ is the probability that this subgraph connects the sources, evaluated by all $2^{13}$ activity masks, rather than by multiplying single-link means. The selected-corridor weight is $P_\mathrm{sel}$; $P_\mathrm{alt} = P_\mathrm{any} - P_\mathrm{sel}$ counts connected configurations without the complete selected corridor, and $P_\mathrm{no} = 1 - P_\mathrm{any}$. The projector $P_\mathrm{ff}(v)$ additionally requires exactly one matter particle at $v$ and no active incident link; $P_\mathrm{both}$ requires this at both sources simultaneously. These are configuration weights, not meson counts or pair-production rates.

| $m$ | $\Delta_\mathrm{src}$ | $\Delta_\mathrm{vac}$ | $P_\mathrm{any}$ | $P_\mathrm{ff}(1)$ | $P_\mathrm{ff}(3)$ | $P_\mathrm{both}$ | $\Delta H_0$ |
|---|---|---|---|---|---|---|---|
| 0.25 | 0.8717 | 2.2481 | 0.0145 | 0.7512 | 0.2572 | 0.184 | 2.5753 |
| 0.5 | 0.7567 | 2.6821 | 0.0289 | 0.7244 | 0.2376 | 0.160 | 3.0817 |
| 0.75 | 0.6518 | 3.1295 | 0.0614 | 0.6712 | 0.2136 | 0.130 | 3.5489 |
| 1 | 0.5734 | 3.5835 | 0.1434 | 0.5487 | 0.1877 | 0.0926 | 3.9547 |
| 1.125 | 0.5665 | 3.8121 | 0.2187 | 0.4514 | 0.1719 | 0.0710 | 4.1203 |
| 1.25 | 0.6014 | 4.0416 | 0.3150 | 0.3454 | 0.1525 | 0.0501 | 4.2510 |
| 1.375 | 0.6863 | 4.2718 | 0.4158 | 0.2543 | 0.1311 | 0.0333 | 4.3447 |
| 1.5 | 0.8165 | 4.5027 | 0.5055 | 0.1889 | 0.1109 | 0.0218 | 4.4056 |
| 1.75 | 1.1716 | 4.9664 | 0.6383 | 0.1160 | 0.0802 | 0.0101 | 4.4590 |
| 2 | 1.5982 | 5.4325 | 0.7244 | 0.0808 | 0.0608 | 0.00528 | 4.4595 |
| 2.5 | 2.5417 | 6.3712 | 0.8242 | 0.0478 | 0.0392 | 0.00196 | 4.3941 |
| 3 | 3.5318 | 7.3177 | 0.8780 | 0.0322 | 0.0277 | 0.000914 | 4.3004 |
| 4 | 5.5486 | 9.2306 | 0.9317 | 0.0176 | 0.0160 | 0.000284 | 4.1151 |

**Table III.1:** Exact source and vacuum ground states on C2, $t_\mathrm{h} = 1, g^2 = 2, \lambda = \mu = 20, \kappa = 0$, with source vertices 1 and 3 at graph distance four. $\Delta_\mathrm{src,vac} = (E_1 - E_0)_\mathrm{src,vac}$. $P_\mathrm{any}$ tests any active connection and $P_\mathrm{no} = 1 - P_\mathrm{any}$. $P_\mathrm{ff}(v)$ requires a flux-free vertex with one matter particle; $P_\mathrm{both}$ requires this at both sources. $\Delta H_0 = \langle H_0 \rangle_\mathrm{src} - \langle H_0 \rangle_\mathrm{vac}$. At degree-one source 1, $P_\mathrm{ff}(1) = 1 - 4\langle E_1^2 \rangle/3$. All 46 exact records, including Table III.3, were separately recomputed on regenerated bases using a shared Hamiltonian term list; ground-state residuals are at most $3.1 \times 10^{-13}$. Digits are rounding, not statistical errors.

At $m = 4$, $P_\mathrm{any} = 0.931665093994$, $P_\mathrm{sel} = 0.931665086296$, and $P_\mathrm{alt} = 7.69836 \times 10^{-9}$.

The mean source-subtracted corridor Casimir response is 0.706560160443, close to the active-link value 3/4, while $P_{\mathrm{both}} = 0.000284349301$. These observables identify a string-like ground state. At $m = 0.5$, connectivity falls to 0.028903779213, with $P_{\mathrm{alt}} = 1.69628 \times 10^{-5}$, while $P_{\mathrm{ff}}(1) = 0.724369043404$, $P_{\mathrm{ff}}(3) = 0.237630994687$, and $P_{\mathrm{both}} = 0.159918944625$. The loss of connectivity is therefore not compensated by rerouting within this sector. However, the joint weight near 0.16 does not imply complete screening in every configuration. At the degree-one source, $P_{\mathrm{ff}}(1) = 1 - 4\langle E_1^2\rangle/3$, verified within approximately $10^{-13}$; these two quantities are not independent witnesses.

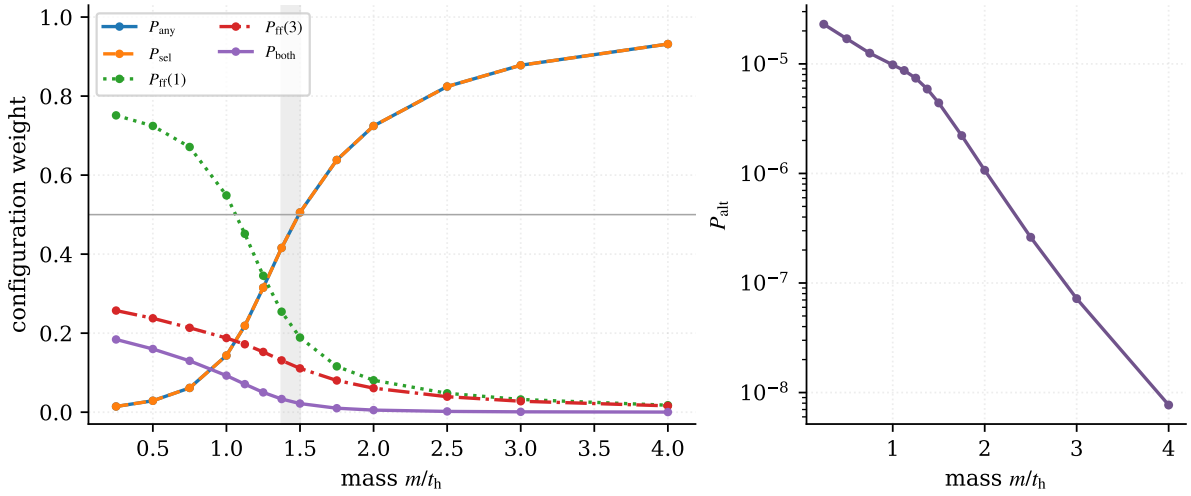



**Figure III.1:** The thirteen-mass C2 source ladder. The all-path and selected-corridor weights nearly coincide; the alternative-path contribution remains small. Source-local flux-free matter support grows as connectivity decreases. The shaded half-weight bracket is a discrete-grid diagnostic, not a critical region or statistical confidence interval. Lines connect sampled values.

**Spatial response and particle–hole dressing.** The electric and matter responses are $\rho_\ell = \langle E_\ell^2\rangle_{\mathrm{src}} - \langle E_\ell^2\rangle_{\mathrm{vac}}$ and $\Delta n_v = \langle n_v\rangle_{\mathrm{src}} - \langle n_v\rangle_{\mathrm{vac}}$. They subtract a fluctuating interacting vacuum, not an empty matter lattice. The staggered reference is $A$-empty and $B$-doubly occupied. At $N = 12$, $N_{\mathrm{pair}} = N_A = \sum_{v\in A} n_v$ and $M = \sum_v \eta_v n_v = 2N_A - 12$. Thus hopping can create particle–hole content relative to that reference without changing total matter number.

At high mass the source and vacuum pair expectations are 0.448023815640 and 0.533404195646; their small negative difference −0.08538038 is reduced dressing, not a negative particle number. At low mass these become 3.830187820885 and 2.851730841905, giving an excess $\Delta\langle N_A\rangle = 0.978456978981$. The source responses $\Delta n_1, \Delta n_3$ change from $(-0.0254124, -0.0222025)$ to $(0.4288789, 0.1923399)$. Positive response also appears at the intervening vertex 2, and negative response at neighbouring $B$ vertices 7–9. The sum over all twelve vertices vanishes, as required by equal total $N$.

Figure III.3 places the same signed responses on the physical graph of Fig. II.1. At high mass the electric response is concentrated on the connecting corridor. At low mass its suppression accompanies a distributed matter rearrangement around the sources and

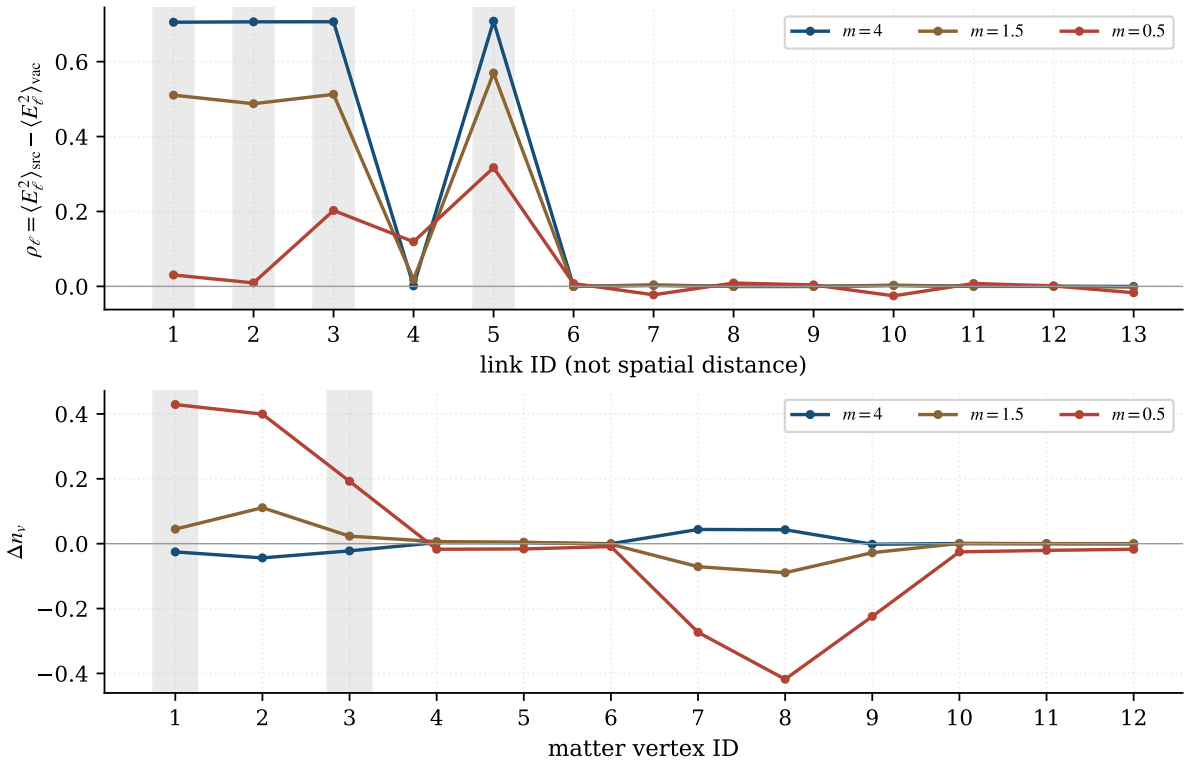



**Figure III.2:** All thirteen link responses and twelve matter responses at three masses. Shading identifies corridor links or source vertices; the link index is not a spatial-distance coordinate. The low-mass state retains short-link electric support and a spatially distributed matter response. The profiles are differences between specified source and vacuum sectors, not colour-resolved field vectors.

neighbouring vertices. Together with the all-path projector, these maps distinguish loss of a connected string from simple rerouting. The colour of a link represents a Casimir difference, not a directed electric-field component; small negative responses are retained.

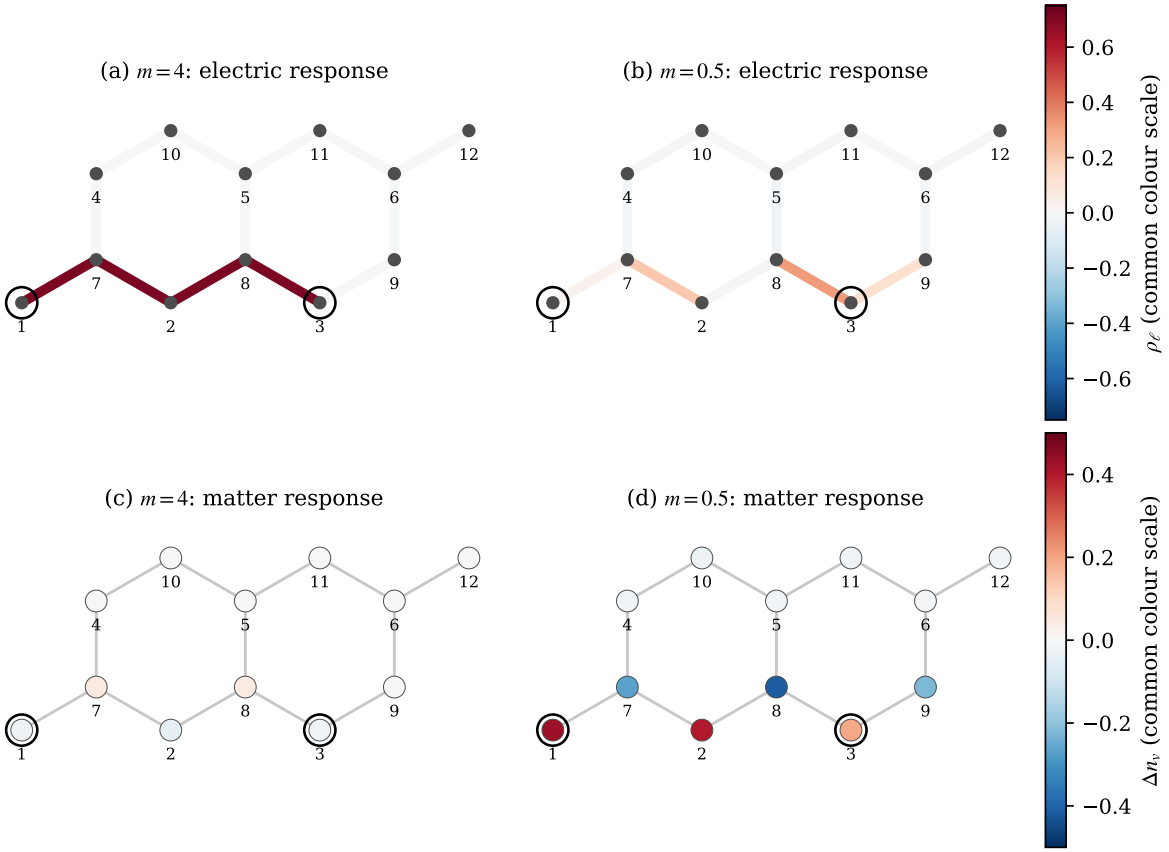



**Figure III.3:** Exact source–vacuum responses at $m = 4$ and $m = 0.5$, in the actual honeycomb embedding. Top: $\rho_\ell$, with one common signed colour scale for both masses. Bottom: $\Delta n_v$, with a separate common signed scale. Black rings mark the imposed source positions; labels are vertex IDs. Identical scales expose the electric suppression without rescaling the low-mass signal. Matter response sums to zero at fixed $N = 12$. These finite-sector maps support the static mechanism; they do not measure a transverse string width or a time-dependent breaking process.

The signed off-corridor fraction is

$$\epsilon_{\text{off}} = 1 - \frac{\sum_{\ell \in \mathcal{C}} \rho_\ell}{\sum_\ell \rho_\ell}, \qquad \sum_\ell \rho_\ell > 0. \tag{III.1}$$

It is neither a transverse width nor a configuration probability; negative link responses retain their signs. The endpoint values are 0.002007684118 and 0.130037321, while the mean corridor response falls from 0.706560160443 to 0.139871624851. At high mass, the largest off-corridor response is 0.002469411606, exceeding the predetermined narrow-tube criterion 0.001. The supported description is therefore a corridor-concentrated string-like response with neighbouring-link support, not a quantitatively resolved narrow tube. At

low mass $\rho_4 = 0.119097348281$ remains, even though this link is outside the chosen corridor. The total source electric energy is 2.423588876213, compared with vacuum 1.780473358051: loss of source-to-source connectivity is not disappearance of all electric flux. Full numerical profiles are given in Appendix A.

**Energy competition.** The physical source energy is

$$V(m) = \langle H_0 \rangle_{\mathrm{src}} - \langle H_0 \rangle_{\mathrm{vac}} = m\Delta\langle M \rangle + \Delta\langle H_E \rangle + \Delta\langle H_{\mathrm{hop}} \rangle \tag{III.2}$$

for the main $\kappa = 0$ ladder. In the verified minimal-Casimir blocks, $V = E_{0,\mathrm{src}}(H_{\mathrm{opt}}) - E_{0,\mathrm{vac}}(H_{\mathrm{opt}}) - 30$. The source penalty cannot be retained in a physical source–vacuum energy. The endpoint budgets are

| $m$ | $m\Delta\langle M \rangle$ | $\Delta\langle H_E \rangle$ | $\Delta\langle H_{\mathrm{hop}} \rangle$ | $V$ |
|---|---|---|---|---|
| 4 | −0.68304304 | 2.83192626 | 1.96625573 | 4.11513894 |
| 0.5 | 0.97845698 | 0.64311552 | 1.46009192 | 3.08166442 |

The electric cost decreases as matter dressing increases. The positive hopping difference is also essential: sources suppress part of the vacuum's kinetic-energy gain. At low mass the electric difference corresponds to 0.857487357549 active links in expectation, not an integer link count for each configuration. Energy components reconstruct $H_0$ within $5.03 \times 10^{-12}$; the mass term and density-derived staggered occupation agree within $9.31 \times 10^{-13}$. The role-wise data satisfy Hellmann–Feynman and concavity bounds,

$$\frac{dE_0}{dm} = \langle M \rangle, \qquad (b-a)\langle M \rangle_b \le E_0(b) - E_0(a) \le (b-a)\langle M \rangle_a.$$

The source–vacuum slopes are −0.170760760013 at $m = 4$ and 1.956913957961 at $m = 0.5$, close to the classical zero-pair and one-pair slopes without implying pure classical configurations.

**Classical configurations and quantum mixing.** At $t_{\mathrm{h}} = \kappa = 0$, exhaustive occupation counting gives the branches in Table III.2, relative to the common source reference $-12m + 30$. For positive mass the source–vacuum lower envelope is

$$V_{\mathrm{cl}}(m) = \min\{3,\ 2m + 3/4,\ 4m\},$$

with crossings at $m = 9/8$ and $3/8$.

The imbalance profile admits only one zero-pair occupation pattern: the selected corridor. Enumeration of all 150 rerouted patterns shows that a connected alternative needs at least four pairs and six active links. Consequently, corridor dominance reflects both sector kinematics and Hamiltonian energetics. It cannot be interpreted as unconstrained route selection by the electric energy alone.

At $t_{\mathrm{h}} = 1$, the smallest sampled source gap is 0.566476607192 at $m = 1.125$, consistent with the first classical competition scale. The half-connectivity bracket is instead $[1.375, 1.5]$, with endpoint weights 0.415795316338 and 0.505536916894. Linear interpolation gives approximately 1.49229, a descriptive value rather than a critical mass. The strongest

| Class | Patterns | Energy | Structure |
|---|---|---|---|
| String | 1 | 3 | Four corridor links, no pairs |
| One-pair | 3 | $2m + 3/4$ | One pair, one active link |
| Two-pair | 3 | $4m$ | Two pairs, no active links |
| Lowest rerouted | 2 | $8m + 9/2$ | Four pairs, six active links |

**Table III.2:** Lowest classical occupation classes in the fixed source sector at $g^2 = 2$, $t_h = \kappa = 0$. Pattern counts are occupation counts, not fusion-resolved Hilbert-space dimensions. The common source reference $-12m + 30$ has been subtracted.

sampled electric and pair-response variation is in $[1.125, 1.25]$, with electric finite-difference slope 2.131284610830; neighbouring intervals have comparable slopes. These distinct diagnostics describe gradual quantum reorganization, not a sharp onset or thermodynamic transition.

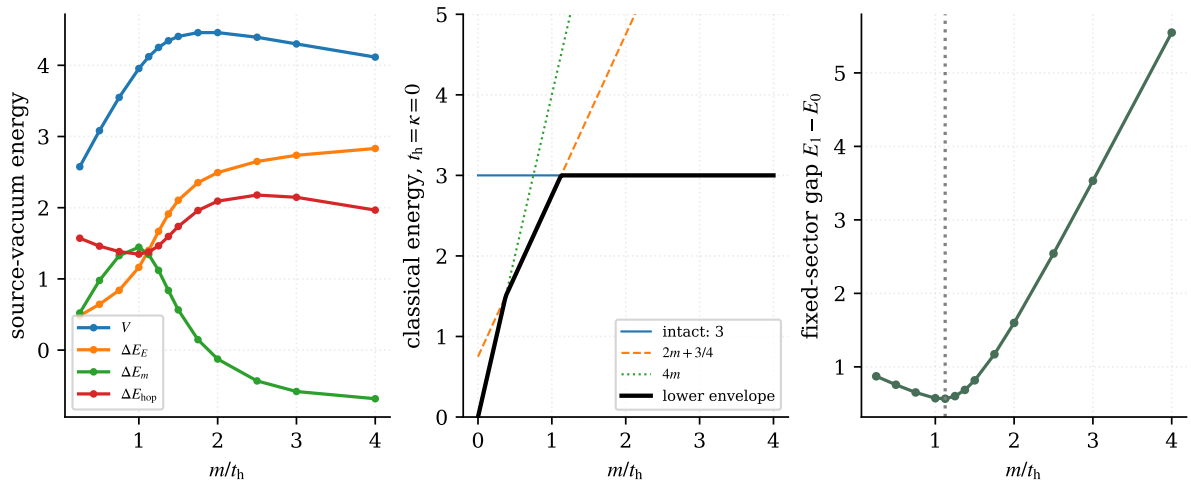



**Figure III.4:** Source–vacuum physical-energy components, the classical lower envelope, and the sampled source gap. Classical branch crossings and the half-connectivity bracket are distinct from the sampled gap minimum. The exact $t_h = 1$ values include kinetic dressing; the classical branches use $t_h = 0$. No critical-point fit is made.

A separate four-level calculation at seven masses further constrains the interpretation. At $m = 1.125$, $P_{\text{any}}(E_0, \ldots, E_3)$ is approximately $(0.2187, 0.1241, 0.0149, 0.1037)$: none of these levels is predominantly string-like. A simple two-level string/screened avoided-crossing picture is therefore not established. At $m = 4$, the three excited levels lie approximately 5.549–5.967 above the ground state and have $P_{\text{any}} \leq 0.002381$, consistent with the classical one-pair excitation scale $2m - 9/4 = 5.75$. Even at $m = 0.25$, where the classical two-pair branch is lowest, the quantum excess pair number is only about 1.042 and $P_{\text{both}} \simeq 0.184$; classical branch labels are not identities of the dressed ground states.

**Plaquette and colour-diagonal controls.** The five-mass controls in Table III.3 retain the same patch and sector definitions. Adding $\kappa = 0.5$ preserves both endpoint characters;

the largest sampled changes in $P_{\text{any}}$, $P_{\text{both}}$ and $V$ are approximately 0.00297244, 0.001226 and 0.02629. This is a two-coupling comparison, not a plaquette phase diagram. The colour-diagonal control retains only $a = b$ hopping. It conserves $G^3$, but changes the full SU(2) dynamics and occupies a nonminimal Gauss block: $\langle H_G\rangle_{\text{src}} - 30 = 0.104823$–$0.112884$ and $\langle H_G\rangle_{\text{vac}} = 0.129294$–$0.137549$. Its differences include finite-$\lambda$ effects as well as removed colour mixing. It is not an independently matched U(1) theory. Moreover, $E_\ell^2 = 3(1 - \mathcal{E}_\ell^2)/4$, where $\mathcal{E}_\ell = (n_{\ell,L} - n_{\ell,R})/2$; scalar occupation and connectivity alone do not establish a uniquely non-Abelian mechanism.

| | $P_{\text{any}}$ | | | $(E_1 - E_0)_{\text{src}}$ | | |
|---|---|---|---|---|---|---|
| $m$ | SU(2) | Cartan | $\kappa = 0.5$ | SU(2) | Cartan | $\kappa = 0.5$ |
| 0.5 | 0.0289 | 0.00584 | 0.0286 | 0.7567 | 0.2901 | 0.7012 |
| 1 | 0.1434 | 0.0832 | 0.1464 | 0.5734 | 0.1670 | 0.5477 |
| 1.5 | 0.5055 | 0.8322 | 0.5084 | 0.8165 | 0.7318 | 0.8144 |
| 2 | 0.7244 | 0.9152 | 0.7245 | 1.5982 | 1.7054 | 1.6005 |
| 4 | 0.9317 | 0.9813 | 0.9316 | 5.5486 | 5.7125 | 5.5493 |

| | $\Delta H_0$ | | |
|---|---|---|---|
| $m$ | SU(2) | Cartan | $\kappa = 0.5$ |
| 0.5 | 3.0817 | 2.3722 | 3.1080 |
| 1 | 3.9547 | 3.3352 | 3.9689 |
| 1.5 | 4.4056 | 3.5635 | 4.4079 |
| 2 | 4.4595 | 3.5239 | 4.4595 |
| 4 | 4.1151 | 3.3582 | 4.1151 |

**Table III.3:** Same-patch controls at the parameters of Table III.1. Cartan retains only colour-diagonal hopping: it conserves $G^3$, breaks $G^{1,2}$ and has non-minimal Gauss penalty, $\langle H_G\rangle_{\text{src}} - 30 = 0.1048$–$0.1129$ and $\langle H_G\rangle_{\text{vac}} = 0.1293$–$0.1375$. The $\kappa = 0.5$ column adds a plaquette term to the full-SU(2) Hamiltonian. The remaining columns use $\kappa = 0$.

**Source-relocation control.** Moving the sources to vertices $(4, 6)$ preserves graph distance four and uses the corridor $4 \to 10 \to 5 \to 11 \to 6$, with links $[6, 9, 8, 12]$. The source degrees become $(2, 3)$; the new placement is not bulk, since vertex 6 has a degree-one neighbour. Its imbalance sector is rederived from the same geometric balance prescription, yielding 6 513 874 kinematic states. Keeping the original imbalance labels would give an empty intersection with the new source targets.

The criteria fixed before data generation were $P_{\text{any}} > P_{\text{no}}$ at $m = 4$, $P_{\text{any}} < P_{\text{no}}$ at $m = 0.5$, and $P_{\text{alt}} < P_{\text{sel}}$ at both endpoints. Both placements satisfy them. The relocated connectivity is 0.931495103145 at high mass and 0.030225535698 at low mass. At low mass, $P_{\text{ff}}(4) = 0.598457337256$, $P_{\text{ff}}(6) = 0.229431925325$, and $P_{\text{both}} = 0.125389271021$; the high-mass mean corridor response is 0.706505447348. The qualitative endpoint change therefore survives removal of the degree-one source, while local screening weights remain placement-dependent.

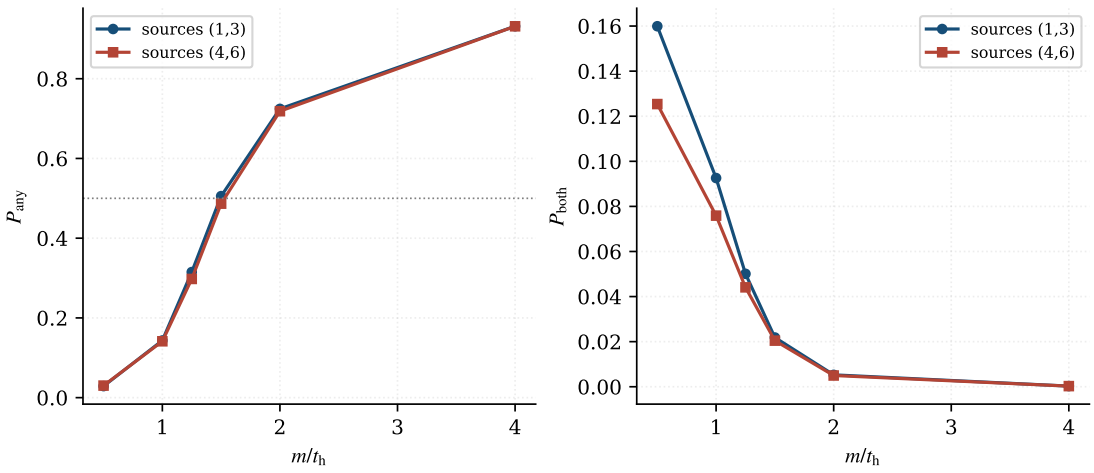



**Figure III.5:** Original and relocated sources at four-link graph separation. The six shared masses bracket half-connectivity by $[1.25, 1.5]$ and $[1.5, 2]$, respectively. These brackets have the same sampling resolution; the denser original thirteen-mass ladder is not used to imply equally fine resolution for the relocated sources, which the post-closure supplement described in the text resolves on the full grid. Lines are guides between measured masses.

The shared-grid interpolants 1.49273364 and 1.52892253 differ by 0.03619, but this is not a continuously resolved shift with statistical error bars. The four intermediate relocated masses were independently reconstructed and solved; 388 numerical checks give worst recorded residual, energy and profile differences of $2.354 \times 10^{-13}$, $1.283 \times 10^{-13}$, and $1.439 \times 10^{-13}$. The control supports persistence of the endpoint characters on this finite patch, not general boundary independence: source position and coordination change together. A post-closure supplement solves the relocated sector with the unchanged solver at the seven remaining masses of the thirteen-mass grid; its recomputed $m = 2$ anchor reproduces the closure record exactly. On the full grid the relocated half-connectivity bracket is $[1.5, 1.75]$, with linear interpolant 1.524, against $[1.375, 1.5]$ and 1.492 for the original sources. The relocated ladder is the exact comparator of Sec. III.1.4.

### III.1.2 Static Potential and Finite-Distance Force

**Separation scan.** The separation scan places the sources at vertices $(1, 2)$, $(1, 3)$ and $(1, 6)$, at graph separations $R = 2, 4, 6$, in the registered sector family $D_R$ of Sec. II.5.2, and evaluates the two endpoint masses $m = 4$ and $m = 0.5$. Each source record is compared with the vacuum of the same mass through the penalty-free potential and finite-distance force of Eq. (II.17). This is a comparison along one registered sector family, not a minimization over all superselection sectors. Table III.4 and Fig. III.6 collect the ground-state data, and Table III.5 gives the secants with their energy budget.

**Heavy mass: confinement-like separation cost.** At $m = 4$ the potential rises in nearly equal steps of graph distance, $V_D = 2.063$, 4.115 and 6.177 for $R = 2, 4, 6$. The finite-distance force is 1.026100 and $1.031033\, t_{\mathrm{h}}/a$ per graph link on the two intervals; the

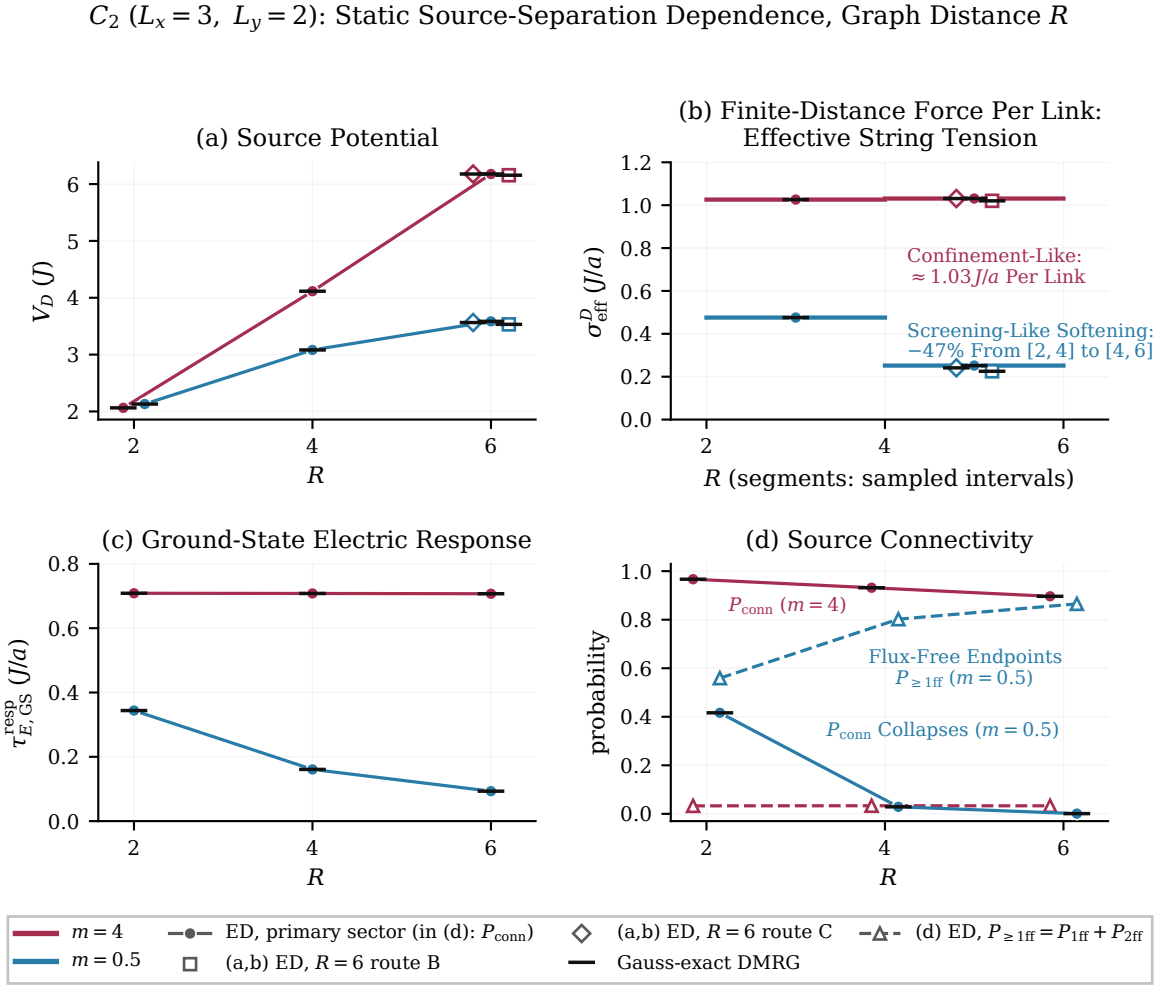



**Figure III.6:** Static separation dependence at $m = 4$ (crimson) and $m = 0.5$ (blue), from exact diagonalization of each registered sector; $J \equiv t_{\mathrm{h}}$ and $P_{\mathrm{conn}} \equiv P_{\mathrm{any}}$. (a) Source potential $V_D$. (b) Finite-distance force per graph link $\sigma^D_{\mathrm{eff}}$, titled effective string tension; that name applies at $m = 4$, where $V_D$ rises linearly. Segments mark the sampled intervals $[2, 4]$ and $[4, 6]$, with symbols at the midpoints. (c) Ground-state electric response. (d) Connection probability (filled circles, solid) and flux-free endpoint support $P_{\geq 1\mathrm{ff}}$ (open triangles, dashed). In (a,b) open squares and diamonds are the alternative $R = 6$ routes B and C, offset horizontally by $+0.2$ and $-0.2$; in (a) the two masses at $R = 2$ are offset by $\mp 0.12$, and in (d) all points by $\mp 0.15$. Black bars are the independent Gauss-exact DMRG values; they agree with the exact values within $1.3 \times 10^{-10}$ in the respective units.

| | $m = 4$ | | | | $m = 0.5$ | | | |
|---|---|---|---|---|---|---|---|---|
| $R$ | $V_D$ | $\tau^{\mathrm{resp}}_{E,\mathrm{GS}}$ | $P_{\mathrm{any}}$ | $P_{\geq 1\mathrm{ff}}$ | $V_D$ | $\tau^{\mathrm{resp}}_{E,\mathrm{GS}}$ | $P_{\mathrm{any}}$ | $P_{\geq 1\mathrm{ff}}$ |
| 2 | 2.063 | 0.709 | 0.967 | 0.033 | 2.130 | 0.344 | 0.416 | 0.559 |
| 4 | 4.115 | 0.708 | 0.932 | 0.033 | 3.082 | 0.161 | 0.029 | 0.802 |
| 6 | 6.177 | 0.707 | 0.896 | 0.033 | 3.585 | 0.093 | 0.0009 | 0.865 |

**Table III.4:** Exact ground-state data of the primary sectors along the registered separation family. $V_D$ is in units of $t_{\mathrm{h}}$ and $\tau^{\mathrm{resp}}_{E,\mathrm{GS}}$ in $t_{\mathrm{h}}/a$; probabilities are dimensionless. The $R = 4$ row is the endpoint of the fixed-separation ladder of Sec. III.1.1.

0.48% change is not significant against the sector systematic. The three registered $R = 6$ sectors give $V_D(6; 4) = 6.177204$, 6.177146 and 6.155685 for the primary sector and routes C and B, so that $\sigma^D_{\mathrm{eff}}[4, 6]$ spans 1.0203–1.0310 across them. The effective string tension of the patch is therefore

$$\sigma^D_{\mathrm{eff}}(m = 4) = 1.03 \pm 0.01\ t_{\mathrm{h}}/a \text{ per graph link}, \tag{III.3}$$

approximately constant over $2 \leq R \leq 6$; the uncertainty is the envelope of the sector spread, not a statistical error. About 69% of the cost is electric ($\sigma^D_E = 0.707366$ and 0.705375), the ground-state electric response is nearly independent of $R$ ($\tau^{\mathrm{resp}}_{E,\mathrm{GS}} = 0.707$–0.709), and 90–97% of the source weight is connected, with flux-free endpoint support near 0.033. Together these observables form the confinement-like signature of the patch: a connected electric flux string whose energy grows approximately linearly with graph separation.

Second-order strong-coupling perturbation theory in the hopping, a heavy-mass expansion, gives without free parameter the cost of a source-connecting link,

$$\sigma_{\mathrm{sc}}(m) = \frac{1}{a}\Big[\frac{3g^2}{8} + \frac{4t_{\mathrm{h}}^2}{2m + 3g^2/8} - \frac{t_{\mathrm{h}}^2}{2m - 3g^2/8}\Big]. \tag{III.4}$$

The first term is the bare $j = 1/2$ electric energy of an active link. In the heavy-mass vacuum no link carries flux, and virtual pair creation across each link, with intermediate energy $2m + 3g^2/8$ and total weight $4t_{\mathrm{h}}^2$, lowers the vacuum energy; a string link forgoes this fluctuation and instead admits a virtual process that removes its flux, with intermediate energy $2m - 3g^2/8$ and weight $t_{\mathrm{h}}^2$. At $m = 4$ the estimate is $1.069\, t_{\mathrm{h}}/a$ (electric 0.697, mass $-0.266$, hopping $+0.638$), compared with the measured 1.026–1.031 (electric 0.705–0.707, mass $-0.163$ to $-0.171$, hopping $+0.482$ to $+0.497$). The predicted correction $+0.32$ to the bare link energy $3g^2/8 = 0.75$ compares with the measured $+0.28$, so that the totals agree to about 4%; the remaining difference is attributed to higher orders in the hopping. The $[4, 6]$ interval also contains an endpoint-degree change that this per-link estimate does not describe.

**Light mass: screening-like softening.** At $m = 0.5$ the finite-distance force falls from 0.475783 to $0.251805\, t_{\mathrm{h}}/a$, by 47%, and by 47–53% across the three $R = 6$ sectors ($\sigma^D_{\mathrm{eff}}[4, 6] = 0.2252$–0.2518). The electric contribution is negative on both intervals,

| $m$ | Interval | $\sigma^D_{\text{eff}}$ | $\sigma^D_E$ | $\sigma^D_m$ | $\sigma^D_{\text{hop}}$ |
|---|---|---|---|---|---|
| 4 | $[2, 4]$ | 1.026100 | 0.707366 | −0.162969 | 0.481702 |
| 4 | $[4, 6]$ | 1.031033 | 0.705375 | −0.171438 | 0.497095 |
| 0.5 | $[2, 4]$ | 0.475783 | −0.022312 | 0.355432 | 0.142664 |
| 0.5 | $[4, 6]$ | 0.251805 | −0.042268 | 0.083543 | 0.210530 |

**Table III.5:** Finite-distance force per graph link (at $m = 4$ the effective string tension of the patch) and its electric, mass and hopping contributions, in units of $t_{\text{h}}/a$, for the primary sectors. Across the three $R = 6$ sectors the total on $[4, 6]$ spans 1.0203–1.0310 at $m = 4$ and 0.2252–0.2518 at $m = 0.5$. The hopping contribution was obtained from the energy balance and is reproduced to the printed digits by a named contraction.

$\sigma^D_E = -0.022312$ and $-0.042268$: separating the sources lowers the source–vacuum electric energy, while mass and hopping contributions keep the total cost positive. Connected flux collapses with distance, $P_{\text{any}} = 0.416$, 0.029 and 0.0009, while flux-free endpoint support rises to 0.559, 0.802 and 0.865, and the electric response falls to $0.093\, t_{\text{h}}/a$ at $R = 6$. For $R \geq 4$ connected-flux support is strongly suppressed and the sources are accompanied by dynamical matter: a deconfinement-like, screening-like static response.

The remaining positive cost does not indicate a residual string tension. Besides the two SU(2) sources, the registered $D_R$ sectors fix the local U(1) labels $D_v$ along the source corridor, and these labels are conserved by $H$. While the connected flux disappears, alternate corridor links remain predominantly in the reversed-rishon configuration, with rishon number 1.72–1.91 at the $A$ end against 0.16–0.31 in the vacuum, on which the vacuum pair fluctuation is kinematically excluded. The continued rise of $V_D$ after the connected flux has disappeared is attributed to these conserved labels; a comparison with other corridor-label families, which would test this attribution, has not been carried out. A plateau of the potential or a string-breaking distance cannot appear within this family and is not claimed; the light-mass potential is conditional on the registered corridor sectors. With dynamical matter a screened branch can become the ground state at large separation [5], so its absence here is a property of the sector family rather than evidence against string breaking.

**Energy budget and electric response.** At heavy mass the electric contribution stays near $0.706\, t_{\text{h}}/a$ and the total near $1.03\, t_{\text{h}}/a$: separating the sources maintains a flux-supported energy cost. At light mass the electric excess decreases with separation, while mass and hopping contributions retain a positive total. The negative heavy-mass mass contribution is the slope of a source–vacuum energy difference, not a negative bare mass, and the hopping contribution is the kinetic energy of the quantum state, not a separately measured classical flow of electric energy. The ground-state electric response is neither the electric derivative $\sigma^D_E$ nor the full force: at $m = 0.5$ both electric slopes are negative although both total forces are positive. At $R = 4$, lowering the mass from 4 to 0.5 reduces $\tau^{\text{resp}}_{E,\text{GS}}$ from 0.707982 to 0.160779 and $P_{\text{any}}$ from 0.931665 to 0.028904. Connectivity is a configuration probability, whereas the force is an energy cost between different separations; their concurrence supports string-like versus screening-like finite-patch character without

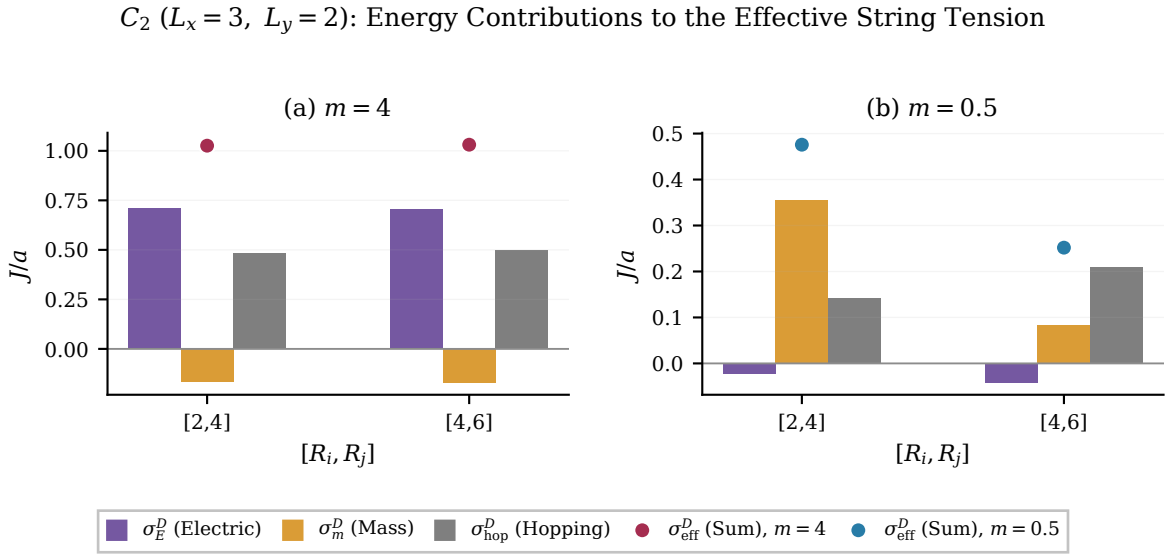



**Figure III.7:** Energy budget of the finite-distance force (titled effective string tension, the heavy-mass name), $\sigma_{\mathrm{eff}}^D = \sigma_E^D + \sigma_m^D + \sigma_{\mathrm{hop}}^D$ (bars; $J \equiv t_{\mathrm{h}}$). Circles in the mass colours are the sums. (a) At $m = 4$ the electric part stays near $0.706\, t_{\mathrm{h}}/a$ on both intervals: a flux-supported, confinement-like cost. (b) At $m = 0.5$ the electric part is negative, while mass and hopping carry the positive total within the registered corridor sector. The $[4, 6]$ bars use the primary $R = 6$ sector; across the three $R = 6$ sectors every component keeps its sign, while the light-mass split varies ($\sigma_E^D$ from $-0.076$ to $-0.042$, $\sigma_m^D$ from 0.018 to 0.084, $\sigma_{\mathrm{hop}}^D$ from 0.211 to 0.284).

establishing singlet screening.

**Geometry of the scan.** At $R = 2$ and 4 the source vertices have degrees $(1, 2)$; at $R = 6$ they have degrees $(1, 3)$, and the source direction rotates by 19.1°. The second interval therefore combines separation and endpoint-geometry effects. Since $R$ counts graph links, $\sigma_{\mathrm{eff}}^D$ is a force per link. A normalization per Euclidean length is not constant: at $m = 4$ it gives 1.185 and $1.844\, t_{\mathrm{h}}$ per bond length for the two intervals. At $R = 2$ and 4 a single shortest path joins the sources, and the heavy-mass geodesic control equals the all-link response within $0.002\, t_{\mathrm{h}}/a$. At $R = 6$ three shortest paths join them; the heavy-mass electric response occupies one of them, whose geodesic weights sum to 5/9, so the path average falls to $0.393 \approx (5/9) \times 0.707$. This reflects path multiplicity, not electric weakening. Effective-string analyses of $(2 + 1)$D pure SU(2) Yang–Mills theory show how finite-distance corrections enter such secants [9]; that theory supplies methodological context, not a numerical string tension for the present quantum-link model. An asymptotic string tension requires additional separations and control of source coordination and orientation; three C2 separations establish a finite-distance force along the registered family.

**Numerical reliability.** All separation-scan values of this subsection are exact sector results, re-derived by the independent first-principles implementation of Sec. II.5.2, which reproduces every potential, secant, budget entry, electric response and connec-

tion probability to the printed digits; recorded and re-derived secants differ by at most $1.6 \times 10^{-12}\, t_{\mathrm{h}}/a$. Gauss-exact DMRG, started from randomized states, reproduces every $V_D$ within $7.9 \times 10^{-11}$, every secant within $4.1 \times 10^{-11}\, t_{\mathrm{h}}/a$ (Table III.6) and every connection probability within $1.3 \times 10^{-10}$.

| Interval | $\sigma^D_{\mathrm{ED}}(m = 4)$ | DMRG−ED | $\sigma^D_{\mathrm{ED}}(m = 0.5)$ | DMRG−ED |
|---|---|---|---|---|
| $[2, 4]$ | 1.0260995059 | $-3.1 \times 10^{-13}$ | 0.4757834773 | $-7.1 \times 10^{-15}$ |
| $[4, 6]$ primary | 1.0310327537 | $-9.8 \times 10^{-13}$ | 0.2518049050 | $+1.1 \times 10^{-14}$ |
| $[4, 6]$ route B | 1.0202731346 | $-5.2 \times 10^{-13}$ | 0.2251810054 | $-1.8 \times 10^{-15}$ |
| $[4, 6]$ route C | 1.0310037752 | $+4.0 \times 10^{-11}$ | 0.2411846125 | $+5.3 \times 10^{-15}$ |

**Table III.6:** First-principles exact finite-distance secants per graph link and signed differences of the independently optimized Gauss-exact DMRG, in units of $t_{\mathrm{h}}/a$.

**Scope of the result.** The fixed-separation ladder establishes a string-like to screened-like reorganization of the specified C2 ground state. The separation scan adds an approximately constant heavy-mass separation cost of $1.03 \pm 0.01\, t_{\mathrm{h}}/a$ per graph link, in agreement with the strong-coupling estimate, and a screening-like light-mass softening with negative electric slope. Numerical residuals describe solver accuracy, not uncertainty from finite volume, sector choice or discrete sampling. No asymptotic string tension, light-mass potential independent of the registered corridor sectors, singlet screening, breaking distance or infinite-volume phase transition follows from these finite-patch data. Real-time breaking and formation are treated in Sec. IV.1. Section III.1.4 compares the heavy- and intermediate-mass response with variational states on the wider C4 and C8 patches; no C4 or C8 value enters the statements above.

**Thermal reference states.** The exact C2 endpoint states at $R = 4$ also supply the zero-temperature references for Sec. V.2. Thermal holds start from the exact vectors. The static connectivity contrast does not itself determine finite-temperature stability or a relaxation time.

**Honeycomb meson scope.** A path-resolved meson operator contracts the matter colours with an ordered product of the link transporters. Its path, orientation and fermionic registration must be held fixed in the source and reference measurements. Different paths on a loop are distinct operators. The all-path projector is an occupation observable, not such a coherent Wilson-line correlator. No honeycomb screening length is inferred from it, and no large-distance meson fit is inferred from the connectivity measurements. Length extraction requires a resolved correlator over several separations, a stable fit window and geometry convergence.

### III.1.3 Coupling Dependence on C2

**Scope and registration.** The C2 results above use $g^2 = 2$ and $\kappa = 0$. Two further exact scans vary one coupling at a time on the same patch, with the same registered source

and vacuum sectors, the $R = 2$ sector and the three $R = 6$ sectors of Sec. III.1.2. The grids, gates and expectations were fixed and sealed before the first point was computed; the plaquette-branch classification used for the $\kappa$ scan below was introduced after the data had been seen, for presentation only. Each point is an exact sector ground state obtained with the solver and gates of Sec. II.5.2; a point that fails a gate or has a degenerate ground level is excluded from every derived quantity. The Gauss penalty is raised with the couplings, $\lambda = 20\max(1, g^2/2, \kappa)$, so that the separation of Casimir blocks grows with the electric and magnetic energy scales; by Eq. (II.22) the physical results do not depend on $\lambda$ once the ground state lies in the minimal block, and two registered checks at $\lambda = 80$ test this premise. Six anchor points reproduce the earlier $g^2 = 2$ records, including the $\kappa = 0.5$ control, within $3.8 \times 10^{-12}$, and the two $\lambda = 80$ checks, at $(g^2, m) = (4, 3)$ and $(\kappa, m) = (2, 1.5)$, reproduce the corresponding $\lambda = 40$ points within $8.4 \times 10^{-14}$. Of the 306 new points, six fail one gate, the identity between two evaluations of the flux-free endpoint probability, with differences of $(1.0\text{–}2.1) \times 10^{-12}$ against a bar of $10^{-12}$; they are excluded, and none of them enters a registered outcome.

**Electric coupling.** The electric coupling takes the values $g^2 \in \{0.5, 1, 1.5, 3, 4, 6, 8\}$ at $\kappa = 0$, on the scaled-mass grid $\mu = 2m/g^2$ of the thirteen-mass ladder, so that $g^2 = 2$ reproduces the original masses. At $t_h = \kappa = 0$ the string, one-pair and two-pair branches of Table III.2 have energies $3g^2/2$, $2m + 3g^2/8$ and $4m$ relative to the common source reference, so the first classical crossing lies at $m_1^* = 9g^2/16$, that is at $\mu_1^* = 9/8$ for every coupling. Expectations fixed before the calculations were that the half-connectivity mass $m_{1/2}$ lies at or above $m_1^*$ for every $g^2 \geq 1$ with $(m_{1/2} - m_1^*)/g^2$ decreasing with $g^2$, and that the heavy-mass force at $\mu = 4$ approaches the strong-coupling estimate of Eq. (III.4) as $g^2$ grows.

| $g^2$ | $\lambda$ | $m_{1/2}$ bracket | $m_{1/2}$ | $m_1^*$ | $\sigma_{\text{eff}}^D[2,4]$ | $[4,6]$ routes | $\sigma_{\text{sc}}$ | $\sigma(\mu{=}\frac{1}{2})/\sigma(\mu{=}4)$ |
|---|---|---|---|---|---|---|---|---|
| 0.5 | 20 | $< \frac{1}{2}$ on grid | – | 0.28125 | 0.620 | 0.530–0.549 | – | 0.352 |
| 1 | 20 | [1, 1.25] | 1.212 | 0.5625 | 0.784 | 0.762–0.795 | 1.013 | 0.386 |
| 1.5 | 20 | [1.3125, 1.5] | 1.349 | 0.84375 | 0.898 | 0.887–0.907 | 0.988 | 0.434 |
| 2 | 20 | [1.375, 1.5] | 1.492 | 1.125 | 1.026 | 1.020–1.031 | 1.069 | 0.464 |
| 3 | 30 | [1.6875, 1.875] | 1.842 | 1.6875 | 1.324 | 1.322–1.325 | 1.338 | 0.472 |
| 4 | 40 | [2.25, 2.5] | 2.309 | 2.25 | 1.653 | 1.652–1.654 | 1.660 | 0.441 |
| 6 | 60 | [3, 3.375] | 3.343 | 3.375 | 2.354 | 2.354–2.355 | 2.356 | 0.359 |
| 8 | 80 | [4, 4.5] | 4.428 | 4.5 | 3.079 | 3.079–3.079 | 3.080 | 0.298 |

**Table III.7:** Exact C2 response at $\kappa = 0$ for each electric coupling $g^2$, on the scaled-mass grid $\mu = 2m/g^2$ of the thirteen-mass ladder. The half-connectivity bracket consists of the two sampled masses between which $P_{\text{any}} - \frac{1}{2}$ changes sign; $m_{1/2}$ is its linear interpolant, a descriptive grid value rather than a critical mass, and $m_1^* = 9g^2/16$ is the classical string/one-pair crossing. The finite-distance forces per graph link $(t_h/a)$ are taken at $\mu = 4$ $(m = 2g^2)$; the $[4,6]$ entry spans the three registered $R = 6$ routes and is a sector spread, not an error bar. $\sigma_{\text{sc}}$ is the second-order strong-coupling estimate, omitted for $g^2 < 1$. The last column is the ratio of the $[2,4]$ forces at $\mu = \frac{1}{2}$ and $\mu = 4$. $\lambda$ is the Gauss penalty used; physical values are $\lambda$-independent (Sec. II.5.2).

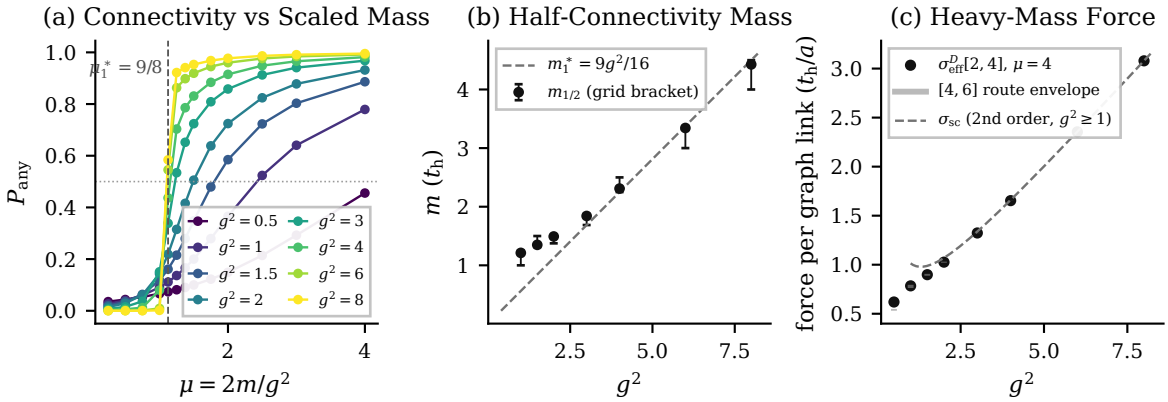



**Figure III.8:** Exact C2 response at $R = 4$ and $\kappa = 0$ for eight electric couplings. (a) Connection probability against the scaled mass $\mu = 2m/g^2$; the dashed line marks the classical crossing $\mu_1^* = 9/8$. (b) Half-connectivity mass, the linear interpolant within the bracketing grid masses (error bars span the bracket; no value exists for $g^2 = 0.5$), against $m_1^* = 9g^2/16$. (c) Finite-distance force per graph link at $\mu = 4$ with the envelope of the three $R = 6$ routes, against the second-order strong-coupling estimate, drawn for $g^2 \geq 1$.

The connectivity change moves with the coupling as the classical energetics suggests, but the registered inequality fails at strong coupling (Table III.7, Fig. III.8). The distance $(m_{1/2} - m_1^*)/g^2$ decreases monotonically, 0.650, 0.337, 0.184, 0.052, 0.015, $-0.005$ and $-0.009$ for $g^2 = 1$ to 8. At $g^2 = 6$ and 8, however, the connection probability already exceeds one half at the classical crossing, $P_{any}(\mu_1^*) = 0.545$ and 0.584, so the half-connectivity mass lies just below $m_1^*$ and the registered inequality $m_{1/2} \geq m_1^*$ fails; the expectation is not supported. The change sharpens with the coupling: at $g^2 = 8$ the connection probability rises from 0.002 at $\mu = 1$ to 0.922 at $\mu = 1.25$, whereas at $g^2 = 0.5$ it stays below one half on the whole grid (0.456 at $\mu = 4$). The mean spread of $P_{any}(\mu)$ across couplings is 0.074 for $g^2 \in \{4, 6, 8\}$ and 0.415 for all $g^2 \geq 1$; since the second set contains the first, this registered comparison is descriptive. An equal-size comparison made after the data were seen gives 0.179 for $g^2 \in \{1, 1.5, 2\}$ against 0.074 for $\{4, 6, 8\}$: the curves lie closer together at strong coupling but do not coincide.

The heavy-mass force approaches the strong-coupling estimate of Eq. (III.4): its relative deviation decreases monotonically, 0.040, 0.011, 0.0038, 0.0008 and 0.0003 for $g^2 = 2$, 3, 4, 6 and 8, as expected before the calculations. The $[4, 6]$ secants of the three $R = 6$ routes deviate from the $[2, 4]$ value by at most 0.089, 0.022, 0.011 and 0.006 for $g^2 = 0.5$, 1, 1.5 and 2, and by at most 0.0021 for $g^2 \geq 3$, so the heavy-mass cost per link becomes approximately distance independent across the sampled window as the coupling grows. The ratio of the light- and heavy-mass forces, $\sigma^D_{eff}[2, 4]$ at $\mu = \frac{1}{2}$ over $\mu = 4$, lies between 0.298 and 0.472 for all couplings. These are statements about the registered sectors of this patch; the connectivity change remains a smooth change on this finite patch in every case.

**Plaquette coupling.** The plaquette coupling takes the values $\kappa \in \{1, 2, 4\}$ at $g^2 = 2$, on the masses $m \in \{0.5, 1, 1.25, 1.375, 1.5, 2, 4\}$ for the $R = 4$ sectors and at $m \in \{4, 0.5\}$ for

the separation sectors, extending the $\kappa = 0.5$ control of Table III.3. The expectations fixed before the calculations were a quadratic small-$\kappa$ response, $(X(1)-X(0))/(X(0.5)-X(0)) \in [3,5]$ for the registered quantities, and a wording rule: the heavy-mass signature is called insensitive to the plaquette coupling up to the largest $\kappa$ for which the $[2,4]$ force at $m = 4$ changes by less than 5% and the $R = 4$ connection probability $P_{\text{any}}(4)$ by less than 0.01.

| | | | $m = 4$ | | | | | $m = \frac{1}{2}$ | | |
|---|---|---|---|---|---|---|---|---|---|---|
| $\kappa$ | $\lambda$ | $m_{1/2}$ bracket | Branch | $P_{\text{any}}$ | $\Delta P_{\text{any}}$ | $\sigma^D_{\text{eff}}[2,4]$ | Change | Branch | $P_{\text{any}}$ | $\sigma^D_{\text{eff}}[2,4]$ |
| 0 | 20 | [1.375, 1.5] | EEE | 0.9317 | 0 | 1.026 | 0 | EEE | 0.0289 | 0.476 |
| 0.5 | 20 | [1, 1.5] | – | 0.9316 | $-1.8 \times 10^{-5}$ | – | – | – | 0.0286 | – |
| 1 | 20 | [1.375, 1.5] | EEE | 0.9316 | $-4.6 \times 10^{-5}$ | 1.026 | −0.01% | EPE | 0.1100 | – |
| 2 | 40 | [2, 4]* | EEE | 0.9315 | $-1.4 \times 10^{-4}$ | 1.026 | −0.04% | PPP | 0.0609 | – |
| 4 | 80 | none ($m \geq 1.375$) | PPE | 0.3180 | −0.614 | – | – | P?P | – | – |

**Table III.8:** Exact C2 response at $g^2 = 2$ for plaquette couplings $\kappa$. Branch (a post hoc classification) lists the $R = 2$ source, $R = 4$ source and vacuum ground states: E if $\langle H_\kappa \rangle \geq -\kappa$, P if the state is plaquette dominated ($\langle H_\kappa \rangle < -\kappa$). Forces per graph link ($t_{\text{h}}/a$) are tabulated only when both source states of the $[2,4]$ secant lie on the E branch. The half-connectivity bracket uses the grid $m \in \{0.5, 1, 1.25, 1.375, 1.5, 2, 4\}$ (for the earlier $\kappa = 0.5$ control $m \in \{0.5, 1, 1.5, 2, 4\}$); * marks a bracket whose two masses lie on different branches, ? a state excluded by a gate, and – an entry not computed (the earlier $\kappa = 0.5$ control has no $R = 2$ sector, hence no force, and no per-sector plaquette energy). Masses excluded by a gate are omitted from the bracket search. $\Delta P_{\text{any}}$ and the force change are taken relative to $\kappa = 0$; $\lambda$ is the Gauss penalty used.

By this registered criterion the heavy-mass response is insensitive to the plaquette coupling up to $\kappa = 2$: at $m = 4$ the $R = 4$ connection probability changes by at most $1.42 \times 10^{-4}$ and the $[2,4]$ force by at most 0.043% (Table III.8). The criterion does not include the $R = 6$ sectors, whose source states at $m = 4$ also remain on the electric branch up to $\kappa = 2$ but respond more strongly. At $\kappa = 2$ and $m = 4$ the $[4,6]$ secants of the primary, route-B and route-C sectors fall from 1.0310, 1.0203 and 1.0310 to 0.9780, 0.9059 and $0.9777\, t_{\text{h}}/a$, by 5.1, 11.2 and 5.2%, and their connection probabilities by 0.015–0.032; at $\kappa = 1$ the secants fall by 1.0–2.3% and the connection probabilities by at most 0.006. The quadratic small-$\kappa$ expectation is not supported; only 3 of the 19 registered ratios lie in $[3,5]$. Four ratios at $m = 0.5$ involve the $\kappa = 1$ source state, which lies on the plaquette-dominated branch described below. The fifteen ratios of states that remain on the electric branch lie between −0.59 and 4.27, nine of them between 1.6 and 2.4, near the value 2 of a response linear in $\kappa$ (a post hoc observation).

The computed ground states fall into two well-separated classes (Fig. III.9(b)): an electric branch, with $\langle H_\kappa \rangle$ within $0.32\kappa$ of zero (within $0.09\kappa$ for the $R = 4$ source and vacuum), and a plaquette-dominated branch, with $\langle H_\kappa \rangle$ between $-9.5\kappa$ and $-14.7\kappa$. The source state is on the plaquette branch at $m = 0.5$ for $\kappa = 1$, up to $m = 2$ for $\kappa = 2$ and up to $m = 4$ for $\kappa = 4$; the vacuum switches at lower masses (at $m \leq 1$ for $\kappa = 2$ and $m \leq 2$ for $\kappa = 4$). Where the source and vacuum states lie on different branches, their energy difference compares two different kinds of ground state and is not a static potential. Where both states lie on the plaquette branch the difference is also negative: −2.78 and −4.12 at $\kappa = 2$, $m = 0.5$ and 1, and −7.9 to −8.4 at $\kappa = 4$, $m = 1.375$–2. At $\kappa = 2$ it

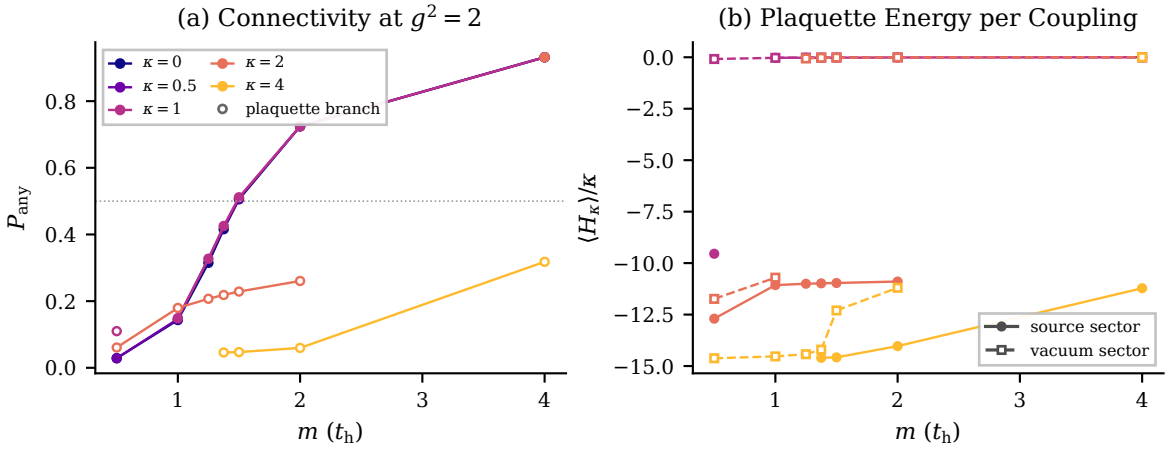



**Figure III.9:** Exact C2 response at $R = 4$ and $g^2 = 2$ for plaquette couplings $\kappa \leq 4$. (a) Connection probability; open symbols mark source states on the plaquette-dominated branch ($\langle H_\kappa \rangle < -\kappa$, a post hoc classification), and lines join only masses on the same branch. (b) Plaquette energy per coupling of the source (filled) and vacuum (open) ground states for $\kappa = 1$, 2 and 4. Points failing a gate are omitted.

stays negative up to $m = 1.5$ and is $+1.23$ at $m = 2$. Like every light-mass $V_D$, these values are conditional on the registered $D_v$ sectors. The table gives $[2, 4]$ forces only where both source states are on the electric branch (a presentation rule); the withheld secants are $-0.776$ ($\kappa = 2$, $m = \frac{1}{2}$), $0.089$ ($\kappa = 1$, $m = \frac{1}{2}$) and $-3.79$ ($\kappa = 4$, $m = 4$). On the plaquette branch the connection probability is low, 0.06–0.26 at $\kappa = 2$ and 0.05–0.32 at $\kappa = 4$. These are ground-state changes in the fixed sectors of this finite patch; they do not establish a magnetic phase or a transition.

**Independent check.** Three registered points, $(g^2, m) = (0.5, 0.375)$ and $(8, 5)$ at $\kappa = 0$ and $(\kappa, m) = (2, 1)$ at $g^2 = 2$, were recomputed with the independent first-principles physical-block implementation of Sec. II.5.2, with a separately written plaquette term. That implementation reproduces the earlier $g^2 = 2$ and $\kappa = 0.5$ records within $2 \times 10^{-12}$, and deliberately mis-oriented or sign-flipped plaquette terms fail this comparison by at least $1.6 \times 10^{-4}$. The registered production values agree with it within $4.7 \times 10^{-12}$ in $\langle H_0 \rangle$ and $1.1 \times 10^{-14}$ in $P_{\mathrm{any}}$, including the plaquette-branch source state at $(\kappa, m) = (2, 1)$.

### III.1.4 Wider-Patch Comparison at Fixed Separation

**Patches and source placement.** Two wider open patches keep $L_x = 3$: C4 ($L_y = 3$; 18 vertices, 21 links, four cycles) and C8 ($L_y = 5$; 30 vertices, 37 links, eight cycles). Their sources sit on the middle row at graph distance four, with degrees $(2, 3)$ and a single shortest four-link corridor, as for the relocated C2 sources $(4, 6)$ of Sec. III.1.1 (Fig. III.10). The comparison therefore uses the exact relocated C2 sector, not the principal $(1, 3)$ ladder. The patches differ in more than width. On C2 the corridor runs along the open edge and its two interior $B$ vertices have degree two; on C4 and C8 every interior corridor

vertex has degree three. From C4 to C8 no vertex within distance one of the corridor changes its degree. A C2–C4 difference thus combines added rows with a changed corridor coordination, and only the C4–C8 step isolates added rows. All differences below are patch differences, not a finite-size scaling.

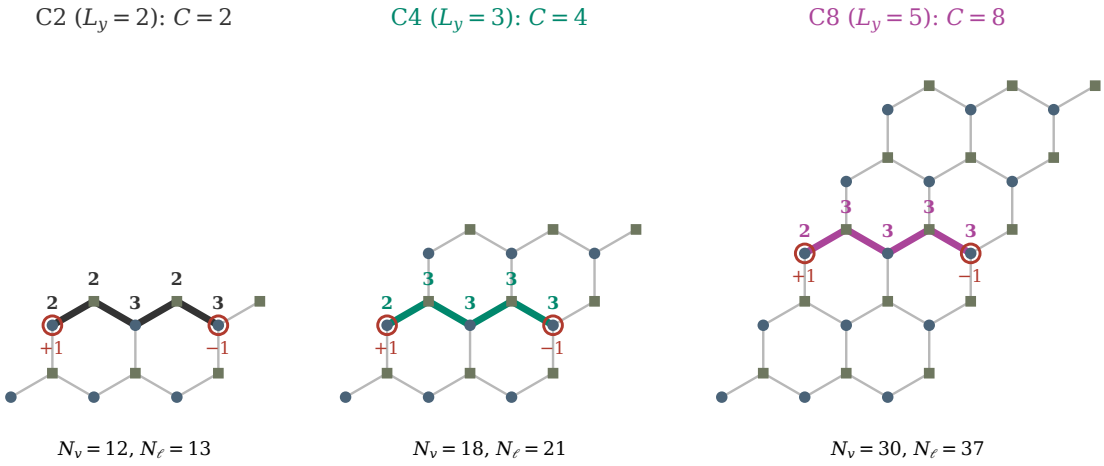



**Figure III.10:** The C2, C4 and C8 patches at a common length scale in the unit-bond embedding of Fig. II.1. Circles and squares are $A$ and $B$ vertices; the coloured four-link corridor joins the circled sources ($q^3 = \pm 1$), which for C2 are the relocated sources $(4, 6)$. Numbers give the degree of each corridor vertex. The two interior $B$ vertices of the corridor have degree two on C2; every interior corridor vertex has degree three on C4 and C8.

**Method and numerical status.** C4 and C8 have 39 and 67 tensor sites and are treated variationally with the penalty-protected two-site DMRG of Sec. II.5.2: the Gauss law and the imbalance labels enter through $H_G$ and $H_D$ with $\lambda = \mu = 20$, the bond dimension $\chi$ is capped at 400–800 and the truncation cutoff is $10^{-13}$. Connection and endpoint-class probabilities are contracted exactly for each stored matrix-product state by a frontier transfer over link-activity configurations, without sampling; they are exact for the stored state, not for the ground state. Two numerical statuses were registered on 7 October 2026, after the selected heavy-mass states had been computed but before any independent cross-check result existed. A *single-state quality pass* requires the final sector criteria (bare source Casimir 3/4 and neutral Gauss residuals within $10^{-6}$; particle-number and imbalance means within $10^{-6}$, variances below $10^{-8}$), an energy spread of at most $10^{-6}$ over the last four sweeps, $\chi$ below its cap, named $H_0$ and $H_{\mathrm{opt}}$ on the stored state with variance per site at most $10^{-5}$, and a consistent reload. *Bounded numerical support* further requires a passing independent cross-check, an independent start or a bond-dimension ladder within $10^{-4}$ in energy, $10^{-5}$ in every $\langle E_\ell^2 \rangle$ and $10^{-4}$ in every $\langle n_v \rangle$, and passing source–vacuum pair criteria. Neither status certifies a ground state. The potential is evaluated as $V_D = E_{\mathrm{opt,src}} - 30 - E_{\mathrm{opt,vac}}$, which equals Eq. (II.17) up to the penalty residuals of the stored states; a difference of two variational energies has no variational bound.

On C2 the same engine reproduced the heavy-mass endpoints ($1 - F \leq 5.8 \times 10^{-10}$),

although other heavy-mass runs stalled below their caps $(1.4–1.5) \times 10^{-4}$ above the exact energies, and it did not reproduce the light-mass source ground state (Sec. II.5.2). At $m = 1–3$ only pilots with $\chi \leq 200$ exist; their source energies lie $1.0 \times 10^{-4}$–$1.2 \times 10^{-3}$ above the exact values, with $|\Delta P_{\text{any}}| \leq 3.4 \times 10^{-4}$. All of these calibration states use the principal sources $(1, 3)$; none exists at $(4, 6)$. C4 and C8 values therefore enter with their status only: campaign statuses at $m = 4$, descriptive values at $m = 3$ and 2, and no values below $m = 2$.

| Patch | $m$ | Status | $\chi$ | $P_{\text{any}}$ | $P_{\geq 1\text{ff}}$ | $V_D$ | $\bar{\rho}_{\mathcal{C}}$ | $\epsilon_{\text{off}}$ |
|---|---|---|---|---|---|---|---|---|
| C2 $(4, 6)$ | 4 | E | – | 0.9315 | 0.0327 | 4.1154 | 0.7065 | 0.0020 |
| C4 | 4 | Q | 773/800 | 0.9364 | 0.0304 | 4.0749 | 0.7091 | 0.0034 |
| C8 | 4 | B | 422/800 | 0.9364 | 0.0304 | 4.0746 | 0.7091 | 0.0034 |
| C2 $(4, 6)^*$ | 3 | E | – | 0.8773 | 0.0576 | 4.3008 | 0.6810 | 0.0039 |
| C4 | 3 | D | 400 (cap) | 0.8911 | 0.0510 | 4.2366 | 0.6873 | 0.0066 |
| C2 $(4, 6)$ | 2 | E | – | 0.7185 | 0.1332 | 4.4570 | 0.6162 | 0.0080 |
| C4$^\dagger$ | 2 | D | 600 (cap) | 0.7750 | 0.1033 | 4.3753 | 0.6395 | 0.0128 |
| C8$^\dagger$ | 2 | D | 600 (cap) | 0.7766 | 0.1024 | – | – | – |

**Table III.9:** Static response at graph separation $R = 4$ for the exact C2 sector with sources $(4, 6)$ and for variational C4 and C8 states with their registered sources (Fig. III.10); $t_{\text{h}} = 1$, $g^2 = 2$, $\lambda = \mu = 20$, $\kappa = 0$. Status: E, exact; B, bounded numerical support; Q, single-state quality pass; D, descriptive; $^*$, post-closure supplement of Sec. III.1.1. $\chi$ gives the largest bond dimension reached by the source state and its cap. $V_D$, the mean corridor response $\bar{\rho}_{\mathcal{C}}$ and $\epsilon_{\text{off}}$ follow Eqs. (II.17) and (II.15); for C4 and C8, $V_D = E_{\text{opt,src}} - 30 - E_{\text{opt,vac}}$. The C4 value of $V_D$ at $m = 4$ depends on the vacuum start path (4.0747–4.0749 over the registered arms). For the cells marked $^\dagger$ the values are those of the lowest-energy stored state; the other measured stored states, from two optimizer start paths (C4: $\chi = 400$–$800$; C8: $\chi = 400$–$600$), differ from it by at most $2.2 \times 10^{-4}$ in $P_{\text{any}}$ and $6.0 \times 10^{-4}$ in $V_D$ on C4 (one vacuum state), and by $2.5 \times 10^{-6}$ in $P_{\text{any}}$ on C8, where no vacuum state exists. Spreads over stored states are not statistical errors.

**Heavy mass.** At $m = 4$ the C8 source and vacuum states carry bounded numerical support (Table III.9). Their single-state criteria pass, a lineage started from an independent random seed agrees with them within the $10^{-4}$ bars at $\chi = 400$, and the pair criteria pass. Continued with cap $\chi = 800$ (reaching 320 and 386), that lineage stays $8.7 \times 10^{-5}$ (vacuum) and $5.8 \times 10^{-5}$ (source) above the selected energies, short of the $10^{-6}$ criterion for an independent-seed reproduction, and a same-sector source state $2.1 \times 10^{-6}$ below the selected one is known. The C4 states pass only the single-state criteria: a fresh random start settled $2.6 \times 10^{-4}$ higher in the vacuum, so the C4 source–vacuum pair is start-dependent ($V_D = 4.0747$–$4.0749$). On both patches the C2 criterion, fixed before the C2 data and carried over unchanged, $P_{\text{any}} > P_{\text{no}}$ with $P_{\text{alt}} \approx 2 \times 10^{-8} \ll P_{\text{sel}}$, holds, and the mean corridor response is $\bar{\rho}_{\mathcal{C}} = 0.709$: the heavy-mass response is string-like on all three patches. Against the exact C2 $(4, 6)$ state, $P_{\text{any}}$ is larger by 0.0049 (0.5%) and $V_D$ is smaller by 1.0% on both wider patches, $\epsilon_{\text{off}}$ rises from 0.0020 to 0.0034 and $P_{\text{both}}$ falls from $2.7 \times 10^{-4}$ to $2.4 \times 10^{-4}$. These are descriptive patch differences. From C4 to C8,

$P_{\text{any}}$ changes by $3.3 \times 10^{-5}$, below the registered mechanism tolerance $10^{-4}$, and $V_D$ by $-3.2 \times 10^{-4}$, comparable to the C4 start dependence ($2.5 \times 10^{-4}$ in $V_D$); neither change is resolved. The largest off-corridor response, $2.2 \times 10^{-3}$ on C2 $(4, 6)$ and $2.4 \times 10^{-3}$ on C4 and C8, exceeds the narrow-tube criterion $10^{-3}$ on all three patches; at $m = 3$ and 2 it is larger still (C2 $(4, 6)$: $4.1 \times 10^{-3}$ and $6.4 \times 10^{-3}$; C4: $4.8 \times 10^{-3}$ and $9.1 \times 10^{-3}$).

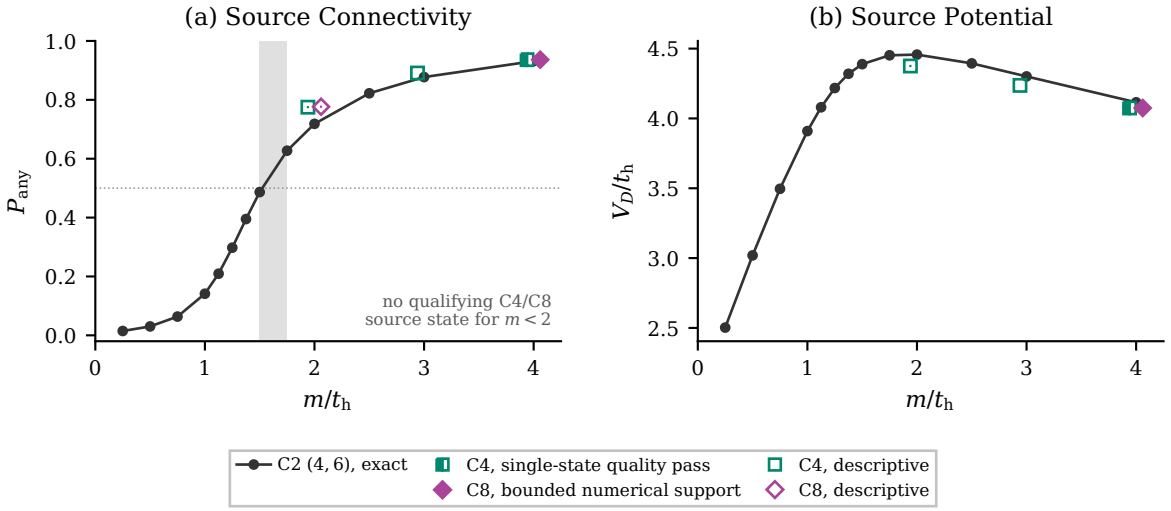



**Figure III.11:** Connection probability (a) and source potential (b) at $R = 4$ against mass. Circles: exact C2 sector with sources $(4, 6)$ at all thirteen ladder masses, seven of them from the post-closure supplement of Sec. III.1.1; lines join the sampled masses. Squares (C4) and diamonds (C8), offset by $\mp 0.06$ in $m$: variational states, with the fill giving the numerical status of Table III.9 (full: bounded numerical support; left half: single-state quality pass; open: descriptive). Spreads over stored states are drawn as vertical bars where they exceed the symbol. No C8 vacuum state exists at $m = 2$. The grey band is the C2 $(4, 6)$ half-connectivity bracket, a grid diagnostic. Below $m = 2$ no C4 or C8 source state meets the sector criteria.

**Intermediate mass.** At $m = 2$ the C4 and C8 source states give $P_{\text{any}} = 0.7750$–$0.7753$ and 0.7766, against 0.7185 for the exact C2 $(4, 6)$ state (Fig. III.11). Each range covers capped states from two optimizer start paths, a four-sweep continuation from $m = 4$ and a fresh random start. On C4 the fresh start stays $1.4 \times 10^{-4}$ above the continuation even at $\chi = 800$, with $\langle E_\ell^2 \rangle$ differences up to $1.3 \times 10^{-4}$, beyond the registered bars $10^{-4}$ and $10^{-5}$. On C8 the two paths agree within $9 \times 10^{-6}$ in energy, but the continuation has only two noise-free final sweeps and no registered energy tail, and both states are capped. The C4 potential, $V_D = 4.3753$–$4.3759$ with a single capped $\chi = 400$ vacuum state (variance per site $1.01 \times 10^{-5}$, just above the $10^{-5}$ bar), lies 1.8% below the C2 value 4.4570; no C8 vacuum state exists at this mass. The connectivity differences of +0.057 (C4) and +0.058 (C8) are therefore reported descriptively. They combine added rows with the changed corridor coordination, whereas the C4–C8 step changes $P_{\text{any}}$ by $(1.3$–$1.6) \times 10^{-3}$. A single capped $\chi = 400$ C4 source–vacuum pair at $m = 3$, whose source state has energy variance per site $3.5 \times 10^{-5}$ above the $10^{-5}$ bar, gives $P_{\text{any}} = 0.8911$ and $V_D = 4.2366$, against 0.8773 and 4.3008 for the exact relocated C2 sector. The C4 intermediate masses were

deferred in the registered scope of the support campaign (9 October 2026). The $m = 2$ continuations finished on 4 and 5 October and the fresh-start stages between 10 and 17 October; the registered $m = 2$ start-path comparison, evaluated on 19 October, gives no conclusion on either patch (on C4 the fresh start lies higher, on C8 the continuation has no converged energy tail). The $m = 3$ values and the vacuum-based $m = 2$ quantities are post hoc diagnostics.

**Light mass and open points.** No C4 or C8 source state at $m \leq 1.5$ meets the sector criteria. The saved light-mass source states sit at their caps $\chi = 400$–$600$, are still descending (last-four spans $1.9 \times 10^{-4}$–$1.3 \times 10^{-2}$) and violate the particle-number variance bar; at $m = 1.5$ no C4 source state was saved (Table C.1). All four registered $m = 0.5$ endpoints remain open, so the C4/C8 character comparison of the support campaign is inconclusive, and below $m = 2$ the mass dependence is known only for C2. Light-mass C4 and C8 source states meeting these criteria remain open. The descriptive $m = 2$ difference leaves open whether a wider-patch half-connectivity bracket differs from the C2 $(4, 6)$ bracket $[1.5, 1.75]$.

**Representation cost.** For $m \geq 2$, where the C4 and C8 states pass the sector criteria, the C4 values of $P_{\text{any}}$, $\bar{\rho}_{\mathcal{C}}$ and $V_D$, and the C8 values of $P_{\text{any}}$, change with $m$ in the same direction as the exact C2 $(4, 6)$ values, while the largest bond entropy of the stored C4 source states rises from 1.37 at $m = 4$ through 1.63 at $m = 3$ to 2.14 at $m = 2$ (Appendix C). The light-mass C4 source states carry 3.5–3.7, reach the bond-dimension cap on 9–12 of 38 bonds and fail the sector criteria, so their observables are not interpreted. These entropies measure the representation cost in the registered tensor ordering, including the gauge-constraint contribution; they are not a physical entanglement measure of the patch.

## III.2 $(2+1)$D Local Response

The link Casimir and matter occupation provide a direct static profile comparison on the loop-bearing C1 patch. The mean corridor Casimir changes from 0.73735 to 0.46814 between $m = 4$ and 0.5, whereas the pair expectation increases from 0.28880 to 2.18689. Neither quantity is a gauge-fixed colour vector. Opposite classical source orientations define a frame for colour-resolved measurements; the colour-summed Casimir itself measures electric representation occupancy.

The nonzero high-mass pair content shows why "pulling particles from the vacuum" requires an explicit reference. The screened state is compared with an already dressed state, not an empty matter Hilbert space. A matched source-free ground state is additionally needed to isolate the part of the mass response caused by the impurities. The same distinction applies to a quench: pair growth after a global mass change contains a bulk preparation response as well as source-induced redistribution.

## III.3 Local Fusion Information

### III.3.1 Vertex, Plaquette, and Loop Information

A trivalent neutral vertex with singly occupied matter can carry four spin-half constituents. Their singlet subspace has multiplicity two, so occupancy labels do not specify a unique colour-fusion state. The measured projector of Eq. (II.19) tests this internal structure without introducing a quantum source absent from the model.

In the pilot ground states the probability of the selected four-spin sector at vertex 5 grows from 0.041923 to 0.083410. Conditional on that sector, the link-1–matter singlet expectation changes from 0.250324 to 0.318772. These are state-dependent local recoupling measurements. They do not establish that a reduced model cannot reproduce the tested observables. That stronger conclusion requires a matched operator-level reduction or control and robustness across geometry. Removing off-diagonal colour hopping changes both recoupling and kinetic channels, so it is not an unqualified comparison with an independently calibrated U(1) gauge theory.

# Chapter IV

# Dynamics

The static C2 states provide controlled initial conditions for a breaking quench and a formation ramp. Source and vacuum states evolve under matched protocols in their own registered sectors, and exact propagation is compared with a Gauss-exact tensor-network representation.

## IV.1 $(2+1)$D Honeycomb Dynamics

### IV.1.1 Protocols and Observables

The real-time calculations use the C2 patch at fixed graph separation $R = 4$, with the sources at vertices 1 and 3 and the sectors of Eq. (II.21). Source and vacuum states evolve under the same protocol in their separately registered sectors. A mass change modifies only the diagonal of the Hamiltonian and preserves every conserved label, so both protocols remain inside these sectors. The sources are present throughout: the string is dynamically prepared or broken within a pre-existing source sector, and the unitary evolution does not create the sources.

- *Breaking quench*: the $m = 4$ source and vacuum ground states evolve with $m = 0.5$. Exact propagation extends to $t = 20$; the tensor-network reproduction covers $0 \le t \le 2$.
- *Formation ramp*: the $m = 0.5$ ground states evolve with $m(t) = 0.5 + 3.5\,t/4$ for $0 \le t \le 4$; the tensor-network reproduction covers the whole ramp. A subsequent hold at $m = 4$ is propagated exactly up to $t = 24$.

Times are in units of $\hbar/t_{\mathrm{h}}$, with $\hbar = 1$. *Exact* denotes full registered-sector propagation without truncation, using exponential actions for the fixed-Hamiltonian quench and the fourth-order commutator-free Magnus scheme of Ref. [4] with step 0.025 for the ramp. It does not denote zero numerical error; the accuracy of every exact curve is established by an independent re-propagation (Sec. IV.1.4). The evolved states are nonequilibrium states rather than separately optimized ground states.

The observables are the electric response $\tau_E^{\mathrm{resp}}(t)$ of Eq. (II.18) with its geodesic control and the four exclusive source classes of Eq. (II.16). Two criteria are applied throughout.

A response is *string-breaking-like* if $\tau_E^{\mathrm{resp}}$ and $P_{\mathrm{any}}$ are suppressed together with growth of the flux-free endpoint classes, and *flux-tube-formation-like* if $\tau_E^{\mathrm{resp}}$ and $P_{\mathrm{any}}$ grow together. String breaking in $(2+1)$D non-Abelian calculations and in a $(2+1)$D Rydberg gauge simulator provides the reference phenomenology [10, 21]. Figure IV.1 shows both protocols, and Table IV.1 lists characteristic values.

| Protocol | $t$ | $\tau_E^{\mathrm{resp}}$ | $P_{\mathrm{any}}$ | $P_{1\mathrm{ff}}$ | $P_{2\mathrm{ff}}$ | $P_{0\mathrm{ff}}$ |
|---|---|---|---|---|---|---|
| Quench | 0 | 0.707982 | 0.931665 | 0.033067 | 0.000284 | 0.034984 |
| Quench | 1.6 | 0.105739 | 0.037104 | 0.587288 | 0.152210 | 0.223397 |
| Quench | 1.75 | 0.135757 | 0.034270 | 0.593573 | 0.150512 | 0.221645 |
| Ramp | 0 | 0.160779 | 0.028904 | 0.642162 | 0.159919 | 0.169015 |
| Ramp | 1.5 | 0.346122 | 0.311539 | 0.557886 | 0.044049 | 0.086526 |
| Ramp | 4 | 0.453518 | 0.477103 | 0.476362 | 0.016219 | 0.030317 |

**Table IV.1:** Characteristic values of the exact source–vacuum evolution at $R = 4$. $t$ is in units of $\hbar/t_{\mathrm{h}}$ and $\tau_E^{\mathrm{resp}}$ in $t_{\mathrm{h}}/a$; the class probabilities refer to the source state and require no vacuum subtraction.

### IV.1.2 Breaking Quench: String-Breaking-Like Rearrangement

The initial state carries the confinement-like string of Sec. III.1.2. Because the electric operator does not depend on the mass, $\tau_E^{\mathrm{resp}}(0) = 0.707981564\, t_{\mathrm{h}}/a$ equals the ground-state response $\tau_{E,\mathrm{GS}}^{\mathrm{resp}}(4; m = 4)$.

After the mass is lowered, the electric response falls to 0.105739 at $t = 1.6$ ($-85\%$) and the connection probability to 0.034270 at $t = 1.75$ ($-96\%$). These are the global minima on the output grid of spacing 0.025; parabolic refinement places them at $t \simeq 1.609$ and 1.739. At $t = 1.75$ the flux-free endpoint classes have grown to $P_{1\mathrm{ff}} = 0.593573$ and $P_{2\mathrm{ff}} = 0.150512$, from 0.033067 and 0.000284 initially, so that 74% of the weight is in configurations in which at least one source holds matter without incident flux, while $P_{0\mathrm{ff}} = 0.221645$ remains disconnected without either endpoint event. Connected flux has thus been strongly suppressed and the probability transferred predominantly to flux-free endpoint matter.

The source–vacuum difference of the total energy $\langle H_0 \rangle$ is conserved with a recorded range of $2.162 \times 10^{-11}$. Electric reduction is therefore redistribution among electric, mass and hopping contributions rather than loss of total energy, consistent with the pair-creation picture of string breaking in $(2+1)$D gauge dynamics [10]. Because $\tau_E^{\mathrm{resp}}$ is a source–vacuum difference, its reduction combines a decrease of the source electric energy ($3.22 \to 2.85$ at $t = 1.6$) with electric energy generated by pair creation in the quenched vacuum ($0.39 \to 2.43$); the class probabilities of the source state, which involve no subtraction, show the rearrangement directly.

The quench injects energy well above the final ground states. With $H_f = H_i + (m_f - m_i)M$,

$$W_{\mathrm{excess}} = \langle H_f \rangle_{\psi_i} - E_{0,f}, \qquad W_{\mathrm{injected}} = \langle H_f - H_i \rangle_{\psi_i}, \tag{IV.1}$$

the source and vacuum give $(W_{\mathrm{excess}}, W_{\mathrm{injected}}, \Delta E_f) = (6.913410,\ 38.863833,\ 4.546184)$ and

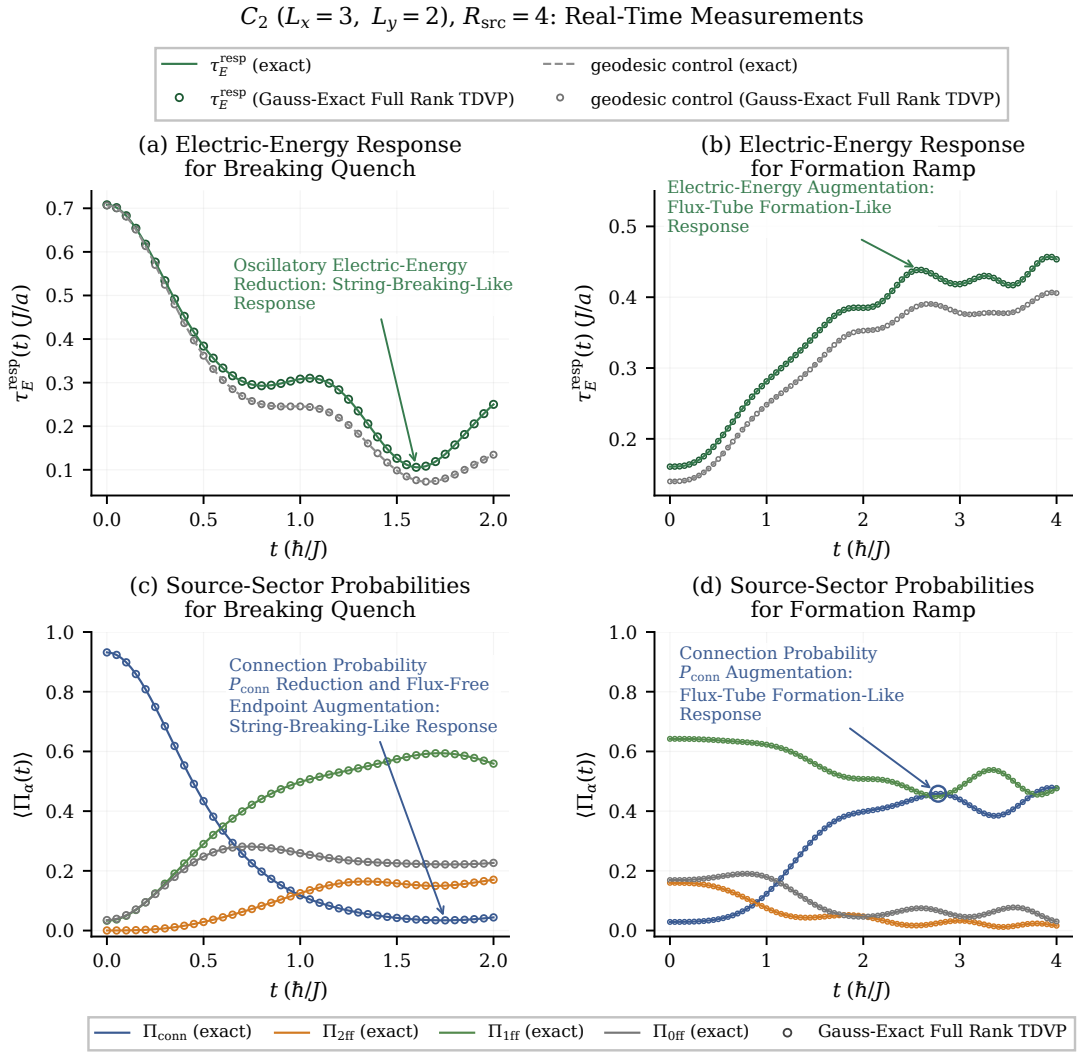



**Figure IV.1:** Real-time response at $R = 4$; $J \equiv t_h$ and $\Pi_{conn} \equiv \Pi_{any}$. Left (a,c): breaking quench $m = 4 \to 0.5$. Right (b,d): formation ramp $m(t) = 0.5 + 3.5t/4$. Top: electric response $\tau_E^{resp}$ (dark green) and geodesic control $\tau_E^{geo}$ (grey dashed). Bottom: the four exclusive source classes $\Pi_{any}$ (blue), $\Pi_{1ff}$ (green), $\Pi_{2ff}$ (orange) and $\Pi_{0ff}$ (grey). Lines are exact propagation. Circles are the independent Gauss-exact tensor-network evolution (full-rank one-site TDVP, fourth order, $\Delta t = 0.025$, sampled every 0.05); compared with exact propagation using the same step, they agree within $5.7 \times 10^{-11}\, t_h/a$ in $\tau_E^{resp}$, $1.2 \times 10^{-10}\, t_h/a$ in the geodesic control and $4.3 \times 10^{-10}$ in the classes. Arrows mark the global minima of the exact quench curves ($t = 1.6$ for $\tau_E^{resp}$, $t = 1.75$ for $P_{any}$) and the first local maxima at $t \geq 2$ of the exact ramp curves ($t = 2.575$ and 2.775); the ring in (d) marks the region where $P_{any}$ and $P_{1ff}$ differ by less than 0.01, between their crossings near $t = 2.66$ and 2.86.

(5.282273, 38.266171, 4.927858), where $\Delta E_f$ is the energy spread under $H_f$. The excess-to-gap ratios, approximately 9.136 and 1.969, caution against assuming a small-excitation linear-response regime.

Over the exactly propagated window $2 \le t \le 20$ the connected flux stays suppressed, $P_{\text{any}} = 0.042$–$0.126$, while $\tau_E^{\text{resp}}$ oscillates between 0.118 and $0.329\, t_{\text{h}}/a$; the averages over $10 \le t \le 20$ are $P_{\text{any}} \simeq 0.075559$ and $P_{\text{2ff}} \simeq 0.128232$. The configuration reached resembles the static light-mass ground state at $R = 4$ ($P_{\text{any}} = 0.029$, $P_{\ge 1\text{ff}} = 0.80$) qualitatively, but closed evolution preserves the post-quench energy and continues to oscillate; the low-mass ground state is not the late-time state. The response is an oscillatory string-breaking-like, screening-like rearrangement. The remaining disconnected class excludes an interpretation as complete two-endpoint screening, and flux-free matter support is not an external-source–matter singlet projector.

### IV.1.3 Formation Ramp: Flux-Tube-Formation-Like Growth

As the mass is raised from the screening-like light-mass ground state, the electric response and the connectivity grow together: $\tau_E^{\text{resp}} = 0.160779 \to 0.453518$ (a factor 2.8) and $P_{\text{any}} = 0.028904 \to 0.477103$ by $t = 4$, while the two-endpoint flux-free class is suppressed, $P_{\text{2ff}} = 0.159919 \to 0.016219$. This is the flux-tube-formation-like signature. At $t = 4$ connected and one-endpoint flux-free configurations carry nearly equal weight, $P_{\text{any}} \simeq P_{\text{1ff}} \simeq 0.48$: formation is supported as a response, not certified as a pure flux tube.

After the ramp the population of the confinement-like $m = 4$ ground state, $F_{\text{GS}} = |\langle \psi_{\text{GS}}(m = 4) | \psi(t = 4) \rangle|^2$, is 0.467 for the source state (vacuum 0.982). During the subsequent exact hold at $m = 4$ the Hamiltonian is constant, so every eigenstate weight is conserved: $F_{\text{GS}}$ stays at 0.467 within $4 \times 10^{-13}$ up to $t = 24$. The observables continue to oscillate. $P_{\text{any}}$ varies between 0.376 and 0.510 with mean 0.442, to which the ground-state component alone contributes $0.467 \times 0.932 = 0.435$, and lies below the static midpoint $(0.931665 + 0.028904)/2 = 0.4803$ in 553 of the 601 samples with $9 \le t \le 24$ (minimum 0.3762 at $t = 18.1$). $\tau_E^{\text{resp}}$ stays between 0.398 and $0.476\, t_{\text{h}}/a$, below its $m = 4$ ground-state value 0.708. The ground-state population is set by the $4\, \hbar/t_{\text{h}}$ ramp; persistence of a fully formed configuration is not supported on this patch at this ramp speed.

On the smaller loop-bearing C1 patch, exact linear ramps $0.5 \to 4$ raise the squared overlap with the high-mass source ground state from 0.672 to 0.999 as the ramp duration increases from 1 to $64\, \hbar/t_{\text{h}}$. The C2 ground-state population is therefore expected to depend strongly on the ramp speed; a slow-ramp series on C2 has not been carried out.

### IV.1.4 Exact Propagation and Tensor-Network Validation

**Exact layer.** The exact curves were re-propagated in the gauge-invariant physical block, with the hopping blocks cross-checked in Sec. II.5.2 and separate integrators: a Lanczos exponential with an a-posteriori error estimate below $10^{-15}$ per step for fixed Hamiltonians and an eighth-order Runge–Kutta (DOP853) integrator with relative tolerance $10^{-13}$ for the ramp. The breaking quench agrees within $8.4 \times 10^{-14}\, t_{\text{h}}/a$ in $\tau_E^{\text{resp}}$, $8.7 \times 10^{-14}$ for every link Casimir and $1.3 \times 10^{-14}$ for the classes over $0 \le t \le 20$. The ramp agrees within $1.5 \times 10^{-9}\, t_{\text{h}}/a$; this residual is the step error of the recorded fourth-order Magnus scheme,

and repeating that scheme reproduces the record within $4 \times 10^{-14}$. Halving the Magnus step from 0.05 to 0.025 changes the electric response by at most $2.29 \times 10^{-8}$. The hold to $t = 24$ agrees within $1.8 \times 10^{-8}$. At $t = 19$ and 20 the propagated state vectors agree with the stored ones to $1 - F \leq 5.2 \times 10^{-14}$, with weight outside the gauge-invariant block below $5.2 \times 10^{-14}$.

**Gauss-exact time evolution.** The tensor-network evolution uses the Gauss-exact representation of Sec. II.5.2. The quench starts from the exact $m = 4$ ground states and the ramp from the exact $m = 0.5$ ground states, each converted to a matrix-product state without truncation (discarded weight below $2 \times 10^{-15}$). The ramp uses the same commutator-free fourth-order scheme as the exact reference: each step is a product of two exponentials of weighted combinations of the Hamiltonian at the two Gauss quadrature nodes, which for the linear mass ramp are half steps with the masses at $t_n + \Delta t/6$ and $t_n + 5\Delta t/6$. Exact propagation in the same basis with the identical scheme and Lanczos exponentials agrees with the certified curves within $7.9 \times 10^{-14}$ in $\tau_E^{\mathrm{resp}}$ and $5.7 \times 10^{-14}$ in the classes, so that the comparison isolates the tensor-network error.

The certified runs use a full-rank one-site TDVP integrator [23]. At every cut and in every prefix sector the bond keeps the maximal Schmidt rank of the constrained space; zero Schmidt values are retained, and their singular vectors complete the bases. For the source the bond dimensions are $[2, 14, 26, 174, 480, 496, 289, 133, 70, 14, 6]$. The variational manifold is then the whole sector and the tangent-space projector is the identity, so the projector-splitting integrator has no projection error; every sample confirms that the bonds still equal these bounds, and the cumulative discarded weight is zero. What remains is the time-splitting error of the fourth-order integrator, a triple-jump composition of symmetric second-order sweeps, the local Krylov tolerance $10^{-14}$ and floating-point roundoff. Halving $\Delta t$ reduces the largest quench deviation in $\tau_E^{\mathrm{resp}}$ by a factor 15.9, as expected for fourth order. In the ramp the deviations at both step sizes stay at the Krylov and floating-point floor (Table IV.2, Fig. IV.2). Because no bond truncation is involved, these runs validate the Gauss-exact operator construction and the integrator on C2; they do not test tensor-network compression on larger patches.

| Run | Samples | $t_{\max}$ | $\lvert\Delta\tau_E^{\mathrm{resp}}\rvert$ | $\lvert\Delta P_\alpha\rvert$ | $\lvert\Delta E_\ell^2\rvert$, $\lvert\Delta n_v\rvert$ | $1 - F$ | Leakage | Energy |
|---|---|---|---|---|---|---|---|---|
| Quench, $\Delta t = 0.025$ | 41 | 2.00 | $5.7\times10^{-11}$ | $4.3\times10^{-10}$ | $6.1\times10^{-10}$ | $7.6\times10^{-14}$ | $2.9\times10^{-15}$ | $2.0\times10^{-14}$ |
| Quench, $\Delta t = 0.0125$ | 41 | 2.00 | $3.6\times10^{-12}$ | $2.5\times10^{-11}$ | $4.3\times10^{-11}$ | $1.6\times10^{-15}$ | $2.9\times10^{-15}$ | $2.6\times10^{-14}$ |
| Ramp, $\Delta t = 0.025$ | 81 | 4.00 | $3.6\times10^{-13}$ | $4.5\times10^{-13}$ | $1.1\times10^{-12}$ | $3.6\times10^{-15}$ | $2.3\times10^{-15}$ | $2.3\times10^{-12}$ |
| Ramp, $\Delta t = 0.0125$ | 81 | 4.00 | $5.5\times10^{-13}$ | $4.7\times10^{-13}$ | $1.9\times10^{-12}$ | $1.6\times10^{-15}$ | $2.5\times10^{-15}$ | $4.1\times10^{-12}$ |

**Table IV.2:** Gauss-exact full-rank TDVP minus exact propagation of the same initial state with the same step. Entries are maxima over source and vacuum, rounded upward to two significant digits. $1 - F$ is the state infidelity after lifting the MPS to the physical block, and leakage the weight outside the constrained configuration space. The energy column gives the relative drift for the quench and $\lvert E_{\mathrm{TN}} - E_{\mathrm{exact}}\rvert$ for the ramp. Norms are conserved within $5.6 \times 10^{-13}$.

Halving $\Delta t$ changes $\tau_E^{\mathrm{resp}}$ and the classes by at most $4.1 \times 10^{-10}$ (quench) and $1.5 \times 10^{-9}$ (ramp), against the predefined threshold $10^{-7}$. For the ramp this change is the step error of the commutator-free scheme itself: the exact references at the two step sizes differ by at most $1.5 \times 10^{-9}$, while the tensor-network error changes by at most $2.7 \times 10^{-13}$.

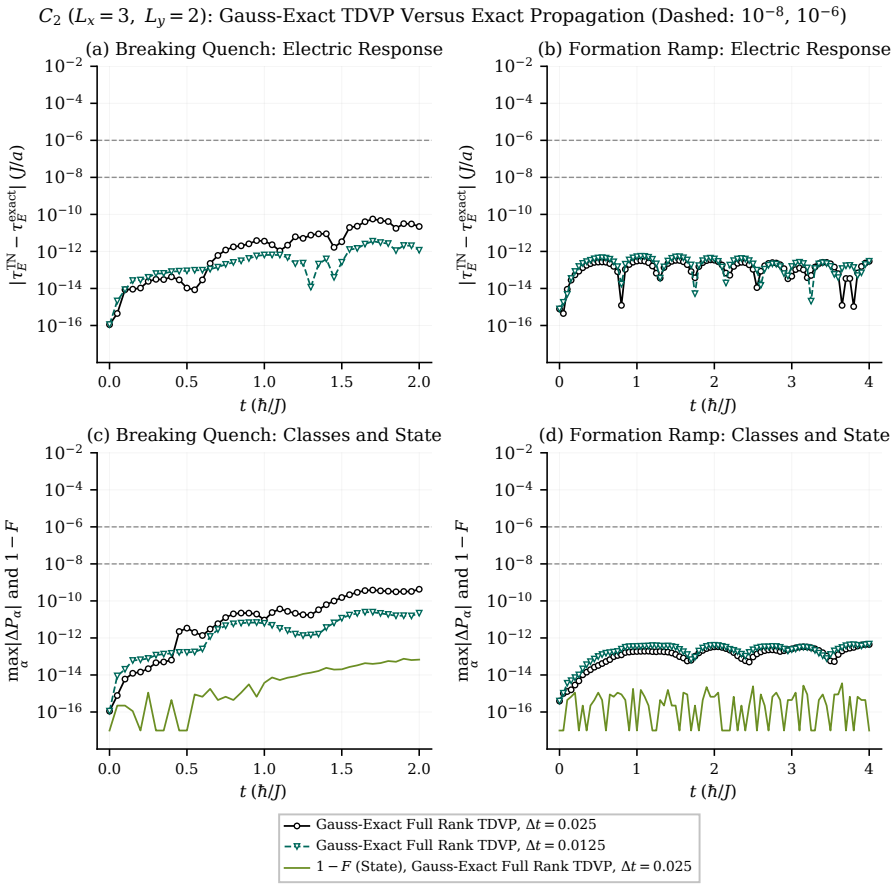



**Figure IV.2:** Gauss-exact TDVP minus exact propagation of the same initial state with the same step, on a logarithmic scale. Dashed lines mark the predefined thresholds $10^{-6}$ for observables and $10^{-8}$ for $1 - F$. (a,b) $|\tau_E^{\mathrm{resp,TN}} - \tau_E^{\mathrm{resp,exact}}|$ for the quench and the ramp. (c,d) Largest class deviation $\max_\alpha |\Delta P_\alpha|$ and, as the olive line, the state infidelity $1 - F$. Black circles and teal triangles: $\Delta t = 0.025$ and 0.0125. Values below $10^{-17}$ are drawn at $10^{-17}$.

Without a truncation parameter a bond-dimension control would be bit-identical, so it is replaced by the structural check of the full-rank bonds. All evaluated observable, fidelity, norm, structural-constraint and time-step checks satisfy their predefined tolerances over $0 \le t \le 2$ (quench) and $0 \le t \le 4$ (ramp). Two checks foreseen in the acceptance protocol have not yet been carried out: an energy balance of the ramp, comparing the energy change with the integrated work $\int \langle \partial_t H \rangle \, dt$, and a reload of the stored checkpoints in a fresh process. The exact quench beyond $t = 2$ and the hold rest on the two exact layers.

**Two-site TDVP and projection error.** A two-site TDVP run specified in advance ($\chi_{\max} = 512$, $\Delta t = 0.025$, fourth order, cutoff $10^{-14}$) fails for the quench. Its state infidelity exceeds $10^{-8}$ from $t = 0.35$ and reaches $3.4 \times 10^{-3}$ (source) and $1.6 \times 10^{-3}$ (vacuum) at $t = 2$, with $|\Delta\tau_E^{\mathrm{resp}}|$ up to $2.1 \times 10^{-4}$. The energy stays at the exact value within $6.7 \times 10^{-11}$, and the bonds stay below the cap, at most 487. Conserved energy together with a degraded state and an unsaturated cap excludes bond truncation as the dominant error and is consistent with tangent-space projection error: in the vertex ordering several hoppings are not nearest neighbours, so the two-site tangent space misses directions generated by $H$. A control with subspace expansion or another ordering was not performed. The two-site ramp run completed with $1 - F \le 4.3 \times 10^{-11}$ and $|\Delta\tau_E^{\mathrm{resp}}| \le 5.6 \times 10^{-11}$ and is retained only as a supporting run. An earlier two-site evolution in the penalty-based representation agreed with exact propagation of its own initial state within $10^{-8}$ in $\tau_E^{\mathrm{resp}}$ for $t \le 0.15$, yet its neutral-vertex Gauss residual grew from $6.7 \times 10^{-11}$ at $t = 0$ to $1.86 \times 10^{-7}$ at $t = 0.225$ with converged local solves and bonds below the cap. Because exact evolution conserves the Casimir $C_{12} = \boldsymbol{G}_{12}^2$ of the neutral leaf vertex 12, with $\|C_{12}\| = 2$, the leak bounds the state error from below, $\|\psi_{\mathrm{TDVP}}(t) - \psi_{\mathrm{exact}}(t)\| \ge (\langle C_{12}\rangle_t^{1/2} - \langle C_{12}\rangle_0^{1/2})/\sqrt{2} = 3.2 \times 10^{-3}$ at $t = 0.5$, although the electric response still agreed within $2.3 \times 10^{-5}$. Agreement of selected observables therefore does not establish sector preservation; the Gauss-exact basis removes this failure mode by construction.

## IV.1.5 Thermal Initial States and Non-Abelian Scope

The exact breaking trajectory supplies candidate initial states for thermal contact. At its $t = 20$ endpoint the source has $P_{\mathrm{any}} \simeq 0.110068$ and $P_{\mathrm{2ff}} \simeq 0.122192$. These checkpoints were selected after inspecting the closed evolution. A thermal comparison must retain that selection history, transfer each state into the declared physical representation and start a matched isolated continuation (Sec. V.2). No such checkpoint has yet supplied a converged thermal ensemble, and the closed-system observables above are not a thermal response.

Real-time string breaking also occurs in a $(2 + 1)$D Rydberg gauge simulator [21]; the present breaking is therefore not a uniquely non-Abelian signature. An SU(2)-specific dynamical claim requires matched Abelian or centre-reduced controls with common geometry and accessible charge content, the same protocol, fusion or multiplet resolution and a numerically converged difference. The colour-diagonal comparison of Table III.3 is a static control, not a substitute for such matched dynamical controls.

### IV.1.6 Physical Interpretation and Experimental Implications

On the C2 patch the static and real-time layers point in the same direction: the quench suppresses the string at conserved total energy, and the ramp rebuilds flux and connectivity together without forming a persistent tube. These finite-patch proxies at one separation do not establish irreversible relaxation, thermalization, a thermodynamic deconfinement transition or a uniquely non-Abelian breaking signal; static, dynamic and thermal evidence remain separate layers.

An experimental realization requires the specified low- and high-mass source states, the complete $m(t)$ schedule and dimensionless duration, and a calibrated conversion to laboratory time. Because the four source classes are projector expectations of joint matter and link occupations, with $E^2 = 3n_L n_R/4$, they are natural observables for site-resolved snapshots. The exact curves and their tensor-network reproduction provide reference targets for Chapter VI; they do not by themselves establish laboratory feasibility.

# Chapter V

# Gauge-Aware Relaxation and Thermalization Criteria

Closed-system response, equilibrium ensembles and reservoir-induced dynamics are treated as distinct evidence layers. The C2 study establishes complete physical sectors, independent Gibbs references, algebraic accessibility and bounded finite-temperature action tests without inferring a mixing time or a flux-tube lifetime.

## V.1 Closed-System Honeycomb Reference

The exact calculations of Sec. IV.1 distinguish changes of connectivity, electric support and matter screening during isolated formation and breaking. Unitary relaxation of selected observables does not change the global entropy of a pure state and need not approach a canonical ensemble. The thermal question is different: does sector-preserving energy exchange alter these structures relative to matched isolated evolution? The C2 study answers the representation, equilibrium-reference and algebraic parts of this question, and supplies initial numerical tests of the open dynamics.

## V.2 Gauge-Aware Honeycomb Thermal Dynamics

### V.2.1 Physical Sectors and Canonical Equilibrium

The thermal calculation uses the open $3 \times 2$, two-cycle C2 patch: twelve four-state matter vertices, thirteen six-state two-rishon links and $N = 12$. The parameters are $t_{\mathrm{h}} = 1$, $g^2 = 2$, $\kappa = 0$ and $\lambda = \mu = 20$. The two classical sources $q^3 = +1, -1$ occupy vertices 1 and 3, both on sublattice $A$; $R_{\mathrm{graph}} = 4$ and $R_x = 2$. The selected path is $1{-}7{-}2{-}8{-}3$, with link indices $(1, 3, 2, 5)$. High and low mass mean $m = 4$ and 0.5, not high and low temperature.

The determinant-free Hamiltonian retains the additional local Abelian charges $D_v$. Thus the ensemble fixes the matter number, every $D_v$, the Cartan targets and the full local Gauss representations; it is not a trace over all sectors of a pure SU(2) theory. The

source and vacuum patterns are

$$\begin{aligned} \boldsymbol{d}_{\mathrm{src}} &= (-1,-2,-1,0,0,0,-2,-2,-2,-2,-2,0), \\ \boldsymbol{d}_{\mathrm{vac}} &= (0,0,0,0,0,0,-4,-4,-2,-2,-2,0). \end{aligned} \tag{V.1}$$

Neutral vertices have $j = 0$; source vertices have $j = 1/2$ and $G^3 = q^3/2$. No quantum external-doublet degree of freedom is added.

Let $\mathcal{U}_s$ embed the complete physical space of role $s$ into its fixed-$G^3, D, N$ product basis. With the penalty-free Hamiltonian $H_0$, define

$$\begin{aligned} P_s &= \mathcal{U}_s\mathcal{U}_s^\dagger, & H_s &= \mathcal{U}_s^\dagger H_0 \mathcal{U}_s, \\ Z_s(\beta) &= \mathrm{Tr}_{\mathcal{H}_s}\, e^{-\beta H_s}, & \rho_{\beta,s} &= Z_s(\beta)^{-1} e^{-\beta H_s}. \end{aligned} \tag{V.2}$$

Here $\hbar = k_B = 1$, and $\beta$ is in units $t_{\mathrm{h}}^{-1}$. The physical dimensions are $d_{\mathrm{src}} = 27{,}580$ and $d_{\mathrm{vac}} = 16{,}502$, compared with product-basis dimensions 6,713,288 and 4,498,399. This projection removes symmetry blocks, not high-energy states. In the embedding space the same reference is

$$\rho_{\beta,Q} = \mathcal{U}_s \rho_{\beta,s} \mathcal{U}_s^\dagger = \frac{P_s e^{-\beta H_0} P_s}{\mathrm{Tr}(P_s e^{-\beta H_0})}, \qquad [P_s, H_0] = 0. \tag{V.3}$$

The earlier notation $P_Q$ denotes this same physical-sector projector. At $\beta = 0$, $Z_s = d_s$, $\log Z_s = \log d_s$ and $\rho_{\beta,Q} = P_s/d_s$.

Within these physical blocks, $H_{\mathrm{opt},s} = H_s + c_s I_s$, where $c_{\mathrm{src}} = 30$, $c_{\mathrm{vac}} = 0$ and the imbalance penalty vanishes. The constant cancels from normalized Gibbs states and Bohr frequencies, but $Z_{\mathrm{opt},s} = e^{-\beta c_s} Z_s$ and $F_{\mathrm{opt},s} = F_s + c_s$. Energy, heat, work and source–vacuum free energies below use $H_0$. This subtraction is not valid as a constant on arbitrary leaking states. A finite-penalty Gibbs ensemble, a fixed-Cartan ensemble and the complete physical-block ensemble are distinct.

Four exact endpoint vectors supply the initial states. Gauss-exact DMRG reproduces all four (Table II.2), but the thermal holds use the exact vectors directly and do not depend on an MPS approximation. The physical transfer gives a maximum $H_0$ residual $1.26004 \times 10^{-13}$, round-trip error $1.92234 \times 10^{-14}$, and initial-observable discrepancy $7.28306 \times 10^{-14}$. These are transfer measurements, not an open-system convergence result.

### V.2.2 Sector-Preserving Contacts and the Davies Reference

The canonical contact set is fixed without rescaling:

$$\{S_a\}_{a=1}^{25} = \{n_1, \ldots, n_{12}, E_1^2, \ldots, E_{13}^2\}. \tag{V.4}$$

These Hermitian gauge scalars preserve $N, D_v, G_v^3$ and the declared Gauss representations. They generally do not commute with hopping and can therefore exchange energy. Equal rates $\gamma_a = 0.02$ do not give equal channel strengths because the contact norms differ. Ambient local-algebra checks are necessary: projecting first would make excluded-sector leakage invisible. The classical-source penalty breaks transverse generator commutation at

a source while retaining its Casimir and $G^3$; it does not justify imposing neutral Gauss targets there.

For a stationary Hamiltonian the Davies construction provides a fully spectral reference [15]. Its GKSL generator is [22, 27, 29]

$$\begin{aligned}\dot{\rho} &= -i[H_s, \rho] + \mathcal{D}(\rho),\\ \mathcal{D}(\rho) &= \sum_{a,\omega} \Gamma_a(\omega)\left(A_a(\omega)\rho A_a^\dagger(\omega) - \tfrac{1}{2}\{A_a^\dagger(\omega)A_a(\omega), \rho\}\right).\end{aligned} \tag{V.5}$$

Complete degenerate energy projectors define

$$A_a(\omega) = \sum_{E_j - E_i = \omega} P_i S_a P_j, \qquad [H_s, A_a(\omega)] = -\omega A_a(\omega), \tag{V.6}$$

so positive $\omega$ lowers energy and

$$\Gamma_a(-\omega) = e^{-\beta\omega}\Gamma_a(\omega), \qquad \omega > 0. \tag{V.7}$$

Equal Bohr frequencies are grouped consistently. The star reference uses $\Gamma_a(\omega) = \gamma_a e^{-\omega/\omega_c}$ for $\omega > 0$, $\omega_c = 10t_\mathrm{h}$, $\Gamma_a(0) = \gamma_a$, and no Lamb shift. This phenomenological reference assumes weak coupling, short reservoir memory and adequate secular separation; the combined rates of all contacts matter. Local $S_a$ do not imply local energy-resolved jumps. No complete C2 Davies Liouvillian or microscopic laboratory reservoir is constructed.

Gibbs stationarity and uniqueness are separate conditions [3]. On the $d = 10, 12$ vacuum/source star blocks, sixteen points satisfy the 30 original numerical checks, and each has one resolved stationary mode. Appendix B retains these tests separately from the two KMS constructions below.

Two controls also remain distinct. A coupling commuting with $H_s$ dephases without changing energy populations. Bare Hermitian contacts instead give the unital $\beta = 0$ generator

$$\mathcal{D}_0(\rho) = -\frac{1}{2}\sum_a \gamma_a [S_a, [S_a, \rho]], \tag{V.8}$$

which can exchange energy and has $I_s/d_s$ as a stationary state. Changing $\gamma$ does not turn this infinite-temperature contact into a finite-temperature reservoir.

### V.2.3 Algorithmic KMS Thermostats

An alternative to resolving every Bohr frequency is to construct a Gibbs sampler satisfying KMS detailed balance [17, 13]. Here the frequency convention is $\nu = E_i - E_j = -\omega$, and $\mathrm{ad}_{H_s}(X) = [H_s, X]$. The Gaussian-filtered jumps are

$$L_a = \sqrt{\gamma_a}\, f_\beta(\mathrm{ad}_{H_s})(S_a), \qquad f_\beta(\nu) = e^{-\beta^2\nu^2/8 - \beta\nu/4}. \tag{V.9}$$

A finite-spectrum modular alternative sets the even prefactor of the filter to one:

$$L_a = \sqrt{\gamma_a}\, e^{-\beta H_s/4} S_a e^{\beta H_s/4}. \tag{V.10}$$

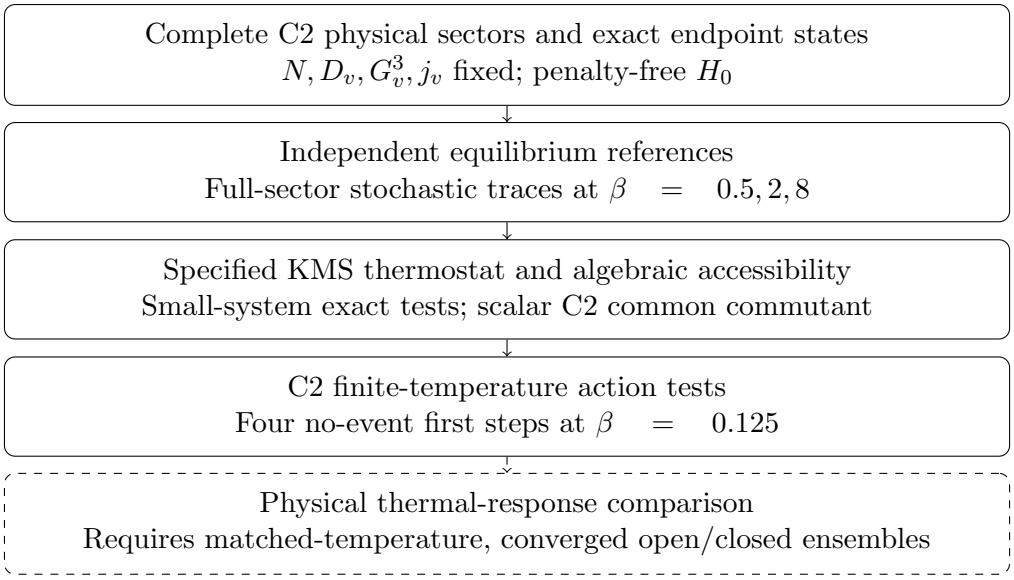



**Figure V.1:** Logical dependencies of the C2 thermal analysis. Solid boxes describe completed constructions or bounded calculations, not interchangeable convergence statements. The dashed box is the unmeasured ensemble-response question. The equilibrium and first-step temperatures differ.

For either choice, the thermostat includes the coherent correction

$$K = \sum_a L_a^\dagger L_a, \qquad G_{\mathrm{KMS}} = \frac{i}{2}\tanh\left(\frac{\beta\ \mathrm{ad}_{H_s}}{4}\right)(K),$$

$$\mathcal{M}_\beta(\rho) = -i[G_{\mathrm{KMS}}, \rho] + \sum_a \left(L_a \rho L_a^\dagger - \tfrac{1}{2}\{L_a^\dagger L_a, \rho\}\right), \tag{V.11}$$

$$\dot{\rho} = -i[H_s, \rho] + \mathcal{M}_\beta(\rho). \tag{V.12}$$

The correction is part of the algorithmic thermostat, not a redefinition of physical energy or a derived microscopic Lamb shift. Omitting it destroys Gibbs stationarity in general: the modular small-system control reaches a fixed-point residual $9.691 \times 10^{-4}$ without it.

The Heisenberg thermostat $\mathcal{M}_\beta^\dagger$ is self-adjoint in $\langle A, B\rangle_{\beta,s} = \mathrm{Tr}(\rho_{\beta,s}^{1/2} A^\dagger \rho_{\beta,s}^{1/2} B)$. The added physical drift $i[H_s, \cdot]$ is skew-adjoint in this product. The two thermostat families share an equilibrium target, not necessarily a relaxation rate or transient trajectory. Their filter, normalization and contact rates define the dynamics. The modular finite-spectrum identity does not inherit a complexity bound requiring a decaying, integrable Fourier filter. The C2 finite-temperature steps use the modular construction only. Its numerical conditioning restriction $\beta(E_{\mathrm{upper}} - E_{\mathrm{lower}}) \leq 24$ is an implementation guard on the stored Hamiltonian, not a mathematical restriction on the exact finite-dimensional construction. Integral formulas and truncation errors are given in Appendix B.2.

### V.2.4 Accessibility and the Conditional Uniqueness Result

Stationarity alone leaves open additional invariant blocks and their initial weights. Contact-signature degeneracies also prevent ordinary Hamiltonian graph connectivity from proving

irreducibility. For the stored, Hermitian-symmetrized C2 Hamiltonian and all 25 canonical contacts, a constructive exact-coefficient witness gives

$$\{H_s, S_1, \ldots, S_{25}\}' = \mathbb{C} I_s \tag{V.13}$$

in both sectors. It covers all 27,580 source and 16,502 vacuum basis directions. Exact rational determinants of the stored binary coefficients resolve the non-singleton contact-signature blocks. Since the mass and electric terms lie in the contact span, the witness applies at both endpoint masses.

For finite $\beta$, a faithful Gibbs state, strictly positive rates and the exact KMS construction, this scalar common commutant implies a unique stationary Gibbs density operator. The proof, including the Hamiltonian drift, appears in Appendix B.3; it uses stationary-state algebra rather than extrapolation from a small residual [19, 44]. The result concerns the declared stored operators. It supplies neither a Liouvillian decay gap nor a fixed-point bound for finite quadrature, and it does not measure mixing time. It must be reconsidered for omitted contacts, zero rates or the nonfaithful zero-temperature limit. Thus C2 accessibility is established conditionally, whereas numerical approach to equilibrium remains a separate question.

### V.2.5 Independent Full-Sector Gibbs Estimates

Equilibrium observables are estimated without using the open trajectories. For independent complex Gaussian probes with $\mathbb{E}[g_r g_r^\dagger] = I_s$, let $v_r = e^{-\beta(H_s - E_{\mathrm{sh}} I_s)/2} g_r$, without normalizing $v_r$. The estimators are

$$\widehat{\langle O \rangle}_{\beta,s} = \frac{\sum_{r=1}^{R_{\mathrm{p}}} v_r^\dagger O_s v_r}{\sum_{r=1}^{R_{\mathrm{p}}} v_r^\dagger v_r},$$
$$\widehat{Z}_s = e^{-\beta E_{\mathrm{sh}}} \frac{1}{R_{\mathrm{p}}} \sum_{r=1}^{R_{\mathrm{p}}} v_r^\dagger v_r. \tag{V.14}$$

The shift improves numerical scaling; it neither removes states nor changes the ensemble. Averaging normalized probes with equal weights would implement a different estimator. This construction is related to canonical thermal pure-state methods [39], with finite-sample ratio and logarithm biases assessed separately.

Table V.1 contains twelve estimates, each using two independent batches of 64 probes. The recorded errors are approximate statistical standard errors, not confidence intervals or total error bounds. Shared 32- and 64-probe subsets within a batch are not independent repeats. Imaginary-time action errors, ratio bias and model/representation uncertainty remain distinct from the displayed sampling errors.

For the high-mass source, connectivity changes from approximately 0.708 at $\beta = 0.5$ to 0.932 at $\beta = 2, 8$. For the low-mass source, the $\beta = 0.5$ value of about 0.066 exceeds the cold value near 0.0289. These estimates resolve canonical redistribution of structural weight, not bath-induced formation. Figure V.2 shows the sampled points without a phase-boundary interpolation. At cold high mass the stochastic error can approach floating-point resolution; this does not establish correspondingly many accurate physical digits. In

| Sector, mass | $\beta$ | $\langle H_0 \rangle$ | $P_{\text{any}}$ |
|---|---|---|---|
| Source, 4 | 0.5 | $-46.713185 \pm 0.142048$ | $0.707634 \pm 0.013530$ |
| | 2 | $-49.349562 \pm 0.000017$ | $0.931625 \pm 0.000029$ |
| | 8 | $-49.349755^{\dagger}$ | $0.931665^{\dagger}$ |
| Source, 0.5 | 0.5 | $-9.960159 \pm 0.030791$ | $0.065958 \pm 0.000957$ |
| | 2 | $-17.025532 \pm 0.027317$ | $0.030944 \pm 0.001860$ |
| | 8 | $-17.397844 \pm 0.000187$ | $0.028891 \pm 0.000163$ |
| Vacuum, 4 | 0.5 | $-52.510122 \pm 0.073214$ | $(6.948 \pm 0.509)10^{-7}$ |
| | 2 | $-53.464893783 \pm 5.03\,10^{-8}$ | $(3.75354 \pm 0.00027)10^{-7}$ |
| | 8 | $-53.464894314^{\dagger}$ | $(3.75337\,10^{-7})^{\dagger}$ |
| Vacuum, 0.5 | 0.5 | $-13.236054 \pm 0.073328$ | $0.00131097 \pm 0.00003609$ |
| | 2 | $-20.424459 \pm 0.004875$ | $0.00029838 \pm 0.00000269$ |
| | 8 | $-20.480996677 \pm 2.37\,10^{-10}$ | $0.000297007 \pm 6.65\,10^{-10}$ |

**Table V.1:** Independent equilibrium estimates at $\beta = 0.5, 2, 8$, in units $t_{\text{h}} = 1$. Errors are approximate statistical standard errors from $2 \times 64$ probes, excluding an established systematic floor. †: cold high-mass sampling errors are exceptionally small and are not interpreted as physical precision; their scales are stated in Appendix B.4. Vacuum connectivity uses the same geometric probe vertices and path projectors without external charges. These are equilibrium references, not measured reservoir-induced dynamics.

particular, the vacuum $m = 4, \beta = 8$ variance diagnostic $-4.54747 \times 10^{-13}$ is a cancellation residual, not negative physical variance. A mixed thermal state normally has nonzero energy variance, so a ground-state variance threshold is inappropriate.

### V.2.6 Finite-Temperature Trajectory Diagnostics

Sparse Monte Carlo wave-function evolution avoids storing a full density matrix [14]. Between two $H_s$ half-steps, the normalized state gives event probabilities $q_a = dt\,\|L_a\psi\|^2$. The sum is checked against $q_{\text{step}} \leq 0.05$ before sampling. A no-event branch uses $\widetilde{\psi} = \psi - dt(iG_{\text{KMS}}\psi + K\psi/2)$ and is normalized; a sampled event uses $L_a\psi/\|L_a\psi\|$. This state-dependent finite-step rule is weakly first order, not an exact linear completely positive trace-preserving (CPTP) map. CPTP evolution follows for the exact GKSL semigroup, not for arbitrary discretizations of a normalized trajectory rule.

Four C2 first steps were computed at $\beta = 0.125$, $\gamma = 0.02$ and $dt = 0.0025$, starting from the exact endpoints (Table V.2). Their coherent actions compare 16 and 32 quadrature nodes on the same initial half-step state. Absolute differences lie between $1.21 \times 10^{-9}$ and $1.75 \times 10^{-9}$; these are state-action comparisons, not operator norm or Gibbs-stationarity bounds. Every sampled branch has no event. The calculations therefore do not sample a C2 jump ensemble.

Restoring the recorded pre-step state and both random-number states reproduces each first step exactly in the tested environment. The closed-control differences range from $2.14 \times 10^{-16}$ to $1.82 \times 10^{-15}$. Replay of one step is not an interrupted long-trajectory continuation test. Likewise, the high-mass source first-step open-minus-closed change $P_{\text{any}} \simeq -7.78283 \times 10^{-6}$ is conditional on no event and is not an ensemble thermal-response

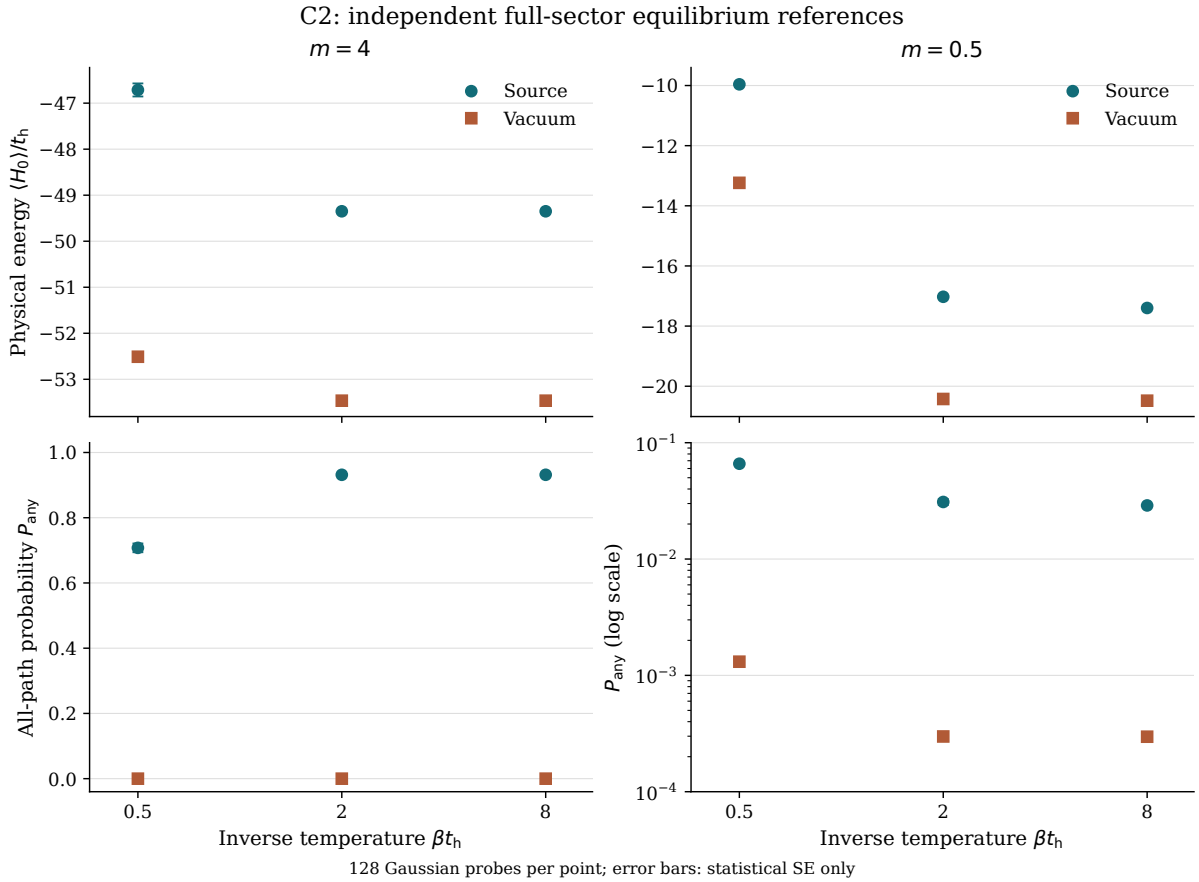



**Figure V.2:** Full-sector stochastic Gibbs estimates for source and vacuum at the two endpoint masses. Points show $\beta = 0.5, 2, 8$; vertical bars denote approximate statistical standard errors only. Invisible bars do not imply zero total error. No curve or transition temperature is fitted. The logarithmic low-mass connectivity panel resolves the vacuum reference. These temperatures differ from the $\beta = 0.125$ first-step calculation.

point. The Gibbs calculations at $\beta = 0.5, 2, 8$ cannot be used as equilibrium targets for this $\beta = 0.125$ pilot without a matched-temperature reference.

### V.2.7 Heat, Work and Ensemble Entropy

Physical energy is $E(t) = \mathrm{Tr}[\rho H_s(t)]$. For a justified time-dependent thermostat,

$$\frac{dE}{dt} = \underbrace{\mathrm{Tr}[\rho \dot{H}_s]}_{\dot{W}} + \underbrace{\mathrm{Tr}[H_s \mathcal{M}_t(\rho)]}_{\dot{Q}}, \qquad \Delta E = W + Q + r_{\mathrm{num}}. \tag{V.15}$$

Heat is positive into the system. The KMS coherent correction contributes:

$$\dot{Q} = \sum_a \mathrm{Tr}[H_s \mathcal{D}[L_a](\rho)] - i\,\mathrm{Tr}[H_s[G_{\mathrm{KMS}}, \rho]]. \tag{V.16}$$

The coherent contribution vanishes for an initial $H_s$ eigenstate up to numerical precision, but need not vanish later. The rates in Table V.2 are therefore not recovered from $\Delta E/dt$

| Sector, mass | $q_{\mathrm{step}}$ | Relative $G$ difference | $\dot{Q}/t_{\mathrm{h}}^2$ | Time (s) |
|---|---|---|---|---|
| Source, 4 | 0.001253639 | $3.5774\,10^{-7}$ | 0.114584083 | 109.35 |
| Source, 0.5 | 0.000971662 | $3.3927\,10^{-7}$ | 0.346097356 | 76.65 |
| Vacuum, 4 | 0.001133959 | $3.6443\,10^{-7}$ | 0.138563833 | 54.70 |
| Vacuum, 0.5 | 0.000995272 | $3.4222\,10^{-7}$ | 0.365283182 | 36.90 |

**Table V.2:** Four no-event steps at $\beta = 0.125$, $\gamma_a = 0.02$, $dt = 0.0025$, with the unrescaled 25-contact set. The $G$ column compares 16 and 32 nodes on the same state. Heat rates are unconditional generator expectations at the initial state, not the realized energy change of the sampled branch. Times refer to one measured step, not the complete diagnostic calculation. These values establish bounded action tests, not ensemble equilibration.

of a single no-event path. Algorithmic heat accounting also does not establish laboratory calorimetry.

For a faithful stationary Gibbs reference, Spohn's relation gives [37]

$$\sigma_{\mathrm{prod}} = -\operatorname{Tr}\{\mathcal{M}_\beta(\rho)[\log\rho - \log\rho_{\beta,s}]\} = \dot{S}_{\mathrm{vN}} - \beta\dot{Q} = -\frac{d}{dt}D(\rho\|\rho_{\beta,s}) \geq 0. \tag{V.17}$$

Here $S_{\mathrm{vN}} = -\operatorname{Tr}\rho\log\rho$ is the ensemble entropy. Averaging pure-trajectory entropies gives zero, not this quantity; bipartite entanglement is different again. Rank-deficient initial states require the appropriate full-rank limit and can have a singular initial entropy derivative. No C2 entropy-production or relative-entropy time curve is inferred from the present first steps.

If the instantaneous Gibbs state is justified during a drive, then at fixed $\beta > 0$

$$\frac{d}{dt}D(\rho\|\rho_{\beta,s}(t)) = -\sigma_{\mathrm{prod}} + \beta(\dot{W} - \dot{F}_{\mathrm{eq},s}), \qquad F_{\mathrm{eq},s} = -\beta^{-1}\log Z_s. \tag{V.18}$$

For fixed $\beta$, $\dot{F}_{\mathrm{eq},s} = \operatorname{Tr}(\rho_{\beta,s}\dot{H}_s)$. If $\beta$ varies, the formula instead uses this fixed-$\beta$ Hamiltonian derivative and adds $\dot{\beta}(E - E_{\mathrm{eq},s})$. A moving-target relative entropy is thus not generally monotone. The sector-conditioned internal- and free-energy differences,

$$\Delta U_\beta = \langle H_0\rangle_{\beta,\mathrm{src}} - \langle H_0\rangle_{\beta,\mathrm{vac}}, \qquad \Delta F_\beta = -\beta^{-1}\log\frac{Z_{\mathrm{src}}}{Z_{\mathrm{vac}}}, \quad \beta > 0, \tag{V.19}$$

are distinct, use a common energy-zero convention and inherit partition estimator uncertainty. At a single source separation neither is a string tension or a bulk deconfinement diagnostic.

### V.2.8 Structural Response and Its Interpretation

The primary operators are link Casimirs, matter occupations and configuration-level path projectors. With $F_\ell = 4E_\ell^2/3$, the all-path probability is not obtained by thresholding a mean flux image. The selected-path, alternative-only and no-path probabilities obey $P_{\mathrm{alt}} = P_{\mathrm{any}} - P_{\mathrm{sel}}$ and $P_{\mathrm{no}} = 1 - P_{\mathrm{any}}$. A shortest-path-only projector is a different observable and must be named as such. At C2 source 1, $P_{\mathrm{ff}} = 1 - 4E_1^2/3$ is a degree-one

identity, not an independent screening witness. Joint support $P_{\mathrm{both}}$ is the expectation of a joint projector, not the product of two means or an external-source–matter singlet probability.

The staggered pair expectation $N_{\mathrm{pair}} = (M + 12)/2$ measures particle–hole dressing, not individual creation events. Source-minus-vacuum link and matter profiles use the two declared $D_v$ patterns and matched Hamiltonian parameters. The signed off-corridor ratio of Eq. (III.1) is interpreted only when its total connected denominator is reliably positive; an unresolved denominator does not define a width by regularization.

The corridor observable is

$$C_E(t;\beta,\gamma) = \frac{1}{|\mathcal{C}|}\sum_{\ell\in\mathcal{C}}\left(\langle E_\ell^2\rangle_{\mathrm{src}} - \langle E_\ell^2\rangle_{\mathrm{vac}}\right). \tag{V.20}$$

To isolate a reservoir effect rather than a quench response, define

$$\begin{aligned}\delta_{\mathrm{bath}}O_s(t) &= \langle O\rangle_{\mathrm{open},s}(t) - \langle O\rangle_{\mathrm{closed},s}(t),\\ \Delta_{\mathrm{bath}}O(t) &= \delta_{\mathrm{bath}}O_{\mathrm{src}}(t) - \delta_{\mathrm{bath}}O_{\mathrm{vac}}(t).\end{aligned} \tag{V.21}$$

The notation $\delta_O^{\mathrm{bath}}$ denotes the same within-sector contrast. Initial states, time grids and Hamiltonian protocols are matched. Covariance enters only for genuinely paired samples; independent samples require independent-variance propagation. A change relative to the initial state alone is not a reservoir effect.

Corridor depletion, alternative-path redistribution and matter screening answer different questions and must be interpreted jointly. Gauge preservation does not guarantee structural survival. Scalar contacts in an SU(2) model also do not establish a uniquely non-Abelian enhancement: that claim requires matched reduced-model controls and fusion-sensitive observables.

### V.2.9 Established Results and Physical Scope

The contribution is a sector-resolved bridge from exact ground states to independent equilibrium references and explicitly defined thermal dynamics. A matched short-ensemble calculation at the same temperature is the next physical comparison, followed by controlled extension of its observation window. The present results do not determine thermal melting, a flux-tube lifetime or a QCD critical temperature.

## V.3 Thermal-Response Protocols and Limits

### V.3.1 Stationary, Post-Preparation and Driven Contact

Stationary holds start from each exact endpoint under its own Hamiltonian. Post-preparation contact instead uses a specified closed checkpoint and matched isolated continuation. Breaking contact follows the sudden $m : 4 \to 0.5$ quench; formation contact accompanies the $m : 0.5 \to 4$ ramp and its moving equilibrium target. These are four distinct protocols.

| Calculation | Established result and limit |
|---|---|
| Physical representation | Complete C2 sectors and four transferred exact endpoints; no energy cutoff |
| Exact small-system tests | Davies, Gaussian-KMS and modular-KMS references with distinct parameter ranges |
| C2 common commutant | Scalar for the stored operators; conditional uniqueness of the ideal faithful Gibbs state |
| Canonical structure | Twelve independent finite-temperature estimates with statistical, not total, errors |
| C2 sparse evolution | Four finite-temperature no-event first steps; bounded action and replay comparisons |
| Reservoir response | Ensemble equilibration, temperature/rate response and structural lifetimes remain unmeasured |
| Experimental transfer | Requires a separately derived reservoir and calibrated coupling and clock |

**Table V.3:** Representation, equilibrium and algebraic results are separated from numerical action tests and physical open-system response. Neither static order nor one sampled trajectory branch establishes thermal survival.

Exact-state holds are independent of MPS approximations and dynamical preparation. Each closed checkpoint requires consistent transfer of its representation, norm, charges, Hamiltonian action and observables. The C2 $t = 20$ choices are retrospective, not preregistered. Their conversion and thermal response are separate from the closed real-time protocols of Sec. IV.1.

A fixed final-Hamiltonian bath is not an instantaneous ramp bath. Adiabatic Davies dynamics requires reservoir-memory, driving, eigenbasis and secular timescale justification [2]; a collision model requires a microscopic derivation. An instantaneous KMS sampler is instead an algorithmic moving-target thermostat, not either physical reservoir. Bath switching, ramp completion and post-preparation contact have distinct start times.

### V.3.2 Controls, Convergence and Statistical Decisions

The comparison requires a closed $\gamma = 0$ control, energy-preserving dephasing, gauge-aware energy exchange, matched source/vacuum evolution, temperature and contact-rate variation, and an independent small-system reference. Deliberately gauge-breaking contacts must be tested in an enlarged space; projection cannot serve as a leakage test. Changing $\gamma$ to increase event frequency changes the kinetics and does not produce additional samples of the same physical experiment.

Independent trajectories, not their correlated output times, are the replication units. At least two independent seed blocks and $N_{\mathrm{traj}}$ versus $2N_{\mathrm{traj}}$ comparisons accompany $dt$ refinement, action/quadrature refinement and a common measurement grid. Tensor-network variants additionally require bond and cutoff controls. Missing, duplicate and failed trajectories must remain accounted for rather than disappearing from an ensemble average. The distinct numerical and statistical requirements are collected in Appendix B.9.

A confidence interval containing zero is inconclusive unless the interval, including the systematic allowance, lies inside an independently chosen physical-equivalence margin. An unresolved estimate, a controlled physical null and a failed numerical calculation

are different outcomes. A structural lifetime requires a specified observable, threshold, persistence rule and crossing convention. Oscillatory dips or revivals do not determine a unique decay time. If no persistent crossing is resolved, the result is right-censored beyond the resolved window $t_{\mathrm{resolved}}$, which can be shorter than the last simulated time. No such C2 lifetime is reported.

# Chapter VI

# Cold-Atom Implementation

This chapter asks which microscopic controls could reproduce the two-rishon honeycomb Hamiltonian and which observables would distinguish source-supported electric response from matter screening. Operator reduction, state preparation, readout and laboratory calibration are kept as separate requirements.

## VI.1 Honeycomb Encoding and Implementation Tests

The experimental target in this section is the honeycomb Hamiltonian of Eq. (II.8). It has four-state fermionic matter sites, six-state links with exactly two rishons, classical $q^3 = +1, -1$ sources and the independent coefficients $t_{\mathrm{h}}, g^2, m, \lambda, \mu$; the present parameter choice is $\kappa = 0$. The two-rishon link representation fixes the local operator algebra and the microscopic couplings. A one-rishon encoding defines a different model and does not determine the coefficients or laboratory time scale of this construction. Quantum-link proposals motivate the finite-mode strategy [6, 42, 24, 16]; they do not by themselves establish the controls assumed in the finite-patch construction.

Figure VI.1 summarizes the encoding, reduction and measurement map of the two-rishon honeycomb realization.

### VI.1.1 Exact Matter and Link Encoding

The matter site uses two fermionic modes, $\psi_\uparrow, \psi_\downarrow$, with Fock basis

$$|0\rangle, \quad |\uparrow\rangle, \quad |\downarrow\rangle, \quad |\uparrow\downarrow\rangle. \tag{VI.1}$$

A full oriented link uses the four modes $(c_{L\uparrow}, c_{L\downarrow}, c_{R\uparrow}, c_{R\downarrow})$. Restricting their total occupation to two gives

$$\dim \mathcal{H}_\ell = \binom{4}{2} = 6. \tag{VI.2}$$

With the least-significant occupation bit assigned to $L\uparrow$, the six target columns select bit strings with integer labels

$$(12, 5, 9, 6, 10, 3) \quad \text{and phases} \quad (+, +, -, -, +, +). \tag{VI.3}$$

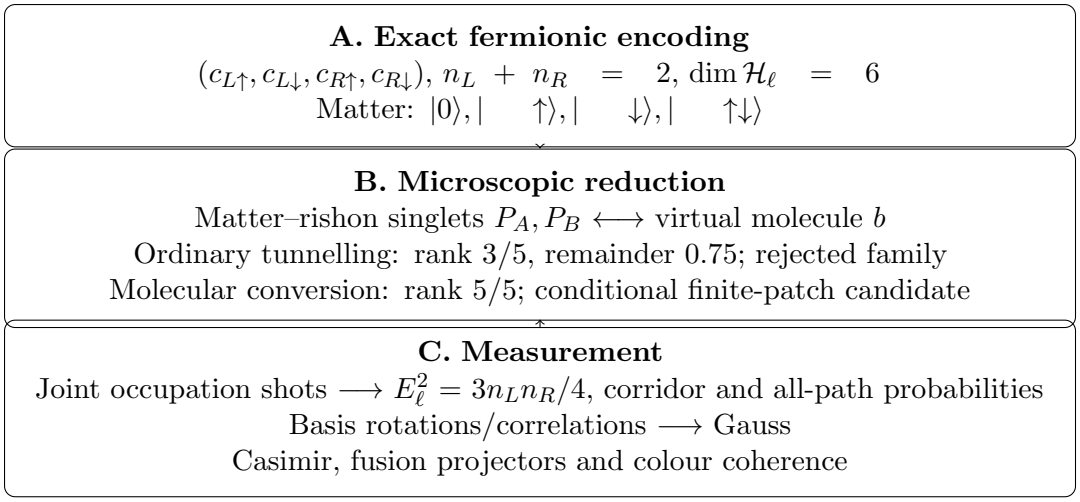



**Figure VI.1:** The present two-rishon route. The molecular reduction assumes independently available conversion, counterterm and target-coefficient controls. Occupation readout and coherent colour/fusion readout require different measurements.

These phases are not cosmetic: changing either negative sign produces a Frobenius residual of 2 in the link transporter.

Let $V_m$ and $V_\ell$ denote the matter and link isometries. The atomic operators are projected according to

$$O_{\text{target}} = V^\dagger O_{\text{atom}} V, \qquad V^\dagger V = \mathbb{1}. \tag{VI.4}$$

All 19 local operator matrices were reconstructed from occupation-bit fermions: matter annihilators and parity, $Q^a$, link-end occupations, $L^a, R^a$, $E^2$, and all four $U^{ab}$. Across 82 local relations the maximum residual was zero in the arithmetic used. The transporter was reconstructed as

$$U^{ab} = \epsilon_{ac}\epsilon_{bd} c^\dagger_{L,c} c_{R,d}, \tag{VI.5}$$

including its orientation and fermionic signs. Reversing $U$ gives the same adversarial residual of 2. The exact occupation identity $E^2 = \frac{3}{4} n_L n_R$ was recovered together with the fixed condition $n_L + n_R = 2$. This establishes the local encoding, rather than a statement about how a laboratory generates the hopping operator.

The local constraints were then assembled without importing a vertex operator from the target implementation. Degree-one, degree-two and trivalent stars were tested for both link-end orientations and for $q^3 = 0, +1, -1$, giving 18 cases. The maximum residual among the Gauss algebra, $D_v$ commutators, penalty expansion and neutral transverse-commutation tests was $2.0 \times 10^{-15}$. A neutral star has zero Gauss-penalty floor. A classical source selects the colour axis, retains $G^3$, and has the expected local floor 3/4; its transverse commutator is nonzero. Thus the sources are c-number terms in $(\boldsymbol{G} - \boldsymbol{q})^2$, not quantum impurity doublets.

## VI.1.2 Microscopic Control Families and Effective Reduction

The target hopping must be generated rather than inserted after projection. The first tested microscopic family contained only colour-preserving matter tunnelling $J_m$, rishon

tunnelling $J_r$, and a common imbalance penalty $\Delta(D_A^2 + D_B^2)$. In the 96-dimensional fixed-link space its selected $D_A = D_B = 0$ subspace has dimension 18 and $PVP = 0$. The second-order Schrieffer–Wolff operator was evaluated as

$$H_{\mathrm{SW}}^{(2)} = -PVQ\,(QH_0Q)^{-1}QVP, \tag{VI.6}$$

and compared with the exact low band using the polar des Cloizeaux map. For $\Delta = 8, 16, 32, 64$, the exact-to-second-order Hamiltonian error falls with fitted power $-2.965$, while the spectral leakage bound falls with power $-1.945$; the largest absolute log-fit residual is 0.028. These exponents describe this four-point finite-patch ladder, not a universal scaling law.

The reduction nevertheless fails as a target realization. After fitting the target operator basis and an identity shift, the non-target Frobenius remainder is 0.75 of the centered effective Hamiltonian at every tested $\Delta$. More fundamentally, the finite-difference coefficient Jacobian has rank three for the five target coefficients. The three controls cannot independently generate all five target coefficients. This obstruction applies to the specified tunnelling family; assisted processes require their own coefficient map. The obstruction is insufficient independent controls, not a no-go theorem for every atomic platform.

A separate audit applies the colour-independent matter-to-incident-rishon tunnelling proposed in Ref. [6] to the registered two-rishon isometry. On one oriented edge its perturbation is

$$V_{\mathrm{vert}} = J_A \sum_a (c_{L,a}^\dagger \psi_{A,a} + \mathrm{h.c.}) + J_B \sum_a (c_{R,a}^\dagger \psi_{B,a} + \mathrm{h.c.}). \tag{VI.7}$$

This operator preserves the local SU(2) Gauss algebra exactly. In the present orientation, however, it changes the registered $D_A, D_B$ labels. Its second-order projection contains same-vertex self energies and the opposite transporter orientation, while the fitted coefficient of the target hopping in Eq. (II.8) is zero. Even after granting four independent diagonal controls for $g^2, m, \lambda, \mu$, the coefficient Jacobian therefore has rank four rather than five. This excludes the literal vertex-tunnelling family for the current isometry; it does not exclude an atomic species or a staggered particle–hole encoding. The latter would be a new model realization and would require a new 19-matrix isometry certificate, fermionic-sign tests and patch validation.

A second family introduces a detuned bosonic molecular mediator for each link. At one end define the colour-singlet pair annihilator

$$P_A = \sum_{a,c} \epsilon_{ac} c_{L,c} \psi_{A,a}, \qquad P_B = \sum_{b,d} \epsilon_{bd} c_{R,d} \psi_{B,b}. \tag{VI.8}$$

For a mediator mode $b$, the microscopic edge Hamiltonian is

$$\begin{aligned} H_{\mathrm{med}} = {} & \Delta b^\dagger b + g_A(b^\dagger P_A + P_A^\dagger b) + g_B(b^\dagger P_B + P_B^\dagger b) \\ & + \frac{g_A^2}{\Delta} P_A^\dagger P_A + \frac{g_B^2}{\Delta} P_B^\dagger P_B + H_E + H_m + H_G + H_D. \end{aligned} \tag{VI.9}$$

Eliminating the molecular mode gives

$$H_{\mathrm{hop}}^{(2)} = -\frac{g_A g_B}{\Delta}(P_A^\dagger P_B + P_B^\dagger P_A), \tag{VI.10}$$

which equals the target $-t_{\mathrm{h}}(\psi_A^\dagger U \psi_B + \mathrm{h.c.})$ element by element. The displayed singlet-density counterterms cancel the two self-energy contributions. Omitting them or changing the relative conversion phase is detected by dedicated negative controls; neither ingredient is hidden in a post-hoc identity shift.

The exact conserved edge sector is $n_{\mathrm{rishon}}+n_{\mathrm{molecule}} = 2$, including molecular occupations 0, 1 and 2. It has dimension 176, of which 96 states form the target subspace. The conversion commutes exactly with the local Gauss and $D_v$ generators. It conserves matter-plus-molecule number rather than bare matter number, so virtual molecular occupation is a physical leakage channel. At fixed target $t_{\mathrm{h}} = g_A g_B/\Delta = 1$, exact reductions at $\Delta = 512, 1024, 2048, 4096$ give a finite-detuning Hamiltonian error proportional to $\Delta^{-0.9994}$, with maximum log-fit residual $1.7 \times 10^{-4}$. With independent bare controls for $g^2, m, \lambda, \mu$, the scaled Jacobian has rank five and condition number $\sqrt{3}$.

This favorable rank is conditional. It assumes coherent singlet pair conversion, separately adjustable singlet-density counterterms, electric and staggered offsets, spin exchange, $D_v$-density interactions, classical-source addressing, and the required occupation and coherence readout. Those controls have not been demonstrated together for a selected platform. The calculation establishes a microscopic Hamiltonian candidate and its controlled finite-patch reduction, not a laboratory working point.

### VI.1.3 Star, Loop, and Source-Pair Validation

The operator comparison was extended beyond one edge. First, an independent global occupation-bit construction was compared with the local-matrix exact diagonalization engine for the eight-mode matter–link hopping and on a closed hexagon. The neutral and classical-source sectors have dimensions 2442 and 1995. In both sectors the full Hamiltonian, two low eigenvalues and short-time states agree within a maximum residual of $5.2 \times 10^{-11}$. These tests include the full fermionic ordering, the source constants, $D_v$, $N$, and $\kappa = 0$.

The mediator model was then enumerated on a trivalent star and on neutral and source-bearing hexagons. Table VI.1 summarizes the largest calculation. The exact second-order operator agrees with the target on each complete declared sector. Higher-order dynamics was tested at $\Delta = 128, 256, 512$ for $t = 0, 0.05, 0.10$ using three fixed probes: the target ground state, a seeded random state and a product-basis state. At the largest detuning and latest time, the largest virtual-sector leakage across all patches and probes is 0.053 and the largest projected-state error is 0.047. These are fixed engineering checks on three probes and a short window; they are not worst-case many-body channel bounds.

On the neutral star the probe restriction can be removed exactly. Propagating all 126 target inputs over the same short-time grid gives, at $t = 0.1$, a maximum projected-channel operator norm 0.0624 and a maximum worst-case molecular/fixed-rishon leakage 0.0369 across $\Delta = 128, 256, 512$. This is a sector-wide bound for the star only; the 2442- and 1995-state hexagon sectors retain their three-probe scope.

| Patch | Target dim. | Micro dim. | $\lvert\Delta E_0\rvert$ | $\lvert\langle\psi_0\rvert P\lvert\Psi_0\rangle\rvert^2$ |
|---|---|---|---|---|
| Trivalent star | 126 | 636 | 0.0240 | 0.9800 |
| Neutral hexagon | 2442 | 38372 | 0.0285 | 0.9646 |
| Source hexagon | 1995 | 31352 | 0.0297 | 0.9630 |

**Table VI.1:** Exact sparse-patch comparison at $\Delta = 512$, with $g_A = g_B = \sqrt{\Delta}$ and target $t_{\mathrm{h}} = 1$. The overlap is reported because the relevant target ground state is isolated; microscopic eigenpair residuals are at most $2.0 \times 10^{-11}$. The numbers validate the stated mediator model on finite patches and do not constitute experimental fidelities.

### VI.1.4 Preparation, Readout, and Error Separation

Preparation has three logically distinct stages: fixing two rishons per link and the intended $N, D_v$ labels; selecting the neutral or classical-source Gauss block; and preparing the interacting endpoint or protocol state. A simple product state in the correct sector is not assumed to be a high- or low-mass ground state. In a formation experiment the classical sources are present from the start and the mass is ramped from low to high. A unitary vacuum evolution does not create a different source sector.

The data-independent preparation test makes this separation operational on a neutral trivalent star. The lexicographically fixed product initializer has the correct two-rishon, $N$, $D_v$, and $G^3$ labels, but its overlap with the three-dimensional zero-hopping ground space is zero. Sector registration is therefore not gauge or endpoint preparation. A separate smooth conversion ramp was run in the 126-state target and 636-state microscopic sectors at $\Delta = 128$, using both a zero-hopping interacting ground-state probe and the product probe. For durations $0.25 \leq T t_{\mathrm{h}} \leq 2$, the largest projected microscopic–target state error is 0.0618 and the largest molecular leakage is 0.0676. The target final-ground-space fidelity reaches only 0.215 for the interacting probe and 0.033 for the product probe over this ladder. This characterizes an executable local primitive and supplies a negative preparation control; without the upstream endpoint states and gaps it does not certify an endpoint or formation schedule.

Occupation snapshots directly give matter $n_\uparrow, n_\downarrow$, link-end $n_L, n_R$, total matter number and $D_v$. Within the fixed two-rishon space they give

$$E_\ell^2 = \frac{3}{4} \mathbb{1}_{n_L = n_R = 1}. \tag{VI.11}$$

The diagonal decoder for $n_L, n_R, E^2, L^3, R^3$ reproduces the target matrices exactly. Each image supplies a joint $n_L n_R$ sample: $\langle n_L n_R \rangle$ cannot be replaced by the product of separate means $\langle n_L \rangle \langle n_R \rangle$. The full Gauss Casimir requires rotated-basis matter–link and link–link correlations; fusion projectors and off-diagonal $U$ coherence likewise require tomography or specified basis rotations. Consequently an $E^2$ image alone cannot establish non-Abelian coherence.

An explicitly synthetic detector test used independent 1% loss and 2% bit misclassification on 20000 snapshots of a fixed six-state distribution. The true $E^2$ mean was 0.525; the raw decoder returned an expectation 0.4759, a bias of $-0.0491$, and standard error

0.00255. Fixed-number postselection retained 90.7% of snapshots and reduced the bias to $-9.3 \times 10^{-5}$. These values validate the error-model pipeline, not an experimental detector. Raw retention and postselected estimates must both be reported because postselection changes the estimator as well as its sample size.

Adding 0.5% link-correlated loss and 0.5% end cross-talk gives raw bias $-0.0515$, 90.1% retention and residual postselected bias $+7.0 \times 10^{-4}$. The correlated model leaves a signal of 0.05 bias limited; under the independent model that signal would require about $6.18 \times 10^5$ shots. These synthetic single-estimator tests cannot replace a calibrated covariance model for paired source-minus-vacuum data. The tested signal scales 0.01 and 0.025 are already limited by raw bias. Single-shot occupation samples must be distinguished from corridor probabilities and source–vacuum differences, which are ensemble quantities.

The physical error budget remains decomposed into upstream numerical uncertainty, effective-Hamiltonian remainder, coefficient calibration, unwanted operators, fixed-rishon and Gauss/$D_v$ leakage, ramp resolution, inhomogeneity, decoherence and loss, and readout statistics. The finite-patch engineering tolerances used above are not substituted for a signal-dependent budget, and correlated systematics are not automatically combined in quadrature.

### VI.1.5 Numerical References and Nominal Statistical Planning

The exact C2 endpoints, profiles and dimensionless trajectories are numerical references, not evidence that an apparatus prepares the corresponding states. The static source and vacuum vectors at $m = 4$ and $m = 0.5$ pass their sector and residual checks. Experimental transfer still requires a microscopic preparation protocol and an apparatus-level error budget.

These exact probabilities do determine a nominal ideal-detector planning oracle. For equal numbers $N_{\mathrm{arm}}$ of independent source and vacuum shots and a Bernoulli observable with known true probabilities $p_s, p_v$, the asymptotic equal-arm count at standardized signal-to-noise ratio six is

$$N_{\mathrm{arm}} = \left\lceil 36 \frac{p_s(1-p_s) + p_v(1-p_v)}{(p_s - p_v)^2} \right\rceil . \tag{VI.12}$$

Table VI.2 applies this formula to the exact endpoint projectors. It assumes an ideal detector, independent arms and no systematic floor. It supplies neither finite-sample significance and power nor multiple-testing control; the entries are nominal oracle counts per arm, not laboratory shot budgets. Applying the same calculation to $a_\ell = 4E_\ell^2/3$, the four high-mass corridor links each require three nominal shots per arm, whereas the low-mass corridor values range from 76 to 73,604 because their source–vacuum signals are strongly link dependent. These are marginal link-by-link calculations; they do not set the joint budget for a complete profile or a multiple-link decision. Moreover, $P_{\mathrm{both}} = 0$ in the vacuum sector is structural. Its entries quantify sampling of a source-sector event and are not an independent matched-null test of the breaking mechanism.

The exact C2 breaking and formation curves provide the same kind of dimensionless planning reference. At the final breaking time $t = 20$, the source–vacuum signals are 0.109409 for $P_{\mathrm{any}}$ and 0.122192 for $P_{\mathrm{both}}$, corresponding nominally to 297 and 259 shots per

| $m$ | Projector | $p_s - p_v$ | $N_{\mathrm{arm}}^{(\mathrm{SNR}=6)}$ | Scope |
|---|---|---|---|---|
| 4 | $P_{\mathrm{any}}$ | 0.931664719 | 3 | high-mass endpoint |
| 4 | $P_{\mathrm{both}}$ | 0.000284349 | 126,569 | high-mass endpoint |
| 0.5 | $P_{\mathrm{any}}$ | 0.028606772 | 1,248 | low-mass endpoint |
| 0.5 | $P_{\mathrm{both}}$ | 0.159918945 | 190 | low-mass endpoint |

**Table VI.2:** Asymptotic source–vacuum SNR oracle obtained from exact C2 endpoint probabilities. The calculation omits finite-sample power, loss, misclassification, correlated drift, postselection and calibration uncertainty; these counts cannot be interpreted as experimental resources or confidence guarantees.

arm. At the end of the duration-four formation ramp the source values are 0.477103 and 0.016219, corresponding nominally to 40 and 2,184 shots per arm. These are descriptive exact-cluster numbers. The independent persistence criterion for formation fails: its minimum lower margin is $-0.104063$ over the duration-20 hold. Larger-patch convergence, a calibrated waveform and the laboratory time map remain absent; the C2 ramp itself is reproduced by full-rank Gauss-exact TDVP (Sec. IV.1.4), which validates the C2 curves but not wider patches, so these numerical references do not certify a laboratory formation mechanism or dynamical transfer.

The thermal results provide independent finite-$\beta$ equilibrium estimates and bounded first-step action tests, but no converged finite-time ensemble, mixing rate, structural lifetime or experimental reservoir. Consequently no atomic endpoint preparation, laboratory waveform, physical clock or thermal coupling is supplied by these numerical references.

## VI.2 QGP-Motivated Correlation Readouts

The measurement correspondence connects source-supported electric flux to matter screening and the lifetime of the resulting structure. On the honeycomb graph, source-minus-vacuum path observables require specified orientations, path ensembles, distances and time windows. Matter response complements the Casimir image; rotated-basis fusion measurements test colour recoupling. A thermal-survival measurement additionally requires the bath or equilibrium target of Sec. V.2.

Temperature- and contact-dependent honeycomb measurements must retain the sector and operator controls of Chapter V. The QGP connection and its continuum limitations are developed in the Conclusion.

## VI.3 Experimental Realization and Scope

The honeycomb analysis establishes an exact local encoding, excludes an ordinary-tunnelling family with three independent controls and identifies a conditional molecular-mediated Hamiltonian. Its physical realization requires coherent conversion, singlet-density counterterms and independent electric, mass, Gauss and imbalance couplings.

The atomic candidates include $^{173}$Yb, $^{171}$Yb and $^{87}$Sr. Coherent orbital Feshbach

molecules, tunable spin exchange and 99.96(1)% single-atom imaging of $^{173}$Yb provide relevant experimental ingredients [8, 34, 1]. Spin-resolved fermionic quantum-gas microscopy of $^{87}$Sr now resolves all ten nuclear-spin occupations in a single realization [20]. Coherent mapping among nuclear-spin, clock and Rydberg qubits in $^{171}$Yb, including many-body entanglement and loss-detectable readout, strengthens a distinct digital route [35]. These demonstrations do not combine the four-mode link, oriented conversion, source addressing and independent counterterms in one apparatus. Two-dimensional U(1) architectures offer related approaches to plaquette dynamics and gauge protection [11, 33, 32]; the SU(2) fusion structure requires additional encoding and readout. Tensor-network resource scaling is also an independent limitation beyond one dimension [28].

### VI.3.1 Operator Validation and Physical Interpretation

Consistency of the model specifications was checked across 23 geometry representations and 18 deliberately inconsistent inputs. Together with the local algebra, star and hexagon comparisons, these tests establish the finite-patch encoding and reduction. Table VI.3 separates those results from the physical assumptions needed in an experiment.

| Component | Quantitative basis | Physical interpretation |
|---|---|---|
| Local encoding | 19 matrices; 82 relations | Exact representation of local operators |
| Star and loop tests | 18 stars; two loop sectors | Agreement on the tested geometries |
| Ordinary tunnelling | Rank 3/5; remainder 0.75 | Insufficient independent couplings |
| Incident-rishon tunnelling | SU(2) exact; target $t = 0$; rank 4/5 | Current-isometry/$D_v$ obstruction; not species-wide |
| Molecular mediator | Rank 5/5 | Conditional low-energy construction |
| Star dynamics | All 126 initial inputs | Short-time sector-wide error bound |
| Preparation | Product and interacting probes | Sector initialization differs from ground-state preparation |
| Readout | Exact decoder; synthetic detector errors | Calibration and correlated noise remain essential |
| Static targets | Exact C2 endpoints and profiles | Reference accepted; formal four-role transfer remains incomplete |
| Driven protocols | Exact C2 descriptive curves | Formation persistence fails and no production transfer is licensed |
| Thermal contact | Engine and zero-time scalar oracles | No finite-time thermal response or reservoir transfer is certified |
| Laboratory scales | Dimensionless model parameters | Require calibrated couplings and coherence times |

**Table VI.3:** Operator identities, finite-patch reductions and experimental assumptions. Numerical overlaps and synthetic detector tests are not experimental fidelities or calibration measurements.

### VI.3.2 Equilibrium Preparation and Driven Protocols

An equilibrium comparison uses source and vacuum states at both high and low mass. Their energies, physical variances, Gauss and $D_v$ diagnostics, link-resolved $E^2$, matter density, corridor profiles and excitation gaps specify the target. Exact C2 vectors provide these references, and Gauss-exact DMRG reproduces them at both masses (Table II.2); the exact real-time curves of Sec. IV.1 provide the corresponding dynamical targets.

A microscopic preparation must additionally reproduce the relevant state or ground subspace. Its accuracy is measured through energy shifts, overlap, observable deviations and leakage relative to the source–vacuum signal. Coefficient tolerances, inhomogeneity and coherence time enter separately. Exact target states therefore establish a numerical reference, not the fidelity with which an atomic apparatus prepares it.

Formation begins with sources already present and ramps the low-mass state toward high mass. Breaking starts from a high-mass source state and quenches to low mass, with a matched vacuum evolution. In both cases target and microscopic trajectories must be compared through connected electric response, matter redistribution, path connectivity and

a colour or fusion channel. Ramp fidelity, leakage and gauge drift distinguish preparation error from subsequent dynamics.

The exact results of Sec. IV.1 provide dimensionless benchmarks. A laboratory prediction also requires a complete waveform, calibrated coupling scale and an error budget for the microscopic evolution. If $J_{\mathrm{cal}}$ is the energy unit used to nondimensionalize a protocol, then

$$t_{\mathrm{lab}} = \frac{\hbar\tau}{J_{\mathrm{cal}}}. \tag{VI.13}$$

This calibrated scale is not automatically any bare tunnelling or conversion amplitude of the microscopic model. Values of $J_{\mathrm{cal}}$, $J_{\mathrm{cal}}/h$ and $J_{\mathrm{cal}}/\hbar$ require distinct conventions. No millisecond duration follows without calibration. Errors must be assessed relative to the intended signal; when that signal vanishes, an absolute equivalence tolerance is needed.

### VI.3.3 Thermal Coupling and Calibrated Readout

The stationary Davies construction in Sec. V.2 is a phenomenological reference, whereas the Gaussian and modular KMS constructions are algorithmic thermostats. Neither is a reservoir already realized by the preparation hardware. Experimental transfer requires a separate derivation of the effective system–reservoir operators, spectrum and coherent correction, with independent calibration of charge preservation, loss and leakage. Exact encoding, canonical contacts and KMS identities alone do not determine a laboratory temperature or relaxation clock. Structural readout can be compared quantitatively only after the ensemble, coupling normalization and laboratory time map are specified.

Temperature and coupling dependence, heat exchange, physical leakage and observable durability must then be evaluated together. The physical target is a window in $(T, \gamma, t)$ over which formation and breaking remain distinguishable. Such a window is not determined by gauge preservation alone: a reservoir preserving only the SU(2) Gauss block can still change $D_v$ labels, while even the fully $D_v$-preserving contacts used here can change the flux profile. The accessible ensemble must be specified in either case.

Readout requires colour-resolved detection matrices, number resolution and calibration of correlated loss, cross-talk and shared source/vacuum errors. Raw and postselected estimators, retention, bias correction and paired-difference covariance determine the shot count at a specified confidence level. The synthetic tests above quantify estimator behaviour under stated noise models; they do not replace detector calibration. Fusion or colour-coherence readout is needed in addition to an $E^2$ image to probe non-Abelian information.

### VI.3.4 Atomic Platforms and Experimental Outlook

The complementary demonstrations in $^{173}$Yb make this species the strongest current analog follow-up candidate for the molecular route. This is a prioritization for further design, not a species selection. The $^{87}$Sr microscope strengthens spin-resolved readout and the $^{171}$Yb result strengthens a digital/tweezer route, but neither implements the present fermionic occupation encoding. The integrated-platform gate therefore remains closed. The central experimental challenge is to combine the four-mode link, oriented singlet

conversion, two independent counterterms, five target coefficients, classical source control and occupation/fusion readout in one geometry. The present calculation does not specify an integrated apparatus or a calibrated parameter set in Hz.

A quantitative implementation must identify internal states, trapping geometry, $g_A, g_B, \Delta$, counterterm generation, $g^2, m, \lambda, \mu$, loss and decoherence, the effective time scale and the measurement sequence consistently. These requirements delimit the experimental interpretation without weakening the operator result: the local fermionic encoding is exact, the ordinary-tunnelling family is insufficient, and the molecular mediator realizes the desired low-energy interaction under the stated control assumptions.

# Chapter VII

# Conclusion

The thesis establishes a controlled finite-patch account of static and driven source response in a two-rishon SU(2) quantum-link model on an open honeycomb lattice. Its central quantitative result is the C2 separation scan in pre-specified conserved-label sectors. At $m = 4$, the potential rises nearly linearly over graph distances $R = 2, 4, 6$, giving a finite-distance force $1.03 \pm 0.01\, t_{\mathrm{h}}/a$ per graph link, approximately 69% electric and within about 4% of the stated second-order strong-coupling estimate. At $m = 0.5$, the force softens by 47–53%, its electric contribution becomes negative, connected flux collapses and flux-free endpoint matter grows. The remaining light-mass cost is conditional on the registered corridor labels; it is not an asymptotic string tension or evidence against screening. On the wider C4 and C8 patches, whose corridor coordination differs from C2, heavy-mass variational states differ from the exact connectivity and source potential of the relocated C2 sources by $+0.5\%$ and $-1.0\%$, with bounded numerical support on C8. At $m = 2$ the capped C4 and C8 source states give larger connectivity, 0.775–0.777, than the relocated C2 sources, 0.718; this difference is reported descriptively. Exact coupling scans on C2 bring the half-connectivity mass close to the classical string/one-pair crossing at strong electric coupling, slightly below it at $g^2 = 6$ and 8 against the registered expectation, and the heavy-mass force towards the strong-coupling estimate. By the registered criterion, which uses the $[2, 4]$ force and the $R = 4$ connectivity, the heavy-mass response is insensitive to the plaquette coupling up to $\kappa = 2$; the $[4, 6]$ secants of the $R = 6$ sectors are not covered by it and fall by 5–11% at $\kappa = 2$. At larger plaquette coupling, or lighter mass, plaquette-dominated ground states appear, and a source–vacuum difference across the two branches is not a potential.

The real-time results connect this static response to controlled preparation. The $m = 4 \to 0.5$ quench suppresses electric support and source connectivity while transferring weight to flux-free endpoint classes. The response remains oscillatory and energy conserving. The inverse ramp increases connectivity and electric response, but the post-ramp ground-state population is incomplete and the hold does not support persistent formation at the tested speed. Exact physical-sector propagation and full-rank Gauss-exact TDVP agree at the stated precision. This validates C2 operators and integrators, not efficient compression or wider-patch convergence.

The non-Abelian content is represented explicitly through SU(2) link matrices, Gauss intertwiners and local fusion projectors. The measured recoupling populations are state-

dependent structural information, but the breaking signal itself is not uniquely non-Abelian. A model-exclusive dynamical claim would require a matched reduced-model control and fusion-resolved evolution.

The thermal contribution supplies complete physical sectors, independent Gibbs estimates, Davies and KMS reference tests, and a conditional uniqueness result for the declared contact algebra. Four finite-temperature first steps test bounded numerical actions. No converged open ensemble, structural lifetime or thermal deconfinement transition is measured. Likewise, the QGP connection is methodological: the model resolves colour-field support, screening and gauge-aware relaxation, but it is neither SU(3) QCD nor a continuum plasma simulation.

The microscopic analysis gives an exact local fermionic encoding and rejects two insufficient tunnelling families for the current isometry. A molecular singlet-pair mediator produces the target hopping conditionally when the counterterms and remaining couplings are independently controllable. Star and hexagon tests quantify finite-patch errors, while preparation and synthetic detector studies expose the remaining experimental requirements. The result is a conditional Hamiltonian construction, not a completed apparatus.

The next decisive steps are C4 and C8 source states that meet the sector and single-state criteria at and below the C2 half-connectivity bracket, slower and gap-informed formation ramps, real-time protocols at nonzero plaquette coupling, matched reduced-model dynamics, multi-step finite-temperature ensembles at matched temperature, and an apparatus-specific preparation and readout design. These extensions follow directly from the limits stated with each present result.

# Appendix A

# Supplementary C2 Results and Numerical Accuracy

This appendix supplies the detailed finite-sector data underlying Secs. III.1.1 and III.1.2, together with the physical-subspace representation and the numerical certification of the C2 calculations. The fixed-separation data come from the C2 static campaign of 6 October 2026 (data release 1.2); the separation-scan records were computed on 6–8 October 2026, and their independent re-derivation and Gauss-exact certification were completed on 12–13 October 2026. Displayed digits are rounded numerical values, not statistical uncertainties. All energies use $t_{\mathrm{h}} = 1$; the main ladder has $g^2 = 2$, $\lambda = \mu = 20$, $\kappa = 0$, and the sectors of Eq. (II.21). The results concern this exact C2 problem and do not extend to C4 or C8.

## A.1 Operator Conventions and Physical Coordinates

The matter basis is $(|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle)$, with $n = \mathrm{diag}(0,1,1,2)$ and parity $F_v = \mathrm{diag}(1,-1,-1,1)$. Its local annihilators are

$$c_\uparrow = |0\rangle\langle\uparrow| + |\downarrow\rangle\langle\uparrow\downarrow|, \qquad c_\downarrow = |0\rangle\langle\downarrow| - |\uparrow\rangle\langle\uparrow\downarrow|.$$

The ordered link basis is $(|0,2\rangle, |L\uparrow, R\uparrow\rangle, |L\uparrow, R\downarrow\rangle, |L\downarrow, R\uparrow\rangle, |L\downarrow, R\downarrow\rangle, |2,0\rangle)$. Numbering these states 1–6, the transporter is

$$\begin{aligned} U_{11} &= |4\rangle\langle 1| - |6\rangle\langle 3|, \quad U_{12} = -|5\rangle\langle 1| - |6\rangle\langle 2|, \\ U_{21} &= |2\rangle\langle 1| + |6\rangle\langle 5|, \quad U_{22} = -|3\rangle\langle 1| + |6\rangle\langle 4|. \end{aligned}$$

It obeys $[n_L, U_{ab}] = U_{ab}$, $[n_R, U_{ab}] = -U_{ab}$, with $E^2 = \mathrm{diag}(0, 3/4, 3/4, 3/4, 3/4, 0)$. The link spin blocks use $(-S_x, -S_y, S_z)$, whereas the matter block uses $(S_x, S_y, S_z)$; intertwiner construction respects this basis rotation. In colour-matrix notation $(U^\dagger)_{ab} = (U_{ba})^\dagger$, distinct from taking the adjoint of a single component $U_{ab}$. The latter enters $(\psi_A^\dagger U \psi_B)^\dagger = \sum_{ab} \psi_{B,b}^\dagger (U_{ab})^\dagger \psi_{A,a}$. Backward plaquette traversal uses the colour-transposed adjoint, and the plaquette term has coefficient $-\kappa$.

Jordan–Wigner strings include intervening matter parity and the appropriate endpoint factors. Link parity is even because the link contains two rishons, but this does not remove matter-fermion signs. On the tested fragment, both the signed and unsigned hopping

constructions obey gauge covariance, yet their Frobenius difference is 11.313708498984761. Statistics must therefore be specified independently of the Gauss algebra. The MPS ordering represents two-dimensional adjacency through nonlocal terms, not a physical reduction to one spatial dimension.

The isometries $\mathcal{U}_s$, distinct from link transporters, map the physical source and vacuum spaces of dimensions 27 580 and 16 502 into the corresponding product sectors:

$$\mathcal{U}_s^\dagger \mathcal{U}_s = I, \qquad P_s = \mathcal{U}_s \mathcal{U}_s^\dagger, \qquad H_{0,s} = \mathcal{U}_s^\dagger H_0 \mathcal{U}_s.$$

This is the complete retained physical sector, not a low-energy window. An independent coordinate-transfer calculation checks all 11,211,687 product-basis rows and all 18,437,415 nonzero isometry entries for number, imbalance, Cartan and diagonal-observable consistency. Relative Frobenius errors of $\mathcal{U}_s^\dagger \mathcal{U}_s - I$ are $1.3452 \times 10^{-14}$ and $2.3226 \times 10^{-16}$. The four exact vectors return to the original basis with relative errors at most $2.4706 \times 10^{-15}$. Fresh physical-coordinate $H_0$ residuals are

| | $m = 4$ | $m = 0.5$ |
|---|---|---|
| source | $4.2882 \times 10^{-14}$ | $1.3405 \times 10^{-13}$ |
| vacuum | $1.5333 \times 10^{-13}$ | $8.8203 \times 10^{-14}$ |

The full orthonormality test is distinct from sampled canonical operator checks: sixteen selected columns at twelve vertices give worst Casimir residual $4.7236 \times 10^{-15}$, and twenty reconstructed Hamiltonian entries differ by at most $2.1317 \times 10^{-14}$. The 222 numerical checks support transfer of exact static vectors. They do not establish finite-temperature Gibbs convergence or dissipative lifetimes.

## A.2 Detailed Mass Ladders and Spatial Responses

**Table A.1:** Ground energies of the optimization Hamiltonian, including the source Gauss constant 30, and same-sector first excitation gaps.

| $m$ | $E_{0,\mathrm{src}}$ | $E_{0,\mathrm{vac}}$ | $\Delta_{\mathrm{src}}$ | $\Delta_{\mathrm{vac}}$ |
|---|---|---|---|---|
| 0.25 | 13.54643 | −19.028855 | 0.87168632 | 2.2480761 |
| 0.5 | 12.600668 | −20.480997 | 0.75674361 | 2.6821173 |
| 0.75 | 11.389136 | −22.159727 | 0.65180356 | 3.1295119 |
| 1 | 9.9333938 | −24.021355 | 0.5733768 | 3.583506 |
| 1.125 | 9.1097606 | −25.010537 | 0.56647661 | 3.8121146 |
| 1.25 | 8.216883 | −26.034166 | 0.60143059 | 4.0415743 |
| 1.375 | 7.2553239 | −27.089346 | 0.68625042 | 4.2717891 |
| 1.5 | 6.2321281 | −28.173473 | 0.81654036 | 4.5026948 |
| 1.75 | 4.0396403 | −30.419381 | 1.171634 | 4.9664136 |
| 2 | 1.7039292 | −32.755573 | 1.598155 | 5.4324958 |
| 2.5 | −3.2530033 | −37.64711 | 2.5416809 | 6.3711653 |
| 3 | −8.4646411 | −42.765049 | 3.5317512 | 7.317697 |
| 4 | −19.349755 | −53.464894 | 5.5485509 | 9.2305657 |

**Table A.2:** Configuration-projector expectations in the original source sector. Local flux-free support requires exactly one matter particle.

| $m$ | $P_{\mathrm{any}}$ | $P_{\mathrm{sel}}$ | $P_{\mathrm{alt}}$ | $P_{\mathrm{ff}}(1)$ | $P_{\mathrm{ff}}(3)$ | $P_{\mathrm{both}}$ |
|---|---|---|---|---|---|---|
| 0.25 | 0.014469757 | 0.014446745 | $2.3012338 \times 10^{-5}$ | 0.75117943 | 0.2571877 | 0.18405594 |
| 0.5 | 0.028903779 | 0.028886816 | $1.6962808 \times 10^{-5}$ | 0.72436904 | 0.237631 | 0.15991895 |
| 0.75 | 0.061403403 | 0.061390874 | $1.2528876 \times 10^{-5}$ | 0.67115732 | 0.21357969 | 0.12994132 |
| 1 | 0.14339623 | 0.14338642 | $9.805135 \times 10^{-6}$ | 0.54869989 | 0.18769289 | 0.09258069 |
| 1.125 | 0.21871283 | 0.21870413 | $8.7017215 \times 10^{-6}$ | 0.45141687 | 0.17188271 | 0.070966306 |
| 1.25 | 0.3150386 | 0.31503119 | $7.4104481 \times 10^{-6}$ | 0.34538904 | 0.15248577 | 0.050088713 |
| 1.375 | 0.41579532 | 0.41578942 | $5.8962165 \times 10^{-6}$ | 0.25428753 | 0.13106125 | 0.033324579 |
| 1.5 | 0.50553692 | 0.50553251 | $4.4046207 \times 10^{-6}$ | 0.18885643 | 0.11086302 | 0.021842285 |
| 1.75 | 0.6383008 | 0.63829859 | $2.2143703 \times 10^{-6}$ | 0.11595339 | 0.08022268 | 0.010052661 |
| 2 | 0.72439944 | 0.72439837 | $1.0692453 \times 10^{-6}$ | 0.080770937 | 0.060823032 | 0.0052799673 |
| 2.5 | 0.82418507 | 0.8241848 | $2.6141317 \times 10^{-7}$ | 0.04781917 | 0.039221536 | 0.0019558381 |
| 3 | 0.87803908 | 0.878039 | $7.2013533 \times 10^{-8}$ | 0.032221172 | 0.027718528 | 0.00091363463 |
| 4 | 0.93166509 | 0.93166509 | $7.6983583 \times 10^{-9}$ | 0.017622049 | 0.016013622 | 0.0002843493 |

**Table A.3:** Physical energy, signed off-corridor fraction, mean corridor response, largest off-corridor absolute response, and source-site matter redistribution.

| $m$ | $V$ | $\epsilon_{\mathrm{off}}$ | $\overline{\rho}_{\mathcal{C}}$ | $\max_{\ell \notin \mathcal{C}} \lvert\rho_\ell\rvert$ | $\Delta n_1$ | $\Delta n_3$ |
|---|---|---|---|---|---|---|
| 0.25 | 2.5752846 | 0.1415016 | 0.10361284 | 0.10933884 | 0.37785359 | 0.21657576 |
| 0.5 | 3.0816644 | 0.13003732 | 0.13987162 | 0.11909735 | 0.42887889 | 0.19233986 |

Continued on the next page

Table A.3 (continued)

| $m$ | $V$ | $\epsilon_{\mathrm{off}}$ | $\overline{\rho}_{\mathcal{C}}$ | $\max_{\ell\notin\mathcal{C}}\lvert\rho_\ell\rvert$ | $\Delta n_1$ | $\Delta n_3$ |
|---|---|---|---|---|---|---|
| 0.75 | 3.5488625 | 0.10339484 | 0.18819013 | 0.11610469 | 0.4302727 | 0.15937136 |
| 1 | 3.9547488 | 0.058556194 | 0.27353921 | 0.08954217 | 0.34845701 | 0.12042581 |
| 1.125 | 4.1202971 | 0.036868015 | 0.33725268 | 0.06709892 | 0.26782732 | 0.096603476 |
| 1.25 | 4.2510486 | 0.022091029 | 0.40755835 | 0.04473341 | 0.17654324 | 0.070113808 |
| 1.375 | 4.3446698 | 0.014581189 | 0.47095529 | 0.028290751 | 0.098567248 | 0.044471336 |
| 1.5 | 4.4056013 | 0.011420789 | 0.52011237 | 0.018613285 | 0.044875906 | 0.023211825 |
| 1.75 | 4.4590214 | 0.0092676547 | 0.58252697 | 0.0104102 | −0.0079831342 | −0.0039665765 |
| 2 | 4.4595026 | 0.0080264069 | 0.61838104 | 0.0083409687 | −0.026788423 | −0.017255502 |
| 2.5 | 4.3941071 | 0.0057725805 | 0.6585053 | 0.0067881596 | −0.034932096 | −0.026111478 |
| 3 | 4.3004075 | 0.0040364486 | 0.68124543 | 0.0049028919 | −0.033016541 | −0.02661708 |
| 4 | 4.1151389 | 0.0020076841 | 0.70656016 | 0.0024694116 | −0.025412435 | −0.022202493 |

**Table A.4:** Separate source and vacuum energy expectations. Tiny negative vacuum Gauss expectations are floating-point roundoff around zero.

| $m$ | Role | $\langle mM\rangle$ | $\langle H_E\rangle$ | $\langle H_{\mathrm{hop}}\rangle$ | $\langle H_G\rangle$ | $\langle H_0\rangle$ |
|---|---|---|---|---|---|---|
| 4 | Source | −44.415809 | 3.2223905 | −8.1563364 | 30 | −49.349755 |
| 4 | Vacuum | −43.732766 | 0.39046427 | −10.122592 | $-1.9014372\times10^{-17}$ | −53.464894 |
| 0.5 | Source | −2.1698122 | 2.4235889 | −17.653109 | 30 | −17.399332 |
| 0.5 | Vacuum | −3.1482692 | 1.7804734 | −19.113201 | $-3.1437181\times10^{-17}$ | −20.480997 |

**Table A.5:** Source–vacuum energy components and excess particle–hole content. The last column is an expectation difference in active-link units, not a configuration-wise integer count.

| $m$ | $m\Delta\langle M\rangle$ | $\Delta\langle H_E\rangle$ | $\Delta\langle H_{\mathrm{hop}}\rangle$ | $V$ | $\Delta\langle N_A\rangle$ | $\frac{4}{3}\Delta\langle H_E\rangle$ |
|---|---|---|---|---|---|---|
| 0.25 | 0.5210054 | 0.48276309 | 1.5715161 | 2.5752846 | 1.0420108 | 0.64368412 |
| 0.5 | 0.97845698 | 0.64311552 | 1.4600919 | 3.0816644 | 0.97845698 | 0.85748736 |
| 0.75 | 1.3259106 | 0.83956745 | 1.3833844 | 3.5488625 | 0.88394043 | 1.1194233 |
| 1 | 1.4451036 | 1.1622115 | 1.3474337 | 3.9547488 | 0.72255179 | 1.5496154 |
| 1.125 | 1.3421324 | 1.4006499 | 1.3775148 | 4.1202971 | 0.5965033 | 1.8675332 |
| 1.25 | 1.1195081 | 1.6670605 | 1.4644799 | 4.2510486 | 0.44780326 | 2.2227473 |
| 1.375 | 0.83725667 | 1.911696 | 1.5957171 | 4.3446698 | 0.30445697 | 2.548928 |
| 1.5 | 0.56485399 | 2.1044844 | 1.736263 | 4.4056013 | 0.18828466 | 2.8059791 |
| 1.75 | 0.1465477 | 2.3519045 | 1.9605691 | 4.4590214 | 0.04187077 | 3.1358727 |
| 2 | −0.12625302 | 2.4935383 | 2.0922173 | 4.4595026 | −0.031563255 | 3.3247178 |
| 2.5 | −0.43313659 | 2.6493146 | 2.1779291 | 4.3941071 | −0.086627319 | 3.5324195 |
| 3 | −0.58177348 | 2.7360255 | 2.1461554 | 4.3004075 | −0.096962246 | 3.648034 |
| 4 | −0.68304304 | 2.8319263 | 1.9662557 | 4.1151389 | −0.08538038 | 3.7759017 |

**Table A.6:** Electric Casimir responses on all thirteen links. Corridor order along the source path is [1,3,2,5]; depth labels are discrete graph-shell labels, not continuum transverse distances.

| Link | Corridor | Depth | $\rho_\ell(4)$ | $\rho_\ell(1.5)$ | $\rho_\ell(0.5)$ |
|---|---|---|---|---|---|
| 1 | Yes | 0 | 0.70531935 | 0.51052537 | 0.03046662 |
| 2 | Yes | 0 | 0.70630646 | 0.48767239 | 0.0093773588 |
| 3 | Yes | 0 | 0.7065391 | 0.51270971 | 0.20266432 |
| 4 | No | 1 | 0.0013261548 | 0.018613285 | 0.11909735 |
| 5 | Yes | 0 | 0.70807574 | 0.56954202 | 0.3169782 |
| 6 | No | 2 | −0.00012005789 | −0.00039060458 | 0.0075523717 |
| 7 | No | 1 | 0.0024694116 | 0.0041361946 | −0.022564224 |
| 8 | No | 2 | −0.00010769821 | −0.00022425495 | 0.0089359441 |
| 9 | No | 2 | −0.00010050153 | $-8.623069 \times 10^{-5}$ | 0.0038888805 |
| 10 | No | 1 | 0.0022687368 | 0.0031367416 | −0.02528686 |
| 11 | No | 3 | $3.0411047 \times 10^{-6}$ | 0.00035923418 | 0.0077136164 |
| 12 | No | 3 | $8.5033409 \times 10^{-6}$ | 0.0003221597 | 0.0013189 |
| 13 | No | 2 | $-6.1976641 \times 10^{-5}$ | −0.0018316539 | −0.017026958 |

**Table A.7:** Matter redistribution at every vertex, with sources at vertices 1 and 3. Summed response vanishes within numerical precision before printed rounding.

| Vertex | Sublattice | $\Delta n_v(4)$ | $\Delta n_v(1.5)$ | $\Delta n_v(0.5)$ |
|---|---|---|---|---|
| 1 | A | −0.025412435 | 0.044875906 | 0.42887889 |
| 2 | A | −0.043877224 | 0.11094927 | 0.39944461 |
| 3 | A | −0.022202493 | 0.023211825 | 0.19233986 |
| 4 | A | 0.0032984462 | 0.0063627825 | −0.017188806 |
| 5 | A | 0.0028838983 | 0.0046227109 | −0.016206216 |
| 6 | A | $-7.0572185 \times 10^{-5}$ | −0.0017378325 | −0.0088113511 |
| 7 | B | 0.043796704 | −0.070962451 | −0.27313116 |
| 8 | B | 0.042918034 | −0.089544004 | −0.4178993 |
| 9 | B | −0.0017784038 | −0.027754911 | −0.22424711 |
| 10 | B | 0.00030940109 | 0.00073616148 | −0.025279719 |
| 11 | B | 0.00013891844 | −0.00016702824 | −0.020625141 |
| 12 | B | $-4.2736489 \times 10^{-6}$ | −0.00059242871 | −0.017274554 |

**Table A.8:** Full-colour main and plaquette-control ladders and colour-diagonal hopping comparison. The latter is not a matched independent U(1) theory.

| | $P_{\text{any}}$ | | | $P_{\text{both}}$ | | |
|---|---|---|---|---|---|---|
| $m$ | Main | $\kappa = .5$ | Diagonal | Main | $\kappa = .5$ | Diagonal |
| 0.5 | 0.028903779 | 0.02862484 | 0.0058446131 | 0.15991895 | 0.1592082 | 0.21791448 |
| 1 | 0.14339623 | 0.14636867 | 0.083247926 | 0.09258069 | 0.091354797 | 0.094917084 |
| 1.5 | 0.50553692 | 0.50840344 | 0.83223183 | 0.021842285 | 0.021302485 | 0.0020066261 |
| 2 | 0.72439944 | 0.72449847 | 0.91515969 | 0.0052799673 | 0.0052368244 | 0.00045410688 |
| 4 | 0.93166509 | 0.93164698 | 0.98131364 | 0.0002843493 | 0.00028435568 | $2.2015082 \times 10^{-5}$ |

**Table A.9:** Finite-penalty departures from the minimal full-colour blocks. These contributions prevent attributing all main/control differences solely to colour mixing.

| $m$ | $\langle H_G\rangle_{\mathrm{src}} - 30$ | $\langle H_G\rangle_{\mathrm{vac}}$ |
|---|---|---|
| 0.5 | 0.10749806 | 0.13056996 |
| 1 | 0.10732007 | 0.13594522 |
| 1.5 | 0.11288401 | 0.1375491 |
| 2 | 0.11259269 | 0.13719577 |
| 4 | 0.10482302 | 0.12929377 |

**Table A.10:** Original sources (1,3) and relocated sources (4,6), both at graph distance four. The relocated sector is separately derived from the same geometric balance rule.

| | $P_{\mathrm{any}}$ | | $\Delta_{\mathrm{src}}$ | | $V$ | |
|---|---|---|---|---|---|---|
| $m$ | Original | Relocated | Original | Relocated | Original | Relocated |
| 0.5 | 0.028903779 | 0.030225536 | 0.75674361 | 0.70265132 | 3.0816644 | 3.0193692 |
| 1 | 0.14339623 | 0.14136187 | 0.5733768 | 0.60798362 | 3.9547488 | 3.9093129 |
| 1.25 | 0.3150386 | 0.29776151 | 0.60143059 | 0.66498308 | 4.2510486 | 4.2176572 |
| 1.5 | 0.50553692 | 0.48658689 | 0.81654036 | 0.8689249 | 4.4056013 | 4.3885589 |
| 2 | 0.72439944 | 0.71846692 | 1.598155 | 1.6254091 | 4.4595026 | 4.4570257 |
| 4 | 0.93166509 | 0.9314951 | 5.5485509 | 5.5584743 | 4.1151389 | 4.1154253 |

**Table A.11:** Independent solutions at Gauss strengths 20 and 40. Shift error subtracts the expected source shift 30 or vacuum shift zero. Negative infidelity at the roundoff scale is not a physical negative probability.

| Role | $m$ | Shift error | $1-F$ | $\max\lvert\delta E^2\rvert$ | $\max\lvert\delta n\rvert$ | Projector error |
|---|---|---|---|---|---|---|
| Source | .5 | $-2\times10^{-13}$ | $2.7\times10^{-15}$ | $2.6\times10^{-14}$ | $4\times10^{-14}$ | $3.5\times10^{-14}$ |
| Source | 4 | $-6.4\times10^{-14}$ | $-4.4\times10^{-14}$ | $3.4\times10^{-14}$ | $8.9\times10^{-14}$ | $1.2\times10^{-15}$ |
| Vacuum | .5 | $6.8\times10^{-14}$ | $2\times10^{-15}$ | $1.9\times10^{-15}$ | $5.6\times10^{-15}$ | $1.2\times10^{-17}$ |
| Vacuum | 4 | $-2.8\times10^{-13}$ | $-1.3\times10^{-14}$ | $8.4\times10^{-16}$ | $2.5\times10^{-14}$ | $1\times10^{-20}$ |

## A.3 Gauss-Exact Static Certification

The Gauss-exact DMRG of Sec. II.5.2 was compared with the first-principles exact ground states for all five source sectors ($R = 2$, 4 and the three $R = 6$ sectors), each relative to the vacuum of the same mass. Figure A.1 shows the potentials and finite-distance secants together with their signed differences. All differences lie far inside the predefined threshold $10^{-8}$; the largest, $+7.8\times10^{-11}\,t_{\mathrm{h}}$ in $V_D(6)$ and $+4.0\times10^{-11}\,t_{\mathrm{h}}/a$ in $\sigma^D_{\mathrm{eff}}[4,6]$, both for route C at $m = 4$, belong to the state with the largest energy error in Table II.2. The recorded production values differ from this reference by at most $3.6\times10^{-12}$ in $V_D$ and $1.6\times10^{-12}\,t_{\mathrm{h}}/a$ in $\sigma^D_{\mathrm{eff}}$. For the $R = 4$ source sectors the Gauss-exact DMRG also reproduces every $\langle E^2_\ell\rangle$ and $\langle n_v\rangle$ of the exact states within $2.0\times10^{-12}$ ($m = 4$) and $2.2\times10^{-13}$ ($m = 0.5$).

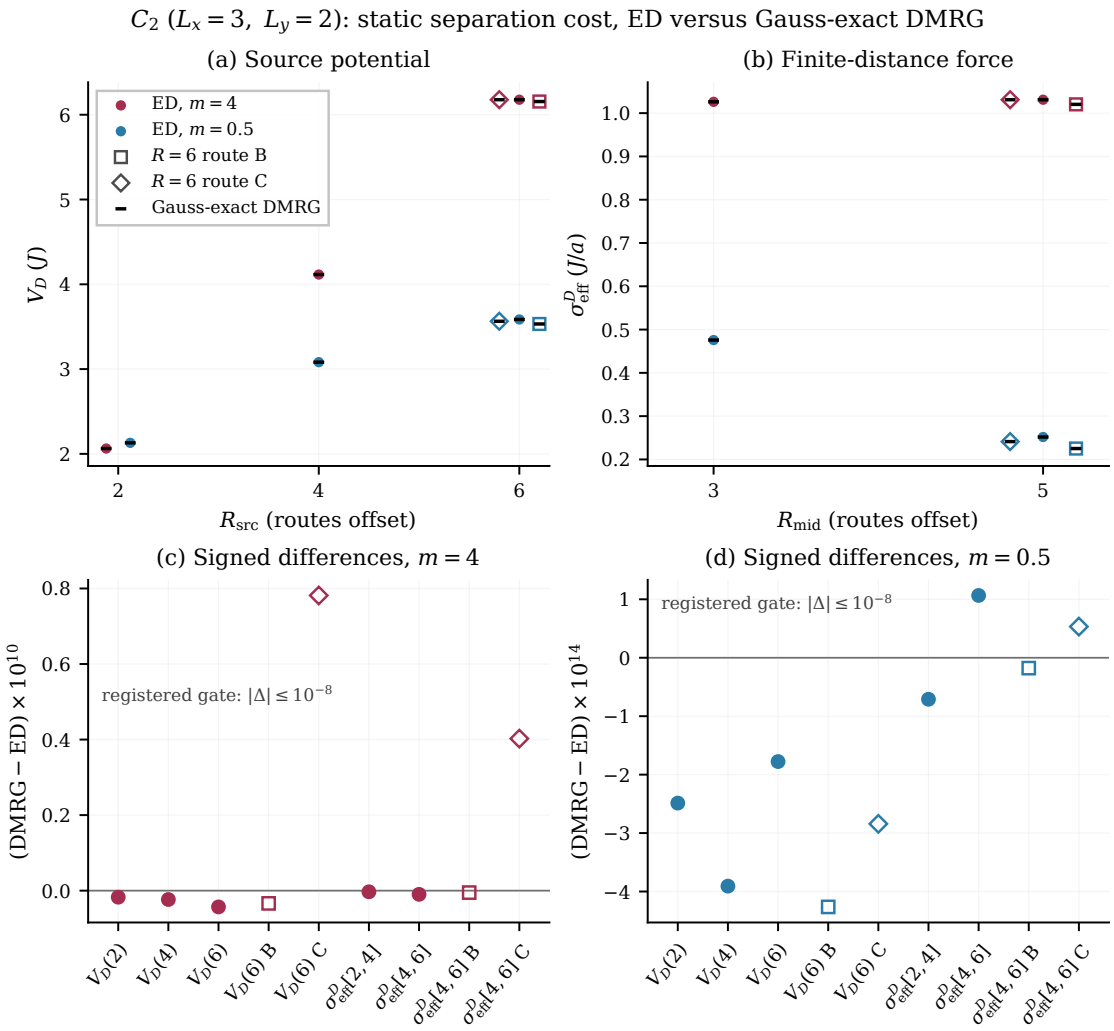



**Figure A.1:** Static separation cost from first-principles exact diagonalization (ED) and Gauss-exact DMRG for the five source sectors at $m = 4$ and $m = 0.5$ (mass colours; $J \equiv t_h$). (a) $V_D(R)$; open squares and diamonds are the $R = 6$ routes B and C, offset by $+0.20$ and $-0.20$, and the two masses at $R = 2$ are offset by $\mp 0.12$. (b) Secants $\sigma^D_{eff}$ per graph link at the interval midpoints 3 and 5. Black bars are the DMRG values. (c,d) Signed differences DMRG−ED at $m = 4$, in units of $10^{-10}$, and at $m = 0.5$, in units of $10^{-14}$ ($t_h$ for $V_D$, $t_h/a$ for $\sigma^D_{eff}$); tick names without a route letter refer to the primary sector.

## A.4 Numerical Reliability and Reproducibility

The local Fock reconstruction, small-system comparisons, sampled production MPO elements, exact-vector recomputation, the independent first-principles re-derivation and the Gauss-exact tensor networks test different parts of the calculation. Their counts are not added as if they were independent physical experiments. The fixed-separation data re-evaluation covers 46 ground-state records, 23 source–vacuum responses, six placement records and two initial quench-energy calculations. Largest record and response differences are approximately $2.023 \times 10^{-12}$ and $2.527 \times 10^{-13}$. The mechanism calculation independently reproduces 78 quantities within $9.3 \times 10^{-12}$. These checks complement, rather than replace, the explicit Hamiltonian and sector definitions.

The definitions of the electric response, the full-energy ratio and the force secants were checked against 10 static records, four force intervals and 1,444 real-time rows reconstructed from their primary measurements; deliberately substituting $H_0/R$ for the electric response is rejected for both static and dynamic tables. The C2 static-response and real-time figures reproduced in this thesis were redrawn from the certified records; their plotted data equal the source records point by point, apart from declared horizontal display offsets of overlapping markers and the declared $10^{-17}$ display floor of the logarithmic deviation plots. In the full campaign documentation every plotted series equals its source record within $6.7 \times 10^{-16}$ (static) and $1.3 \times 10^{-15}$ (real time), and every drawn primitive lies within $1.1 \times 10^{-6}$ pt of its source value. Figure inputs and outputs are bound by SHA-256 hashes.

The sparse C2 implementation uses approximately 5.2–6.2 GiB for source states and 3.2–4.3 GiB for vacuum states. Recorded main-ladder computations take about 9–27 minutes per point; the slowest colour-diagonal control takes approximately 75 minutes. Each production task constructs its own working representation. These are implementation-specific measurements, not scaling claims for other geometries or MPS contractions.

Reproducibility requires binding a state to its basis order, local matrices, fermionic convention, geometry, parameters and measurement Hamiltonian. Residual calculations retain the state independently of subsequent reporting so that a measurement failure does not erase a valid optimization result. Original records and corrections are versioned rather than silently overwritten. Hashes identify the data used, but are neither digital signatures nor independent proofs of the physical model. The acceptance criteria of the Gauss-exact calculations were fixed before any production data; two amendments, each fixed before the data it governs, are recorded with the run files.

# Appendix B

# Sector-Resolved Thermal Methods and Numerical Controls

This appendix gives the algebraic and numerical details underlying Sec. V.2. It distinguishes complete physical-sector calculations from spectral approximations, exact generators from numerical actions, and equilibrium estimates from trajectory statistics. The numerical results refer to the C2 representation and computations reported on 6 October 2026; earlier star and $C = 1$ results retain their own geometries.

## B.1 Physical Representation and State Transfer

The C2 source and vacuum product spaces contain 6,713,288 and 4,498,399 states. Projecting onto all prescribed local Casimirs gives physical dimensions 27,580 and 16,502. Every retained energy level remains present. The isometry $\mathcal{U}_s$ is distinct from the link transporter $U_\ell^{ab}$; the projector $P_s$ is distinct from an energy-window projector. The extra $D_v$ conservation follows from the determinant-free Hamiltonian and gives a U(2)-type local algebraic structure. Restricting to a $D_v$ pattern is additional conditioning, not the removal of non-Abelian fusion multiplicity.

Three Hamiltonian conventions require separation:

$$H_{\mathrm{opt}} = H_0 + H_G + H_D, \qquad H_{\mathrm{prot}} = H_0 + H_G, \qquad H_0 = H_{\mathrm{hop}} + H_m + H_E + H_p.$$

Some earlier numerical records name $H_{\mathrm{prot}}$ physical. All thermal energies here instead use the penalty-free $H_0$. At $\lambda = 20$ the retained source contribution is 30 and the vacuum contribution is zero. Normalized Gibbs states are unchanged by those constants, but raw partition functions and free energies are not. Outside the retained block the protection energy is state dependent.

The four exact input vectors contain 1,411,136 bytes of complex128 state payload. Their maximum physical-energy residual is $1.26004 \times 10^{-13}$, back-transfer error $1.92234 \times 10^{-14}$, initial-observable discrepancy $7.28306 \times 10^{-14}$, and discrepancy against the independent earlier initial-observable calculation $9.59233 \times 10^{-14}$. The preparation took 3.9165 s wall time, 2.79465 s CPU time and 0.770794 GiB peak resident memory; a 20% allowance gives 0.924953 GiB. Casimir action checks sampled eight columns in each sector. They are

not full-matrix Casimir tests. These transfer measurements are separate from the static certification of Appendix A.

An earlier independent infinite-temperature review covered 222 numerical checks and 81 input/operator bindings in 12.62 s with 0.807 GiB peak memory. Its representation analysis included all 11,211,687 ambient product-basis rows and 18,437,415 nonzero isometry coefficients. Source and vacuum orthogonality residuals were $1.3452 \times 10^{-14}$ and $2.3226 \times 10^{-16}$, respectively. The sampled Casimir-action residual was $4.6451 \times 10^{-15}$ on sixteen columns, and twenty hopping-action checks gave a maximum $2.1317 \times 10^{-14}$. These scopes must not be conflated: complete row/coefficient coverage does not turn sampled operator actions into a full operator-norm bound.

## B.2 Filtered Actions and Numerical Conditioning

For $\beta > 0$, the coherent correction of Eq. (V.11) has the paired integral representation

$$G_{\mathrm{KMS}} = \int_0^\infty \frac{e^{iH_s u} K e^{-iH_s u} - e^{-iH_s u} K e^{iH_s u}}{\beta \sinh(2\pi u/\beta)}\, du. \tag{B.1}$$

Pairing cancels the apparent singularity: the integrand tends to $i[H_s, K]/\pi$ at $u = 0$. At $\beta = 0$, $G_{\mathrm{KMS}} = 0$ follows directly from the algebraic formula rather than division by a vanishing inverse temperature. For a time cutoff $T_c > 0$,

$$\|G_{\mathrm{KMS}} - G_{\mathrm{KMS}}^{(T_c)}\| \le \frac{\|K\|}{\pi} \log \coth\left(\frac{\pi T_c}{\beta}\right). \tag{B.2}$$

For the Gaussian jump filter,

$$\|L_a - L_a^{(T_c)}\| \le \sqrt{\gamma_a}\|S_a\| e^{1/8} \operatorname{erfc}\left(\frac{\sqrt{2}T_c}{\beta}\right). \tag{B.3}$$

These are cutoff-tail bounds only. Quadrature, exponential-action and approximate-$K$ errors require separate control. Refining $G\psi$ on one vector does not bound the full operator or establish Gibbs stationarity. Refinements on different vectors cannot be concatenated into a convergence sequence.

For the modular alternative, write $a_\beta = \beta/4$. With a constant energy shift suppressed,

$$B_\beta = \sum_a \gamma_a S_a e^{-2a_\beta H_s} S_a, \qquad K = e^{a_\beta H_s} B_\beta e^{a_\beta H_s}. \tag{B.4}$$

Shared outer factors and bounded contact blocks reduce repeated action costs; the same factorization enters the $G_{\mathrm{KMS}}$ action. An outward-padded Gershgorin interval of the stored Hermitian sparse Hamiltonian supplies $E_{\mathrm{lower}}, E_{\mathrm{upper}}$. It is not an interval proof for an independently reconstructed ambient Hamiltonian. The guard $\beta(E_{\mathrm{upper}} - E_{\mathrm{lower}}) \le 24$ limits amplification in the actual implementation. Exact finite-dimensional identities remain valid outside that guard, but numerical conditioning and jump rates can worsen severely. The finite-temperature C2 steps at $\beta = 0.125$ do not demonstrate tractability at $\beta = 0.5, 2, 8$.

| Witness quantity | Source | Vacuum |
|---|---|---|
| Physical directions covered | 27,580/27,580 | 16,502/16,502 |
| Distinct contact signatures | 20,071 | 16,452 |
| Largest initial signature block | 8 | 3 |
| Exact determinant witnesses | 11,503 | 49 |
| Propagation rounds | 5 | 1 |
| Smallest recorded singular diagnostic | 0.613645 | $\simeq 2$ |
| Hermitian correction, Frobenius norm | $4.889\,10^{-12}$ | $2.160\,10^{-12}$ |

**Table B.1:** Coverage of the source and vacuum physical spaces for the stored, Hermitian-symmetrized $H_s$ and canonical contacts. Rational determinant witnesses establish the block propagation; the singular diagnostics are not the exactness criterion. This is not an independent ambient reconstruction.

The Gaussian construction follows the filtered KMS framework [17]. Its operator identities and equilibrium target do not imply identical kinetics to the modular alternative. The latter's constant even prefactor is admissible on a finite spectrum; an efficient Fourier implementation theorem with integrable decaying filters is a different statement.

## B.3 Constructive Accessibility and Uniqueness

The 25 contacts are diagonal in the chosen physical basis. Assign to each basis direction its complete vector of contact eigenvalues. If an operator $X$ commutes with every $S_a$, then $X$ is block diagonal with respect to these signatures. Degenerate signatures can support nonscalar blocks even when the Hamiltonian's ordinary connectivity graph is connected.

Singleton signatures and nonzero Hamiltonian links first propagate a common scalar $c$. For a remaining signature block $B$, commutation with $H_s$ implies

$$(X_B - cI_B)\, H_s[B, \mathcal{K}] = 0, \tag{B.5}$$

where $\mathcal{K}$ contains directions already fixed to that scalar. An invertible square minor of $H_s[B, \mathcal{K}]$ fixes $X_B = cI_B$. Iterating this argument reaches every coordinate. Determinants of the stored binary floating-point coefficients are evaluated as exact rational numbers, rather than declared nonzero from a floating-point singular-value threshold. The numerical singular values are supplementary conditioning diagnostics.

The mass and electric diagonal terms belong to the contact span, so changing the endpoint mass leaves the common-commutant argument unchanged.

For completeness, consider the exact finite-dimensional generator at finite $\beta$, with all $\gamma_a > 0$ and faithful $\rho_{\beta,s}$. Let $X$ be a fixed observable of $\mathcal{M}_\beta^\dagger + i\,\mathrm{ad}_{H_s}$. KMS self-adjointness and nonpositivity of the thermostat, together with skew-adjointness of the Hamiltonian drift, imply

$$\mathrm{Re}\langle X, (\mathcal{M}_\beta^\dagger + i\,\mathrm{ad}_{H_s})X\rangle_{\beta,s} = 0 \quad \Longrightarrow \quad \mathcal{M}_\beta^\dagger X = 0, \quad [H_s, X] = 0.$$

The stationary-state Kadison identity then gives

$$\sum_a \mathrm{Tr}\big(\rho_{\beta,s}[L_a, X]^\dagger [L_a, X]\big) = 0.$$

Faithfulness makes every commutator vanish. Since $X$ commutes with $H_s$, invertible modular conjugation yields $[S_a, X] = 0$. For the Gaussian filter the same implication follows because its multiplier is strictly positive at each finite Bohr frequency. Equation (V.13) therefore leaves only $X = cI_s$; the stationary density is unique. This application of stationary-state algebra retains the faithful-state and exact-generator hypotheses [19, 44].

The argument gives neither a numerical relaxation rate nor a finite-step convergence bound. Additional exact invariant blocks, if present for another contact set, would preserve initial weights. With mixing inside such blocks and decay of inter-block coherences, the accessible reference would instead be $\sum_\alpha p_\alpha \rho_{\beta,\alpha}$. The scalar C2 result rules out this particular obstruction for the declared stored algebra, not for all gauge-aware reservoirs. Near-zero modes at finite numerical precision must not be promoted to exact conservation laws.

## B.4 Equilibrium Estimators and Uncertainty

For Eq. (V.14), set $x_r = v_r^\dagger O_s v_r$, $y_r = v_r^\dagger v_r$, $\bar{x} = R_{\mathrm{p}}^{-1} \sum_r x_r$ and $\bar{y} = R_{\mathrm{p}}^{-1} \sum_r y_r$. The trace estimates are unbiased before taking ratios, whereas $\bar{x}/\bar{y}$ and $\log \bar{y} - \beta E_{\mathrm{sh}}$ have finite-sample bias. The covariance contribution in a delta-method variance estimate is

$$\widehat{\mathrm{Var}}(\bar{x}/\bar{y}) \simeq \frac{s_x^2 - 2\widehat{O}\, s_{xy} + \widehat{O}^{\,2} s_y^2}{R_{\mathrm{p}} \bar{y}^{\,2}}, \qquad \widehat{O} = \bar{x}/\bar{y}. \tag{B.6}$$

Student intervals and jackknife bias estimates complement this diagnostic; an effective weight count is not an exact confidence guarantee. The stored sufficient statistics include $H_0$, $H_0^2$ and 35 diagonal observables, without requiring a saved bank of propagated probes.

At infinite temperature, independent exact trace identities are

$$\langle O \rangle_{0,s} = \frac{\operatorname{Tr} O_s}{d_s}, \qquad \langle H_s^2 \rangle_{0,s} = \frac{\|H_s\|_F^2}{d_s}. \tag{B.7}$$

They have no Monte Carlo sampling error, apart from arithmetic error in the stored representation. At finite $\beta$, comparing one exponential action with two half-actions tests composition but is not a rigorous global error bound. Using an accepted ground-state energy as $E_{\mathrm{sh}}$ does not certify omitted levels; the present estimator does not omit levels at all.

Each of the twelve reference points uses two independent PCG64 batches, 64 probes per batch. At $\beta = 0.5$, the maximum composition discrepancy is $1.526 \times 10^{-15}$; effective weight counts are 94.4, 127.5, 74.0 and 125.8 for source high mass, source low mass, vacuum high mass and vacuum low mass, respectively. Nested 32-to-64 comparisons reuse probes.

The unrounded cold connectivity estimates are 0.931665094 for source $m = 4, \beta = 8$ and $3.753374241 \times 10^{-7}$ for the corresponding vacuum. Their approximate sampling errors are $9.29 \times 10^{-13}$ and $5.44 \times 10^{-23}$, while high-mass energy errors are of order $10^{-14}$. These scales indicate nearly collinear cold propagated probes, not an independent systematic accuracy of twenty-three decimal places. For comparison, the source $m = 0.5, \beta = 8$ energy variance is about 0.001147, and the vacuum $m = 0.5, \beta = 2$ variance is about

| Construction | Diagnostic | Value |
|---|---|---|
| Davies | Gibbs fixed-point residual | $2.58797\,10^{-14}$ |
| | Generator trace residual | $1.17161\,10^{-16}$ |
| | Integrated trace error | $2.6244\,10^{-15}$ |
| | Hermiticity error | $1.0929\,10^{-14}$ |
| | Minimum density eigenvalue | $-1.49802\,10^{-16}$ |
| | Gauss expectation residual | $7.33\,10^{-15}$ |
| | $G^3, D, N$ target-square residual | 0 |
| | Matrix-action/integrator difference | $1.01\,10^{-14}$ |
| | Smallest resolved decay gap | $2.95795\,10^{-4}$ |
| Gaussian KMS | Thermostat Gibbs residual | $2.08\,10^{-16}$ |
| | Gibbs residual including $H_s$ | $1.83356\,10^{-14}$ |
| | Weighted KMS residual | $1.77204\,10^{-16}$ |
| | Spohn/first-law consistency error | $1.53350\,10^{-14}$ |
| | Minimum Choi eigenvalue | $-4.26805\,10^{-15}$ |
| Modular KMS | Gibbs fixed-point residual | $4.25752\,10^{-15}$ |
| | Weighted KMS residual | $6.35426\,10^{-17}$ |
| | Residual with $G_{\mathrm{KMS}}$ omitted | $9.691\,10^{-4}$ |

**Table B.2:** Separate maxima or minima over sixteen Davies cases, sixteen Gaussian cases and eight modular cases. Davies and Gaussian sets use $\beta = 0, 0.5, 2, 8$; the modular set uses $\beta = 0, 0.125$, each at both masses and source/vacuum roles. Negative eigenvalues shown are at roundoff scale. These are tractable-system tests, not C2 finite-step CPTP or many-body mixing bounds.

0.169788. Both are legitimate thermal fluctuations. The cold vacuum high-mass variance $-4.54747 \times 10^{-13}$ is numerical cancellation.

For numerical traceability, the standard errors rounded in Table V.1 are as follows. At source high mass, $\beta = 2$, the energy and connectivity errors are $1.724 \times 10^{-5}$ and $2.874 \times 10^{-5}$; at source low mass, $\beta = 8$, they are $1.8695 \times 10^{-4}$ and $1.6297 \times 10^{-4}$. The vacuum high-mass connectivity at $\beta = 0.5$ is $(6.94774 \pm 0.50871) \times 10^{-7}$, and at $\beta = 2$ it is $(3.753538 \pm 0.000267) \times 10^{-7}$. The vacuum low-mass $\beta = 2$ value is $0.000298377 \pm 0.000002691$. Reporting these retained digits does not remove the systematic qualifications.

## B.5 Distinct Small-System Reference Calculations

The Davies star calculation, Gaussian KMS calculation and modular KMS calculation are separate numerical references. Their inverse-temperature sets and residuals cannot be pooled into a single C2 accuracy bound. Table B.2 retains their principal measurements.

The Davies physical star spaces have dimensions 10 and 12; retaining only $G^3, D, N$ would instead give 40 and 32. Its tests use $\gamma = 0.02$, $t = 0, 0.25, 1, 3$ and a high-energy initial eigenstate. At $t = 3$, energy changes range from $-0.234$ to $-0.0556$. All sixteen cases have one stationary and one nondecaying mode. The original acceptance bounds are

$10^{-10}$ for fixed-point, trace, Hermiticity, Gauss and positivity deviations and $10^{-8}$ for the matrix-action comparison. The resolved decay gap belongs to the star calculation alone.

The analytic two-level reference checks the sign of detailed balance, $\beta = 0$, cold-state flow and $\gamma = 0$. A degenerate four-level reference is covariant under rotations within energy eigenspaces. A commuting contact dephases without energy flow and loses uniqueness. A bare-fermion annihilation contact produces detectable $N, D, G^3$ leakage in a space where leakage can occur. Density-matrix continuation reproduces the corresponding uninterrupted star evolution; it is not evidence for C2 MCWF continuation.

The Gaussian reference has up to six numerically unresolved near-zero modes at cold points. They are not established extra conservation laws and cannot be used to infer rapid mixing. In the modular small-system reference, coherent-action errors at quadrature orders 8, 16 and 32 are $7.8674 \times 10^{-6}$, $9.74328 \times 10^{-10}$ and $2.87132 \times 10^{-13}$, respectively. This ladder is distinct from the C2 random-vector and ground-state comparisons below.

## B.6 Stochastic Evolution, Replay and Error Separation

For the infinite-temperature Hermitian contact, one splitting step is

$$e^{-iH_s dt/2} e^{-i\sum_a \sqrt{\gamma_a} S_a \Delta W_a} e^{-iH_s dt/2}, \qquad \Delta W_a \sim \mathcal{N}(0, dt). \tag{B.8}$$

The ideal finite-step ensemble is a mixture of unitary channels and hence CPTP. This does not prove time-step convergence or a finite-temperature target. Its earlier C2 closed control has energy drift at most $3.20 \times 10^{-13}$ and phase error at most $3.90 \times 10^{-16}$. Interrupted continuation differs from the uninterrupted result by at most $1.504 \times 10^{-16}$ over 32 Brownian draws. Coarse/fine differences $2.65$–$3.15 \times 10^{-5}$ lack a predefined physical-equivalence margin and do not close a time-step error budget. These continuation tests belong to that older random-unitary method, not to the finite-$\beta$ MCWF implementation.

The modular MCWF step described in Sec. V.2.6 uses one PCG64 uniform event draw after checking the probability cap. At $\gamma = 0$ it applies only the coherent physical-Hamiltonian action and does not consume an event draw. A separate auxiliary NumPy MT19937 stream is used by exponential-action norm estimation. First-step replay restores the actual recorded pre-step state of both streams, not merely an initial seed. The four replay differences are zero in the tested arithmetic; the source high/low and vacuum high/low closed-control discrepancies are, respectively,

$$2.21 \times 10^{-16}, \quad 4.44 \times 10^{-16}, \quad 1.82 \times 10^{-15}, \quad 2.14 \times 10^{-16}.$$

This establishes reproducibility of one bounded operation, not universal bitwise reproducibility across numerical environments.

Each measured step uses 1,560 exponential-action calls, with a limit of 20,000 calls and 600 s for the bounded calculation. All sampled branches are no-event branches. The coherent contribution to the initial heat rate is of order $10^{-15}$ because the inputs are energy eigenstates. For a general pure state it is

$$-i\langle [H_s, G_{\mathrm{KMS}}] \rangle = 2\ \mathrm{Im}\langle H_s\psi, G_{\mathrm{KMS}}\psi \rangle$$

| Measured calculation | Wall time (s) | Peak GiB | CPU/wall |
|---|---|---|---|
| Gibbs references, $\beta = 0.5$ | 6.17–19.70 | 0.102–0.143 | 0.993–0.994 |
| Gibbs references, $\beta = 2, 8$ | 21.47–361.65 | 0.101–1.376 | — |
| Complete first-step diagnostics | 128.26–368.90 | 0.186–0.329 | 0.983–0.992 |
| One MCWF step | 36.90–109.35 | — | — |

**Table B.3:** Reference estimation, complete diagnostic tasks and individual steps are different workloads. Complete tasks include preparation, action refinement and replay. A dash denotes an unreported quantity, not zero. The larger reference peak requires its 2 GiB allocation rather than a 1 GiB assumption.

and need not vanish. The realized conditional energy changes can be of order $10^{-9}$ even though $dt\,\dot{Q}$ from the unconditional generator is much larger. This difference is expected from conditional sampling and cannot diagnose energy-balance failure by itself.

Small-system comparison with exact Liouvillian propagation gives local errors proportional to $dt^2$ and fixed-time errors proportional to $dt$. It supports first-order weak accuracy, not a C2 time-step ladder or an exact linear finite-step channel. The frozen numerical development suite contains 125 tests; that count does not imply complete physical validation.

Long-trajectory continuation additionally requires a stored physical state, time, Hamiltonian and bath specification, source sector, contact order and rates, both random streams, integrity checks and an actual interrupted-versus- uninterrupted comparison. Measurement must leave the state unchanged. Model error, state-transfer error, Gibbs sampling error, filtered-action error, time-step error, trajectory-ensemble error, tensor truncation and physical effect size are eight distinct layers. Passing one does not set the others to zero.

## B.7 Memory and Measured Action Costs

A complex128 vector, density matrix and dense Liouvillian require, respectively, $16d_s$, $16d_s^2$ and $16d_s^4$ bytes. For C2, density matrices alone require approximately 12.17 GB (11.33 GiB) in the source sector and 4.36 GB (4.06 GiB) in vacuum, before work buffers. Matrix-free Liouvillian application avoids the $d_s^4$ operator, not the $d_s^2$ density state. Pure-state trajectories avoid the latter but retain potentially expensive global filtered jumps.

The earlier $C = 1$, $2 \times 2$ source block has $d = 17{,}431$, not either C2 dimension. Its closed quench at $t = 0, 0.25, 0.5$ gives ground-state residual $1.55 \times 10^{-13}$, norm drift $1.27 \times 10^{-14}$, energy drift $4.33 \times 10^{-13}$ and variance drift $5.03 \times 10^{-13}$. Its sparse Hamiltonian has about 6.28 MB raw storage; a dense density matrix would take 4.86 GB and a dense Liouvillian about 1.48 EB. These estimates motivate sparse methods without supplying a $C = 1$ dissipative trajectory.

A direct Gaussian implementation with 25 contacts and sixteen jump and coherent quadrature nodes would require 102,464 exponential calls for one coherent action. This is a cost estimate, not a completed full C2 action. Earlier single-contact Gaussian timings at $\beta = 0.5$, for eight and sixteen nodes, were 7.917 and 13.239 s in the low-mass source sector, and 4.787 and 8.236 s in vacuum. Modular random-vector eight/sixteen-node

comparisons took 12.94/25.59 s in the source sector and 6.11/12.50 s in vacuum, with relative differences 0.00267 and 0.00268. These vectors differ from the eigenstate inputs of the sixteen/thirty-two-node first-step measurements. The two sets do not form a common refinement curve.

The measured near-single-core use supports parallelism over independent probes or trajectories rather than assuming unmeasured within-action speedup. Long ensembles require a measured multi-step throughput including observables and checkpoints. The small first-step event probabilities also make rare-event resolution a separate cost. Memory allowance and checkpoint overhead must accompany any extrapolation from these bounded timings.

## B.8 Spectral and Tensor-Network Alternatives

An energy-window approximation would require an omitted-state bound that the complete-sector calculation does not need. Suppose $r$ retained levels have known partition sum $Z_r > 0$, and every omitted energy is certified to exceed $E_\star$. Then

$$B_{\mathrm{tail}} = (d_s - r)e^{-\beta E_\star}, \qquad w_{\mathrm{tail}} \le \frac{B_{\mathrm{tail}}}{Z_r + B_{\mathrm{tail}}}, \qquad |\delta\langle O\rangle| \le 2\|O\| w_{\mathrm{tail}}. \tag{B.9}$$

At $\beta = 0$, the omitted weight is $(d_s - r)/d_s$. A few small Ritz residuals do not lower-bound every omitted energy; agreement between two energy windows is not this certificate. Transient dynamics additionally needs bounds on state projection and omitted jump couplings. No such truncation is used for the twelve Gibbs estimates.

MPS trajectories can reduce state storage but must approximate global filtered jumps and control truncation errors. Matrix-product density operators require independent positivity checks. Locally purified representations preserve positivity by construction while retaining bond and ancillary/Kraus truncation errors [40]. Purification-based finite-temperature tensor-network calculations supply a related equilibrium method [18]. METTS samples equilibrium states; it does not specify reservoir kinetics [38]. These are possible extensions, not additional C2 results.

## B.9 Ensemble Entropy and Response Design

For independent normalized, equally weighted trajectories, $\rho_{N_{\mathrm{traj}}} = N_{\mathrm{traj}}^{-1} \sum_r |\psi_r\rangle\langle\psi_r|$ is positive with unit trace. Its nonzero eigenvalues equal those of the Gram matrix

$$\mathsf{C}_{rs} = \frac{\langle\psi_r|\psi_s\rangle}{N_{\mathrm{traj}}}.$$

This offers an entropy estimator without explicitly storing a $d_s^2$ density matrix, but finite-ensemble bias and the overlap cost remain. For $N_{\mathrm{traj}} \ge 2$, an unbiased purity estimator is

$$\widehat{\mathrm{Tr}\,\rho^2} = \frac{1}{N_{\mathrm{traj}}(N_{\mathrm{traj}} - 1)} \sum_{r \ne s} |\langle\psi_r|\psi_s\rangle|^2. \tag{B.10}$$

Its negative logarithm is a biased estimator of second-Renyi entropy, not von Neumann entropy. These are prospective estimators; no C2 entropy or relative-entropy time series is available.

A physical response study fixes geometry, sectors, initial vectors, $H(t)$, bath family, contact normalization, rates, temperature, switch times and observation window. Numerical controls specify a time-step ladder, quadrature/cutoff refinement, exponential tolerances and any representation truncation. Stochastic controls specify independent seed blocks, ensemble-size refinement and a genuine pairing rule. Primary observables, temporal summaries, uncertainty procedure, physical-equivalence margin and lifetime persistence criterion must be declared independently of the response being tested. Exploratory checkpoint selection is identified as such rather than retrospectively labelled prospective.

The existing proposed sector tolerances are

$$\begin{aligned} \max_v |\langle G_v^3(t)\rangle - \langle G_v^3(0)\rangle| &\le 10^{-6}, \\ \max_v |\langle D_v(t)\rangle - d_v| &\le 10^{-6}, \qquad \max_v \mathrm{Var}(D_v) \le 10^{-8}, \\ |\langle N(t)\rangle - \langle N(0)\rangle| &\le 10^{-6}, \qquad \mathrm{Var}(N) \le 10^{-8}. \end{aligned} \tag{B.11}$$

They accompany, rather than replace, the correct local Casimir targets. In a projected representation the corresponding residuals can be kinematically zero, which is why ambient operator checks remain necessary.

For a dynamically prepared input, the transfer compares norm loss, round trip, ambient/physical $H_0$ action and initial observables. Its energy spread $\|(H_0 - \langle H_0\rangle)\psi\|$ can be physical and is not required to vanish as for an eigenstate. Any renormalization must retain its pre-correction discrepancy. A matched closed continuation starts from the identical checkpoint; successful conversion does not establish a producer's physical breaking or persistence mechanism.

The next controlled comparison requires a $\beta = 0.125$ equilibrium reference, multi-step and sampled-jump tests, actual interrupted continuation, and independent short ensembles with a matched closed control. Forcing a jump into a selected path is a diagnostic, not unbiased sampling. Equilibrium accuracy, numerical dynamics, approach to equilibrium, structural response and persistent lifetimes have distinct acceptance conditions. Several initial states and a resolved time window are needed to characterize approach; measured observables do not determine the entire density matrix. Gauge conservation does not establish flux survival.

Independent reproduction should challenge sector swaps, omitted $G_{\mathrm{KMS}}$, reversed temperature/rate conventions, altered contact normalization, missing or duplicated trajectories, invalid random states, probability-cap violations and state-changing measurement. These are requirements for future ensemble validation, not a claim that every adverse case has already been tested. Positive response, controlled null response and statistically unresolved response remain equally admissible scientific outcomes.

# Appendix C

# Wider-Patch Numerical Status

This appendix lists every recorded C4 and C8 optimization behind Sec. III.1.4 and the bond entropies of the stored states. The heavy-mass statuses are those of the registered support campaign, whose final manifest (11 October 2026) keeps the overall verdict open: the C4 heavy-mass pair has no passing independent cross-check, and no light-mass endpoint was selected. The intermediate masses were deferred in that registered scope; their values in Sec. III.1.4 are post hoc diagnostics.

Table C.1 separates three situations. At $m = 4$ the selected endpoints pass the final sector criteria below their bond-dimension caps. At $m = 3$ and 2 the stored states pass the sector criteria but sit at their caps, so they describe the stored approximants and are used descriptively. At $m \leq 1.5$ no source state passes the sector criteria; the light-mass source optimizations were still descending when they stopped, and no observable from them enters the main text.

The bond entropies of Table C.2 were computed from the stored states by exact singular-value decompositions of the orthogonalized matrix-product states, without truncation. Rows with different caps are not comparable in their capped-bond counts, and an entropy of a capped state is a property of the stored approximant only.

| Patch | Sector | $m$ | Records (saved) | $\chi$/cap | $\Delta E_4$ | Sector bars | Use |
|---|---|---|---|---|---|---|---|
| C4 | source | 4 | 14 (8) | 773/800 | $6.8 \times 10^{-10}$ | pass | endpoint (Q) |
| C4 | source | 3 | 2 (1) | 400/400 | $5.6 \times 10^{-10}$ | pass | descriptive |
| C4 | source | 2 | 8 (6) | 600/600 | $7.2 \times 10^{-5}$ | pass | descriptive |
| C4 | source | 1.5 | 4 (0) | – | – | – | no saved state |
| C4 | source | 1 | 6 (1) | 600/600 | $3.5 \times 10^{-3}$ | fail | not used (bars) |
| C4 | source | 0.5 | 5 (1) | 400/400 | $1.3 \times 10^{-2}$ | fail | open endpoint |
| C4 | vacuum | 4 | 14 (10) | 381/800 | $3.6 \times 10^{-11}$ | pass | endpoint (Q) |
| C4 | vacuum | 3 | 2 (1) | 400/400 | $1.9 \times 10^{-10}$ | pass | descriptive |
| C4 | vacuum | 2 | 3 (1) | 400/400 | $5.5 \times 10^{-10}$ | pass | descriptive |
| C4 | vacuum | 1.5 | 4 (1) | 400/400 | $7.2 \times 10^{-10}$ | pass | not used (no source) |
| C4 | vacuum | 1 | 4 (2) | 400/400 | $2.5 \times 10^{-6}$ | fail | not used (bars) |
| C4 | vacuum | 0.5 | 11 (9) | 800/800 | $1.9 \times 10^{-5}$ | fail | open endpoint |
| C8 | source | 4 | 8 (7) | 422/800 | $4.5 \times 10^{-10}$ | pass | endpoint (B) |
| C8 | source | 2 | 4 (4) | 600/600 | $1.6 \times 10^{-7}$ | pass | descriptive, no vacuum |
| C8 | source | 1 | 1 (1) | 600/600 | $2.1 \times 10^{-3}$ | fail | not used (bars) |
| C8 | source | 0.5 | 3 (1) | 400/400 | $1.9 \times 10^{-4}$ | fail | open endpoint |
| C8 | vacuum | 4 | 9 (7) | 469/800 | $2.8 \times 10^{-12}$ | pass | endpoint (B) |
| C8 | vacuum | 0.5 | 5 (4) | 400/400 | $2.6 \times 10^{-3}$ | fail | open endpoint |

**Table C.1:** Numerical status of all recorded C4 and C8 optimizations ($R = 4$, $\kappa = 0$, $g^2 = 2$). For the four heavy-mass targets the selected endpoint of the registered support campaign is shown, otherwise the lowest-energy saved state. $\Delta E_4$ is the energy spread of the last four completed sweeps; sector bars are the final particle-number and imbalance bars (means $10^{-6}$, variances $10^{-8}$). Use in Sec. III.1.4: endpoint (Q) or (B), selected endpoint with a single-state quality pass or bounded numerical support; descriptive, quoted descriptively; open endpoint, registered light-mass endpoint without a selected state; not used (bars), sector criteria fail; not used (no source), no source state at that mass. A state at its bond-dimension cap describes only the stored approximant.

| Patch | Sector | $m$ | State | $\chi$/cap | $S_{\max}$ | Bonds at cap | $w_{\min}$ at $S_{\max}$ |
|---|---|---|---|---|---|---|---|
| C4 | source | 4 | S038 | 773/800 | 1.37 | 0/38 | $1 \times 10^{-13}$ |
| C4 | source | 3 | S023 | 400/400 | 1.63 | 7/38 | $1 \times 10^{-11}$ |
| C4 | source | 2 | S021 | 800/800 | 2.14 | 2/38 | $5 \times 10^{-11}$ |
| C4 | source | 2 | S022 | 600/600 | 2.14 | 4/38 | $1 \times 10^{-9}$ |
| C4 | source | 1 | S010 | 600/600 | 3.46 | 9/38 | $2 \times 10^{-12}$ |
| C4 | source | 0.5 | S000 | 400/400 | 3.74 | 12/38 | $2 \times 10^{-10}$ |
| C4 | vacuum | 4 | S076 | 381/800 | 0.91 | 0/38 | $6 \times 10^{-15}$ |
| C4 | vacuum | 2 | S058 | 400/400 | 1.65 | 8/38 | $2 \times 10^{-11}$ |
| C4 | vacuum | 0.5 | S046 | 800/800 | 2.63 | 4/38 | $4 \times 10^{-12}$ |
| C4 | vacuum | 0.5 | S049 | 800/800 | 2.64 | 2/38 | $1 \times 10^{-12}$ |
| C8 | source | 4 | S092 | 422/800 | 1.37 | 0/66 | $9 \times 10^{-13}$ |
| C8 | source | 2 | S083 | 600/600 | 2.06 | 3/66 | $2 \times 10^{-11}$ |
| C8 | source | 1 | S080 | 600/600 | 2.83 | 7/66 | $9 \times 10^{-11}$ |
| C8 | source | 0.5 | S077 | 400/400 | 2.82 | 29/66 | $1 \times 10^{-9}$ |
| C8 | vacuum | 4 | S106 | 469/800 | 0.91 | 0/66 | $1 \times 10^{-14}$ |
| C8 | vacuum | 0.5 | S093 | 400/400 | 2.69 | 21/66 | $3 \times 10^{-10}$ |

**Table C.2:** Largest bond entanglement entropy $S_{\max}$ of stored C4 and C8 states in the registered tensor ordering, the number of bonds whose Schmidt rank equals the bond-dimension cap, and the smallest retained Schmidt weight on the bond of $S_{\max}$. The entropies include the gauge-constraint contribution and are a representation cost of this ordering, not a physical entanglement measure; for capped states they describe only the stored approximant. Caps differ between rows, so the capped-bond counts are not comparable across rows.

# Bibliography


[1] O. Abdel Karim, A. Muzi Falconi, R. Panza, W. Liu, and F. Scazza. Single-atom imaging of $^{173}$Yb in optical tweezers loaded by a five-beam magneto-optical trap. *Quantum Science and Technology*, 10:045019, 2025.

[2] Tameem Albash, Sergio Boixo, Daniel A. Lidar, and Paolo Zanardi. Quantum adiabatic Markovian master equations. *New Journal of Physics*, 14:123016, 2012.

[3] Robert Alicki. On the detailed balance condition for non-hamiltonian systems. *Reports on Mathematical Physics*, 10(2):249–258, 1976.

[4] Andreas Alvermann and Holger Fehske. High-order commutator-free exponential time-propagation of driven quantum systems. *Journal of Computational Physics*, 230:5930, 2011.

[5] Gunnar S. Bali, Hartmut Neff, Thomas Düssel, Thomas Lippert, and Klaus Schilling. Observation of string breaking in QCD. *Physical Review D*, 71(11):114513, 2005.

[6] Debasish Banerjee et al. Atomic quantum simulation of U(N) and SU(N) non-Abelian lattice gauge theories. *Physical Review Letters*, 110:125303, 2013.

[7] Debasish Banerjee, F.-J. Jiang, T. Z. Olesen, Peter Orland, and Uwe-Jens Wiese. From the SU(2) quantum link model on the honeycomb lattice to the quantum dimer model on the Kagomé lattice: Phase transition and fractionalized flux strings. *Physical Review B*, 97:205108, 2018. arXiv:1712.08300.

[8] G. Cappellini, L. F. Livi, L. Franchi, D. Tusi, D. Benedicto Orenes, M. Inguscio, J. Catani, and L. Fallani. Coherent manipulation of orbital Feshbach molecules of two-electron atoms. *Physical Review X*, 9:011028, 2019.

[9] Michele Caselle, Michele Pepe, and Antonio Rago. Static quark potential and effective string corrections in the (2+1)-d SU(2) Yang–Mills theory. *Journal of High Energy Physics*, 2004(10):005, 2004.

[10] Giovanni Cataldi, Simone Orlando, and Jad C. Halimeh. Real-time string dynamics in a (2+1)D non-Abelian lattice gauge theory: String breaking, glueball formation, baryon blockade, and tension reduction, 2025. arXiv:2509.08868.

[11] Alessio Celi, Benoît Vermersch, Oscar Viyuela, Hannes Pichler, Mikhail D. Lukin, and Peter Zoller. Emerging two-dimensional gauge theories in Rydberg configurable arrays. *Physical Review X*, 10:021057, 2020.

[12] Shailesh Chandrasekharan and Uwe-Jens Wiese. Quantum link models: A discrete approach to gauge theories. *Nuclear Physics B*, 492:455–471, 1997.

[13] Chi-Fang Chen, Michael Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Efficient quantum thermal simulation. *Nature*, 646:561–566, 2025.

[14] Jean Dalibard, Yvan Castin, and Klaus Mølmer. Wave-function approach to dissipative processes in quantum optics. *Physical Review Letters*, 68(5):580–583, 1992.

[15] E. B. Davies. Markovian master equations. *Communications in Mathematical Physics*, 39:91–110, 1974.

[16] Gaia De Paciani, Lukas Homeier, Jad C. Halimeh, Monika Aidelsburger, and Fabian Grusdt. Quantum simulation of fermionic non-Abelian lattice gauge theories in (2 + 1)D with built-in gauge protection. *Physical Review A*, 113:063327, 2026.

[17] Zhiyan Ding, Bowen Li, and Lin Lin. Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition. *Communications in Mathematical Physics*, 406:67, 2025.

[18] Adrian E. Feiguin and Steven R. White. Finite-temperature density matrix renormalization using an enlarged Hilbert space. *Physical Review B*, 72:220401, 2005.

[19] Alberto Frigerio. Stationary states of quantum dynamical semigroups. *Communications in Mathematical Physics*, 63:269–276, 1978.

[20] Carlos Gas-Ferrer, Antonio Rubio-Abadal, Sandra Buob, Leonardo Bezzo, Jonatan Höschele, and Leticia Tarruell. Spin-resolved microscopy of $^{87}$Sr SU(N) Fermi–Hubbard systems. *Physical Review X*, 2026. Accepted 15 September 2026.

[21] Daniel González-Cuadra, Majd Hamdan, Torsten V. Zache, Boris Braverman, Milan Kornjača, Alexander Lukin, Sergio H. Cantú, Fangli Liu, Sheng-Tao Wang, Alexander Keesling, Mikhail D. Lukin, Peter Zoller, and Alexei Bylinskii. Observation of string breaking on a (2 + 1)D Rydberg quantum simulator. *Nature*, 642:321–326, 2025.

[22] Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan. Completely positive dynamical semigroups of N-level systems. *Journal of Mathematical Physics*, 17:821–825, 1976.

[23] Jutho Haegeman, Christian Lubich, Ivan Oseledets, Bart Vandereycken, and Frank Verstraete. Unifying time evolution and optimization with matrix product states. *Physical Review B*, 94:165116, 2016.

[24] Jad C. Halimeh, Lukas Homeier, Annabelle Bohrdt, and Fabian Grusdt. Spin exchange-enabled quantum simulator for large-scale non-Abelian gauge theories. *PRX Quantum*, 5:030358, 2024.

[25] Daniel Horn. Finite matrix models with continuous local gauge invariance. *Physics Letters B*, 100:149–151, 1981.

[26] John Kogut and Leonard Susskind. Hamiltonian formulation of Wilson's lattice gauge theories. *Physical Review D*, 11:395–408, 1975.

[27] G. Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48:119–130, 1976.

[28] Giuseppe Magnifico, Giovanni Cataldi, Marco Rigobello, Peter Majcen, Daniel Jaschke, Pietro Silvi, Simone Montangero, et al. Tensor networks for lattice gauge theories beyond one dimension. *Communications Physics*, 8:322, 2025.

[29] Daniel Manzano. A short introduction to the Lindblad master equation. *AIP Advances*, 10:025106, 2020.

[30] Berndt Müller and Xiaojun Yao. Simple Hamiltonian for quantum simulation of strongly coupled (2+1)D SU(2) lattice gauge theory on a honeycomb lattice. *Physical Review D*, 108:094505, 2023.

[31] Peter Orland and Daniel Rohrlich. Lattice gauge magnets: Local isospin from spin. *Nuclear Physics B*, 338:647–672, 1990.

[32] Jesse J. Osborne, Ian P. McCulloch, Bing Yang, Philipp Hauke, and Jad C. Halimeh. Large-scale (2 + 1)D U(1) gauge theory with dynamical matter in a cold-atom quantum simulator. *Communications Physics*, 8:273, 2025.

[33] Robert Ott, Torsten V. Zache, Fred Jendrzejewski, and Jürgen Berges. Scalable cold-atom quantum simulator for two-dimensional QED. *Physical Review Letters*, 127:130504, 2021.

[34] L. Riegger, N. Darkwah Oppong, M. Höfer, D. R. Fernandes, I. Bloch, and S. Fölling. Localized magnetic moments with tunable spin exchange in a gas of ultracold fermions. *Physical Review Letters*, 120:143601, 2018.

[35] Aruku Senoo, Alexander Baumgärtner, Joanna W. Lis, Gaurav M. Vaidya, Zhongda Zeng, Giuliano Giudici, Hannes Pichler, and Adam M. Kaufman. High-fidelity entanglement and coherent multi-qubit mapping in an atom array. *Nature Physics*, 22:903–909, 2026.

[36] Rainer Sommer. A new way to set the energy scale in lattice gauge theories and its application to the static force and $\alpha_s$ in SU(2) Yang–Mills theory. *Nuclear Physics B*, 411:839–854, 1994.

[37] Herbert Spohn. Entropy production for quantum dynamical semigroups. *Journal of Mathematical Physics*, 19(5):1227–1230, 1978.

[38] E. Miles Stoudenmire and Steven R. White. Minimally entangled typical thermal state algorithms. *New Journal of Physics*, 12:055026, 2010.

[39] Sho Sugiura and Akira Shimizu. Canonical thermal pure quantum state. *Physical Review Letters*, 111:010401, 2013.

[40] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero. Positive tensor network approach for simulating open quantum many-body systems. *Physical Review Letters*, 116:237201, 2016.

[41] Steven R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69:2863–2866, 1992.

[42] Uwe-Jens Wiese. Ultracold quantum gases and lattice systems: quantum simulation of lattice gauge theories. *Annalen der Physik*, 525:777–796, 2013.

[43] Kenneth G. Wilson. Confinement of quarks. *Physical Review D*, 10:2445–2459, 1974.

[44] Yikang Zhang and Thomas Barthel. Criteria for Davies irreducibility of Markovian quantum dynamics. *Journal of Physics A: Mathematical and Theoretical*, 57:115301, 2024.

# Declaration

I hereby declare that this thesis is my own work, and that I have not used any sources and aids other than those stated in the thesis.

Munich,
Date of submission

Author's signature