

Master in Quantum Science and Technology Barcelona

Master's Thesis

Master in Quantum Science and Technology

Linear Anyon Model with Rydberg Atoms

Aditya Banerjee

Supervisor: Prof. Alessio Celi

Universitat de Barcelona

Host institution: Universitat Autònoma de Barcelona

Barcelona, 3 September 2026



Linear Anyon Model with Rydberg Atoms

Aditya Banerjee

Supervised by: Prof. Alessio Celi

Departament de Física, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain
3 September 2026

We propose a scheme for quantum-simulating the one-dimensional anyon-Hubbard model with Rydberg atoms. Starting from the generalized Jordan–Wigner transformation, which trades anyonic exchange statistics for a density-dependent Peierls phase, we encode the truncated local occupations $\{|0\rangle, |1\rangle, |2\rangle\}$ in the internal levels of Rydberg atoms driven by Raman lasers with site-dependent phases. Adiabatically eliminating the far-detuned intermediate states in the dispersive regime yields an effective bond Hamiltonian whose four tunnelling amplitudes we obtain in second-order perturbation theory. First, we show that the Raman scheme allows one to simulate pseudo-fermions with bosons, but not generic anyons: Floquet driving is necessary to engineer a generic statistical phase. In fact, we find that single-tone Floquet driving leads to density-dependent amplitudes, but it is not sufficient to achieve the desired density-dependent tunnelling. This calculation suggests that to achieve the phases one needs a drive with several tones, each resonant with one occupation-resolved process and each carrying its own phase, a mechanism similar to the one employed in the recent few-particle realization of the anyon-Hubbard model with ultracold neutral atoms in optical lattices.

Keywords: quantum simulators, anyon-Hubbard model, Rydberg atoms, optical lattice

Acknowledgements

I would like to thank Prof. Alessio Celi for being my supervisor, for supporting me throughout this work, and for pointing out this interesting problem and introducing me to the area of quantum simulators. I would also like to thank Prof. Celi for letting me work remotely. Finally, I would like to thank the course on quantum simulators with ultracold atoms that I attended in 2024, which kept me motivated to work on this interesting problem.

Contents

1 Introduction **4**
- 1.1 Anyonic statistics and the generalized Jordan–Wigner transformation 7
- 1.2 Density-dependent tunnelling and the anyon-Hubbard Hamiltonian 7
- 1.3 Truncated local Hilbert space and the bond Hamiltonian 8

2 Rydberg-atom implementation **9**
- 2.1 Encoding the local states and the Raman coupling 9
- 2.2 Rydberg interactions and the dispersive regime 10
- 2.3 Effective exchange from second-order perturbation theory 13
- 2.4 Effective Hamiltonian and the flux-locking restriction 13

3 Floquet engineering of the statistical phase **14**
- 3.1 The driven two-site Hamiltonian 17
- 3.2 Rotating frame 17
- 3.3 The Floquet factor 18
- 3.4 Magnus expansion and Fourier mode decomposition 19
- 3.5 Secular approximation and commutators 20
- 3.6 Final second-order effective Hamiltonian 20
- 3.7 Resolving the spectrum by particle number 21
- 3.8 Time evolution 23
  - 3.8.1 Analytic solution within each manifold 24
  - 3.8.2 From the bond to the chain: spacetime dynamics 25

4 Matching the effective Hamiltonian to the anyon-Hubbard model **28**
- 4.1 Diagonal sector: chemical potential and interaction 29
- 4.2 Effective chemical potential, interaction, and their ratio 29
- 4.3 Off-diagonal sector: target versus engineered amplitudes 33
- 4.4 Amplitude conditions 33
- 4.5 Gauge freedom and the uniformity of the induced phase 33
- 4.6 Phase conditions and the locking of the statistical angle 35

5 Conclusions and outlook **37**

Bibliography **40**

A Numerical results and experimental parameters **43**
- A.1 Two-site building block: spectrum and dynamics 43
- A.2 The density-dependent phase 44

B A concrete optical-tweezer experiment **45**

C Full derivation of the anyon-Hubbard model **47**

D The four second-order transitions **52**
- D.1 Transition $|0,1\rangle \to |1,0\rangle$ 52
- D.2 Transition $|0,2\rangle \to |1,1\rangle$ 54
- D.3 Transition $|1,1\rangle \to |2,0\rangle$ 56
- D.4 Transition $|1,2\rangle \to |2,1\rangle$ 58

1 Introduction

Quantum statistics is usually introduced as a postulate: the wavefunction of a system of identical particles is required to be either symmetric or antisymmetric under the interchange of any two of them, giving bosons and fermions. Leinaas and Myrheim showed that this restriction is not an independent axiom but a consequence of the topology of the configuration space [1]. Because configurations that differ by a permutation of the particle labels describe the same physical state, the classical configuration space of N identical particles is not the Cartesian product of the one-particle spaces but the identification space obtained by identifying such points, and this space has a different global topological structure. Quantizing directly on it, no symmetrization postulate need be imposed: the admissible statistics follow from the topology itself.

The number of possibilities depends on the dimension. The essential observation is that exchanging two particles twice is equivalent to transporting one of them in a closed loop around the other [2]. In three or more dimensions such a loop can always be contracted to a point, so a double exchange must act trivially and a single exchange can only produce ± 1: bosons and fermions are the only options, corresponding to the two one-dimensional representations of the permutation group. Equivalently, a rotation by 4π is topologically equivalent to the identity in $SO(3)$, so the eigenvalues of a 2π rotation can only be $+1$ for integer and -1 for half-integer spin—the topological content of the spin–statistics connection [3]. In two dimensions a loop encircling another particle cannot be shrunk without passing through it, the winding number is well defined, and trajectories are instead classified by the braid group, which is infinite where the permutation group has $N!$ elements. A double exchange need no longer return the system to its initial state, so a single counter-clockwise exchange may multiply the wavefunction by an arbitrary phase,

$$\psi(\mathbf{r}_1, \mathbf{r}_2) \longrightarrow e^{i\phi}\, \psi(\mathbf{r}_1, \mathbf{r}_2), \tag{1}$$

with $\phi = 0$ and $\phi = \pi$ recovering bosons and fermions. Particles with intermediate values of the statistical angle are called *anyons*, a name coined by Wilczek while showing that composites of a charge and a magnetic flux tube realize this behaviour: the Aharonov–Bohm phase accumulated when one composite encircles another interpolates continuously between the bosonic and fermionic cases, and the composite carries a correspondingly fractional spin [4]. This *flux-attachment* picture remains the standard way of thinking about anyons and underlies the Chern–Simons description of fractional quantum Hall states [1, 2]. When the states carrying a given set of particle positions are degenerate, braiding acts as a matrix rather than a phase; such *non-Abelian* anyons are the basis of proposals for topological quantum computation, in which information is stored non-locally and manipulated by braiding, making it intrinsically robust against local perturbations [2]. Anyonic physics of this kind is no longer confined to solid-state settings. Since the first proposals for computing with individually addressed, laser-controlled atomic systems [5], ultracold gases have matured into a general-purpose platform for quantum simulation [3, 6], offering the control over band structure and lattice geometry that such engineering demands [7], and a small fractional quantum Hall state—the setting in which anyons were first predicted—has recently been assembled atom by atom [8].

One dimension is a special case, and it is the one relevant here. Two particles on a line cannot be interchanged without passing through one another, so statistics cannot be defined by braiding at all; a hard-core repulsion alone already makes bosons equivalent to fermions

[2]. Statistics nevertheless survives as a genuine degree of freedom. In the continuum treatment of Ref. [1] the coincidence point becomes a boundary of the configuration space, and requiring only that probability be conserved there admits a one-parameter family of boundary conditions $\partial_z\psi|_{z=0} = \eta\,\psi|_{z=0}$ in the relative coordinate $z = |x_1 - x_2|$, with $\eta = 0$ describing bosons, $\eta^{-1} = 0$ fermions, and a continuum of intermediate values describing particles that are neither. On a lattice the same freedom reappears in a form that can be *engineered*. The paradigmatic lattice model is the anyon-Hubbard model (AHM), in which the statistical phase is traded, through a generalized Jordan–Wigner transformation, for a *density-dependent Peierls phase* in an equivalent bosonic model [9, 10]. A boson then picks up a phase equal to the statistical phase ϕ times the occupation of the site it tunnels from or to, so that the statistics is encoded in the tunnelling itself. This is more than a change of variables: varying ϕ drives statistically induced quantum phase transitions [9], and the ground-state phase diagram hosts pair superfluid, dimer, and partially paired phases together with asymmetric momentum distributions [10]. Dynamically, the model displays asymmetric expansion and supports chiral solitons. Remarkably, both features have been observed [11] in the chiral BF theory [12], which is the continuum limit of the anyon-Hubbard model [13, 14]. A complementary route to one-dimensional anyons, in which the statistical phase is imprinted on a continuum Bose gas rather than on a lattice, has also been demonstrated [15].

The anyon-Hubbard model belongs to a wider family of Hubbard Hamiltonians that cold atoms can realize. The most general such Hamiltonian contains, besides the usual tunnelling, on-site and extended interaction terms, an *interaction-assisted tunnelling* contribution in which the tunnelling matrix element itself depends on the occupations of the sites involved; this term is normally discarded in the Wannier expansion, and its reinstatement defines the class of Hubbard models with occupation-dependent parameters [3]. The anyon-Hubbard model belongs to this class, with the particular feature that what the statistics makes occupation dependent is not the modulus of the tunnelling but its *phase*. This distinction matters throughout: a linear coupling can readily be made to give an amplitude that depends on the density, but not a phase.

Realizing the AHM therefore amounts to engineering an occupation-dependent tunnelling phase. The routes proposed with ultracold atoms all work in the same way: the interactions split the occupation-resolved tunnelling processes apart in energy, and a drive that is resonant with one of them at a time imprints a phase on that process alone. The original proposal uses Raman-assisted tunnelling in a tilted lattice, where the internal atomic structure supplies the occupation selectivity [9]. Greschner and Santos showed that no additional lasers are needed: by shaking a tilted lattice with strong on-site interactions, resonant photon-assisted tunnelling against the tilt acquires a number-dependent Peierls phase [10, 16]. These are not competing alternatives—the recent realization of one-dimensional anyons with *arbitrary* statistical phase in a quantum-gas microscope, where the anyonic Hanbury Brown–Twiss effect and asymmetric transport were observed directly [17], implements this Floquet scheme. A variant modulates the depth of the lattice rather than its position: since the various occupation-dependent hops are energetically distinguishable, a *multicolour* modulation can address each of them with its own frequency component and its own phase [18]. In every case the phase is carried by a drive tone resonant with a single occupation-resolved process. The analysis below shows that this is the ingredient the scheme developed here lacks.

Periodic driving is by now a standard tool for band-structure and gauge-field engineering in optical lattices [19, 20]. Suitably designed forcing renders the tunnelling matrix element complex, with a Peierls phase tunable continuously over the full range $(0, 2\pi)$ [21],

and generalizations produce non-Abelian gauge structures and topological bands [22]. A clear demonstration of this control is the realization of the topological Haldane model, in which circular modulation of a honeycomb lattice induces the complex next-nearest-neighbour tunnelling the model requires, with the imaginary part of that tunnelling—and hence the topology of the band—set by the phase of the modulation [23]. Driving also causes heating: a driven many-body system generically absorbs energy from the drive [24], though at large drive frequency the heating rate is exponentially suppressed, leaving a long-lived prethermal window in which the engineered Hamiltonian governs the dynamics [25].

Rydberg atoms offer a complementary platform. Their strong, distance- and state-dependent interactions provide the large, configuration-resolved energy scales needed to make a given tunnelling process resonant while detuning all others, and individual atoms can be positioned and addressed with single-site resolution [26, 27]. The platform this points to is the optical-tweezer array, whose geometry is programmable and whose sites are prepared and read out one by one: where an optical lattice fixes the interaction energy once the lattice depth is chosen, here that role is played by the separation and orientation of the traps, which can be set independently of the coupling strengths. Using this interaction as a resource goes back to the proposal of Jaksch *et al.* for a fast two-qubit gate between neutral atoms [28], which established the two features the present construction inherits. The first is that the dipole–dipole energy of a pair of laser-excited Rydberg atoms can be made so much larger than the mechanical scales of the trap that the atoms may be manipulated as if pinned in place, the internal dynamics decoupling from the external one. The second is the hierarchy $u \gg \Omega$ between that interaction and the Rabi frequency of the exciting laser—the blockade regime—in which the interacting pair state is detuned out of resonance by the interaction itself and is therefore only virtually populated, yet still controls the phase acquired by the pair. Both recur below: the interaction energies supply the offsets that put the transitions between two Rydberg states of a single atom off resonance, Eq. (13), and super-exchange transitions in which two neighbouring atoms exchange their states simultaneously occur at second order in the Rabi frequencies. This process can be interpreted as physical tunnelling by mapping the internal levels of the atoms to particle occupation. That multi-level encodings of this kind can realize lattice models has recently been demonstrated: encoding two spin states and a hole in three Rydberg levels of ^{87}Rb realizes a doped quantum antiferromagnet—a hard-core bosonic t–J–V model—in which the dipolar tunnelling, the spin exchange and the hole–hole interaction have different distance and angular dependences and can therefore be tuned separately through the geometry of the array [29]. This suggests encoding the local occupations of the AHM not in the motional state of atoms in a lattice, but in the internal levels of Rydberg atoms, with the tunnelling generated as an effective second-order (super-exchange) process between them.

This thesis develops such a scheme. We encode the truncated local Hilbert space $\{|0\rangle, |1\rangle, |2\rangle\}$ of the AHM in three internal levels of each of two Rydberg atoms, drive them with Raman lasers carrying site-dependent phases, and eliminate the far-detuned intermediate states to obtain an effective bond Hamiltonian. Matching this to the AHM shows that a purely *static* protocol is restricted: the phases of the four independent tunnelling amplitudes are not independent, and the statistical phase is locked to the values $\phi \in \{0, \pi\}$. The value $\phi = \pi$ is not trivial. It describes pseudofermions—particles that are antisymmetric under exchange between different sites while still allowing multiple occupancy [16, 30]—and in a random-walk experiment they would display the antibunching characteristic of fermions. In other words, the static Raman scheme can realize pseud-

ofermions, but not arbitrary one-dimensional anyons. We then add a periodic drive on the doubly-occupied levels and derive the resulting Floquet effective Hamiltonian. Matching it to the anyon-Hubbard bond shows that the restriction is not lifted: the phases available—two laser phases together with one gauge shift—are occupation independent, whereas the statistical angle enters only those processes that leave a doubly-occupied site behind, so the matching conditions force $\phi = 0$. What the drive delivers instead is a renormalized Bose-Hubbard bond with widely tunable parameters.

Outline. The remainder of this section introduces the anyon-Hubbard model and the mapping that underlies the whole construction. Section 2 presents the Rydberg encoding, the second-order perturbation theory, and the flux-locking restriction. Section 3 develops the Floquet treatment of the driven bond. Section 4 compares the resulting effective Hamiltonian with the anyon-Hubbard model. Section 5 critically summarizes the results of this thesis and discusses the research direction it opens. The interaction geometry and driving parameters assumed throughout are collected in Table 1; plausible experimental parameters chosen to illustrate the spectrum and the dynamics of the effective Hamiltonian, the full derivations, and the experimental considerations are given in the appendices.

1.1 Anyonic statistics and the generalized Jordan–Wigner transformation

One-dimensional lattice anyons are described by operators a_j obeying the deformed commutation relations

$$a_j a_k^\dagger - e^{\,i\phi\,\mathrm{sgn}(j-k)}\, a_k^\dagger a_j = \delta_{jk}, \tag{2}$$

$$a_j a_k \;-\; e^{-i\phi\,\mathrm{sgn}(j-k)}\, a_k a_j = 0, \tag{3}$$

where ϕ is the statistical angle. The sign function makes the algebra non-local: on a single site ($j = k$) the phase factor is unity and the operators are ordinary bosons, $[a_j, a_j^\dagger] = 1$, so that several anyons may occupy the same site, whereas on different sites the exchange produces $e^{\pm i\phi}$. Lattice anyons at $\phi = \pi$ are therefore *pseudofermions*: they are antisymmetric under exchange between different sites but still allow multiple occupancy [16, 30].

The anyons are obtained from bosonic operators b_j, with $n_j = b_j^\dagger b_j$, through the generalized Jordan–Wigner transformation

$$a_j = b_j\, \exp\Big(i\phi \sum_{l<j} n_l\Big), \qquad a_j^\dagger = b_j^\dagger\, \exp\Big(-i\phi \sum_{l<j} n_l\Big), \tag{4}$$

which attaches to each particle a string counting the total occupation to its left. Because the string runs only over sites $l < j$ it commutes with b_j, and the local occupation is left invariant, $n_j = a_j^\dagger a_j = b_j^\dagger b_j$. The full verification that (4) reproduces (2)–(3) is given in Appendix C.

1.2 Density-dependent tunnelling and the anyon-Hubbard Hamiltonian

The anyon-Hubbard model is the Hubbard model written in terms of these anyonic operators,

$$H = -J \sum_j \left(a_j^\dagger a_{j+1} + \mathrm{H.c.}\right) + \frac{U}{2} \sum_j n_j(n_j - 1). \tag{5}$$

Inserting (4), the strings on neighbouring sites cancel up to the occupation of the site being left, $\sum_{m<j+1} n_m - \sum_{l<j} n_l = n_j$, so that the tunnelling becomes $a_j^\dagger a_{j+1} = b_j^\dagger\, e^{i\phi n_j}\, b_{j+1}$,

while the interaction is untouched. The model is thus equivalent to a Bose–Hubbard model with a *density-dependent Peierls phase*,

$$H = -J \sum_j \left(b_j^\dagger \, e^{i\phi n_j} \, b_{j+1} + \text{H.c.} \right) + \frac{U}{2} \sum_j n_j (n_j - 1). \tag{6}$$

A boson tunnelling out of site j acquires a phase $\phi\, n_j$ set by the occupation of that site: the exchange statistics has been converted into an occupation-dependent tunnelling phase, and it is this bosonic form that is engineered in the laboratory [9, 10, 16, 17]. Note that a boson tunnelling to the left instead acquires $-\phi\, n_{j-1}$, so the model breaks spatial inversion symmetry. This asymmetry shows up in the momentum distribution [10, 17], but it is not on its own a signature of fractional statistics: a static gauge field with a phase ϕj linear in position shifts the momentum of the ground state away from zero and makes the momentum distribution asymmetric in exactly the same way. What distinguishes the anyonic case is that the phase is set by the *occupation* rather than by the position. For $\phi = 0$ one recovers the Bose–Hubbard model, and for $\phi = \pi$ the pseudofermionic limit.

1.3 Truncated local Hilbert space and the bond Hamiltonian

Everything that follows is derived on a single bond. This is not a restriction: the Hamiltonian (6) is a sum of identical nearest-neighbour terms, and the implementation of Sec. 2 is translationally invariant, so the effective Hamiltonian of the whole chain is obtained by repeating the one derived for the bond $(j, j+1)$ on every bond. It is therefore sufficient to work out a single bond in full.

The scheme developed below realizes a single bond of (6) with two atoms, each carrying at most two particles, $\{|0\rangle, |1\rangle, |2\rangle\}$, where the bosonic operators truncate to

$$b = |0\rangle\langle 1| + \sqrt{2}\, |1\rangle\langle 2|, \qquad n = |1\rangle\langle 1| + 2\, |2\rangle\langle 2|, \qquad b^3 = 0. \tag{7}$$

Such a two-body hard-core constraint is not merely a convenience: it is an ingredient of the proposed realizations [10] with neutral atoms in optical lattices, where triply occupied sites are strongly affected by three-body losses. It is nevertheless a genuine physical restriction rather than a purely technical one. The unconstrained model, in which arbitrarily many particles may share a site, has been studied separately and hosts additional structure, including an emergent statistically induced exclusion that stabilizes the density even at vanishing interaction, and a two-component partially paired phase near the pseudofermionic limit [31]. On a bond $(j, j+1)$ the density-dependent tunnelling connects $|n_j, n_{j+1}\rangle \to |n_j + 1,\, n_{j+1} - 1\rangle$ with amplitude $-J\, e^{i\phi n_j} \sqrt{n_{j+1}} \sqrt{n_j + 1}$, giving the four independent processes

$$\begin{aligned} |0,1\rangle &\to |1,0\rangle : \; -J, & |0,2\rangle &\to |1,1\rangle : \; -J\sqrt{2}, \\ |1,1\rangle &\to |2,0\rangle : \; -J\sqrt{2}\, e^{i\phi}, & |1,2\rangle &\to |2,1\rangle : \; -2J\, e^{i\phi}. \end{aligned} \tag{8}$$

The statistical phase appears only in those processes that leave a doubly-occupied site behind, which is the structure that any physical implementation must reproduce. Equation (8) is the target of the remainder of this thesis; the explicit 9×9 bond Hamiltonian is given in Appendix C.

2 Rydberg-atom implementation

2.1 Encoding the local states and the Raman coupling

The physical platform is an array of individually trapped atoms. Single atoms are held in optical tweezers imprinted by a spatial light modulator or an acousto-optic deflector; because the traps load stochastically, defect-free arrays are assembled by imaging the loaded configuration and rearranging the atoms into the target geometry, which routinely yields filling fractions above 98% for arrays of up to about a hundred atoms [27]. Each atom can then be addressed and read out individually, which is what makes a single bond $(j, j+1)$ an experimentally meaningful object.

The three local occupations $|0\rangle, |1\rangle, |2\rangle$ of each site are encoded in internal (Rydberg) levels of an atom, and the two sites $j = 0, 1$ of the bond are addressed by Raman lasers. The $0 \leftrightarrow 1$ and $1 \leftrightarrow 2$ transitions are driven with Rabi frequencies Ω_α and Ω_β and site-dependent phases,

$$H_\Omega = \sum_{j=0,1} \Big[\Omega_\alpha \big(e^{i\alpha_j} |1\rangle_j \langle 0|_j + e^{-i\alpha_j} |0\rangle_j \langle 1|_j \big) + \Omega_\beta \big(e^{i\beta_j} |2\rangle_j \langle 1|_j + e^{-i\beta_j} |1\rangle_j \langle 2|_j \big) \Big], \tag{9}$$

with the site-dependent phases $\alpha_j = j\alpha$ and $\beta_j = j\beta$, i.e. $\alpha_0 = \alpha$, $\alpha_1 = 2\alpha$, $\beta_0 = \beta$, $\beta_1 = 2\beta$ for the two atoms of the bond. In the taxonomy of Hubbard Hamiltonians realizable with cold atoms, (9) is the on-site Rabi term: the part that describes coherent transitions between the internal states of an atom, driven by laser-induced Raman fields, as distinct from the tunnelling term that moves atoms between sites [3]. A Raman coupling is essential here, and a microwave coupling would not do: the wavevector of a microwave field is of order 1/mm, so on the micrometre scale of the array it would give $\alpha_j = \beta_j = 0$ and no position dependence at all.

The reason for choosing a Raman process, rather than any other coupling between the internal levels, is that a Raman process imprints a phase that advances *linearly* with position. Each site j therefore carries $\alpha_j = j\alpha$ and $\beta_j = j\beta$, and in a super-exchange process built from two opposite Raman transitions on neighbouring sites the two phases do not cancel: what survives is a phase difference proportional to the separation of the sites, and hence a non-zero Peierls phase on the bond. This is the same mechanism that produces a constant synthetic flux in synthetic-dimension lattices, where a one-dimensional optical lattice dressed by Raman transitions becomes a synthetic two-dimensional lattice pierced by a uniform magnetic field, the tunnelling along the synthetic direction carrying a phase $e^{ij\gamma}$ that advances linearly with the spatial index j [32]. Through the encoding, the truncated anyon-Hubbard model can likewise be viewed as a three-leg ladder, though with an unusual pattern of tunnelling amplitudes.

Two properties of Rydberg states make this semi-synthetic ladder a natural one to build. Transition dipole moments between neighbouring levels scale as n^2, so that Rydberg states are coupled to one another by microwave or two-photon fields with Rabi frequencies of many megahertz at modest power, while radiative lifetimes grow as n^3 and reach hundreds of microseconds at $n \approx 60$ [26]. The phases α_j and β_j are then properties of the driving fields rather than of the atoms, and are set directly by the experimenter. Which three levels are used is constrained by the requirement, imposed in Sec. 2.2 below, that the Rydberg–Rydberg interaction be diagonal in the occupation basis. What the construction needs is that the initial and final states of a super-exchange process differ in energy only by the difference of their on-site Hubbard interactions, which is of the order of the second-order tunnelling amplitude itself and, by assumption, far smaller than the Rydberg–Rydberg interactions; measured against the latter the process is then effectively resonant. Two

encoded levels that are *resonantly dipole coupled*, by contrast, exchange their internal states already at first order, which would swamp the second-order process engineered here. The encoding must therefore be built from levels for which that first-order exchange is off-resonant. Multi-level Rydberg encodings of this kind, with one internal state per local occupation, are already in use—two S levels and one P level of ^{87}Rb realize a t–J–V model in a tweezer array [29].

Written out, the eight elementary processes are

$$H_\Omega = \Omega_\alpha e^{i\alpha}\,|1\rangle_0\langle 0|_0 + \Omega_\alpha e^{-i\alpha}\,|0\rangle_0\langle 1|_0 + \Omega_\beta e^{i\beta}\,|2\rangle_0\langle 1|_0 + \Omega_\beta e^{-i\beta}\,|1\rangle_0\langle 2|_0 \\ +\Omega_\alpha e^{i2\alpha}\,|1\rangle_1\langle 0|_1 + \Omega_\alpha e^{-i2\alpha}\,|0\rangle_1\langle 1|_1 + \Omega_\beta e^{i2\beta}\,|2\rangle_1\langle 1|_1 + \Omega_\beta e^{-i2\beta}\,|1\rangle_1\langle 2|_1 \,. \tag{10}$$

2.2 Rydberg interactions and the dispersive regime

Before the derivation begins it is useful to fix, once and for all, the parameter set used in every numerical evaluation and every figure of this thesis. It is collected in Table 1. These values are not derived from a specific set of ^{87}Rb levels: they are a *representative* set, chosen to be of the order of magnitude that dipole–dipole couplings between Rydberg states with principal quantum number $n \approx 50$–60 take at micrometre separations [26, 27], and constrained by the one relation the construction actually requires, the resonance condition $2C_{11} = C_{02}$ of (13). A quantitative determination of the C_{ij} for a named triple of levels would require a full pair-interaction calculation including the Förster defects and the angular dependence of each channel, which is beyond the scope of this thesis; what matters for the argument developed below is the *hierarchy* of scales, $C_{ij} \gg \Omega_{\alpha,\beta}$, and the resonance condition, not the individual numbers. Every energy quoted in megahertz therefore inherits its scale from Table 1, and rescaling all C_{ij} by a common factor rescales all second-order amplitudes by its inverse.

Symbol	Value	Unit	Meaning
C_{00}	0	MHz	interaction shift of $\|0\rangle\,\|0\rangle$ (reference)
C_{01}	20	MHz	interaction shift of $\|0\rangle\,\|1\rangle$
C_{11}	15	MHz	interaction shift of $\|1\rangle\,\|1\rangle$
C_{02}	30	MHz	interaction shift of $\|0\rangle\,\|2\rangle$; set by $C_{02} = 2C_{11}$
C_{12}	45	MHz	interaction shift of $\|1\rangle\,\|2\rangle$
C_{22}	55	MHz	interaction shift of $\|2\rangle\,\|2\rangle$
Ω_α	1	MHz	Rabi frequency of the $0 \leftrightarrow 1$ Raman coupling
Ω_β	1	MHz	Rabi frequency of the $1 \leftrightarrow 2$ Raman coupling
ω'	40	MHz	Floquet drive frequency (Sec. 3)
κ	1.2	—	modulation index of the drive (Sec. 3)

Table 1: Reference parameters assumed throughout this thesis, unless a figure states otherwise. The derived combination $2C_{11} - C_{01} = 10\,\text{MHz}$ sets the scale of the second-order denominators. See App. A for the values that follow from them, and Sec. 2.2 for the physical assumptions behind the hierarchy $C_{ij} \gg \Omega_{\alpha,\beta}$.

The two sites interact through configuration-dependent energy shifts,

$$H_c = \sum_{i,\,j\le i} C_{ij}\big(\,|i\rangle_0\,|j\rangle_1\langle i|_0\langle j|_1 + |j\rangle_0\,|i\rangle_1\langle j|_0\langle i|_1\,\big), \tag{11}$$

so that $H = H_\Omega + H_c$. Written out over the nine two-site configurations this reads

$$\begin{aligned} H_c = {} & 2C_{00}\ |0\rangle_0\,|0\rangle_1\langle 0|_0\langle 0|_1 \\ & + C_{01}\big(\,|1\rangle_0\,|0\rangle_1\langle 1|_0\langle 0|_1 + |0\rangle_0\,|1\rangle_1\langle 0|_0\langle 1|_1\,\big) \\ & + 2C_{11}\ |1\rangle_0\,|1\rangle_1\langle 1|_0\langle 1|_1 \\ & + C_{02}\big(\,|2\rangle_0\,|0\rangle_1\langle 2|_0\langle 0|_1 + |0\rangle_0\,|2\rangle_1\langle 0|_0\langle 2|_1\,\big) \\ & + C_{12}\big(\,|2\rangle_0\,|1\rangle_1\langle 2|_0\langle 1|_1 + |1\rangle_0\,|2\rangle_1\langle 1|_0\langle 2|_1\,\big) \\ & + 2C_{22}\ |2\rangle_0\,|2\rangle_1\langle 2|_0\langle 2|_1\,. \end{aligned} \tag{12}$$

H_c is therefore diagonal in the occupation basis, assigning each two-site state $|n_0, n_1\rangle$ its interaction energy: $E_{01} = C_{01}$, $E_{11} = 2C_{11}$, $E_{00} = 2C_{00}$, $E_{02} = C_{02}$, $E_{12} = C_{12}$, $E_{22} = 2C_{22}$. We work in the dispersive regime

$$|C_{ij} - C_{mn}| \gg \Omega \qquad (ij \neq mn), \tag{13}$$

so that single-hop (first-order) processes are strongly off-resonant and the dynamics within each degenerate manifold is governed by second-order (super-exchange) processes [26, 27]. Super-exchange of this kind is the standard route to magnetic couplings in optical lattices, where it has been resolved directly through the short-range spin correlations of fermions in a tunable lattice [33].

Condition (13) is the same hierarchy on which the original Rydberg gate rests, and the contrast between its two limits fixes what the second-order treatment of Sec. 2.3 is and is not doing. Reference [28] considers two atoms whose ground states are laser-coupled to a Rydberg level with Rabi frequency Ω, the doubly excited pair state being shifted by the dipole–dipole energy u. When $\Omega \gg u$ the laser simply transfers the population to the Rydberg level, the pair is held there for a time Δt, and the shift is read out as the dynamical phase $\varphi = u\,\Delta t$: the atoms spend that time actually excited and actually interacting, and pay for it with the mechanical force the interaction exerts on them. When instead $u \gg \Omega$ the excitation of the second atom is blocked while the first is excited, so the doubly excited state is populated only virtually; the conditional dynamics then arises because the pulse sequence acts differently on blocked and unblocked pairs, and in the adiabatic version of that gate the interaction enters only through the dispersive shift $\sim \Omega^2/u$ of the dressed level connected to the doubly excited configuration [28]. The scheme developed here works throughout in this second regime, of which (13) is the statement, but puts the virtual excursion to a different use. There it produces a phase on a pair state that does not change; here it must transport an excitation across the bond, so the object of interest is the *off-diagonal* second-order amplitude (15) rather than a diagonal shift—and, as Sec. 2.4 will show, it is this change of role that makes the accompanying phase so much harder to control.

The shifts C_{ij} originate in the electric dipole–dipole interaction $V_{\rm dd} \sim d_1 d_2/4\pi\epsilon_0 R^3$ between the two atoms, and take two qualitatively different forms depending on which pair of levels is involved [27]. When both atoms occupy the *same* Rydberg level, $V_{\rm dd}$ has no effect at first order, because such a state carries no average electric dipole moment; the interaction appears only at second order through virtual transitions to opposite-parity pair states detuned by Δ, giving a van der Waals shift C_6/R^6 with $C_6 \sim d^4/\Delta$. When the two atoms occupy *different*, dipole-coupled levels, for instance $|nS\rangle$ and $|nP\rangle$, the pair state is instead coupled directly to the exchanged configuration already at first order; this produces a resonant dipolar shift $\pm C_3/R^3$ together with the coherent exchange of the internal states, and is anisotropic, varying as $1 - 3\cos^2\Theta$ with Θ the angle between the internuclear axis and the quantization axis; this angular structure is common to all dipolar systems [34].

Both scales are large—van der Waals coefficients grow as n^{11}, and interaction energies of order 1–10 MHz are reached at micrometre separations for principal quantum numbers $n \approx 50$—which is what allows the splittings between different two-site configurations to satisfy (13) while the Rabi frequencies remain comparatively small. These two regimes are what underlie the standard mappings of Rydberg arrays onto quantum spin models: same-state van der Waals interactions realize Ising Hamiltonians, whereas dipole-coupled states realize XY models whose exchange term $\propto (\sigma_+^i\sigma_-^j + \sigma_-^i\sigma_+^j)$ is the spin analogue of a particle hop [27]. The density-dependent exchange engineered below is a generalization of this second case, in which the amplitude of the hop is made to depend on the occupations and to carry a tunable phase.

A third arrangement, the one adopted in Ref. [28], removes the distinction between the two cases by applying a constant electric field $\mathcal{E}$ along the quantization axis. The linear Stark effect splits a given manifold n into states $|n, q, m\rangle$ with $\Delta E_{nqm} = 3nqea_0\mathcal{E}/2$, and these Stark eigenstates carry a *permanent* dipole moment $\mu = 3nqea_0/2$ along the field, so that two atoms interact through V_{dd} already at first order even when they occupy the same level. For the outermost Stark state of the hydrogenic manifold the shift is $u(R) = -9[n(n-1)]^2(a_0/R)^3e^2/8\pi\epsilon_0 a_0$, growing as n^4 rather than as the n^{11} of the van der Waals case, and in alkalis n is replaced by the effective quantum number ν [28]. A field of this kind is a useful additional knob in the present context, because the six shifts of (12) are not free parameters: besides being mutually well separated in the sense of (13), they must satisfy the resonance condition $2C_{11} = C_{02}$ used throughout Sec. 2.3, and the field strength, the choice of Stark states and the geometry of the pair together supply the freedom to impose it. The price is that the C_{ij} then inherit the stability of the applied field, and that the nondiagonal part of V_{dd}, which couples adjacent m manifolds, must be suppressed by an appropriate choice of Stark eigenstate [28].

Throughout this thesis the two atoms are treated as pinned at a fixed separation R, so that the C_{ij} are numbers rather than operators and the problem reduces to one of internal levels alone. The conditions under which this is legitimate were set out in Ref. [28]. The first is that the extent a of the atomic wavepacket be small compared with the separation, $\eta = a/R \ll 1$, so that the interaction may be expanded about R; the leading correction is a state-dependent force, $F = 3u/R$ for a $1/R^3$ shift and $6u/R$ for a van der Waals one, which kicks the atoms whenever the corresponding configuration is populated. The second is that the protocol be short compared with the period of the motion in the trap, so that internal and motional degrees of freedom never become entangled.

The second condition is not met by a second-order protocol of the present kind: the exchange amplitudes computed below are of order 10^{-2}–10^{-1} MHz, so the bond dynamics unfolds over tens to hundreds of microseconds, far longer than a trap period. What matters here, however, is not the absolute force but its *difference* across the states that are actually populated, since a force common to a degenerate manifold displaces the atoms without distinguishing their internal states. In the $N = 1$ manifold $\{|0,1\rangle, |1,0\rangle\}$ the two states are related by exchange and carry the same shift C_{01}, so no differential force arises at all. The $N = 2$ manifold is more delicate. Its degeneracy is not a symmetry but the tuned condition $2C_{11} = C_{02}$, imposed at one separation, and if the two shifts have different distance dependences—as they do when $|1\rangle|1\rangle$ is a same-state pair interacting through van der Waals while $|0\rangle|2\rangle$ is dipole coupled—then it is broken as soon as the atoms move,

$$\delta(2C_{11} - C_{02}) \;\simeq\; (-6+3)\,C_{02}\,\frac{\delta R}{R} \;=\; -3\,C_{02}\,\frac{\delta R}{R}. \tag{14}$$

With $C_{02} \sim 1$ MHz a relative displacement of one per cent already detunes the resonance by $\sim$30 kHz, which is comparable to the exchange amplitudes themselves. Tight confinement

in the Lamb–Dicke sense of Ref. [28] is therefore not a technicality of the readout but a precondition for the resonance condition on which the whole construction rests, and the same estimate applies to the site-to-site uniformity of $2C_{11} - C_{02}$ required when the bond is extended to a chain.

2.3 Effective exchange from second-order perturbation theory

For a degenerate transition the second-order amplitude is

$$A_{i\to f} = \sum_m \frac{\langle f|\, H_\Omega\, |m\rangle\, \langle m|\, H_\Omega\, |i\rangle}{E_i - E_m}, \tag{15}$$

and for a non-degenerate transition ($E_i \neq E_f$) we use the symmetrized form

$$A_{i\to f} = \frac{1}{2}\sum_m \langle f|\, H_\Omega\, |m\rangle\, \langle m|\, H_\Omega\, |i\rangle \left(\frac{1}{E_i - E_m} + \frac{1}{E_f - E_m}\right). \tag{16}$$

Both expressions are instances of the standard construction of an effective low-energy Hamiltonian in the strong-coupling limit, in which a unitary $e^{i\hat{S}}$ is chosen to decouple manifolds separated by a large interaction energy and the series is developed to second order in the ratio of tunnelling to that energy [3]. The symmetric combination of the two energy denominators in (16) is the pair of resolvents $\hat{Q}_i + \hat{Q}_f$ that the general second-order formula places between the two tunnelling operators, and it keeps the effective Hamiltonian Hermitian when the initial and final states are not degenerate. Each transition proceeds through two intermediate paths (A and B) built from the same pair of elementary hops in opposite order; they share a common phase and differ only in their energy denominators. We use the resonance condition $2C_{11} = C_{02}$ throughout, and abbreviate

$$\begin{aligned} K_1 &= \frac{1}{C_{01} - C_{02}} + \frac{1}{C_{01} - 2C_{00}}, \qquad K_2 = \frac{1}{C_{02} - C_{12}} + \frac{1}{C_{02} - C_{01}}, \\ K_3 &= \frac{1}{C_{12} - 2C_{22}} + \frac{1}{C_{12} - C_{02}}. \end{aligned} \tag{17}$$

The four processes are worked out one at a time in App. D, where each is resolved into its two virtual paths and its energy denominators; Figs. 10–17 there show the path diagram, the two-site level picture and the photon picture of each. The results are collected in (18) below.

2.4 Effective Hamiltonian and the flux-locking restriction

Collecting the four amplitudes gives the effective second-order Hamiltonian at resonance,

$$\begin{aligned} H^{(2)R} = {} & \Omega_\alpha^2\, e^{-i\alpha}\, K_1 \big(\, |1,0\rangle\langle 0,1| + \text{H.c.}\big) + \Omega_\alpha\Omega_\beta\, e^{i(\alpha-2\beta)}\, K_2\ |1,1\rangle\langle 0,2| \\ & + \Omega_\alpha\Omega_\beta\, e^{i(\beta-2\alpha)}\, K_2\ |2,0\rangle\langle 1,1| + \Omega_\beta^2\, e^{-i\beta}\, K_3 \big(\, |2,1\rangle\langle 1,2| + \text{H.c.}\big) + \text{H.c.}. \end{aligned} \tag{18}$$

Matching each second-order amplitude to the corresponding target amplitude of the anyon-Hubbard bond, Eq. (8), the conditions on the drive parameters become

$$\begin{aligned} \Omega_\alpha^2 e^{-i\alpha} K_1 &= -J, & \Omega_\alpha\Omega_\beta\, e^{i(\alpha-2\beta)} K_2 &= -J\sqrt{2}, \\ \Omega_\alpha\Omega_\beta\, e^{i(\beta-2\alpha)} K_2 &= -J\sqrt{2}\, e^{i\phi}, & \Omega_\beta^2 e^{-i\beta} K_3 &= -2J\, e^{i\phi}. \end{aligned} \tag{19}$$

Dividing the third condition by the second fixes the synthetic flux,

$$e^{i[(\beta-2\alpha)-(\alpha-2\beta)]} = e^{i3(\beta-\alpha)} = e^{i\phi} \quad \Longrightarrow \quad \phi = 3(\beta-\alpha)\ (\mathrm{mod}\ 2\pi). \tag{20}$$

Consistency of the magnitudes follows from the first, second, and fourth conditions, $\Omega_\alpha^2|K_1| = J$, $\Omega_\alpha\Omega_\beta|K_2| = \sqrt{2}\,J$, $\Omega_\beta^2|K_3| = 2J$, which fixes

$$\frac{\Omega_\beta^2|K_3|}{\Omega_\alpha^2|K_1|} = 2 \qquad \text{and} \qquad \frac{(\Omega_\alpha\Omega_\beta|K_2|)^2}{\Omega_\alpha^2|K_1|\cdot\Omega_\beta^2|K_3|} = \frac{2J^2}{2J^2} = 1 \implies |K_2|^2 = |K_1||K_3|, \tag{21}$$

i.e. the interaction energies C_{ij} must be tuned so that $K_2^2 = K_1K_3$ (up to sign), while the Rabi frequencies satisfy $\Omega_\beta/\Omega_\alpha = \sqrt{2|K_1|/|K_3|}$.

The magnitude ratio also fixes

$$\frac{\Omega_\alpha\Omega_\beta|K_2|}{\Omega_\beta^2|K_3|} = \frac{\Omega_\alpha|K_2|}{\Omega_\beta|K_3|} = \frac{1}{\sqrt{2}}, \tag{22}$$

so that dividing the third condition of (19) by the fourth and using this ratio leaves $\mathrm{sgn}(K_2/K_3)\,e^{i2(\beta-\alpha)} = 1$. If K_2, K_3 share a sign (uniform tunnelling, no spurious π shifts) this requires $e^{i2(\beta-\alpha)} = 1$, i.e. $\beta - \alpha = \pi n$. With (20) the flux is therefore quantized,

$$\boxed{\phi = 3\pi n \implies \phi \in \{0, \pi\} \ (\mathrm{mod}\ 2\pi).} \tag{23}$$

A *static* second-order Raman protocol thus locks the statistical phase to the values $\{0, \pi\}$. This means that pseudofermions can be realized experimentally with this scheme, by taking $\alpha = 0$ and $\beta = \pi$. However, we cannot realize anyons with an arbitrary statistical phase ϕ (e.g. $\phi = \pi/2$), which is what is required to simulate chiral dynamics [19, 20]. Whether a time-dependent (Floquet) drive can lift this restriction is the subject of Secs. 3 and 4.

The conditions derived here are stated for a single bond, but they extend directly to every bond of a chain. The Raman phases are linear in position, $\alpha_j = j\alpha$ and $\beta_j = j\beta$, so the second-order process on the bond $(j, j+1)$ depends only on the phase *difference* across it and is therefore the same on every bond: the construction is translationally invariant, and it is enough to derive the effective Hamiltonian on one bond. The phase bookkeeping is summarized in Table 2.

Transition	Site-0 hop (phase)	Site-1 hop (phase)	Net phase
$\lvert 0,1\rangle \to \lvert 1,0\rangle$	$0 \to 1$ $(e^{i\alpha})$	$1 \to 0$ $(e^{-i2\alpha})$	$e^{-i\alpha}$
$\lvert 0,2\rangle \to \lvert 1,1\rangle$	$0 \to 1$ $(e^{i\alpha})$	$2 \to 1$ $(e^{-i2\beta})$	$e^{i(\alpha-2\beta)}$
$\lvert 1,1\rangle \to \lvert 2,0\rangle$	$1 \to 2$ $(e^{i\beta})$	$1 \to 0$ $(e^{-i2\alpha})$	$e^{i(\beta-2\alpha)}$
$\lvert 1,2\rangle \to \lvert 2,1\rangle$	$1 \to 2$ $(e^{i\beta})$	$2 \to 1$ $(e^{-i2\beta})$	$e^{-i\beta}$

Table 2: Phase accumulated in each second-order transition; the two paths share this phase and differ only in their energy denominators.

The whole construction of this section is summarized in Fig. 1, which shows how the off-resonant states are eliminated, which transitions survive inside the resonant manifold, and how each of them is assembled from two virtual paths.

3 Floquet engineering of the statistical phase

We now ask whether the restriction of Sec. 2.4 can be lifted by making the drive time dependent. Both doubly-occupied levels are periodically modulated and the resulting time-dependent problem is solved by transforming to a rotating frame, expanding the accumulated phase with the Jacobi–Anger identity, and applying a Magnus expansion in the secular approximation.

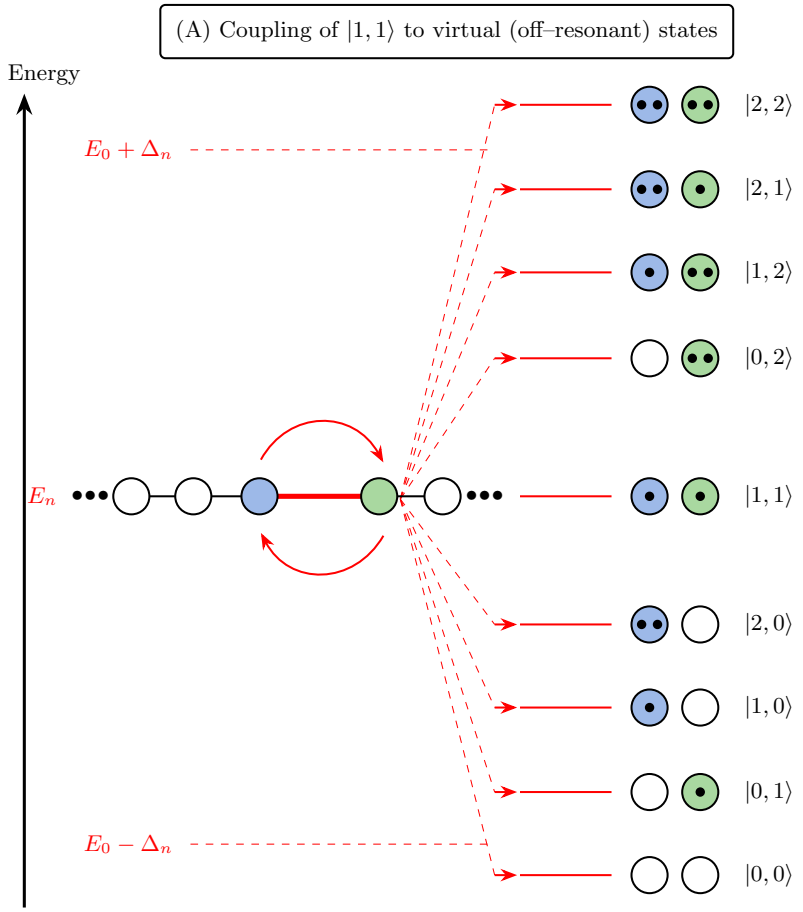


Figure 1: **One-dimensional Rydberg lattice: the active bond** $(j, j+1)$**.** Kets are written $|n_j, n_{j+1}\rangle$, the occupations of the two atoms of the bond. **(A) Encoding.** Each site j is represented by an atom in a tweezer driven between three Rydberg states, labelled 0, 1 and 2 and representing $n = 0, 1, 2$ particles respectively; the Raman beams Ω_α $(0 \leftrightarrow 1)$ and Ω_β $(1 \leftrightarrow 2)$ move one excitation across the bond, imprinting $e^{i\phi}$ on a hop $j \to j+1$ and $e^{-i\phi}$ on the reverse hop. Starting from $|1,1\rangle$ at energy E_n, a single hop reaches every other two-site configuration, but none of them is energy-matched: the interaction energies $C_{nn'}$ put them at $E_0 \pm \Delta_n$, the upper group being the higher virtual states and the lower group the lower ones. Because $|\Delta_n| \gg \Omega$ these states are never populated—they are virtual, entered and left within one drive cycle. Eliminating them to second order in Ω leaves an effective Hamiltonian acting only inside the resonant manifold, which is panel (B). Panels (B) and (C) are continued overleaf.

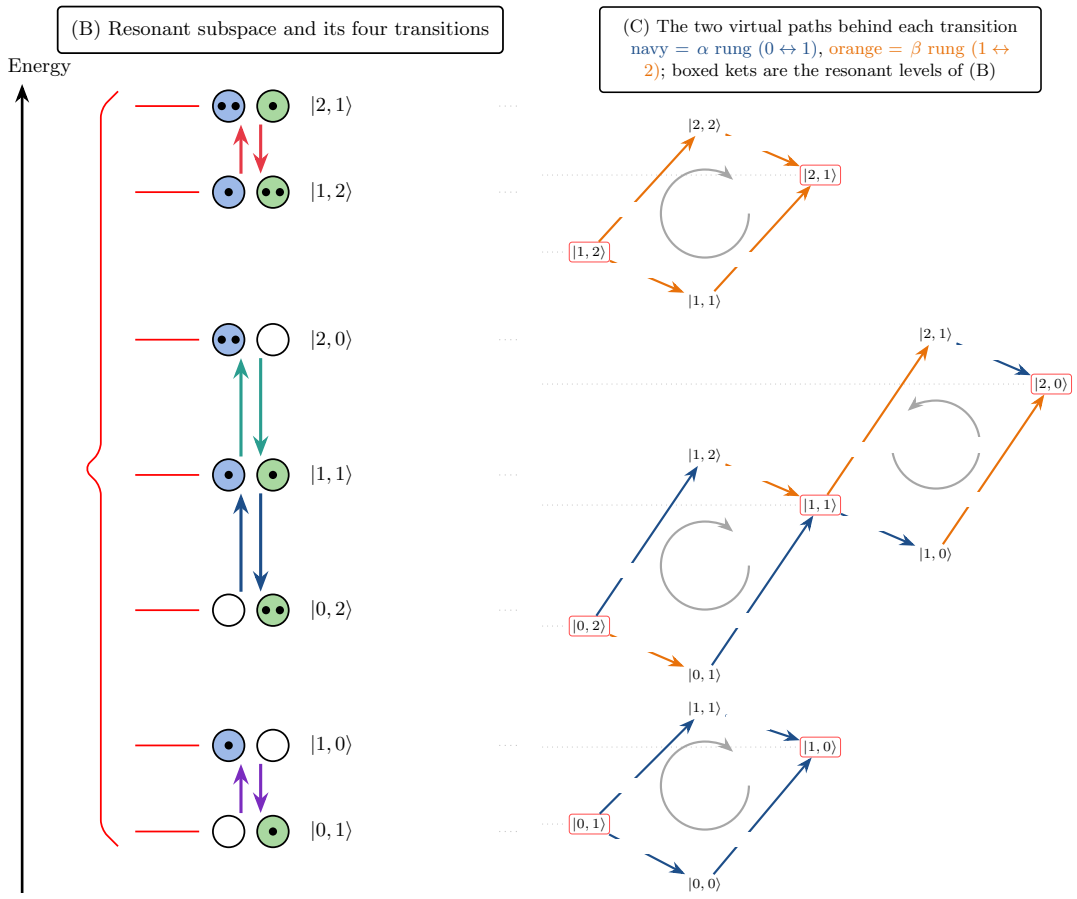


Figure 1: (continued) **(B) Effective Hamiltonian.** The drive conserves the total occupation $N = n_j + n_{j+1}$, so the resonant states, all lying within $\sim E_n$, split into blocks: $N = 3$ $\{|2,1\rangle, |1,2\rangle\}$, $N = 2$ $\{|2,0\rangle, |1,1\rangle, |0,2\rangle\}$ and $N = 1$ $\{|1,0\rangle, |0,1\rangle\}$, separated here by the wider gaps. Only four second-order processes connect states inside a block; every other pair of levels is left untouched. Each pair carries two directed arrows. Reading upward, the four upward transitions acquire $e^{-i\alpha}$ for $|0,1\rangle \to |1,0\rangle$, $e^{i(\alpha-2\beta)}$ for $|0,2\rangle \to |1,1\rangle$, $e^{i(\beta-2\alpha)}$ for $|1,1\rangle \to |2,0\rangle$ and $e^{-i\beta}$ for $|1,2\rangle \to |2,1\rangle$—the four net phases of Table 2—and the downward transitions their complex conjugates, as required by hermiticity. The three $N = 2$ levels are degenerate only when $2C_{11} = C_{02}$; that resonance condition is what selects density-assisted tunnelling and makes the phase occupation dependent. **(C) Second-order processes.** Every arrow of (B) is a two-photon process running through one virtual state above the pair and one below. Boxed kets are the resonant levels of (B) and sit at the same height (dotted guides); unboxed kets are virtual. Each leg is labelled by the phase it imprints and coloured by the rung it uses—navy for α, orange for β. The legs are not labelled individually: a leg acting on site j contributes $e^{+i\varphi}$ and one acting on site $j+1$ contributes $e^{-2i\varphi}$, where φ stands for the beam phase of the rung it uses, since $\varphi_j = \varphi$ and $\varphi_{j+1} = 2\varphi$. The two paths of a process use the same two rungs in opposite order, so they carry the same total phase and differ only in their denominators Δ_1, Δ_2—the interaction differences of (17), read off the pair of virtual states of each loop—giving $A = \Omega\Omega' e^{i\phi}\left(\Delta_1^{-1} + \Delta_2^{-1}\right)$. The circular arrow marks that closed loop and ϕ is the phase it encloses: reading along the drawn circulation, $\phi = 0$ for the $N = 3$ and $N = 1$ loops, whose two legs use the same rung, while $\phi = \alpha - 2\beta$ for $|0,2\rangle \to |1,1\rangle$ and $\phi = -(\beta - 2\alpha)$ for $|1,1\rangle \to |2,0\rangle$. Since the two $N = 2$ processes share $|1,1\rangle$, they merge into a single chain.

Periodic driving replaces the static Hamiltonian by a time-periodic one, $H(t+T) = H(t)$, for which the Floquet theorem, the temporal analogue of Bloch's theorem, guarantees solutions $|\Psi(t)\rangle = e^{-i\varepsilon t}\,|\Phi(t)\rangle$ with $|\Phi(t+T)\rangle = |\Phi(t)\rangle$ and a quasi-energy ε defined only modulo the drive frequency [19, 35]. Expanding both the Hamiltonian and the Floquet states in harmonics $e^{-im\Omega t}$ turns the time-dependent problem into a time-independent eigenvalue problem in an extended Hilbert space, in which the harmonic index m behaves as a coordinate along a fictitious "Floquet direction" and the drive frequency acts as a uniform field along it [35]. The physical content of the drive is then a ladder of sidebands separated by the drive frequency; when that frequency is large the rungs become energetically well separated and a high-frequency expansion—here the Magnus expansion of Sec. 3.4—is controlled [20, 35]. The photon index ℓ introduced below is this Floquet-direction label. The specific mechanism we exploit—that the *phase* of the drive is inherited by the induced tunnelling matrix element—is the same one behind the experimental realization of continuously tunable Peierls phases [21] and of the topological Haldane model, where the modulation phase fixes the imaginary part of the induced next-nearest-neighbour tunnelling [23].

Driven lattice systems fall into two broad classes [36]. In the first the drive frequency is off-resonant with every intrinsic energy scale, as in a uniformly shaken lattice. In the second—the one relevant here—large static energy offsets inhibit the bare tunnelling, and a modulation resonant with those offsets restores it in a controlled way, generating complex tunnelling matrix elements in the process. In the usual realization the offsets are provided by a superlattice or a Wannier–Stark ladder; here they are supplied instead by the Rydberg interaction energies C_{ij}, whose role in the dispersive condition (13) is to suppress the direct hops that the drive then restores. The formal treatment of such resonantly modulated systems proceeds by splitting the Hilbert space into sectors separated by multiples of $\hbar\omega$ and removing the offsets by a time-dependent unitary before expanding in $1/\omega$ [36], which is the route followed in Secs. 3.2–3.5.

3.1 The driven two-site Hamiltonian

The total Hamiltonian is split into a part that is diagonal in the occupation basis and the Raman coupling,

$$H(t) = \underbrace{H_c + H_{\mathrm{Floquet}}(t)}_{H_{\mathrm{diag}}(t)} + H_\Omega, \tag{24}$$

with H_Ω and H_c as in (10) and (12). The drive acts on the doubly-occupied level of each site,

$$H_{\mathrm{Floquet}}(t) = \sum_j \delta_{2j}(t)\;|2\rangle_j\langle 2|_j = \delta_{20}(t)\;|2\rangle_0\langle 2|_0 + \delta_{21}(t)\;|2\rangle_1\langle 2|_1\,, \tag{25}$$

where each modulation consists of a static offset plus a pulsed sinusoid of duty cycle T_j' within the period T_0,

$$\delta_{2j}(t) = \nu_j\hbar\omega_j + A\sum_{n=-\infty}^{\infty}\Theta(t-nT_0)\,\Theta(nT_0+T_j'-t)\,\cos(\omega'(t-nT_0)), \qquad \omega' = \frac{2\pi}{T_2'}. \tag{26}$$

Adding (25) to H_c gives the full diagonal part $H_{\mathrm{diag}}(t)$.

3.2 Rotating frame

We now transform to a frame rotating with the diagonal part of the Hamiltonian, the standard route to effective Hamiltonians for periodically driven lattices [19, 22]. Since

$H_{\rm diag}(t)$ is diagonal at all times, $[H_{\rm diag}(t), H_{\rm diag}(t')] = 0$ and no time ordering is required:

$$U_0(t) = \exp\Big[-i\int_0^t H_{\rm diag}(\tau)\,d\tau\Big] = \sum_{n,m} e^{-i\theta_{nm}(t)}\,|nm\rangle\langle nm|\,, \qquad \Phi_j(t) = \int_0^t \delta_{2j}(\tau)\,d\tau, \tag{27}$$

which written out over the nine configurations reads

$$\begin{aligned} U_0(t) = {} & e^{-i2C_{00}t}\,|00\rangle\langle 00| + e^{-iC_{01}t}\,|01\rangle\langle 01| + e^{-i[C_{02}t+\Phi_1(t)]}\,|02\rangle\langle 02| \\ & + e^{-iC_{01}t}\,|10\rangle\langle 10| + e^{-i2C_{11}t}\,|11\rangle\langle 11| + e^{-i[C_{12}t+\Phi_1(t)]}\,|12\rangle\langle 12| \\ & + e^{-i[C_{02}t+\Phi_0(t)]}\,|20\rangle\langle 20| + e^{-i[C_{12}t+\Phi_0(t)]}\,|21\rangle\langle 21| + e^{-i[2C_{22}t+\Phi_0(t)+\Phi_1(t)]}\,|22\rangle\langle 22|\,. \end{aligned} \tag{28}$$

Transforming the coupling, $H_{\rm rot}(t) = U_0^\dagger(t) H_\Omega U_0(t)$, leaves the α-rung terms carrying only their static phase and the bare detuning,

$$\begin{aligned} H_{\rm rot}^{(\alpha)}(t) = {} & \Omega_\alpha e^{i[\alpha+(C_{01}-2C_{00})t]}\,|10\rangle\langle 00| + \Omega_\alpha e^{i[\alpha+(2C_{11}-C_{01})t]}\,|11\rangle\langle 01| \\ & + \Omega_\alpha e^{i[\alpha+(C_{12}-C_{02})t]}\,|12\rangle\langle 02| + \Omega_\alpha e^{i[2\alpha+(C_{01}-2C_{00})t]}\,|01\rangle\langle 00| \\ & + \Omega_\alpha e^{i[2\alpha+(2C_{11}-C_{01})t]}\,|11\rangle\langle 10| + \Omega_\alpha e^{i[2\alpha+(C_{12}-C_{02})t]}\,|21\rangle\langle 20| + \text{H.c.}, \end{aligned} \tag{29}$$

while every β-rung term, which by construction couples to a doubly-occupied level, is dressed by the accumulated drive phase $\mathcal{F}_j(t) \equiv e^{i\Phi_j(t)}$,

$$\begin{aligned} H_{\rm rot}^{(\beta)}(t) = {} & \Omega_\beta e^{i[\beta+(C_{02}-C_{01})t]}\mathcal{F}_0(t)\,|20\rangle\langle 10| + \Omega_\beta e^{i[\beta+(C_{12}-2C_{11})t]}\mathcal{F}_0(t)\,|21\rangle\langle 11| \\ & + \Omega_\beta e^{i[\beta+(2C_{22}-C_{12})t]}\mathcal{F}_0(t)\,|22\rangle\langle 12| + \Omega_\beta e^{i[2\beta+(C_{02}-C_{01})t]}\mathcal{F}_1(t)\,|02\rangle\langle 01| \\ & + \Omega_\beta e^{i[2\beta+(C_{12}-2C_{11})t]}\mathcal{F}_1(t)\,|12\rangle\langle 11| + \Omega_\beta e^{i[2\beta+(2C_{22}-C_{12})t]}\mathcal{F}_1(t)\,|22\rangle\langle 21| + \text{H.c.}, \end{aligned} \tag{30}$$

with $H_{\rm rot} = H_{\rm rot}^{(\alpha)} + H_{\rm rot}^{(\beta)}$.

3.3 The Floquet factor

Integrating the drive (26) to obtain the accumulated phase $\Phi_j(t) = \int_0^t \delta_{2j}(\tau)\,d\tau$ and expanding the result with the Jacobi–Anger identity gives, in the continuous-drive limit $T_j' = T_0$,

$$\mathcal{F}_j(t) = e^{i\Phi_j(t)} = \sum_{\ell=-\infty}^{\infty} J_\ell(\kappa)\, e^{i(\tilde\nu_j + \ell\omega')t}, \qquad \kappa \equiv \frac{A}{\omega'}, \qquad \tilde\nu_j \equiv \nu_j \hbar\omega_j. \tag{31}$$

The drive therefore splits each β-rung coupling into an infinite ladder of sidebands separated by ω' and weighted by Bessel functions of the modulation index κ. This Bessel renormalization is the same effect that produces *dynamic localization* in shaken lattices, where a sinusoidal force rescales the bare tunnelling matrix element by a Bessel function of the scaled driving amplitude and can suppress it entirely at the zeros of that function [19]. Here the factors $J_\ell(\kappa)$ play the same role for each sideband, and it is this amplitude control—independent of the phase control provided by ϕ_f—that is used in Sec. 3.6 to satisfy the magnitude conditions. The general pulsed case of finite duty cycle, the stroboscopic condition, and the derivation of (31) are given in Appendix E. This continuous-drive limit is assumed from here on, since it is what makes a rigorous Magnus expansion possible.

3.4 Magnus expansion and Fourier mode decomposition

What we compute is the high-frequency approximation to the Floquet Hamiltonian H^F, defined by requiring that $\exp\left(-iTH^F/\hbar\right)$ reproduce the evolution over one driving period, so that stroboscopically the driven system behaves as if governed by a time-independent Hamiltonian [19, 37]. Expanding it in the Magnus series $H_{\text{eff}} = H^{(1)}_{\text{eff}} + H^{(2)}_{\text{eff}}$, the first order vanishes for the same reason as in the static Raman case of Sec. 2.3. There the interaction shifts C_{ij} were chosen, through (13), so that every single-atom transition is off resonance, which is what removes the first-order term and leaves the second-order exchange as the leading process. Here the same choice, made in the rotating frame, means that every mode $\hat{V}_\mu$ of H_{rot} carries a strictly non-zero frequency ω_μ: the drive contributes multiples of ω' and the interactions contribute the offsets C_{ij}, and no combination of them vanishes. The period average therefore has no term to survive it,

$$H^{(1)}_{\text{eff}} = \frac{1}{T}\int_0^T H_{\text{rot}}(t)\,dt = 0, \qquad H^{(2)}_{\text{eff}} = \frac{-i}{2T}\int_0^T dt_1 \int_0^{t_1} dt_2 \left[H_{\text{rot}}(t_1), H_{\text{rot}}(t_2)\right]. \tag{32}$$

To evaluate the time-ordered integrals we decompose H_{rot} into transition modes, separating the time-independent operators from their dynamical phases,

$$H_{\text{rot}}(t) = \sum_\mu \hat{V}_\mu e^{i\omega_\mu t}, \tag{33}$$

where μ indexes a transition pathway and hermiticity guarantees a conjugate mode with $\hat{V}_{-\mu} = \hat{V}^\dagger_\mu$ and $\omega_{-\mu} = -\omega_\mu$. A decomposition of this kind into Fourier harmonics is the standard starting point of the inverse-frequency expansion, in which the n-th order contribution to the Floquet Hamiltonian is of order Ω^{-n} and is built from nested commutators of the harmonics [37].

Two caveats attach to this construction. First, what the Magnus series produces is the *stroboscopic* Floquet Hamiltonian, whose higher-order corrections depend on the Floquet gauge—the choice of initial time within the drive cycle—while the micromotion between stroboscopic times is carried by a separate kick operator; a manifestly gauge-invariant expansion for the effective Hamiltonian is also available [37]. Since we compare populations at stroboscopic times, the stroboscopic form is the appropriate one here. Second, the Magnus series is not guaranteed to converge, and in general it is not known whether it is asymptotic or has a finite radius of convergence; for the finite-time evolution considered here, with a drive frequency large compared with the effective couplings, the short-period expansion is nonetheless well behaved [37].

The two technical tools used in this thesis are closely related. The adiabatic elimination of Sec. 2.3 is a Schrieffer–Wolff transformation, and the inverse-frequency expansion used here is its extension to driven systems, the large interaction splitting (13) and the large drive frequency playing formally parallel roles [37]. The static α-drives form single-frequency modes, for example

$$\hat{V}_{10,00} = \Omega_\alpha e^{i\alpha}\,|10\rangle\langle 00|\,,\ \omega_{10,00} = C_{01} - 2C_{00}; \qquad \hat{V}_{11,10} = \Omega_\alpha e^{i2\alpha}\,|11\rangle\langle 10|\,,\ \omega_{11,10} = 2C_{11} - C_{01}, \tag{34}$$

whereas the Floquet-modulated β-drives expand into an infinite ladder of modes labelled by the photon index ℓ, each rung displaced in quasi-energy by $\ell\omega'$ and weighted by $J_\ell(\kappa)$ — the sideband structure of the extended Floquet space [35],

$$\begin{aligned} \hat{V}^{(\ell)}_{20,10} &= \Omega_\beta J_\ell(\kappa) e^{i\beta}\,|20\rangle\langle 10|\,, & \omega^{(\ell)}_{20,10} &= C_{02} - C_{01} + \tilde{\nu}_0 + \ell\omega', \\ \hat{V}^{(\ell)}_{12,11} &= \Omega_\beta J_\ell(\kappa) e^{i2\beta}\,|12\rangle\langle 11|\,, & \omega^{(\ell)}_{12,11} &= C_{12} - 2C_{11} + \tilde{\nu}_1 + \ell\omega'. \end{aligned} \tag{35}$$

3.5 Secular approximation and commutators

Substituting (33) into (32) and performing the inner integral, $\int_0^{t_1} e^{i\omega_\nu t_2} dt_2 = (e^{i\omega_\nu t_1} - 1)/i\omega_\nu$, gives

$$H_{\text{eff}}^{(2)} = \frac{-1}{2T} \sum_{\mu,\nu} \frac{[\hat{V}_\mu, \hat{V}_\nu]}{\omega_\nu} \int_0^T \Big(e^{i(\omega_\mu+\omega_\nu)t_1} - e^{i\omega_\mu t_1} \Big) dt_1. \tag{36}$$

In the limit of many driving periods $\frac{1}{T}\int_0^T e^{i\Omega t} dt = \delta_{\Omega,0}$, so the first term survives only for $\omega_\mu = -\omega_\nu$ while the second vanishes identically, and the infinite sum collapses to

$$H_{\text{eff}}^{(2)} = \sum_{\mu,\nu} \frac{-[\hat{V}_\mu, \hat{V}_\nu]}{2\omega_\nu} \, \delta(\omega_\mu + \omega_\nu). \tag{37}$$

Two classes of term contribute. *Type I* pairs $\hat{V}_\mu$ and $\hat{V}_\nu$ that are conjugates, which satisfy $\omega_\mu + \omega_\nu = 0$ inherently; using $[|a\rangle\langle b|, |b\rangle\langle a|] = |a\rangle\langle a| - |b\rangle\langle b|$ they generate AC Stark shifts, e.g.

$$\Delta E_{11} = \frac{|\hat{V}_{11,01}|^2}{\omega_{11,01}} + \frac{|\hat{V}_{11,10}|^2}{\omega_{11,10}} - \sum_\ell \frac{|\hat{V}_{21,11}^{(\ell)}|^2}{\omega_{21,11}^{(\ell)}} - \sum_\ell \frac{|\hat{V}_{12,11}^{(\ell)}|^2}{\omega_{12,11}^{(\ell)}}. \tag{38}$$

Type II pairs form a pathway across three states $|a\rangle \to |b\rangle \to |c\rangle$,

$$H_{\text{off-diag}}^{(2)} = \tfrac{1}{2} \hat{V}_{ab} \hat{V}_{bc} \Big(\frac{1}{\omega_{ab}} - \frac{1}{\omega_{bc}} \Big) \delta(\omega_{ab} + \omega_{bc}) \, |a\rangle\langle c| \,, \tag{39}$$

and survive only when initial and final states are bridged by an integer number of drive quanta. For the cross-site $|12\rangle \leftrightarrow |21\rangle$ transition through $|11\rangle$ this imposes

$$\omega_{21,11}^{(\ell)} + \omega_{11,12}^{(\ell')} = 0 \quad \Longrightarrow \quad \ell - \ell' = \frac{\tilde{\nu}_1 - \tilde{\nu}_0}{\omega'} \equiv q, \tag{40}$$

so the term exists only if the commensurability parameter q is an integer, which is enforced below by a Kronecker delta $\delta_{q,\mathbb{Z}}$.

3.6 Final second-order effective Hamiltonian

Assembling all commutators gives $H_{\text{eff}}^{(2)} = H_{\text{diag}}^{(2)} + H_{\text{off-diag}}^{(2)}$. To write the diagonal block compactly we abbreviate the α-type shift and the Bessel-weighted β-type sum by

$$\mathcal{A}_X \equiv \frac{\Omega_\alpha^2}{X}, \qquad \mathcal{S}_j(X) \equiv \Omega_\beta^2 \sum_{\ell=-\infty}^{\infty} \frac{J_\ell^2(\kappa)}{X + \tilde{\nu}_j + \ell\omega'}. \tag{41}$$

The AC Stark shifts are then

$$\begin{aligned}
H_{\text{diag}}^{(2)} = \; & -2\mathcal{A}_{C_{01}-2C_{00}} \, |00\rangle\langle 00| \\
& + \left[\mathcal{A}_{C_{01}-2C_{00}} - \mathcal{A}_{2C_{11}-C_{01}} - \mathcal{S}_1(C_{02}-C_{01})\right] |01\rangle\langle 01| \\
& + \left[\mathcal{S}_1(C_{02}-C_{01}) - \mathcal{A}_{C_{12}-C_{02}}\right] |02\rangle\langle 02| \\
& + \left[\mathcal{A}_{C_{01}-2C_{00}} - \mathcal{A}_{2C_{11}-C_{01}} - \mathcal{S}_0(C_{02}-C_{01})\right] |10\rangle\langle 10| \\
& + \left[2\mathcal{A}_{2C_{11}-C_{01}} - \mathcal{S}_0(C_{12}-2C_{11}) - \mathcal{S}_1(C_{12}-2C_{11})\right] |11\rangle\langle 11| \\
& + \left[\mathcal{A}_{C_{12}-C_{02}} + \mathcal{S}_1(C_{12}-2C_{11}) - \mathcal{S}_0(2C_{22}-C_{12})\right] |12\rangle\langle 12| \\
& + \left[\mathcal{S}_0(C_{02}-C_{01}) - \mathcal{A}_{C_{12}-C_{02}}\right] |20\rangle\langle 20| \\
& + \left[\mathcal{A}_{C_{12}-C_{02}} + \mathcal{S}_0(C_{12}-2C_{11}) - \mathcal{S}_1(2C_{22}-C_{12})\right] |21\rangle\langle 21| \\
& + \left[\mathcal{S}_0(2C_{22}-C_{12}) + \mathcal{S}_1(2C_{22}-C_{12})\right] |22\rangle\langle 22| \,.
\end{aligned} \tag{42}$$

Every off-diagonal term is conditionally gated by a Kronecker delta, guaranteeing conservation of quasi-energy up to Floquet quanta,

$$\begin{aligned} H^{(2)}_{\text{off-diag}} = {} & \Omega_\alpha^2\Big[\tfrac{1}{C_{01}-2C_{00}} - \tfrac{1}{2C_{11}-C_{01}}\Big]\big(e^{-i\alpha}\,|10\rangle\langle 01| + e^{i\alpha}\,|01\rangle\langle 10|\,\big) \\ & + \delta_{\mathcal{N}_0,\mathbb{Z}}\cdot\Omega_\alpha\Omega_\beta\Big[\tfrac{1}{2C_{11}-C_{01}} - \tfrac{1}{C_{12}-C_{02}}\Big]J_{\mathcal{N}_0}(\kappa)\big(e^{i(\beta-2\alpha)}\,|20\rangle\langle 11| + \text{H.c.}\big) \\ & + \delta_{\mathcal{N}_1,\mathbb{Z}}\cdot\Omega_\alpha\Omega_\beta\Big[\tfrac{1}{2C_{11}-C_{01}} - \tfrac{1}{C_{12}-C_{02}}\Big]J_{\mathcal{N}_1}(\kappa)\big(e^{i(\alpha-2\beta)}\,|11\rangle\langle 02| + \text{H.c.}\big) \\ & + \delta_{q,\mathbb{Z}}\cdot\Omega_\beta^2\Big[\sum_{\ell=-\infty}^{\infty} J_\ell(\kappa)J_{\ell-q}(\kappa)\Sigma_\ell\Big]\big(e^{-i\beta}\,|21\rangle\langle 12| + \text{H.c.}\big), \end{aligned} \tag{43}$$

where the first line is the degenerate exchange ($\Delta E = 0$), the second and third are the Floquet-assisted transitions $|11\rangle \leftrightarrow |20\rangle$ and $|02\rangle \leftrightarrow |11\rangle$, and the last is the cross-site commensurate transition $|12\rangle \leftrightarrow |21\rangle$. The integer resonance parameters count the Floquet quanta needed to bridge each gap,

$$\mathcal{N}_0 = \frac{2C_{11} - C_{02} - \tilde{\nu}_0}{\omega'}, \qquad \mathcal{N}_1 = \frac{2C_{11} - C_{02} - \tilde{\nu}_1}{\omega'}, \qquad q = \frac{\tilde{\nu}_1 - \tilde{\nu}_0}{\omega'}, \tag{44}$$

and the virtual-state summation factor for the cross-site coupling is

$$\Sigma_\ell = \frac{1}{C_{12} - 2C_{11} + \tilde{\nu}_0 + \ell\omega'} - \frac{1}{2C_{22} - C_{12} + \tilde{\nu}_0 + \ell\omega'}. \tag{45}$$

Comparing (43) with the target amplitudes (8), the statistical phase now enters through the combinations $e^{i(\beta-2\alpha)}$ and $e^{i(\alpha-2\beta)}$ weighted by the Bessel factors $J_{\mathcal{N}_0}(\kappa)$ and $J_{\mathcal{N}_1}(\kappa)$, whose magnitudes are set independently by the modulation index κ.

3.7 Resolving the spectrum by particle number

Before turning to the dynamics it is worth asking what the spectrum of $H^{(2)}_{\text{eff}}$ looks like once the two-site operator is promoted to a chain, $H = \sum_j H^{(2)}_{\text{eff},\,j,j+1}$. Every term of (42) and (43) moves exactly one quantum between neighbouring sites, so the total particle number

$$\hat{N} = \sum_j \hat{n}_j, \qquad \hat{n}_j = \sum_{n=0}^{2} n\;|n\rangle_j\langle n|_j\,, \tag{46}$$

commutes with $H^{(2)}_{\text{eff}}$. The Hamiltonian is therefore block diagonal in N and every eigenstate carries a definite *integer* N: the pair (E_m, N_m) is unambiguous, and no state has to be assigned a fractional particle number. This makes the exact diagonalization cheap enough to be done in full—each of the $2L+1$ blocks is diagonalized separately and the blocks concatenated, giving all 3^L levels rather than only the low-lying ones, since the largest block has dimension 8953 rather than 59 049 for $L = 10$. This is the standard route to making exact diagonalization tractable: whenever an operator commutes with the Hamiltonian, the Hilbert space splits into sectors of definite eigenvalue and the single large diagonalization is replaced by several small ones, which both reduces the memory requirement and gives access to a far larger number of excited states than a sparse solver would return [3].

Figure 2 shows the complete spectrum for a chain of $L = 10$ sites, i.e. all $3^{10} = 59\,049$ levels, at the parameters used throughout this section. Panel (a) plots E against the eigenvalue of $\hat{N}$; panel (b) is the intensive version, E/L against the average number density

$$\bar{n} = \hat{N}/L, \tag{47}$$

whose eigenvalues are the values N/L, spaced by $1/L$ and bounded above by $\bar{n} = 2$ because the local space is truncated at $|2\rangle$. Dividing *both* axes by L is what makes panel (b) more than a relabelling of panel (a): it is the combination that stays finite as L grows, and hence the form in which chains of different length may be compared.

Two features matter for what follows. The energies below are obtained by diagonalizing the effective Hamiltonian of Sec. 3.6 numerically at the reference parameters of Table 1; they inherit their absolute scale from the C_{ij} and $\Omega_{\alpha,\beta}$ assumed there, and are quoted in megahertz for that reason alone. First, the lower envelope $E_0(N)$ is flat at small filling: the sectors $N = 0, \ldots, 4$ span only -0.900, -0.912, -0.919, -0.919, $-0.910\,\mathrm{MHz}$, a spread of $0.019\,\mathrm{MHz}$ across five different particle numbers, and the global minimum falls at $N = 2$ rather than at unit filling. The lowest six levels of the whole spectrum are distributed over the four sectors $N = 1, 2, 3, 4$, and the global gap $E_1 - E_0 = 1.3 \times 10^{-4}\,\mathrm{MHz}$ connects $N = 2$ to $N = 3$: it is an *inter*-sector spacing, not an excitation at fixed filling. States of different particle number are thus nearly degenerate at the bottom of the band, so which sector holds the global minimum is decided by small changes in the drive parameters, and any statement made about the unconstrained ground state is correspondingly fragile. The remedy is to work at *fixed* filling, which is legitimate precisely because $\hat{N}$ is conserved exactly—restricting to one block is not an approximation—and this is what makes the sector-resolved treatment of the dynamics in Sec. 3.8 well posed. Second, the spectrum is strongly asymmetric in N even though the Hilbert space is particle-hole symmetric in dimension, $\dim(N) = \dim(2L - N)$: the total bandwidth is $1.436\,\mathrm{MHz}$ and the highest level lies in $N = 15$ rather than in the largest sector $N = 10$. That asymmetry comes from the AC-Stark diagonal (42), which is not particle-number neutral, and it is the same feature that flattens the low-filling envelope.

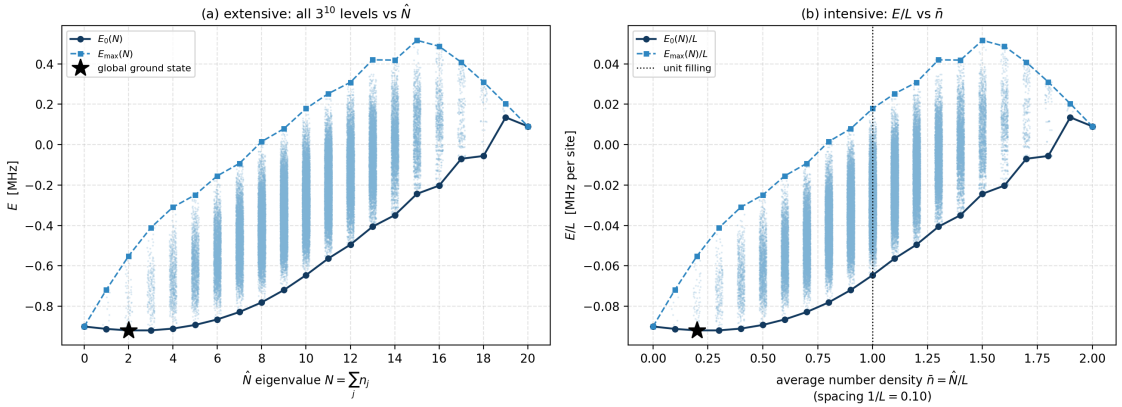


Figure 2: Spectrum of the effective chain $H_{\mathrm{eff}}^{(2)}$ plotted against the conserved particle number, for $L = 10$ sites with open boundaries. All $3^{10} = 59\,049$ eigenvalues are shown (pale points, jittered horizontally inside each sector) together with the sector envelopes $E_0(N)$ and $E_{\max}(N)$; the star marks the global ground state. (a) Extensive view: E versus the eigenvalue N of $\hat{N}$, Eq. (46). (b) Intensive view: E/L versus the average number density $\bar{n} = \hat{N}/L$, Eq. (47), whose eigenvalues are spaced by $1/L = 0.1$ and bounded by $\bar{n} = 2$ through the truncation at $|2\rangle$; the dotted line marks unit filling. Parameters as in Fig. 3, with $\omega' = 40\,\mathrm{MHz}$ and drive phases $\alpha = 0$, $\beta = \theta = 0$. Because $[H_{\mathrm{eff}}^{(2)}, \hat{N}] = 0$ the spectrum splits into $2L + 1$ blocks and each level carries an integer N, so the figure contains the entire spectrum and not a truncation of it. The envelope $E_0(N)$ is flat to within $0.019\,\mathrm{MHz}$ over $N = 0, \ldots, 4$, so the lowest levels come from several different sectors and the global gap is an inter-sector spacing.

3.8 Time evolution

Because $H_{\text{eff}}^{(2)}$ is static, the dynamics follows from the matrix exponential

$$|\psi(t)\rangle = e^{-iH_{\text{eff}}^{(2)}t}\,|\psi(0)\rangle\,. \tag{48}$$

This describes the evolution at stroboscopic times. The rapid micromotion within a driving period is carried by a separate kick operator, and its effect on observables depends on the details of the modulation scheme, so that the effective Hamiltonian alone does not in general capture the full time evolution [36, 37]; the curves below should be read accordingly. Figure 3 shows the resulting population dynamics for each initial configuration of the bond.

The oscillations are the direct signature of the engineered exchange. Because $H_{\text{eff}}^{(2)}$ conserves the total occupation $N = n_0 + n_1$, the dynamics decomposes into the manifolds identified in Fig. 1(B): the $N = 1$ and $N = 3$ blocks are two-level systems in which the population flops completely between $|01\rangle$ and $|10\rangle$, and between $|12\rangle$ and $|21\rangle$, at a rate set by the corresponding off-diagonal element of (43); the $N = 2$ block forms the three-state chain $|02\rangle \leftrightarrow |11\rangle \leftrightarrow |20\rangle$, whose two legs are the Floquet-assisted transitions and whose beat pattern reflects their relative amplitude; and $|22\rangle$ is uncoupled and therefore stationary. Reading off the oscillation periods thus gives the effective tunnelling rates directly, while the asymmetry between the two legs of the $N = 2$ chain is where the statistical phase enters.

Physically these are resonant dipolar exchange oscillations. Two Rydberg atoms prepared in different, dipole-coupled states have their pair state directly coupled to the exchanged configuration by the dipole–dipole interaction, so that an atom pair initially in one arrangement coherently evolves into the other and back—the "flip-flop" oscillation, with a rate scaling as C_3/R^3 in the resonant dipolar case and C_6/R^6 for van der Waals coupled pairs [27]. This is the process that the second-order treatment of Sec. 2.3 produces, dressed here by the drive. Oscillations of this form have been resolved between individual atom pairs in Rydberg tweezer arrays. There the number of atoms never changes; what is followed in real time is the population of each internal level of one atom as the pair coherently exchanges its internal states, and the exchange rate is tuned through the orientation of the pair relative to the quantization axis [29]. The populations themselves are directly observable: both tweezer arrays and quantum-gas microscopes provide single-site resolved readout of the atomic state, which is what makes site-resolved occupations and their correlations experimentally accessible [27].

Two experimental scales should be kept in mind when reading Fig. 3. The interaction energies assumed there (C_{ij} of tens of MHz) are of the order reached between Rydberg atoms at micrometre separations, where interaction strengths of $\gtrsim 1$ MHz are routine for principal quantum numbers $n \approx 50$ [27]. Rydberg lifetimes, however, are only of the order of a few hundred microseconds [27], so while the fast oscillations of the $N = 1$ manifold (periods of a few tens of microseconds) are comfortably resolved within the coherence time, the slower $N = 3$ traces extending beyond a millisecond exceed it; observing those would require either stronger couplings, obtained at smaller separations or higher n, or a longer-lived encoding such as Rydberg dressing. The trade-off is the one identified already in the original gate proposal [28]: the dipolar interaction grows as ν^4 with the effective principal quantum number while spontaneous emission and photoionization by the driving lasers fall as ν^{-3}, so the ratio of useful coupling to loss improves rapidly with ν, and the natural strategy is to push the interaction up rather than the decay down. How far this can be taken is limited by the geometry rather than by the atoms: with $\nu = 18$, $R = 300$ nm and a field of 100 V/cm, Ref. [28] quotes $u = 1.8$ GHz, three orders of magnitude beyond the shifts

assumed here, but at such separations the single-site addressing on which the encoding of Sec. 2.1 depends is no longer available, and the micrometre spacings of a tweezer array are the compromise. One caveat attaches to the scaling argument: the estimates of Ref. [28] neglect blackbody-induced transitions, which is justified there for $\nu \lesssim 20$ but not at the $n \approx 50$ assumed here, where blackbody redistribution contributes appreciably to the finite lifetime quoted above.

3.8.1 Analytic solution within each manifold

Because $H_{\mathrm{eff}}^{(2)}$ is block diagonal in N, each manifold can be solved in closed form and the curves of Fig. 3 admit exact expressions. The two-level blocks give generalized Rabi oscillations, in which the residual detuning between the two levels—itself generated by the AC Stark shifts (42)—reduces the contrast.

Degenerate exchange $|01\rangle \leftrightarrow |10\rangle$. With the coupling read off from the first line of (43),

$$V_{01,10} = \Omega_\alpha^2 \left(\frac{1}{C_{01} - 2C_{00}} - \frac{1}{2C_{11} - C_{01}} \right), \tag{49}$$

the populations are

$$P_{10}(t) = \frac{4V_{01,10}^2}{\Delta_{01,10}^2 + 4V_{01,10}^2} \sin^2\!\left(\frac{t}{2} \sqrt{\Delta_{01,10}^2 + 4V_{01,10}^2} \right), \qquad P_{01}(t) = 1 - P_{10}(t), \tag{50}$$

which manifestly conserve probability within the block. Subtracting the diagonal shift of $|10\rangle$ from that of $|01\rangle$, the Ω_α contributions cancel exactly and the detuning is purely Floquet-assisted,

$$\Delta_{01,10} = \Omega_\beta^2 \sum_{\ell=-\infty}^{\infty} J_\ell^2(\kappa) \left(\frac{1}{C_{02} - C_{01} + \tilde{\nu}_0 + \ell\omega'} - \frac{1}{C_{02} - C_{01} + \tilde{\nu}_1 + \ell\omega'} \right). \tag{51}$$

Cross-site commensurate exchange $|12\rangle \leftrightarrow |21\rangle$. Provided the commensurability condition $\delta_{q,\mathbb{Z}} = 1$ holds, the same structure applies with the coupling of the last line of (43), $V_{12,21} = \Omega_\beta^2 \sum_\ell J_\ell(\kappa) J_{\ell-q}(\kappa) \Sigma_\ell$, giving

$$P_{21}(t) = \frac{4V_{12,21}^2}{\Delta_{12,21}^2 + 4V_{12,21}^2} \sin^2\!\left(\frac{t}{2} \sqrt{\Delta_{12,21}^2 + 4V_{12,21}^2} \right), \qquad P_{12}(t) = 1 - P_{21}(t), \tag{52}$$

with, again after exact cancellation of the Ω_α terms,

$$\begin{aligned} \Delta_{12,21} = \Omega_\beta^2 \sum_{\ell=-\infty}^{\infty} J_\ell^2(\kappa) \Bigg(& \frac{1}{C_{12} - 2C_{11} + \tilde{\nu}_1 + \ell\omega'} - \frac{1}{C_{12} - 2C_{11} + \tilde{\nu}_0 + \ell\omega'} \\ & + \frac{1}{2C_{22} - C_{12} + \tilde{\nu}_1 + \ell\omega'} - \frac{1}{2C_{22} - C_{12} + \tilde{\nu}_0 + \ell\omega'} \Bigg). \end{aligned} \tag{53}$$

Floquet-assisted V-type block $|11\rangle \leftrightarrow \{|02\rangle, |20\rangle\}$. Here the central state is coupled to two others, with amplitudes differing only through the Bessel factors $J_{\mathcal{N}_0}(\kappa)$ and $J_{\mathcal{N}_1}(\kappa)$. Assuming exact multi-photon resonance, so that the detuning vanishes, $|11\rangle$ couples to a single bright superposition of $|02\rangle$ and $|20\rangle$ with effective Rabi frequency

$$\Omega_{\mathrm{eff}} = |\Omega_\alpha \Omega_\beta| \left| \frac{1}{2C_{11} - C_{01}} - \frac{1}{C_{12} - C_{02}} \right| \sqrt{J_{\mathcal{N}_0}^2(\kappa) + J_{\mathcal{N}_1}^2(\kappa)}, \tag{54}$$

and the populations are

$$P_{11}(t) = \cos^2(\Omega_{\text{eff}} t),$$
$$P_{02}(t) = \frac{J_{\mathcal{N}_1}^2(\kappa)}{J_{\mathcal{N}_0}^2(\kappa) + J_{\mathcal{N}_1}^2(\kappa)} \sin^2(\Omega_{\text{eff}} t), \qquad P_{20}(t) = \frac{J_{\mathcal{N}_0}^2(\kappa)}{J_{\mathcal{N}_0}^2(\kappa) + J_{\mathcal{N}_1}^2(\kappa)} \sin^2(\Omega_{\text{eff}} t). \tag{55}$$

The orthogonal (dark) superposition is decoupled and never populated. The branching ratio between $|02\rangle$ and $|20\rangle$ is therefore controlled entirely by the two Bessel factors, so that tuning the modulation index κ redistributes weight between the two legs of the chain—visible in Fig. 3 as the unequal amplitudes of the $|02\rangle$ and $|20\rangle$ curves.

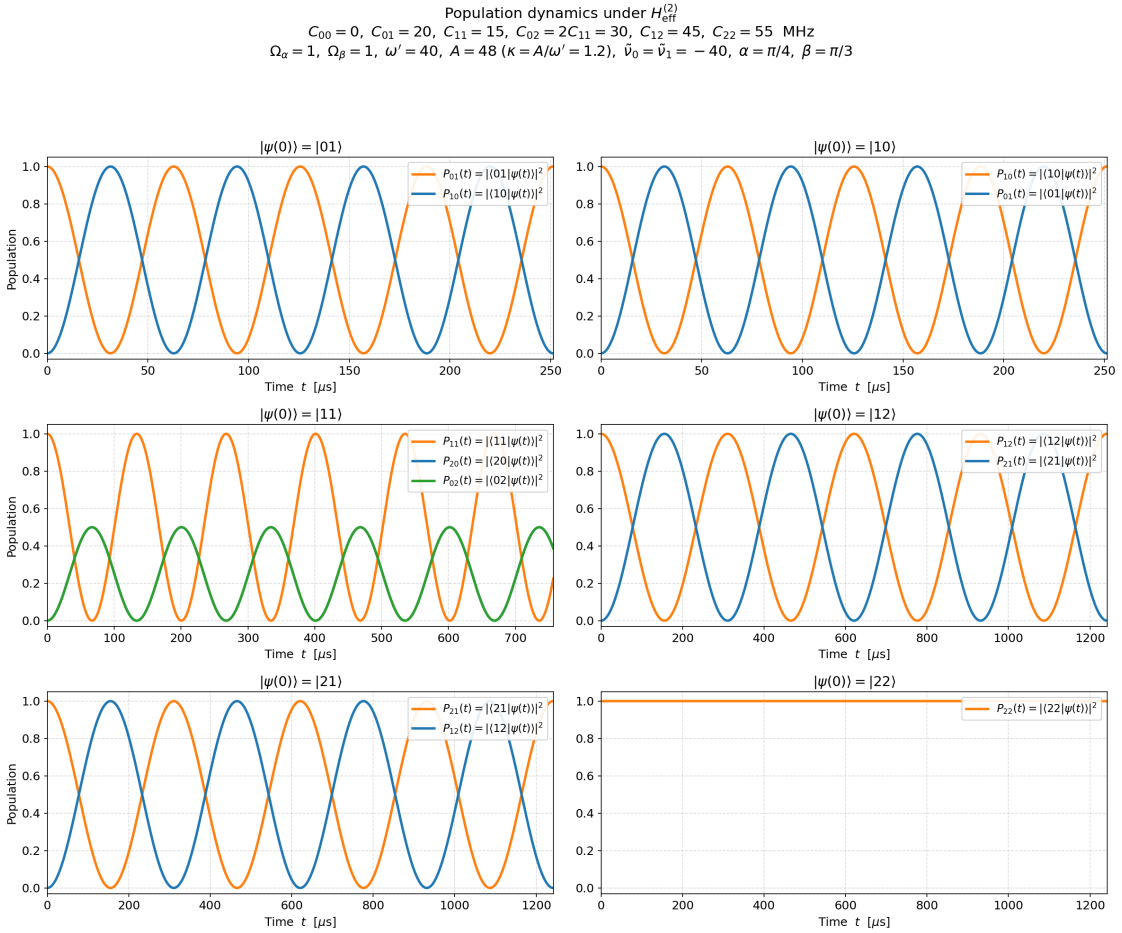


Figure 3: Population dynamics under $H_{\text{eff}}^{(2)}$ for the six initial two-site configurations, obtained from (48). Parameters (in MHz) are $C_{00} = 0$, $C_{01} = 20$, $C_{11} = 15$, $C_{02} = 2C_{11} = 30$, $C_{12} = 45$, $C_{22} = 55$, with $\Omega_\alpha = \Omega_\beta = 1$, $\omega' = 40$, $A = 48$ (so $\kappa = A/\omega' = 1.2$), $\tilde{\nu}_0 = \tilde{\nu}_1 = -40$, $\alpha = \pi/4$ and $\beta = \pi/3$. The $N = 1$ and $N = 3$ manifolds show two-level exchange, the $N = 2$ manifold the three-state chain $|02\rangle \leftrightarrow |11\rangle \leftrightarrow |20\rangle$, while $|22\rangle$ is decoupled and remains stationary.

3.8.2 From the bond to the chain: spacetime dynamics

The bond dynamics above extends directly to a chain, $H = \sum_j H_{\text{eff},\, j,j+1}^{(2)}$. Because $\hat{N}$ is conserved, Eq. (46), the evolution never leaves the sector of the initial state, so (48) can be evaluated exactly by diagonalizing that one block—this is the fixed-filling statement of Sec. 3.7 put to use. Releasing a single doubly-occupied site into an otherwise empty

$L = 10$ chain, $|\psi(0)\rangle = |0, 0, 0, 0, 2, 0, 0, 0, 0, 0\rangle$, confines the dynamics to $N = 2$, a block of dimension 55, and the site-resolved occupation $\langle \hat{n}_j(t) \rangle$ follows without approximation.

Figure 4 shows the result. Panel (b) is a reference: a *single* excitation released from the same site can only use the degenerate-exchange channel A_1 of (43), labelled as in (75), and produces a clean single-particle light cone that reflects off the open ends and reconstructs, with revivals near $t \simeq 340$ and $530\,\mu$s. Panel (a) is the doublon. It stays bound for $t \lesssim 100\,\mu$s—breaking it costs the AC-Stark diagonal (42), while the Floquet-assisted amplitudes that can break it are small, $A_2 = A_3 = 0.017\,$MHz against $A_1 = 0.050\,$MHz—and only then dissociates, the occupation at the initial site falling to a minimum of 0.079 and 1.92 of the two quanta having left it by $t \simeq 577\,\mu$s. The contrast between the two panels is the clearest statement that the doublon propagates as a composite object rather than as two independent particles: it is slower by roughly the ratio A_4/A_1 and it survives long before spreading. The left–right imbalance of panel (a), 0.58 against 1.32 at the end of the window, is a boundary effect of releasing the doublon off centre (site 4 of ten) and not a signature of statistics.

Two caveats apply. First, these are stroboscopic occupations: the micromotion within a driving period is excluded, as noted at the beginning of this subsection. Second, the doublon channels are slow. With $A_4 \simeq 0.01\,$MHz the dissociation unfolds over hundreds of microseconds, which is comparable to Rydberg lifetimes; resolving the late-time structure of panel (a) would need the stronger couplings or longer-lived encoding discussed above, whereas the early bound plateau and the single-particle light cone of panel (b) lie comfortably inside the coherence time.

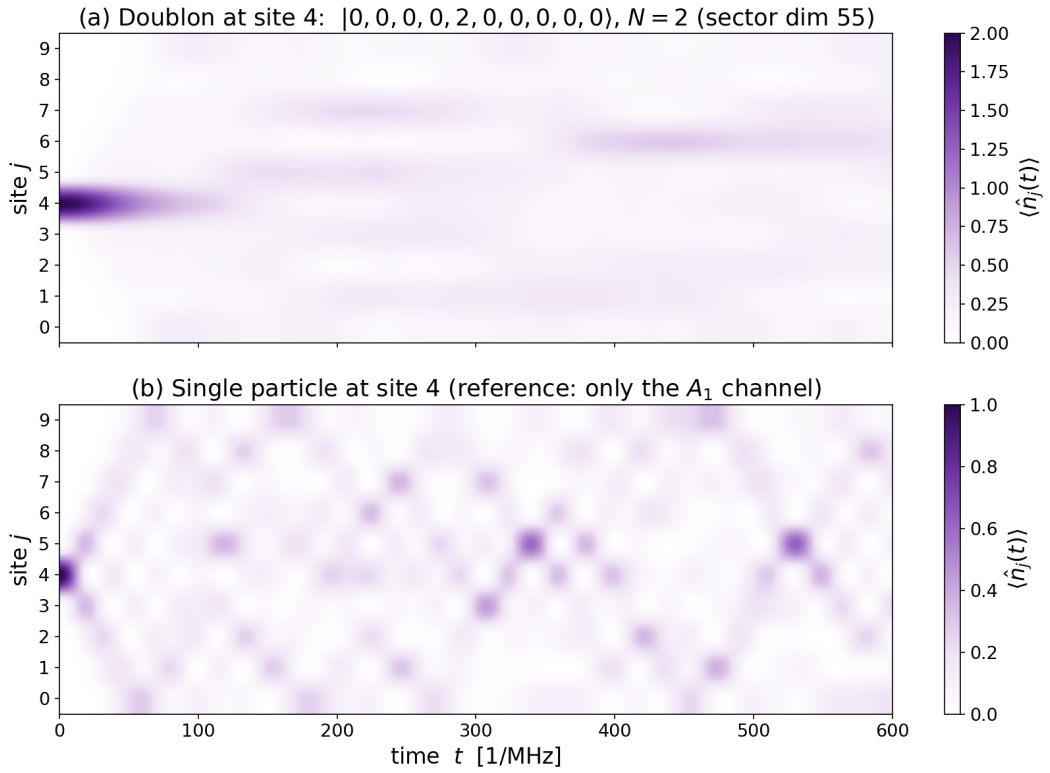


Figure 4: Spacetime maps of the site-resolved occupation $\langle \hat{n}_j(t) \rangle$ under $H_{\text{eff}}^{(2)}$ on a chain of $L = 10$ sites with open boundaries, obtained from (48). Colour runs from white at $\langle \hat{n}_j \rangle = 0$ to deep violet at the maximum. (a) A doublon released at site 4, $|0, 0, 0, 0, 2, 0, 0, 0, 0, 0\rangle$, so $N = 2$ and the evolution is exact inside a block of dimension 55. (b) A single excitation released at the same site, which can only use the A_1 channel and therefore shows the bare single-particle light cone with its reflections. The evolution uses the four amplitudes $A_1, \ldots, A_4$ of (75) with the signed values $A_1 = -0.0500$, $A_2 = A_3 = +0.0166$ and $A_4 = -0.0101\,\text{MHz}$; parameters as in Fig. 3 with $\omega' = 40\,\text{MHz}$ and $\theta = 0$, the value the drive actually realizes (Sec. 4.6). Time is in units of $1/\text{MHz}$, i.e. microseconds.

What a non-zero statistical angle would look like. The matching of Sec. 4 works at the level of the effective Hamiltonian. It is therefore worth asking whether the same conclusion can be reached from the *full* dynamics alone, without analysing the effective Hamiltonian at all—that is, what an experiment measuring site-resolved densities would see if the statistical phase could be engineered. Figure 5 repeats the doublon release with the bond phases set to the pattern $(\vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4) = (0, \theta, -2\theta, -\theta)$ obtained from $\alpha = 0$, $\beta = \theta$, whose flux through one elementary loop of the configuration graph is $-2\vartheta_1 + \vartheta_2 + \vartheta_3 = -\theta$, and which the diagonal gauge with $\gamma_2 - 2\gamma_1 = -2\theta$ carries onto the anyon-Hubbard target $(0, \phi, 0, \phi)$ with $\phi = -\theta$. The angle changes the dynamics substantially: the largest deviation from the $\theta = 0$ map, $\max_{t,j} |\langle \hat{n}_j \rangle_\theta - \langle \hat{n}_j \rangle_0|$, reaches 0.66 out of 2, and it does so at $\theta = 4\pi/3$ rather than at $\theta = \pi$, taking the values 0.435, 0.451, 0.408, 0.657 and 0.361 at $\theta = \pi/3$, $2\pi/3$, π, $4\pi/3$ and $5\pi/3$. At $\theta = 4\pi/3$ the doublon migrates steadily along the chain, reaching site 7 by $t \simeq 500\,\mu\text{s}$, whereas at $\theta = 5\pi/3$ it barely leaves its initial neighbourhood—so a directed, phase-controlled transport of the composite object is the signature to look for.

That the circle must be sampled in full, rather than only on $[0, \pi]$, follows from a symmetry that is easy to misread. Since the diagonal is real and only ϑ_2, ϑ_4 carry the phase, $H(-\theta) = H(\theta)^*$; with a real initial state this gives $|\psi(t; -\theta)\rangle = |\psi(-t; \theta)\rangle^*$ and

hence

$$\langle \hat{n}_j(t) \rangle_{-\theta} = \langle \hat{n}_j(-t) \rangle_{\theta}. \tag{56}$$

This is *not* evenness in θ: that would additionally require the dynamics to be time-reversal symmetric, which it is not. Reversing the statistical angle runs the density backwards in time instead of leaving it unchanged, so θ and $2\pi - \theta$ are different evolutions rather than mirror images—which is why $4\pi/3$ and $5\pi/3$ differ by nearly a factor of two in the table above.

It must be stressed that Fig. 5 is a statement about the target model and not about the drive of Sec. 3, which delivers only $\phi \in \{0, \pi\}$ (Sec. 4.6). Indeed the uniform bond pattern used there cannot be produced by any assignment of laser phases to the sites: $\vartheta_1(j) = \alpha_j - \alpha_{j+1} = 0$ forces α_j to be constant, $\vartheta_4(j) = -\theta$ forces $\beta_j = \beta_0 + j\theta$, and then $\vartheta_2(j) = \beta_j - \alpha_{j+1} = \beta_0 + j\theta - \alpha$ grows linearly with j and cannot equal the constant θ unless $\theta = 0$. This reproduces the restriction of Sec. 4.6 from a different direction—per-bond flux rather than the linear system of Table 3—and shows what is lost by it: the transport asymmetry of Fig. 5 is the observable that an occupation-selective, multi-tone drive would unlock.

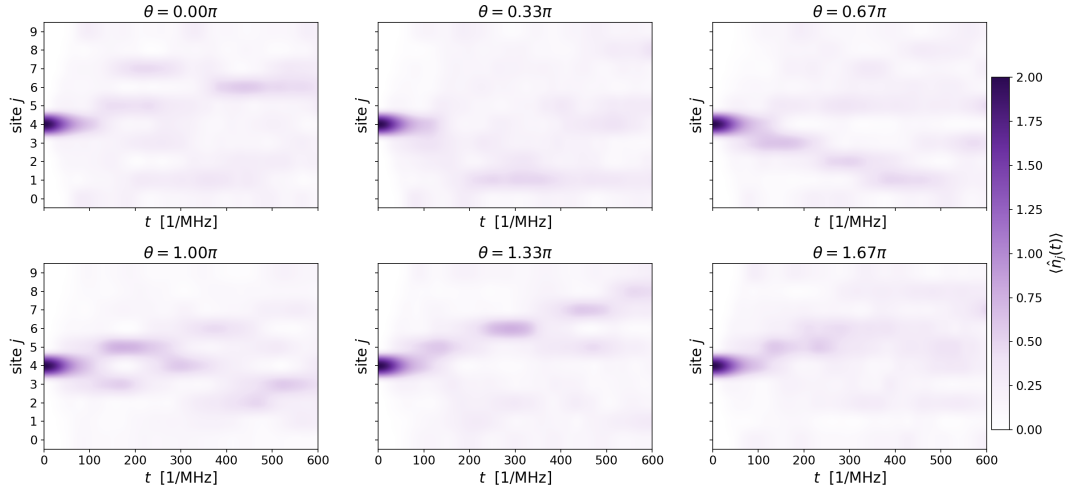


Figure 5: The doublon release of Fig. 4(a) repeated at six values of the drive phase θ covering the full circle, with the bond phases $(0, \theta, -2\theta, -\theta)$ that are gauge-equivalent to the anyon-Hubbard pattern with statistical angle $\phi = -\theta$. The colour scale is common to all panels, white at $\langle \hat{n}_j \rangle = 0$. The angle controls both how far and in which direction the doublon moves: at $\theta = 4\pi/3$ it migrates to site 7 by $t \simeq 500\,\mu$s, while at $\theta = 5\pi/3$ it stays near its initial site. Because of (56) the density is *not* even in θ, so θ and $2\pi - \theta$ are distinct evolutions and the whole circle must be sampled. This figure characterizes the target model: the single-tone drive of Sec. 3 realizes only $\theta = 0$ and $\theta = \pi$, as established in Sec. 4.6.

4 Matching the effective Hamiltonian to the anyon-Hubbard model

Sections 2–3 produced an effective Hamiltonian $H_{\text{eff}}^{(2)} = H_{\text{diag}}^{(2)} + H_{\text{off-diag}}^{(2)}$ for the driven two-atom bond. Section 1.3 fixed the target: the anyon-Hubbard bond, with on-site interaction U and the four tunnelling amplitudes (8). It remains to ask under what conditions the former can be made equal to the latter. We treat the diagonal sector first, where the matching fixes the effective chemical potential and interaction, and then the off-diagonal sector, where the matching turns out to be restricted.

Throughout this section we abbreviate the two families of AC Stark shifts appearing

in (42) as

$$\Delta_1 = \frac{\Omega_\alpha^2}{C_{01} - 2C_{00}}, \qquad \Delta_2 = \frac{\Omega_\alpha^2}{2C_{11} - C_{01}}, \qquad \Delta_3 = \frac{\Omega_\alpha^2}{C_{12} - C_{02}}, \tag{57}$$

$$\begin{aligned} S_1 &= \Omega_\beta^2 \sum_\ell \frac{J_\ell^2(\kappa)}{C_{02} - C_{01} + \tilde{\nu} + \ell\omega'}, \qquad S_2 = \Omega_\beta^2 \sum_\ell \frac{J_\ell^2(\kappa)}{C_{12} - 2C_{11} + \tilde{\nu} + \ell\omega'}, \\ S_3 &= \Omega_\beta^2 \sum_\ell \frac{J_\ell^2(\kappa)}{2C_{22} - C_{12} + \tilde{\nu} + \ell\omega'}, \end{aligned} \tag{58}$$

i.e. $\Delta_i = \mathcal{A}_X$ and $S_i = \mathcal{S}(X)$ of (41) evaluated at the relevant denominators, having anticipated the condition $\tilde{\nu}_0 = \tilde{\nu}_1 \equiv \tilde{\nu}$ that Sec. 4.4 will impose.

4.1 Diagonal sector: chemical potential and interaction

The diagonal elements $D_{n_1,n_2} = \langle n_1, n_2 | H^{(2)}_{\text{diag}} | n_1, n_2 \rangle$ are the quasi-energies of the nine configurations. For a Bose- or anyon-Hubbard bond with local energy E_0, chemical potential μ and interaction U, the target spectrum is

$$\begin{aligned} E_{00} &= E_0, & E_{01} = E_{10} &= E_0 + \mu, & E_{11} &= E_0 + 2\mu, \\ E_{02} = E_{20} &= E_0 + U + 2\mu, & E_{12} = E_{21} &= E_0 + U + 3\mu, & E_{22} &= E_0 + 2U + 4\mu. \end{aligned} \tag{59}$$

Requiring the spacings of $H^{(2)}_{\text{diag}}$ to reproduce this ladder gives a set of consistency conditions. Demanding that adding the first particle cost the same as adding the second, $D_{01} - D_{00} = D_{11} - D_{01}$, yields

$$2\Delta_1 - 2\Delta_2 = S_1 - S_2, \tag{60}$$

in which case the chemical potential is well defined and equal to

$$\mu = 3\Delta_1 - \Delta_2 - S_1. \tag{61}$$

Similarly, requiring the combination $U + \mu$ to be the same for the transitions $01 \to 02$, $11 \to 12$ and $12 \to 22$ gives two further constraints,

$$\begin{aligned} D_{02} - D_{01} = D_{12} - D_{11} &\implies -\Delta_1 + 3\Delta_2 - 2\Delta_3 + 2S_1 - 3S_2 + S_3 = 0, \\ D_{12} - D_{11} = D_{22} - D_{12} &\implies \Delta_2 - \Delta_3 - 2S_2 + 2S_3 = 0, \end{aligned} \tag{62}$$

under which the Hubbard interaction is

$$U = -4\Delta_1 + 2\Delta_2 - \Delta_3 + 3S_1. \tag{63}$$

4.2 Effective chemical potential, interaction, and their ratio

Independently of whether the consistency conditions hold exactly, one may read the low-occupation energies directly off $H^{(2)}_{\text{diag}}$,

$$E_{00} = -2\Delta_1, \qquad E_{01} = \Delta_1 - \Delta_2 - S_1, \qquad E_{02} = S_1 - \Delta_3, \tag{64}$$

and define the effective parameters from the first and second particle addition energies,

$$\mu_{\text{eff}} = E_{01} - E_{00} = 3\Delta_1 - \Delta_2 - S_1, \qquad U_{\text{eff}} = E_{02} - 2E_{01} + E_{00} = -4\Delta_1 + 2\Delta_2 - \Delta_3 + 3S_1, \tag{65}$$

which reproduce (61) and (63). Written out,

$$\mu_{\rm eff} = \Omega_\alpha^2\left[\frac{3}{C_{01}-2C_{00}} - \frac{1}{2C_{11}-C_{01}}\right] - \Omega_\beta^2 \sum_{\ell=-\infty}^{\infty} \frac{J_\ell^2(\kappa)}{C_{02}-C_{01}+\tilde{\nu}+\ell\omega'}, \tag{66}$$

$$U_{\rm eff} = \Omega_\alpha^2\left[\frac{2}{2C_{11}-C_{01}} - \frac{4}{C_{01}-2C_{00}} - \frac{1}{C_{12}-C_{02}}\right] + 3\,\Omega_\beta^2 \sum_{\ell=-\infty}^{\infty} \frac{J_\ell^2(\kappa)}{C_{02}-C_{01}+\tilde{\nu}+\ell\omega'}. \tag{67}$$

The dimensionless ratio that controls the physics of the Hubbard model follows as

$$\frac{\mu_{\rm eff}}{U_{\rm eff}} = \frac{3\Delta_1-\Delta_2-S_1}{3S_1-4\Delta_1+2\Delta_2-\Delta_3} = \frac{\dfrac{3}{C_{01}-2C_{00}} - \dfrac{1}{2C_{11}-C_{01}} - r\,\Sigma}{\dfrac{2}{2C_{11}-C_{01}} - \dfrac{4}{C_{01}-2C_{00}} - \dfrac{1}{C_{12}-C_{02}} + 3r\,\Sigma}, \tag{68}$$

where the overall factor Ω_α^2 cancels between numerator and denominator and we introduced

$$r \equiv \frac{\Omega_\beta^2}{\Omega_\alpha^2}, \qquad \Sigma \equiv \sum_{\ell=-\infty}^{\infty} \frac{J_\ell^2(\kappa)}{C_{02}-C_{01}+\tilde{\nu}+\ell\omega'}. \tag{69}$$

The ratio is therefore controlled by only two experimentally accessible knobs, the Rabi-frequency ratio r and the modulation index κ entering through Σ, at fixed interaction geometry.

In fact the two enter only in one combination. Collecting the geometry factors of (68) into

$$a = \frac{3}{C_{01}-2C_{00}} - \frac{1}{2C_{11}-C_{01}}, \qquad b = \frac{4}{C_{01}-2C_{00}} + \frac{1}{C_{12}-C_{02}} - \frac{2}{2C_{11}-C_{01}}, \tag{70}$$

the ratio becomes a Möbius function of the single product $x \equiv r\,\Sigma(\kappa)$,

$$\frac{\mu_{\rm eff}}{U_{\rm eff}} = \frac{a-x}{3x-b}, \tag{71}$$

with a pole at $x = b/3$, a zero at $x = a$, the value $-a/b$ at $x = 0$ and the limit $-1/3$ as $|x| \to \infty$. For the interaction geometry of Table 1, $a = 0.0500$ and $b = 0.0667\,\mathrm{MHz}^{-1}$, so the pole sits at $x = 0.0222$ and the two limiting values are $-3/4$ and $-1/3$. Two consequences follow. First, r and κ are not independent handles on the ratio: any pair with the same $r\Sigma(\kappa)$ gives the same $\mu_{\rm eff}/U_{\rm eff}$, and the remaining freedom acts only on the amplitudes. Second, since $x > 0$ whenever $\Sigma > 0$, letting x run over $(0, \infty)$ covers $(-\infty, -3/4]$ and $(-1/3, +\infty)$ but never the interval in between:

$$-\tfrac{3}{4} < \frac{\mu_{\rm eff}}{U_{\rm eff}} < -\tfrac{1}{3} \qquad \text{is unreachable for } \Sigma > 0. \tag{72}$$

That window opens only where the Bessel sum changes sign: Σ vanishes at $\kappa = 1.0585$, 3.4910 and 4.2841, and for $\Sigma < 0$ the product x runs over $(-\infty, 0)$ and sweeps the excluded interval. Choosing κ below the first of these zeros is therefore the only way to reach it, which constrains the modulation index rather than the drive powers.

Figure 6 shows the amplitudes of (75) against the ratio. Because $A_1 \propto \Omega_\alpha^2$ is blind to both Ω_β and κ, while $A_4 \propto \Omega_\beta^2$ is blind to Ω_α, no single sweep of two knobs moves all of them; each panel therefore states the sweep it uses, Ω_α for A_1 and Ω_β for $A_2 = A_3$ and A_4. In every panel the curve splits into two branches on either side of the pole of (71), and

the shaded strip is the forbidden window (72)—the $\kappa = 0.8$ curves, for which $\Sigma < 0$, are the ones that cross it. At the reference parameters ($\Omega_\alpha = \Omega_\beta = 1$, $\kappa = 1.2$, starred) the ratio is $\mu_{\text{eff}}/U_{\text{eff}} = -0.924$ with $A_1 = 0.0500$, $A_2 = A_3 = 0.0166$ and $A_4 = 0.0101\,\text{MHz}$, i.e. $A_2/A_1 = 0.332$ and $A_4/A_1 = 0.202$ against the targets $\sqrt{2}$ and 2 of (8). The figure makes clear how that shortfall can be reduced: moving along Ω_β raises A_2 and A_4 relative to the fixed A_1, at the price of driving the ratio toward its $-1/3$ asymptote.

Two features of the curves are set by Bessel zeros rather than by the ratio. $A_2 = A_3$ vanishes at $\kappa = 3.8317$, the first zero of J_1, and A_4 vanishes at $\kappa = 1.3292$, 3.3175, 4.5008, where the sum Σ_4 of (75) changes sign. At such κ the corresponding channel switches off entirely while $\mu_{\text{eff}}/U_{\text{eff}}$ remains perfectly finite. The reference value $\kappa = 1.2$ lies close to the first of these, which is why A_4 comes out an order of magnitude below A_1 there; moving to $\kappa = 0.8$ raises A_4 by a factor of four.

Finally, the amplitude conditions of Sec. 4.4 interact with this ratio. At fixed interaction geometry those three conditions fix Ω_α, Ω_β and κ completely, hence fix x, hence fix the ratio: $\mu_{\text{eff}}/U_{\text{eff}}$ is then no longer a free parameter. Solving them numerically for $\kappa \in [0.3, 6]$ gives six solutions, all requiring a strongly asymmetric drive $\Omega_\beta/\Omega_\alpha$ between 3.9 and 10.6, and all landing in the narrow band

$$\frac{\mu_{\text{eff}}}{U_{\text{eff}}} \in [-0.303,\ -0.263], \tag{73}$$

just above the $-1/3$ asymptote. So while the magnitudes pose no difficulty on their own, matching them does pin the chemical-potential-to-interaction ratio to a small negative value; tuning it elsewhere requires retuning the interaction geometry C_{ij}, which is fixed by the atomic states and the lattice spacing.

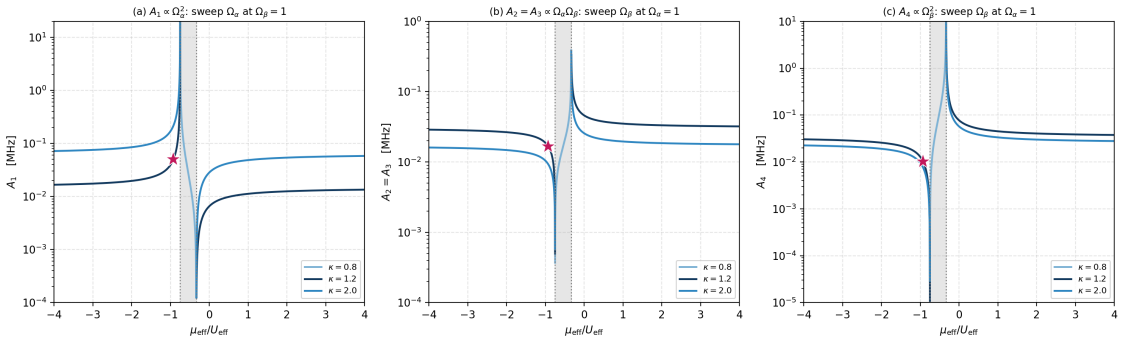


Figure 6: The engineered amplitudes of (75) against the ratio $\mu_{\text{eff}}/U_{\text{eff}}$ of (68), for three modulation indices. Each panel is traced by the sweep that moves the amplitude in question: (a) $A_1 \propto \Omega_\alpha^2$ by sweeping Ω_α at $\Omega_\beta = 1$, (b) $A_2 = A_3 \propto \Omega_\alpha\Omega_\beta$ and (c) $A_4 \propto \Omega_\beta^2$ by sweeping Ω_β at $\Omega_\alpha = 1$. Each curve breaks into two branches at the pole $x = b/3$ of (71); the shaded strip is the window (72), unreachable when $\Sigma > 0$ and crossed only by the $\kappa = 0.8$ curves, for which $\Sigma < 0$. Stars mark the reference parameters $\Omega_\alpha = \Omega_\beta = 1$, $\kappa = 1.2$, where $\mu_{\text{eff}}/U_{\text{eff}} = -0.924$ and $A_1 = 0.0500$, $A_2 = A_3 = 0.0166$, $A_4 = 0.0101\,\text{MHz}$. The sharp dips of A_2 and A_4 are the Bessel zeros discussed in the text, at which a channel switches off while the ratio stays finite. Interaction geometry as in Fig. 3, with $\omega' = 40\,\text{MHz}$ and $\mathcal{N} = 1$.

The same information can be read the other way round. Figure 6 places the ratio $\mu_{\text{eff}}/U_{\text{eff}}$ on the horizontal axis and asks what the amplitudes do; Fig. 7 instead places the dimensionless couplings A_i/U_{eff} there and asks what the chemical potential does. The second reading is the one that connects to the Hubbard literature, because A_1/U_{eff} is the tunnelling-to-interaction ratio J/U in which Mott lobes are conventionally drawn, and

A_2/U_{eff} and A_4/U_{eff} are its density-assisted counterparts. Only three panels are needed: the resonance lock $\tilde{\nu}_0 = \tilde{\nu}_1$ of Sec. 4.6 forces $\mathcal{N}_0 = \mathcal{N}_1 = \mathcal{N}$, so $A_3 = A_2$ identically and the pair shares a panel. The two are equal in magnitude only—in $H^{(2)}_{\text{eff}}$ they act on different bond configurations and carry the phases ϕ and -2ϕ.

The axes behave differently under an overall rescaling of the drive, and the distinction matters when quoting numbers. Every amplitude, μ_{eff} and U_{eff} carries a factor Ω_α^2, so under $(\Omega_\alpha, \Omega_\beta) \to \lambda(\Omega_\alpha, \Omega_\beta)$ at fixed r and κ the horizontal axis is invariant while the vertical one scales as λ^2. The figure is therefore drawn at a stated $\Omega_\alpha = 1\,\text{MHz}$, the value used in the exact diagonalizations of Sec. 3.7 and App. C; changing it slides every curve vertically and does nothing else. That choice also keeps the construction inside its own regime, since the smallest energy denominator appearing in $H^{(2)}_{\text{eff}}$ is $2C_{11} - C_{01} = 10\,\text{MHz}$ and the second-order expansion requires Ω well below it.

Three features are visible immediately. First, all four A_i are positive by construction, so the sign of A_i/U_{eff} is the sign of U_{eff} alone: the left half-plane of each panel is the attractive side of the model, and the vertical asymptote separating the halves is the locus $U_{\text{eff}} = 0$. For $\kappa = 1.2$ this is crossed at $\Omega_\beta/\Omega_\alpha = 1.842$ and for $\kappa = 2.0$ at $\Omega_\beta/\Omega_\alpha = 0.887$, whereas $\kappa = 0.8$ has $\Sigma < 0$ and U_{eff} never changes sign. Second, μ_{eff} stays bounded exactly where A_i/U_{eff} diverges: approaching the pole, $\mu_{\text{eff}} \to \Omega_\alpha^2(a - b/3) = 0.0278\,\text{MHz}$, which is the plateau the curves run into at large $|A_i/U_{\text{eff}}|$ on both sides. The divergence is entirely in the denominator, not in the chemical potential. Third, in the opposite limit of a weak second tone, $\mu_{\text{eff}} \to \Omega_\alpha^2 a = 0.0500\,\text{MHz}$ while $A_1/U_{\text{eff}} \to -3/4$, so the accessible region of each panel is bounded above and the useful tuning range is the steep descent between the two plateaus. At the reference parameters (stars) the model sits at $A_1/U_{\text{eff}} = -1.064$, $A_2/U_{\text{eff}} = A_3/U_{\text{eff}} = -0.353$ and $A_4/U_{\text{eff}} = -0.215$, with $\mu_{\text{eff}} = 0.0435\,\text{MHz}$ and $U_{\text{eff}} = -0.0470\,\text{MHz}$—strongly attractive, and with a tunnelling comparable to the interaction rather than small against it.

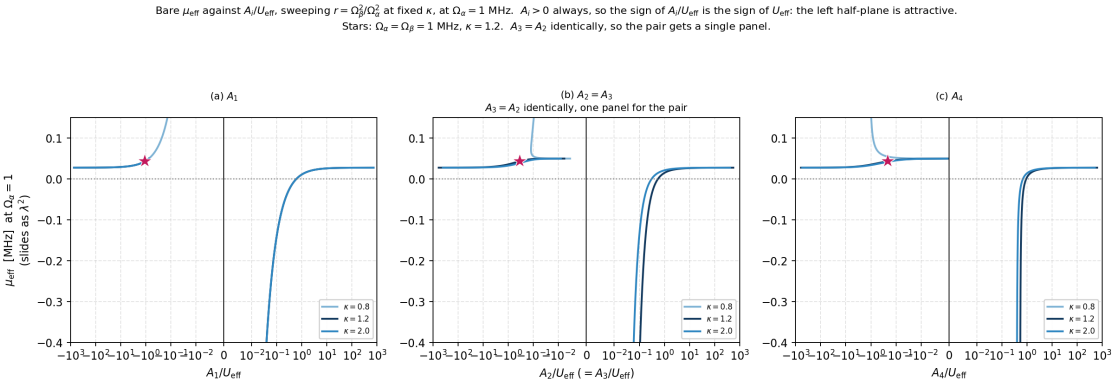


Figure 7: The bare chemical potential μ_{eff} of (66) against the dimensionless couplings A_i/U_{eff}, sweeping $r = \Omega_\beta^2/\Omega_\alpha^2$ at three modulation indices. Panel (b) carries A_2 and A_3 together because $\mathcal{N}_0 = \mathcal{N}_1$ makes them identical in magnitude. The horizontal axes are invariant under $(\Omega_\alpha, \Omega_\beta) \to \lambda(\Omega_\alpha, \Omega_\beta)$; the vertical axis is not, and is drawn at $\Omega_\alpha = 1\,\text{MHz}$. Since $A_i > 0$ always, the sign of the abscissa is the sign of U_{eff} and the left half-plane is attractive; the axes are symmetric-logarithmic so that both signs and the divergence at $U_{\text{eff}} = 0$ are visible. The curves run into a plateau at $\Omega_\alpha^2(a - b/3) = 0.0278\,\text{MHz}$ as the pole is approached, and tend to $\Omega_\alpha^2 a = 0.0500\,\text{MHz}$ in the weak-Ω_β limit. Stars mark $\Omega_\alpha = \Omega_\beta = 1\,\text{MHz}$, $\kappa = 1.2$, where $\mu_{\text{eff}} = 0.0435\,\text{MHz}$ and $U_{\text{eff}} = -0.0470\,\text{MHz}$. Interaction geometry as in Fig. 3, with $\omega' = 40\,\text{MHz}$ and $\mathcal{N} = 1$.

4.3 Off-diagonal sector: target versus engineered amplitudes

Comparing the off-diagonal part (43) with the target amplitudes (8) gives the four correspondences collected in Table 3. Each engineered element carries both a magnitude, set by the Rabi frequencies, the energy denominators and a Bessel factor, and a phase, set by the laser phases α and β.

Transition	Target (anyon-Hubbard)	Engineered element of $H^{(2)}_{\text{off-diag}}$
$\lvert 01\rangle \to \lvert 10\rangle$	$-J$	$\Omega_\alpha^2\left[\frac{1}{C_{01}-2C_{00}}-\frac{1}{2C_{11}-C_{01}}\right]e^{-i\alpha}$
$\lvert 02\rangle \to \lvert 11\rangle$	$-J\sqrt{2}$	$\Omega_\alpha\Omega_\beta\left[\frac{1}{2C_{11}-C_{01}}-\frac{1}{C_{12}-C_{02}}\right]J_{\mathcal{N}_1}(\kappa)\,e^{i(\alpha-2\beta)}$
$\lvert 11\rangle \to \lvert 20\rangle$	$-J\sqrt{2}\,e^{i\phi}$	$\Omega_\alpha\Omega_\beta\left[\frac{1}{2C_{11}-C_{01}}-\frac{1}{C_{12}-C_{02}}\right]J_{\mathcal{N}_0}(\kappa)\,e^{i(\beta-2\alpha)}$
$\lvert 12\rangle \to \lvert 21\rangle$	$-2J\,e^{i\phi}$	$\Omega_\beta^2\left[\sum_\ell J_\ell(\kappa)J_{\ell-q}(\kappa)\Sigma_\ell\right]e^{-i\beta}$

Table 3: Matching of the four independent bond processes. The target column is Eq. (8); the engineered column is read off Eq. (43).

4.4 Amplitude conditions

The two processes $|02\rangle \to |11\rangle$ and $|11\rangle \to |20\rangle$ carry the same target magnitude $J\sqrt{2}$ but are proportional to $J_{\mathcal{N}_1}(\kappa)$ and $J_{\mathcal{N}_0}(\kappa)$ respectively. Equality therefore demands identical Floquet modulation on the two sites,

$$\tilde{\nu}_0 = \tilde{\nu}_1 \equiv \tilde{\nu} \quad\Longrightarrow\quad \mathcal{N}_0 = \mathcal{N}_1 \equiv \mathcal{N}, \qquad q = 0, \tag{74}$$

which also justifies the single $\tilde{\nu}$ used in (58). The magnitudes then reduce to three conditions,

$$\begin{aligned}
A_1 &= \Omega_\alpha^2\left|\frac{1}{C_{01}-2C_{00}}-\frac{1}{2C_{11}-C_{01}}\right| = J,\\
A_2 = A_3 &= \Omega_\alpha\Omega_\beta\left|\frac{1}{2C_{11}-C_{01}}-\frac{1}{C_{12}-C_{02}}\right|\left|J_{\mathcal{N}}(\kappa)\right| = J\sqrt{2},\\
A_4 &= \Omega_\beta^2\left|\sum_{\ell=-\infty}^{\infty} J_\ell^2(\kappa)\left(\frac{1}{C_{12}-2C_{11}+\tilde{\nu}+\ell\omega'}-\frac{1}{2C_{22}-C_{12}+\tilde{\nu}+\ell\omega'}\right)\right| = 2J,
\end{aligned} \tag{75}$$

subject to the resonance requirement

$$\mathcal{N} = \frac{2C_{11}-C_{02}-\tilde{\nu}}{\omega'} \in \mathbb{Z}. \tag{76}$$

These are three equations for the four free quantities Ω_α, Ω_β, κ and the interaction geometry C_{ij}, and can be satisfied simultaneously. The magnitudes therefore pose no difficulty.

4.5 Gauge freedom and the uniformity of the induced phase

The phases in Table 3 do not match the target phases $(0, 0, \phi, \phi)$ as they stand. This is not immediately fatal, because the phase of a basis state is not physical: a local $U(1)$ gauge transformation may be applied without changing any observable.

It is important to be precise about what kind of freedom this is. In a genuine lattice gauge theory the matter fields and the gauge field living on each link transform *together*: a local $U(1)$ rotation acts as $\hat{a}_j \to e^{i\alpha_j}\hat{a}_j$ on the matter and simultaneously as $\hat{U}_j \to e^{-i\alpha_j}\hat{U}_j e^{i\alpha_{j+1}}$ on the link operator, so that the gauge-invariant tunnelling $\hat{a}_{j+1}^\dagger \hat{U}_j \hat{a}_j$

is left unchanged, and the physical Hilbert space is the one annihilated by the Gauss-law generators $\hat{G}_j = \hat{L}_j - \hat{L}_{j-1} - \hat{Q}_j$ [38, 39]. The link variable is what absorbs the site-dependent phase difference there. This structure is not only formal: the elementary building block of such a theory has been assembled experimentally in a two-species atomic mixture, with the gauge field carried by the collective spin of one species, the matter field by the other, and the gauge-invariant coupling $\hat{b}_{\mathrm{p}}^{\dagger}\hat{L}_{-}\hat{b}_{\mathrm{v}}$—the direct analogue of $\hat{a}^{\dagger}\hat{U}\hat{a}$—implemented by heteronuclear spin-changing collisions [40]. Note already here what supplies the gauge structure in that experiment: two-body collisions, not a one-body drive. Comparable constructions extend to non-Abelian gauge groups, where the link variables become matrix-valued and the local constraints correspondingly richer [41]. In the present scheme there is no dynamical link variable: the phase carried by a hop is imposed externally by the lasers and the drive, not by a gauge field with its own dynamics and its own constraint. The only freedom available is therefore the matter-field redefinition, generated by a linear function of the densities,

$$\hat{U} = \exp\bigl[i(\gamma_a \hat{n}_a + \gamma_{0a}) + i(\gamma_b \hat{n}_b + \gamma_{0b})\bigr], \tag{77}$$

with $\gamma_{a,b}$ the gauge phase per particle on each site and $\gamma_{0a,0b}$ constant vacuum offsets. Since $|n_a, n_b\rangle$ is a simultaneous eigenstate of both number operators, the transformation acts directly on the eigenvalues,

$$\hat{U}\,|n_a, n_b\rangle = e^{i\Theta(n_a,n_b)}\,|n_a, n_b\rangle\,, \qquad \Theta(n_a, n_b) = \gamma_a n_a + \gamma_b n_b + \gamma_{0a} + \gamma_{0b}. \tag{78}$$

To see how this affects the tunnelling, transform the field operators. Using $[\hat{a}, \hat{a}^{\dagger}] = 1$ and the Baker–Campbell–Hausdorff lemma,

$$\hat{U}\hat{a}_a\hat{U}^{\dagger} = e^{-i\gamma_a}\hat{a}_a, \qquad \hat{U}\hat{a}_a^{\dagger}\hat{U}^{\dagger} = e^{+i\gamma_a}\hat{a}_a^{\dagger}, \tag{79}$$

and likewise on site b; note that the vacuum offsets γ_{0a}, γ_{0b} cancel identically and so have no physical effect. A single-particle hop from b to a therefore transforms as

$$(\hat{a}_a^{\dagger}\hat{a}_b)' = (e^{i\gamma_a}\hat{a}_a^{\dagger})(e^{-i\gamma_b}\hat{a}_b) = e^{i(\gamma_a-\gamma_b)}\,\hat{a}_a^{\dagger}\hat{a}_b, \tag{80}$$

so that every such hop acquires one and the same phase

$$\Delta\gamma = \gamma_a - \gamma_b. \tag{81}$$

Because (80) is an identity between field operators, it holds irrespective of the occupations of the two sites before or after the hop. *The gauge shift is therefore strictly uniform: it cannot distinguish between processes that differ only in how many particles are present.*

Equation (81) is the lattice gauge transformation in its standard form. For a Hubbard model whose tunnelling carries phases φ_{ij}, the on-site redefinition $b_j \to e^{i\chi_j}b_j$ sends $\varphi_{ij} \to \varphi_{ij} + \chi_i - \chi_j$, the discrete counterpart of $\mathbf{A} \to \mathbf{A} + \nabla\phi$; what survives it is the cumulative phase around a closed loop, the Wilson loop, which is the lattice statement of the Aharonov–Bohm flux [3]. Infinitely many assignments of link phases give the same set of loop fluxes and hence the same physics, so a claim about a tunnelling phase has no content unless it is referred to a closed circuit. This is why the restriction below is stated twice—once as the linear system of Table 3 after gauge fixing, and once in Sec. 3.8.2 as a flux through an elementary loop of the configuration graph—and why the experimental test proposed in App. B must close a loop rather than measure a single bond.

This makes the structure of the restriction transparent. A tunnelling phase that can be absorbed by (77) is pure gauge and carries no physical content; conversely, a phase that does carry physical content must not be removable in this way. The density-dependent Peierls phase $e^{i\phi n_j}$ of (6) is of the latter type—which is why it encodes exchange statistics rather than a choice of basis—and, for the same reason, it cannot be manufactured out of α, β and $\Delta\gamma$, which between them supply only occupation-independent phases.

4.6 Phase conditions and the locking of the statistical angle

All four processes of Table 3 move one particle in the same direction, so each acquires the same shift $\Delta\gamma$. Adding it to the raw phases and equating to the target phases $(0, 0, \phi, \phi)$ gives the linear system

$$\begin{aligned} |01\rangle \to |10\rangle: &\quad -\alpha + \Delta\gamma = 0, & |02\rangle \to |11\rangle: &\quad (\alpha - 2\beta) + \Delta\gamma = 0, \\ |11\rangle \to |20\rangle: &\quad (\beta - 2\alpha) + \Delta\gamma = \phi, & |12\rangle \to |21\rangle: &\quad -\beta + \Delta\gamma = \phi. \end{aligned} \tag{82}$$

All four relations are equalities between phases, and therefore hold only modulo 2π. The first condition fixes $\Delta\gamma = \alpha$. Substituting into the second gives $2(\alpha - \beta) = 0 \pmod{2\pi}$, i.e. $\alpha - \beta = \pi n$ with n an integer—not $\alpha = \beta$. Substituting $\Delta\gamma = \alpha$ into the third and fourth conditions gives

$$\beta - \alpha = \phi \qquad \text{and} \qquad \alpha - \beta = \phi \qquad (\text{mod } 2\pi), \tag{83}$$

which are compatible only if $2\phi = 0 \pmod{2\pi}$. The system is therefore consistent, and its solutions are the ones already found for the static protocol in (23),

$$\boxed{\phi = \pi n \implies \phi \in \{0, \pi\} \ (\text{mod } 2\pi).} \tag{84}$$

A single-tone drive thus neither lifts the restriction found in the static case nor makes it worse: the statistical phase remains locked to $\{0, \pi\}$. As there, $\phi = \pi$ is attainable and describes pseudofermions; what remains out of reach is a generic statistical phase such as $\phi = \pi/2$.

The conclusion follows from the structure of the coupling, not from the particular choice of parameters. The target phases $(0, 0, \phi, \phi)$ are *occupation dependent*: the statistical angle appears only in those processes that leave a doubly-occupied site behind, which is the content of the density-dependent Peierls phase $e^{i\phi n_j}$ of (6). The phases actually available—two laser phases α, β together with one uniform gauge shift $\Delta\gamma$—are occupation *independent*, by (80). No choice of α, β and $\Delta\gamma$ can therefore generate a phase that depends on n_j, and the over-determined system collapses onto the two values $\phi = 0$ and $\phi = \pi$ that a global, occupation-independent phase can already supply.

There is a more physical way to see why this had to happen, which also explains why the failure is one of phases and not of magnitudes. The Raman coupling H_Ω of (9) is a *one-body* operator: each of its terms moves a single atom between two internal levels, and the phase it imprints, α or β, is a property of the laser driving that particular rung—not of how many particles happen to be present. A second-order process is a product of two such factors, so its phase is always a fixed sum of two laser phases, determined solely by which rungs are traversed. What the occupation *does* control is the energy denominator, and hence the magnitude: this is why the four amplitudes $A_1, \ldots, A_4$ of (75) differ from one another. An occupation-dependent *phase*, by contrast, is an intrinsically many-body object—the operator $b_j^\dagger e^{i\phi n_j} b_{j+1}$ of (6) is cubic in the fields—and it cannot be assembled out of linear processes whose phases are fixed by the driving lasers. Linear couplings can produce density-dependent *amplitudes* through their energy denominators, but not density-dependent *phases*. Section 4.5 states this argument formally.

Contrast this with what running laser phases *can* achieve. In the synthetic-dimension construction of Sec. 2.1, site-dependent Raman phases of the form used here do generate a genuine synthetic gauge field—but a *uniform* one, a flux per plaquette fixed by the laser geometry and depending on *where* a particle hops, never on how many particles are present [32]. A one-body coupling yields a one-body gauge field. The anyon-Hubbard model asks instead for a gauge field that responds to the local density, and that is a qualitatively

different object. The distinction matters here: a geometric phase picked up by a particle transported around a closed circuit is not the same thing as the phase produced when one particle is carried around another, and only the latter is statistics [3]. What the construction of this thesis delivers is a phase of the first kind; what the anyon-Hubbard model requires is one of the second.

At $\phi = 0$ the complex target amplitudes $-J\sqrt{2}\,e^{i\phi}$ and $-2J\,e^{i\phi}$ reduce to the real bosonic values $-J\sqrt{2}$ and $-2J$, and the many-body wavefunction remains symmetric under exchange: the scheme then realizes an effective *Bose*-Hubbard bond, with tunnelling amplitudes and an interaction U_{eff} that are renormalized and widely tunable through (75) and (67). At $\phi = \pi$ it realizes the pseudofermionic point of the anyon-Hubbard model, which is a genuine—if not generic—anyonic statistics. What the scheme does not deliver is a *tunable* statistical angle. Generating a genuinely occupation-dependent phase requires an ingredient beyond the present scheme—a drive whose phase itself depends on the local occupation, rather than the uniform shift permitted by (80).

This is also what indicates that a time-dependent drive is the right tool, and what it must do. A periodic drive is not a static one-body coupling with a fixed phase: it carries its own frequency scale, and because the interaction energies split the occupation-resolved processes apart—by the C_{ij} here, and by the on-site interaction in a tilted lattice—*frequency selectivity becomes occupation selectivity.* A drive containing several tones, each quasi-resonant with one process and each carrying its own independently adjustable phase, can therefore imprint a different phase according to the occupation, which is the nonlinearity that the static Raman scheme lacks.

This is the mechanism behind the experiment on which the present work is modelled. There a tilted optical lattice with energy offset E per site is modulated in depth with three tones, resonant respectively with tunnelling from a singly-occupied site to an empty one (at E), from a singly-occupied site to a singly-occupied one (at $E+U_0$), and from a doubly-occupied site to an empty one (at $E-U_0$), where U_0 is the on-site interaction; offsetting the phase of the $E+U_0$ component by θ relative to the other two is what realizes the density-dependent Peierls phase, and hence the statistical angle [17]. The multicolour lattice-depth modulation of Ref. [18] implements the same idea, and the resulting occupation-dependent gauge fields map onto the anyon-Hubbard model.

A second and independent route reaches the same end from the interaction side. A density-dependent gauge field is, by definition, a statement about how the vector potential depends on the matter density, so demanding a particular matter-dependent gauge field is equivalent to demanding a particular form of the matter–matter interaction [42]. Instead of making the drive occupation selective, one may therefore engineer the interactions themselves: in a Raman-coupled Bose–Einstein condensate with imbalanced interstate scattering lengths, the low-energy description acquires the chiral, density-dependent interactions of the chiral BF theory—a one-dimensional reduction of Chern–Simons theory—and hence a density-dependent vector potential [42]. This is consistent with, rather than a counterexample to, the argument above: there the nonlinearity is supplied by the two-body scattering, not by the Raman coupling itself. The same is true of the atomic-mixture realization of local $U(1)$ invariance discussed in Sec. 4.5, where the gauge-invariant matter–gauge coupling is likewise collisional [40]. In every case the density dependence must originate in something genuinely many-body—either two-body collisions [40, 42] or an occupation-selective drive [17, 18].

Seen in this light, the restriction established above is specific to the drive employed in Sec. 3, and is not a limitation of Floquet engineering as such. The single-tone modulation used there offers only one drive phase ϕ_f, shared by every process that touches $|2\rangle$, so

the phases it generates remain locked together and, after gauge fixing, collapse onto the two occupation-independent values $\phi \in \{0, \pi\}$. Driving each of the resonances identified in Sec. 2.3 with its own tone and its own phase—the direct analogue of the three-tone protocol above—is therefore the natural extension of this work, and the analysis of this section shows that some such occupation-selective drive is necessary, not just convenient.

5 Conclusions and outlook

This thesis has asked whether the one-dimensional anyon-Hubbard model can be realized with Rydberg atoms, using internal atomic levels rather than motional states to carry the local occupation.

The starting point was the generalized Jordan–Wigner transformation, which converts anyonic exchange statistics into a density-dependent Peierls phase and so reduces the problem to a concrete engineering target: the four tunnelling amplitudes (8) of a two-site bond, of which the two that leave a doubly-occupied site behind must carry the statistical phase ϕ. The truncated local Hilbert space $\{|0\rangle, |1\rangle, |2\rangle\}$ was encoded in three internal levels of each of two atoms—a synthetic-dimension construction in which the Raman couplings provide the tunnelling along the internal direction and the running laser phases attach the Peierls phases [32]. The configuration-dependent Rydberg interactions C_{ij} were then used to place every single-hop process far off resonance, so that the bond dynamics is governed by second-order exchange.

Three results follow. *First*, a purely static Raman protocol realizes part of the model but not all of it: the four amplitudes are controlled by only two laser phases, and the resulting consistency conditions force $\beta - \alpha = \pi n$, locking the statistical phase to $\phi \in \{0, \pi\}$, Eq. (23). The value $\phi = \pi$ is a genuine result rather than a trivial one: it is the pseudofermionic point of the anyon-Hubbard model, realizable experimentally by taking $\alpha = 0$ and $\beta = \pi$. Thus the scheme allows for pseudofermions, but not for generic anyons; chiral dynamics at, say, $\phi = \pi/2$ is out of reach of any time-independent scheme of this kind. *Second*, adding a periodic drive on the doubly-occupied levels yields a well-controlled effective description. Transforming to the rotating frame, expanding the accumulated drive phase with the Jacobi–Anger identity, and applying a Magnus expansion in the secular approximation gives a static effective Hamiltonian whose diagonal part is a set of AC Stark shifts, Eq. (42), and whose off-diagonal part comprises four exchange processes, each gated by a Kronecker delta enforcing conservation of quasi-energy up to Floquet quanta, Eq. (43). The Bessel factors $J_\ell(\kappa)$ renormalize the couplings—the same mechanism that produces dynamic localization in shaken lattices—and the effective chemical potential and interaction, (66) and (67), are controlled by the Rabi-frequency ratio and the modulation index alone. *Third*, and centrally, matching this effective Hamiltonian to the anyon-Hubbard bond shows that the magnitudes can be matched while the phases cannot: the amplitude conditions (75) are three equations for four free quantities and are simultaneously satisfiable, whereas the phase conditions, once the residual gauge freedom has been fixed, again admit only $\phi \in \{0, \pi\}$. A single-tone drive therefore neither lifts the restriction found in the static case nor makes it worse. Promoting the bond to a chain makes that result measurable rather than merely formal: because $\hat{N}$ is conserved exactly the dynamics can be evaluated without approximation at fixed filling, and the doublon released into an empty register propagates symmetrically at the reachable values of ϕ while acquiring directed transport at generic ϕ (Sec. 3.8.2)—the observable the construction misses, and the one the experiment of App. B is built around.

This last result follows from the structure of the coupling. The Raman coupling is a

one-body operator: each of its terms moves a single atom and imprints a phase belonging to the laser driving that particular transition, not to the occupations of the sites involved. A second-order process is a product of two such factors, so its phase is a fixed sum of two laser phases, while the occupation enters only through the energy denominators. Linear couplings can therefore render tunnelling *amplitudes* density dependent—which is why the four amplitudes of (75) differ from one another—but not tunnelling *phases*. An occupation-dependent phase is an intrinsically many-body object, and the local $U(1)$ transformations available in a matter-only theory shift every hop by one and the same amount. The same contrast is visible in the synthetic-dimension construction used for the encoding, where laser phases of this form generate a genuine but *uniform* synthetic gauge field, one that depends on where a particle hops and never on how many are present [32]. What the present scheme delivers is an effective Bose-Hubbard bond whose tunnelling and on-site interaction are renormalized and widely tunable.

This analysis also indicates what a successful realization must supply, and there are two complementary routes. The first is to make the drive occupation selective. Because the interactions split the occupation-resolved processes apart in energy, frequency selectivity becomes occupation selectivity: a drive with several tones, each quasi-resonant with one process and each carrying its own adjustable phase, can imprint a different phase according to the occupation. This is how the recent experimental realization of one-dimensional anyons with neutral atoms operates, with three tones addressing the three occupation-resolved tunnelling processes of a tilted lattice and the phase of one of them setting the statistical angle [17], and it is the principle behind multicolour lattice-depth modulation [18]. Applied here, this means driving each of the resonances identified in Sec. 2.3 with its own tone and phase—a direct and well-defined extension of the present work. The second route supplies the nonlinearity through the interactions themselves rather than through the drive, as in the Raman-coupled condensate with imbalanced interstate scattering lengths that realizes the chiral BF theory [42]; there too the density dependence originates in genuinely two-body physics, as it does in the atomic mixture where local $U(1)$ invariance is protected by heteronuclear spin-changing collisions [40].

Several further directions remain open. The construction developed here assumes translational invariance, which is what allows the effective Hamiltonian of a chain to be derived from that of a single bond. In a regular array that invariance is obvious; what deserves explicit study is the effect of uncertainties in the tweezer positions. Indeed, as Eq. (14) shows, the resonance condition $2C_{11} = C_{02}$ is sensitive at the tens-of-kilohertz level to displacements of order one per cent of the atomic separation. Relaxing the two-particle truncation would give access to the full phase diagram of the model, including the pair-superfluid and partially paired phases and the statistically induced exclusion found in the unconstrained case [10, 31]. On the experimental side the dominant limitations are the finite Rydberg lifetime, which as noted in Sec. 3.8 already constrains the slower oscillations of Fig. 3—the same competition between interaction strength and decay rate that has set the figure of merit for Rydberg-based schemes since the original gate proposal [28]—and the heating inherent to periodic driving, which is exponentially suppressed at large drive frequency and leaves a long prethermal window in which the engineered Hamiltonian governs the dynamics [25].

Taken together, the analysis identifies where the difficulty in engineering fractional statistics lies. It is not in producing large, tunable, occupation-dependent tunnelling *amplitudes*, which second-order Rydberg exchange supplies readily, but in producing an occupation-dependent tunnelling *phase*, which no arrangement of static one-body couplings can generate. Recognizing this fact converts the negative result into a concrete design re-

quirement, and one that existing multi-tone Floquet techniques are already known to meet.

Bibliography

[1] J. M. Leinaas and J. Myrheim. On the theory of identical particles. *Il Nuovo Cimento B*, 37(1):1–23, 1977. DOI: 10.1007/BF02727953.

[2] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma. Non-Abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80:1083–1159, 2008. DOI: 10.1103/RevModPhys.80.1083.

[3] M. Lewenstein, A. Sanpera, and V. Ahufinger. *Ultracold Atoms in Optical Lattices: Simulating Quantum Many-Body Systems*. Oxford University Press, Oxford, 2012. ISBN 978-0-19-957312-7. DOI: 10.1093/acprof:oso/9780199573127.001.0001.

[4] F. Wilczek. Quantum mechanics of fractional-spin particles. *Phys. Rev. Lett.*, 49: 957–959, 1982. DOI: 10.1103/PhysRevLett.49.957.

[5] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091–4094, 1995. DOI: 10.1103/PhysRevLett.74.4091.

[6] I. Bloch, J. Dalibard, and S. Nascimbène. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8:267–276, 2012. DOI: 10.1038/nphys2259.

[7] L. Tarruell, D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger. Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice. *Nature*, 483: 302–305, 2012. DOI: 10.1038/nature10871.

[8] J. Léonard, S. Kim, J. Kwan, P. Segura, F. Grusdt, C. Repellin, N. Goldman, and M. Greiner. Realization of a fractional quantum Hall state with ultracold atoms. *Nature*, 619:495–499, 2023. DOI: 10.1038/s41586-023-06122-4.

[9] T. Keilmann, S. Lanzmich, I. McCulloch, and M. Roncaglia. Statistically induced phase transitions and anyons in 1D optical lattices. *Nature Communications*, 2:361, 2011. DOI: 10.1038/ncomms1353.

[10] S. Greschner and L. Santos. Anyon Hubbard model in one-dimensional optical lattices. *Phys. Rev. Lett.*, 115:053002, 2015. DOI: 10.1103/PhysRevLett.115.053002.

[11] A. Frölian, C. S. Chisholm, E. Neri, C. R. Cabrera, R. Ramos, A. Celi, and L. Tarruell. Realizing a 1D topological gauge theory in an optically dressed BEC. *Nature*, 608: 293–297, 2022. DOI: 10.1038/s41586-022-04943-3.

[12] U. Aglietti, L. Griguolo, R. Jackiw, S.-Y. Pi, and D. Seminara. Anyons and chiral solitons on a line. *Phys. Rev. Lett.*, 77:4406–4409, 1996. DOI: 10.1103/PhysRevLett.77.4406.

[13] A. Kundu. Exact solution of double δ function Bose gas through an interacting anyon gas. *Phys. Rev. Lett.*, 83:1275–1278, 1999. DOI: 10.1103/PhysRevLett.83.1275.

[14] M. Bonkhoff, K. Jägering, S. Eggert, A. Pelster, M. Thorwart, and T. Posske. Bosonic continuum theory of one-dimensional lattice anyons. *Phys. Rev. Lett.*, 126:163201, 2021. DOI: 10.1103/PhysRevLett.126.163201.

[15] S. Dhar, B. Wang, M. Horvath, A. Vashisht, Y. Zeng, M. B. Zvonarev, N. Goldman, Y. Guo, M. Landini, and H.-C. Nägerl. Observing anyonization of bosons in a quantum gas. *Nature*, 642:53–57, 2025. DOI: 10.1038/s41586-025-09016-9.

[16] C. Sträter, S. C. L. Srivastava, and A. Eckardt. Floquet realization and signatures of one-dimensional anyons in an optical lattice. *Phys. Rev. Lett.*, 117:205303, 2016. DOI: 10.1103/PhysRevLett.117.205303.

[17] J. Kwan, P. Segura, Y. Li, S. Kim, A. V. Gorshkov, A. Eckardt, B. Bakkali-Hassani, and M. Greiner. Realization of one-dimensional anyons with arbitrary statistical phase. *Science*, 384(6696):546–551, 2024. DOI: 10.1126/science.adi3252.

[18] L. Cardarelli, S. Greschner, and L. Santos. Engineering interactions and anyon statis-

tics by multicolor lattice-depth modulations. *Phys. Rev. A*, 94:023615, 2016. DOI: 10.1103/PhysRevA.94.023615.

[19] A. Eckardt. Colloquium: Atomic quantum gases in periodically driven optical lattices. *Rev. Mod. Phys.*, 89:011004, 2017. DOI: 10.1103/RevModPhys.89.011004.

[20] N. Goldman and J. Dalibard. Periodically driven quantum systems: Effective Hamiltonians and engineered gauge fields. *Phys. Rev. X*, 4:031027, 2014. DOI: 10.1103/PhysRevX.4.031027.

[21] J. Struck, C. Ölschläger, M. Weinberg, P. Hauke, J. Simonet, A. Eckardt, M. Lewenstein, K. Sengstock, and P. Windpassinger. Tunable gauge potential for neutral and spinless particles in driven optical lattices. *Phys. Rev. Lett.*, 108:225304, 2012. DOI: 10.1103/PhysRevLett.108.225304.

[22] P. Hauke, O. Tieleman, A. Celi, C. Ölschläger, J. Simonet, J. Struck, M. Weinberg, P. Windpassinger, K. Sengstock, M. Lewenstein, and A. Eckardt. Non-Abelian gauge fields and topological insulators in shaken optical lattices. *Phys. Rev. Lett.*, 109:145301, 2012. DOI: 10.1103/PhysRevLett.109.145301.

[23] G. Jotzu, M. Messer, R. Desbuquois, M. Lebrat, T. Uehlinger, D. Greif, and T. Esslinger. Experimental realization of the topological Haldane model with ultracold fermions. *Nature*, 515:237–240, 2014. DOI: 10.1038/nature13915.

[24] M. Weinberg, C. Ölschläger, C. Sträter, S. Prelle, A. Eckardt, K. Sengstock, and J. Simonet. Multiphoton interband excitations of quantum gases in driven optical lattices. *Phys. Rev. A*, 92:043621, 2015. DOI: 10.1103/PhysRevA.92.043621.

[25] A. Rubio-Abadal, M. Ippoliti, S. Hollerith, D. Wei, J. Rui, S. L. Sondhi, V. Khemani, C. Gross, and I. Bloch. Floquet prethermalization in a Bose-Hubbard system. *Phys. Rev. X*, 10:021044, 2020. DOI: 10.1103/PhysRevX.10.021044.

[26] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82:2313–2363, 2010. DOI: 10.1103/RevModPhys.82.2313.

[27] A. Browaeys and T. Lahaye. Many-body physics with individually controlled Rydberg atoms. *Nature Physics*, 16:132–142, 2020. DOI: 10.1038/s41567-019-0733-z.

[28] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85:2208–2211, 2000. DOI: 10.1103/PhysRevLett.85.2208.

[29] M. Qiao, G. Emperauger, C. Chen, L. Homeier, S. Hollerith, G. Bornet, R. Martin, B. Gély, L. Klein, D. Barredo, S. Geier, N.-C. Chiu, F. Grusdt, A. Bohrdt, T. Lahaye, and A. Browaeys. Realization of a doped quantum antiferromagnet in a Rydberg tweezer array. *arXiv preprint arXiv:2501.08233*, 2025.

[30] A. A. Abrikosov. Electron scattering on magnetic impurities in metals and anomalous resistivity effects. *Physics Physique Fizika*, 2:5–20, 1965. DOI: 10.1103/PhysicsPhysiqueFizika.2.5.

[31] W. Zhang, S. Greschner, E. Fan, T. C. Scott, and Y. Zhang. Ground-state properties of the one-dimensional unconstrained pseudo-anyon Hubbard model. *Phys. Rev. A*, 95:053614, 2017. DOI: 10.1103/PhysRevA.95.053614.

[32] A. Celi, P. Massignan, J. Ruseckas, N. Goldman, I. B. Spielman, G. Juzeliūnas, and M. Lewenstein. Synthetic gauge fields in synthetic dimensions. *Phys. Rev. Lett.*, 112:043001, 2014. DOI: 10.1103/PhysRevLett.112.043001.

[33] D. Greif, T. Uehlinger, G. Jotzu, L. Tarruell, and T. Esslinger. Short-range quantum magnetism of ultracold fermions in an optical lattice. *Science*, 340(6138):1307–1310, 2013. DOI: 10.1126/science.1236362.

[34] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein, and T. Pfau. The physics of dipolar

bosonic quantum gases. *Reports on Progress in Physics*, 72(12):126401, 2009. DOI: 10.1088/0034-4885/72/12/126401.

[35] T. Oka and S. Kitamura. Floquet engineering of quantum materials. *Annu. Rev. Condens. Matter Phys.*, 10:387–408, 2019. DOI: 10.1146/annurev-conmatphys-031218-013423.

[36] N. Goldman, J. Dalibard, M. Aidelsburger, and N. R. Cooper. Periodically driven quantum matter: The case of resonant modulations. *Phys. Rev. A*, 91:033632, 2015. DOI: 10.1103/PhysRevA.91.033632.

[37] M. Bukov, L. D'Alessio, and A. Polkovnikov. Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering. *Advances in Physics*, 64(2):139–226, 2015. DOI: 10.1080/00018732.2015.1055918.

[38] T. Chanda, J. Zakrzewski, M. Lewenstein, and L. Tagliacozzo. Confinement and lack of thermalization after quenches in the bosonic Schwinger model. *Phys. Rev. Lett.*, 124:180602, 2020. DOI: 10.1103/PhysRevLett.124.180602.

[39] M. Aidelsburger, L. Barbiero, A. Bermudez, T. Chanda, A. Dauphin, D. González-Cuadra, P. R. Grzybowski, S. Hands, F. Jendrzejewski, J. Jünemann, G. Juzeliūnas, V. Kasper, A. Piga, S.-J. Ran, M. Rizzi, G. Sierra, L. Tagliacozzo, E. Tirrito, T. V. Zache, J. Zakrzewski, E. Zohar, and M. Lewenstein. Cold atoms meet lattice gauge theory. *Philos. Trans. R. Soc. A*, 380(2216):20210064, 2021. DOI: 10.1098/rsta.2021.0064.

[40] A. Mil, T. V. Zache, A. Hegde, A. Xia, R. P. Bhatt, M. K. Oberthaler, P. Hauke, J. Berges, and F. Jendrzejewski. A scalable realization of local U(1) gauge invariance in cold atomic mixtures. *Science*, 367(6482):1128–1130, 2020. DOI: 10.1126/science.aaz5312.

[41] L. Tagliacozzo, A. Celi, P. Orland, M. W. Mitchell, and M. Lewenstein. Simulation of non-Abelian gauge theories with optical lattices. *Nature Communications*, 4:2615, 2013. DOI: 10.1038/ncomms3615.

[42] C. S. Chisholm, A. Frölian, E. Neri, R. Ramos, L. Tarruell, and A. Celi. Encoding a one-dimensional topological gauge theory in a Raman-coupled Bose-Einstein condensate. *Phys. Rev. Research*, 4:043088, 2022. DOI: 10.1103/PhysRevResearch.4.043088.

[43] G. Evenbly. Tensors.net: Example codes for exact diagonalization. `https://www.tensors.net`, 2019. Accessed 2026.

[44] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Natl. Bur. Stand.*, 45(4):255–282, 1950. DOI: 10.6028/jres.045.026.

[45] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, 1998. DOI: 10.1137/1.9780898719628.

[46] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods*, 17:261–272, 2020. DOI: 10.1038/s41592-019-0686-2.

A Numerical results and experimental parameters

The figures of this thesis are placed where the quantities they show are derived. This section collects what they establish, states it in the units an experiment would use, and reads it against the measurements the construction is modelled on. Everything below refers to the reference interaction geometry of Fig. 3—$C_{00} = 0$, $C_{01} = 20$, $C_{11} = 15$, $C_{02} = 2C_{11} = 30$, $C_{12} = 45$, $C_{22} = 55\,\mathrm{MHz}$—with $\Omega_\alpha = \Omega_\beta = 1\,\mathrm{MHz}$, $\omega' = 40\,\mathrm{MHz}$ and $\kappa = 1.2$ unless stated otherwise.

A.1 Two-site building block: spectrum and dynamics

The driven bond behaves exactly as the effective Hamiltonian (42)–(43) predicts, and the prediction is quantitative. Total occupation is conserved, so the nine configurations split into the blocks $N = 1$, $N = 2$, $N = 3$ and the decoupled $|22\rangle$, and each block is solved in closed form in Sec. 3.8.1: the $N = 1$ and $N = 3$ blocks give generalized Rabi flopping, Eqs. (50) and (52), whose contrast is reduced only by the residual AC-Stark detunings (51) and (53), while the $N = 2$ block is a V system in which $|11\rangle$ couples to a single bright combination of $|02\rangle$ and $|20\rangle$ with the branching ratio (55) fixed entirely by the two Bessel factors. Physically these are the resonant dipolar exchange oscillations of a Rydberg pair, dressed by the drive; oscillations of this form have been resolved between individual atoms in a tweezer array, with the exchange rate tuned through the orientation of the pair [27, 29].

The four amplitudes come out separated by an order of magnitude at the reference point, $A_1 = 0.0500$, $A_2 = A_3 = 0.0166$ and $A_4 = 0.0101\,\mathrm{MHz}$, i.e. oscillation periods from tens to hundreds of microseconds. The hierarchy is set by the Bessel factors rather than by the Rabi frequencies: $A_2 = A_3$ carries $J_{\mathcal{N}}(\kappa)$ and vanishes at $\kappa = 3.8317$, the first zero of J_1, while A_4 carries the sum Σ_4 of (75) and vanishes at $\kappa = 1.3292$, 3.3175 and 4.5008. That the reference value $\kappa = 1.2$ sits just below the first of these is why A_4 is an order of magnitude below A_1 there, and moving to $\kappa = 0.8$ raises it by a factor of four. This switching of a channel by the modulation index alone is the same dynamic localization that suppresses tunnelling at the zeros of a Bessel function in a shaken lattice [19, 21], and it is the sharpest available check that the Floquet description of Sec. 3, and not some residual static coupling, is what governs the bond.

Read as a Hubbard model, the bond is strongly attractive and not weakly coupled. The addition energies give $\mu_{\mathrm{eff}} = 0.0435$ and $U_{\mathrm{eff}} = -0.0470\,\mathrm{MHz}$ at the reference point, so $\mu_{\mathrm{eff}}/U_{\mathrm{eff}} = -0.924$ and the dimensionless couplings are $A_1/U_{\mathrm{eff}} = -1.064$, $A_2/U_{\mathrm{eff}} = -0.353$ and $A_4/U_{\mathrm{eff}} = -0.215$ (Fig. 7): the tunnelling is comparable to the interaction rather than small against it, which places the model outside the regime in which Mott lobes are the useful picture. Two structural constraints on this tuning emerged in Sec. 4.2. The ratio depends on the Rabi frequencies and the modulation index only through the single product $x = r\Sigma(\kappa)$, Eq. (71), so those are not two independent handles; and because $x > 0$ whenever $\Sigma > 0$, the window $-\frac{3}{4} < \mu_{\mathrm{eff}}/U_{\mathrm{eff}} < -\frac{1}{3}$ of (72) is closed unless κ is taken below the first zero of Σ at $\kappa = 1.0585$. Against this, U_{eff} itself passes through zero at $\Omega_\beta/\Omega_\alpha = 1.842$ for $\kappa = 1.2$—a useful knob, since it makes the effective interaction tunable through zero at fixed atomic geometry, something an optical lattice cannot do without a Feshbach resonance.

The magnitudes can be matched to the anyon-Hubbard target, but not comfortably. At the reference point $A_2/A_1 = 0.332$ and $A_4/A_1 = 0.202$ against the required $\sqrt{2}$ and 2 of (8): the two doublon channels come out roughly four and ten times too small. Since (75) is three equations in four unknowns this is not a difficulty, and solving them over $\kappa \in [0.3, 6]$ yields six solutions—but every one of them demands a strongly asymmetric drive, $\Omega_\beta/\Omega_\alpha$

between 3.9 and 10.6, and every one lands in the narrow band $\mu_{\text{eff}}/U_{\text{eff}} \in [-0.303, -0.263]$ just above the $-1/3$ asymptote. Matching the amplitudes therefore spends the freedom that would otherwise tune the chemical-potential-to-interaction ratio, and moving that ratio elsewhere requires retuning the interaction geometry C_{ij} itself, which is fixed by the atomic states and the trap spacing.

A.2 The density-dependent phase

A *tunable* density-dependent phase does not emerge. Matching the phases of Table 3 to the target pattern $(0, 0, \phi, \phi)$ gives an over-determined but consistent linear system whose solutions are $\phi \in \{0, \pi\}$, Eq. (84), and the same conclusion follows independently in Sec. 3.8.2 from the per-bond fluxes: $\vartheta_1(j) = 0$ forces the α_j to be constant, $\vartheta_4(j) = -\theta$ forces $\beta_j = \beta_0 + j\theta$, and $\vartheta_2(j)$ then grows linearly with j and cannot equal a constant unless $\theta = 0$ modulo π. Two derivations that share no algebra reach the same values, and Sec. 4.5 identifies the reason: the available phases are occupation independent, whereas the target phases are not. The scheme therefore reaches the pseudofermionic point $\phi = \pi$, which is the same set of values a static Raman protocol already permits; what the single-tone drive adds is a renormalized *Bose*-Hubbard bond with widely tunable parameters, not a tunable statistical angle.

Promoting the bond to a chain is what makes this failure measurable rather than merely formal. Because every term of $H_{\text{eff}}^{(2)}$ moves one quantum between neighbours, $\hat{N}$ of (46) is conserved and the spectrum splits into $2L + 1$ blocks, so the $L = 10$ diagonalization of Sec. 3.7 returns all $3^{10} = 59\,049$ levels rather than a low-lying truncation. The spectrum is not that of a conventional Hubbard chain: the lower envelope $E_0(N)$ is flat to within 0.019 MHz across $N = 0, \ldots, 4$, the global minimum falls at $N = 2$ rather than at unit filling, and the global gap $E_1 - E_0 = 1.3 \times 10^{-4}$ MHz is an inter-sector spacing rather than an excitation at fixed filling. Any statement about the unconstrained ground state is correspondingly fragile, and the remedy is to work at fixed filling—legitimate because $\hat{N}$ is conserved exactly, so restricting to one block is not an approximation. That is what makes the dynamics of Sec. 3.8.2 exact rather than variational.

At fixed filling the doublon is the observable that carries the statistics. Released into an otherwise empty chain it stays bound for $t \lesssim 100\,\mu$s and only then dissociates, the initial site retaining 0.079 of its two quanta by $t \simeq 577\,\mu$s, whereas a single excitation released from the same site uses only the A_1 channel and produces a clean light cone with revivals near $t \simeq 340$ and $530\,\mu$s (Fig. 4). The composite propagates slower than its constituents by roughly A_4/A_1 and survives long before spreading, which is the direct dynamical counterpart of the amplitude hierarchy measured above. Setting the bond phases to the pattern gauge-equivalent to statistical angle $\phi = -\theta$ changes this substantially: the largest deviation from the $\theta = 0$ map reaches 0.657 out of 2, and it does so at $\theta = 4\pi/3$—where the doublon migrates steadily to site 7 by $t \simeq 500\,\mu$s—rather than at $\theta = \pi$, taking the values 0.435, 0.451, 0.408, 0.657 and 0.361 at $\theta = \pi/3$, $2\pi/3$, π, $4\pi/3$ and $5\pi/3$ (Fig. 5). The whole circle must be sampled, because the reversal identity (56) runs the density backwards in time rather than leaving it unchanged, so θ and $2\pi - \theta$ are distinct evolutions. This directed, phase-controlled transport of a composite object is the real-space counterpart of the asymmetric momentum distribution that is the standard signature of one-dimensional anyons [10, 17], and it is the observable lost when ϕ is restricted to $\{0, \pi\}$.

Finally, the exact diagonalization of the *bare* anyon-Hubbard chain in App. C gives the reference against which such a measurement would be read. Over the plane spanned by J/U and $\phi \in [0, \pi]$ at $L = 10$, the two lowest gaps are largest at $\phi = \pi$, peaking at $0.33\,U$ and $0.36\,U$ near $J/U \simeq 0.7$–0.8, while $E_3 - E_0$ and $E_4 - E_0$ peak instead at $\phi = 0$ (Figs. 8

and 9): the statistical phase pushes the lowest excitations up while lowering the higher ones, by amounts of order a third of the interaction energy. The target model therefore depends on ϕ strongly enough to be resolved spectroscopically, which is what makes the null result established here a falsifiable statement about the drive rather than a statement about the sensitivity of the probe.

B A concrete optical-tweezer experiment

We close by asking what one would actually do in the laboratory, because the result of this thesis is not a dead end for the platform but a prescription. Everything derived above—the encoding, the interaction ladder, the second-order exchange, the Bessel renormalization, the effective Hubbard parameters, and the doublon transport that would carry the statistics—is measurable in a present-day Rydberg tweezer array, and the measurements form a natural sequence whose last step is the multi-tone upgrade that would deliver $\phi \neq 0$.

Register and encoding. The starting point is two atoms, later a short chain of five to ten, held in optical tweezers imprinted by a spatial light modulator or a pair of crossed acousto-optic deflectors. Because the traps load stochastically, the array is imaged and the atoms rearranged into the target geometry, which routinely gives defect-free registers with filling above 98% for arrays of up to about a hundred atoms [27]. The separation R is a few micrometres and the angle Θ between the array axis and the quantization axis is a free parameter; together they are tuned, as in Sec. 2.2, so that the six shifts C_{ij} satisfy the dispersive condition (13) and the resonance $2C_{11} = C_{02}$. The tens-of-megahertz shifts assumed throughout are what principal quantum numbers $n \approx 50$–60 give at micrometre spacings [27, 29]; a static electric field along the quantization axis, as in Ref. [28], supplies the extra freedom needed to impose the resonance exactly. Three Rydberg levels per atom encode $|0\rangle, |1\rangle, |2\rangle$, chosen so that no pair of them is resonantly dipole coupled and H_c is therefore diagonal in the occupation basis, as required in Sec. 2.1. The two rungs are driven at $\Omega_\alpha, \Omega_\beta \approx 1\,\mathrm{MHz}$, an order of magnitude below the smallest energy denominator $2C_{11} - C_{01} = 10\,\mathrm{MHz}$, so that the expansion parameter is $\Omega/\Delta E \approx 0.1$ and the second-order treatment is controlled. The site-dependent phases $\alpha_j = j\alpha$, $\beta_j = j\beta$ are supplied either by the running-wave geometry of the driving beams, as in the synthetic-dimension construction [32], or by individually addressed beams with programmable phase—both standard in tweezer arrays. The Floquet drive is a modulation of the energy of $|2\rangle$ at $\omega'/2\pi = 40\,\mathrm{MHz}$ with modulation index $\kappa = A/\omega'$ of order unity.

One requirement deserves emphasis because it is easily mistaken for a technicality. The $N = 2$ degeneracy is not protected by any symmetry: it is the tuned condition $2C_{11} = C_{02}$, imposed at one separation, and Eq. (14) shows that a relative displacement of one per cent already detunes it by $\sim 30\,\mathrm{kHz}$—comparable to the exchange amplitudes themselves. Tight confinement in the Lamb–Dicke sense is therefore a precondition for the whole construction rather than a refinement of the readout, and the same estimate sets the site-to-site uniformity the traps must achieve once the bond is extended to a chain.

Sequence and readout. Each shot prepares a chosen configuration $|n_0, n_1\rangle$ by state-selective Rydberg excitation, evolves it for a time t under $H_\Omega + H_c$ plus the drive, and reads it out by state-selective de-excitation followed by fluorescence imaging, which returns the occupation of every site [27]. Scanning t reproduces the traces of Fig. 3. The expected rates are set by $\Omega^2/\Delta E \approx 100\,\mathrm{kHz}$, i.e. periods of tens of microseconds, comfortably inside the few-hundred-microsecond Rydberg lifetime; the slower $N = 3$ traces are not,

as already noted in Sec. 3.8, and reaching them would require smaller separations, higher n, or a dressed encoding. Two features of the effective model make this readout well posed. First, $\hat{N}$ is conserved exactly, Eq. (46), so preparing a definite total occupation confines the dynamics to one block for all time and every observable is measured at fixed filling—which Sec. 3.7 showed to be essential, since the sector envelopes are degenerate to within 0.019 MHz at the parameters of Table 1, and the unconstrained ground state is not a robust reference. Second, the populations are stroboscopic: comparison with theory must be made at integer multiples of the drive period, the micromotion being carried by a separate kick operator [36, 37].

The measurements. Four steps validate the construction, in an order in which each supplies the input for the next. *(i) Calibrate the interaction ladder.* Microwave spectroscopy of the pair in each of the nine configurations returns the six C_{ij} directly, verifies (13), and locates the geometry at which $2C_{11} = C_{02}$. *(ii) Measure the amplitude hierarchy.* Preparing $|0,1\rangle$, $|1,1\rangle$ and $|1,2\rangle$ separately and reading off the flop frequencies gives A_1, $A_2 = A_3$ and A_4. This thesis predicts 0.0500, 0.0166 and 0.0101 MHz at $\Omega_\alpha = \Omega_\beta = 1$ MHz and $\kappa = 1.2$, against the anyon-Hubbard target ratio $1 : \sqrt{2} : \sqrt{2} : 2$ of (8)—a sharp question with a sharp answer, and (75) says which $\Omega_\alpha, \Omega_\beta, \kappa$ close the gap, at the cost of a drive asymmetry $\Omega_\beta/\Omega_\alpha$ between 3.9 and 10.6. *(iii) Verify the Bessel renormalization.* Scanning κ at fixed geometry should make the $|11\rangle \leftrightarrow \{|02\rangle, |20\rangle\}$ rate track $|J_{\mathcal{N}}(\kappa)|$ and vanish at $\kappa = 3.8317$, and the $|12\rangle \leftrightarrow |21\rangle$ rate vanish at $\kappa = 1.3292$, 3.3175 and 4.5008—the dynamic-localization signature familiar from shaken lattices [19, 21]. Observing it establishes that the Floquet description of Sec. 3 is the correct one *before* any claim about phases is made. *(iv) Map the effective Hubbard parameters.* The addition energies $E_{01} - E_{00}$ and $E_{02} - 2E_{01} + E_{00}$ are spectroscopically accessible and give μ_{eff} and U_{eff} directly, Eq. (65). Sweeping the Rabi ratio at fixed κ then traces out Fig. 7, including the sign change of U_{eff}—predicted at $\Omega_\beta/\Omega_\alpha = 1.842$ for $\kappa = 1.2$—which is both a falsifiable prediction and useful in its own right, since it makes the effective interaction tunable through zero at fixed atomic geometry.

Testing the flux locking. The central result of Sec. 4 is that a single drive tone collapses the four tunnelling phases onto $\phi = 0$. On two sites this is not directly visible: the phase of a single bond can be removed by the local $U(1)$ redefinition of Sec. 4.5, so a genuine test requires a closed loop, and the chain supplies one. Releasing a doublon into an otherwise empty register and following $\langle \hat{n}_j(t) \rangle$ is the measurement, because Sec. 3.8.2 computed exactly what it must show. At $\phi = 0$ the doublon stays bound for $t \lesssim 100\,\mu$s and then dissociates symmetrically about its release site, up to the boundary asymmetry of an off-centre release; at $\phi \neq 0$ it acquires directed transport, migrating to site 7 of ten by $t \simeq 500\,\mu$s at $\theta = 4\pi/3$, with density differences of up to 0.657 out of 2 quanta. That contrast is large on the scale of the projection noise of site-resolved imaging, which for a single-site occupation falls as $1/\sqrt{M}$ over M repetitions and is therefore at the few-per-cent level after a few hundred shots, so the discrimination is not statistics limited. This thesis predicts the symmetric result for the single-tone drive. Observing it confirms the flux-locking argument; observing directed transport falsifies it. Because the reversal identity (56) makes θ and $2\pi - \theta$ distinct evolutions rather than mirror images, the scan must cover the full circle and not merely $[0, \pi]$.

The upgrade. The final step is the one the analysis points to. Replacing the single drive tone by three, each quasi-resonant with one of the occupation-resolved processes of

Sec. 2.3 and each carrying its own phase, makes the phase imprinted on a hop a function of the occupation left behind. This is the mechanism by which one-dimensional anyons of arbitrary statistical phase were realized in a quantum-gas microscope, where a tilted lattice was modulated with three components—resonant with the tunnelling from a singly-occupied to an empty site, between singly-occupied sites, and from a doubly-occupied site—and offsetting the phase of the middle component set the statistical angle [17]; it is the same principle as multicolour lattice-depth modulation [18]. The Rydberg register has a structural advantage here. What separates the tones in the lattice experiment is the on-site interaction, of order a few hundred hertz, whereas here it is the interaction shifts C_{ij}, tens of megahertz apart—close to five orders of magnitude more. The tones are therefore trivially resolved without compromising the drive strength, and giving each its own phase is a matter of adding channels to a synthesizer rather than a limitation of the platform. The same register, the same readout and the same sequence then measure ϕ itself: scanning the phase of the tone addressing the doubly-occupied channel should rotate the statistical angle continuously between the bosonic and pseudofermionic limits, with the directed doublon transport of Fig. 5 as the observable and the gap structure of App. C—where ϕ moves the lowest excitations by up to a third of U—as the spectroscopic cross-check.

That is the experiment this thesis argues for. Steps (i)–(iv) validate the machinery developed here—the encoding, the interaction ladder, the second-order exchange and the Floquet renormalization—as derived; the doublon-transport measurement confirms the restriction; and only the last ingredient, a drive that is selective in occupation rather than merely in frequency, remains to be added before an array of optical tweezers hosts anyons.

C Full derivation of the anyon-Hubbard model

This appendix derives the one-dimensional anyon-Hubbard model in full: the generalized Jordan–Wigner mapping, the anyonic commutation relations, and the bond Hamiltonian in the truncated occupation basis, in both the bosonic (density-dependent-phase) and the anyonic representations.

Mapping the bosonic anyon-Hubbard model

We start from the Bose–Hubbard model with a density-dependent phase [10],

$$H = -J\sum_j \left(b_j^\dagger e^{i\phi n_j} b_{j+1} + \text{H.c.}\right) + \frac{U}{2}\sum_j n_j(n_j-1), \tag{85}$$

and introduce the generalized Jordan–Wigner transformation

$$a_j = e^{i\phi\sum_{l<j} n_l}\, b_j, \qquad a_j^\dagger = b_j^\dagger\, e^{-i\phi\sum_{l<j} n_l}, \qquad n_l = b_l^\dagger b_l. \tag{86}$$

The phase string commutes with the on-site operators, so the occupation is invariant, $n_j = a_j^\dagger a_j = b_j^\dagger b_j$, and the interaction term $\frac{U}{2}\sum_j n_j(n_j-1)$ is unchanged. For the tunnelling, since n_l commutes with b_{j+1} for all $l \le j$,

$$e^{-i\phi\sum_{l<j} n_l}\, e^{i\phi\sum_{m<j+1} n_m} = e^{i\phi n_j} \qquad \Longrightarrow \qquad a_j^\dagger a_{j+1} = b_j^\dagger e^{i\phi n_j} b_{j+1}, \tag{87}$$

so that, summing over sites and adding the Hermitian conjugate, the model becomes the anyon-Hubbard Hamiltonian

$$H = -J\sum_j \left(a_j^\dagger a_{j+1} + \text{H.c.}\right) + \frac{U}{2}\sum_j n_j(n_j-1). \tag{88}$$

Verifying the anyonic commutation relations

The anyonic operators obey [9]

$$a_j a_k^\dagger - e^{\,i\phi\,\mathrm{sgn}(j-k)}\, a_k^\dagger a_j = \delta_{jk}, \tag{89}$$

$$a_j a_k - e^{-i\phi\,\mathrm{sgn}(j-k)}\, a_k a_j = 0. \tag{90}$$

We verify these using $f(n_j)\, b_j = b_j\, f(n_j - 1)$ and $e^{i\phi n_j}\, b_j^\dagger = e^{i\phi}\, b_j^\dagger\, e^{i\phi n_j}$.

Case 1: $a_j a_k^\dagger$**.** Assume $j > k$, so that n_k appears in the string of a_j. Then

$$a_j a_k^\dagger = (e^{i\phi \sum_{l<j} n_l} b_j)(b_k^\dagger e^{-i\phi \sum_{m<k} n_m}) = e^{i\phi}\, b_k^\dagger\, e^{i\phi \sum_{l<j} n_l}\, e^{-i\phi \sum_{m<k} n_m}\, b_j = e^{i\phi}\, a_k^\dagger a_j, \tag{91}$$

where pushing $b_k^\dagger$ through $e^{i\phi n_k}$ produces the factor $e^{i\phi}$. For $j < k$ one obtains $e^{-i\phi}$, so that $a_j a_k^\dagger = e^{i\phi\,\mathrm{sgn}(j-k)} a_k^\dagger a_j$, i.e. (89).

Case 2: $a_j a_k$**.** Assume $j > k$. Passing b_k through $e^{i\phi n_k}$ gives $b_k e^{i\phi n_k} = e^{i\phi} e^{i\phi n_k} b_k$, hence $a_k a_j = e^{i\phi} a_j a_k$, i.e. $a_j a_k = e^{-i\phi} a_k a_j$. Since $\mathrm{sgn}(j-k) = 1$ for $j > k$, this is (90). (The overall sign of ϕ is a matter of convention; the signs above are those consistent with the mapping (86) and the phase $e^{i\phi n_j}$ in (85).)

Truncated local Hilbert space $\{|0\rangle, |1\rangle, |2\rangle\}$

Restricting each site to at most two particles, with $|0\rangle = (1,0,0)^{\mathsf{T}}$, $|1\rangle = (0,1,0)^{\mathsf{T}}$, $|2\rangle = (0,0,1)^{\mathsf{T}}$, the bosonic operators become

$$b = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & \sqrt{2} \\ 0 & 0 & 0 \end{pmatrix}, \quad b^\dagger = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \sqrt{2} & 0 \end{pmatrix}, \quad n = b^\dagger b = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix}, \tag{92}$$

i.e. $b = |0\rangle\langle 1| + \sqrt{2}\, |1\rangle\langle 2|$ and $b^\dagger = |1\rangle\langle 0| + \sqrt{2}\, |2\rangle\langle 1|$, with $b\,|n\rangle = \sqrt{n}\, |n-1\rangle$ $(n \le 2)$ and the truncation $b^3 = (b^\dagger)^3 = 0$.

Bond Hamiltonian in the occupation basis (bosonic representation)

Here the diagonal phase operator is $e^{i\phi n} = \mathrm{diag}(1,\, e^{i\phi},\, e^{i2\phi})$, and the nonzero tunnelling matrix elements on a bond $(j, j+1)$ are

$$\langle n_j + 1,\, n_{j+1} - 1|\, H_{j,j+1}\, |n_j, n_{j+1}\rangle = -J\, e^{i\phi n_j} \sqrt{n_{j+1}} \sqrt{n_j + 1}, \tag{93}$$

$$\langle n_j - 1,\, n_{j+1} + 1|\, H_{j,j+1}\, |n_j, n_{j+1}\rangle = -J\, e^{-i\phi(n_j - 1)} \sqrt{n_j} \sqrt{n_{j+1} + 1}, \tag{94}$$

with $n_j, n_{j+1} \in \{0, 1, 2\}$. The four single-particle processes carry

$$\begin{aligned} &|0,1\rangle \to |1,0\rangle : \; -J, && |0,2\rangle \to |1,1\rangle : \; -J\sqrt{2}, \\ &|1,1\rangle \to |2,0\rangle : \; -J\sqrt{2}\, e^{i\phi}, && |1,2\rangle \to |2,1\rangle : \; -2J\, e^{i\phi}. \end{aligned} \tag{95}$$

Collecting them gives the bond operator

$$H_{j,j+1} = -J \sum_{n_j, n_{j+1}=0}^{2} \Big[e^{i\phi n_j} \sqrt{n_{j+1}} \sqrt{n_j + 1}\; |n_j + 1, n_{j+1} - 1\rangle\langle n_j, n_{j+1}| + \mathrm{H.c.} \Big], \tag{96}$$

and the full chain of L sites is $H = \sum_{j=1}^{L-1} H_{j,j+1} + U \sum_{j=1}^{L} |2\rangle_j\langle 2|$. In the ordered basis $\{|0,0\rangle, |0,1\rangle, \ldots, |2,2\rangle\}$ (lexicographic in $n_j n_{j+1}$), the bond Hamiltonian, with the on-site interaction U on both sites, is

$$H_{j,j+1} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -J & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & U & 0 & -J\sqrt{2} & 0 & 0 & 0 & 0 \\ 0 & -J & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -J\sqrt{2} & 0 & 0 & 0 & -J\sqrt{2}\, e^{-i\phi} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & U & 0 & -2J\, e^{-i\phi} & 0 \\ 0 & 0 & 0 & 0 & -J\sqrt{2}\, e^{i\phi} & 0 & U & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -2J\, e^{i\phi} & 0 & U & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2U \end{pmatrix}. \tag{97}$$

Bond Hamiltonian in the anyonic representation

In the anyonic representation the tunnelling $-J(a_j^\dagger a_{j+1} + \text{H.c.})$ carries no phases (they have been absorbed into the anyonic operators), so the amplitudes are

$$\begin{aligned} |0,1\rangle \to |1,0\rangle &: -J, & |0,2\rangle \to |1,1\rangle &: -J\sqrt{2}, \\ |1,1\rangle \to |2,0\rangle &: -J\sqrt{2}, & |1,2\rangle \to |2,1\rangle &: -2J, \end{aligned} \tag{98}$$

so that $H_{j,j+1} = -J \sum_{n_j, n_{j+1}=0}^{2} \left[\sqrt{n_{j+1}}\sqrt{n_j+1}\; |n_j+1, n_{j+1}-1\rangle\langle n_j, n_{j+1}| + \text{H.c.}\right]$, and, in the same ordered basis,

$$H_{j,j+1} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -J & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & U & 0 & -J\sqrt{2} & 0 & 0 & 0 & 0 \\ 0 & -J & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -J\sqrt{2} & 0 & 0 & 0 & -J\sqrt{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & U & 0 & -2J & 0 \\ 0 & 0 & 0 & 0 & -J\sqrt{2} & 0 & U & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -2J & 0 & U & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2U \end{pmatrix}. \tag{99}$$

Exact diagonalization of the truncated chain

To characterize the model derived above beyond the two-site building block, we diagonalize the chain of L sites numerically. With the local Hilbert space truncated to $\{|0\rangle, |1\rangle, |2\rangle\}$ the local dimension is $d = 3$ and the many-body space has dimension d^L, which already reaches $3^{10} = 59\,049$ for the chains considered here. Rather than storing H as a matrix, we follow the matrix-free strategy of [43]: the Hamiltonian is represented only as a *linear map* $|\psi\rangle \mapsto H|\psi\rangle$, so the memory cost is that of a single state vector, $O(d^L)$, instead of the $O(d^{2L})$ required by an explicit matrix.

The map is assembled from the two pieces of $H = \sum_{j=1}^{L-1} H_{j,j+1} + U \sum_{j=1}^{L} |2\rangle_j\langle 2|$ separately. The bond term is the $d^2 \times d^2$ operator

$$h = -J\Big[(b^\dagger e^{i\phi n}) \otimes b \;+\; (e^{-i\phi n} b) \otimes b^\dagger\Big], \tag{100}$$

reshaped to a rank-4 tensor $h^{n_j n_{j+1}}_{n'_j n'_{j+1}}$. To apply it to bond $(k, k+1)$ the state is reshaped as $\psi_{(d^k,\, d^2,\, d^{L-2-k})}$ and contracted with h over the two-site index, the result being accumulated over all bonds. Because $e^{i\phi n} = \mathrm{diag}(1, e^{i\phi}, e^{i2\phi})$ is complex for generic ϕ, h is complex Hermitian rather than real symmetric, and the state vector is kept complex throughout. The interaction term is diagonal in the occupation basis, so it is precomputed once as a vector $U \sum_j \delta_{n_j,2}$ over the full Hilbert space and applied by elementwise multiplication; keeping it separate from (100) guarantees that each site is counted exactly once, whereas folding it into the bond operator would double-count every interior site of an open chain.

This linear map is handed to a sparse eigensolver (the implicitly restarted Lanczos algorithm [44, 45] as provided by `scipy.sparse.linalg.eigsh` [46]), requesting the algebraically smallest eigenvalues. Lanczos is the standard choice here because it reduces the Hamiltonian to a tridiagonal form through a three-term recurrence that needs only the action of H on a vector, never its matrix elements, which is what the matrix-free representation above supplies [3]. Since the low-lying levels become nearly degenerate in the strongly interacting regime $J \ll U$, the Krylov subspace is enlarged to $n_{\mathrm{cv}} = 40$ and the tolerance set to 10^{-8}, which is far below the resolution of the figures below; the Lanczos start vector is drawn from a fixed seed so that the sweep is reproducible.

We verified the implementation in three ways. First, for $L = 2$ the assembled Hamiltonian reproduces the 9×9 bond matrix of the previous subsection, including the phased amplitudes $-J\sqrt{2}\, e^{\pm i\phi}$ and $-2J\, e^{\pm i\phi}$, to a maximum absolute deviation of 3×10^{-16}. Second, the assembled operator is Hermitian to machine precision. Third, for $L = 6$ the Lanczos ground energy agrees with a full dense diagonalization of the same operator to $\sim 10^{-14}$, both at $\phi = 0$ (where the model reduces to Bose–Hubbard) and at $\phi = \pi/2$.

Using this machinery we scanned the plane spanned by the statistical phase $\phi \in [0, \pi]$ and the ratio J/U, on a 21×21 grid with $U = 1$ and $L = 10$ sites with open boundaries, computing the five lowest eigenvalues $E_0, \ldots, E_4$ at every point. The 441 solves are mutually independent and were distributed over eight worker processes, the whole sweep taking about 39 minutes. The resulting gaps measured from the ground state are shown in Fig. 8 and the consecutive level spacings in Fig. 9. The two lowest gaps $E_1 - E_0$ and $E_2 - E_0$ are largest at $\phi = \pi$, peaking at $0.33\,U$ and $0.36\,U$ near $J/U \simeq 0.7$–0.8, whereas $E_3 - E_0$ and $E_4 - E_0$ instead peak at $\phi = 0$: the statistical phase pushes the lowest excitations up while lowering the higher ones. The narrow dark curves crossing the panels are genuine near-degeneracies, where the gap drops to $\sim 10^{-5}\, U$. We emphasize that the spectrum is computed in the full Hilbert space without a chemical potential, so these differences are global gaps over all particle-number sectors and the dark curves mark level crossings between states of different total occupation; the gap at fixed filling would require the Lanczos iteration to be restricted to a single number sector.

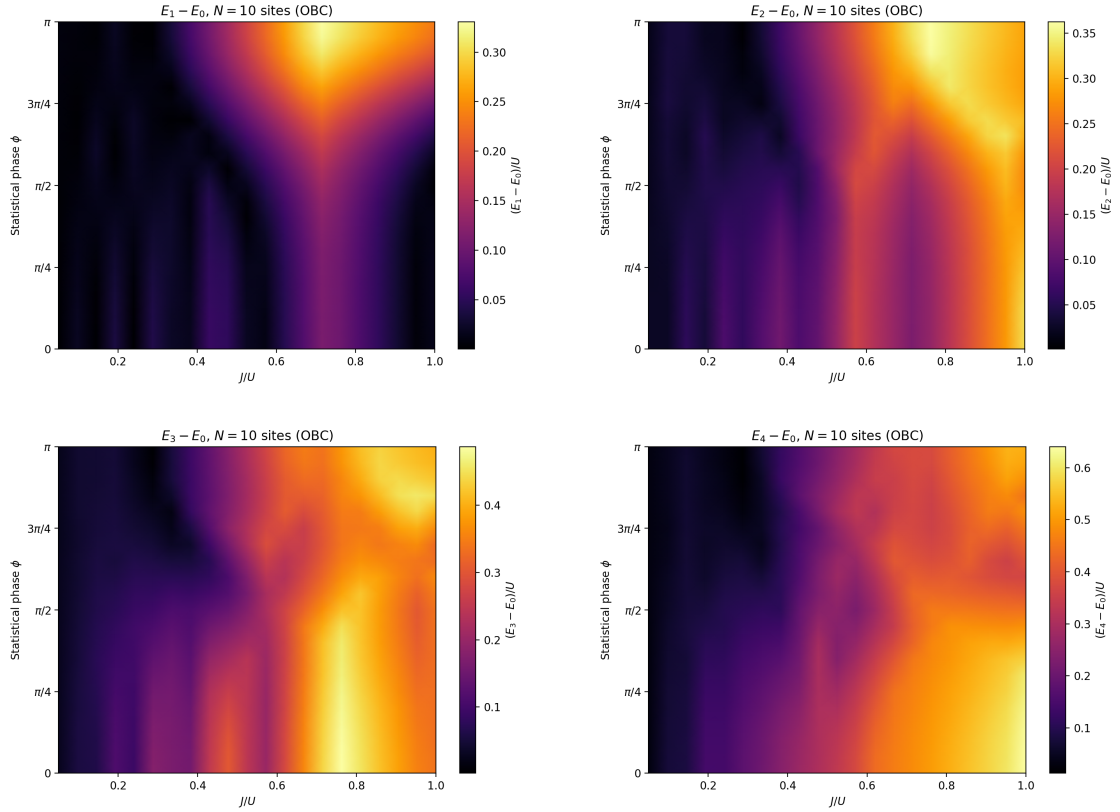


Figure 8: Energy gaps measured from the ground state of the anyon-Hubbard chain, $E_k - E_0$ for $k = 1, \ldots, 4$ (top left to bottom right), as a function of the ratio J/U and the statistical phase ϕ, for $L = 10$ sites with open boundaries. Energies are in units of U. The lowest two gaps are maximal at $\phi = \pi$, while $E_3 - E_0$ and $E_4 - E_0$ are maximal at $\phi = 0$.

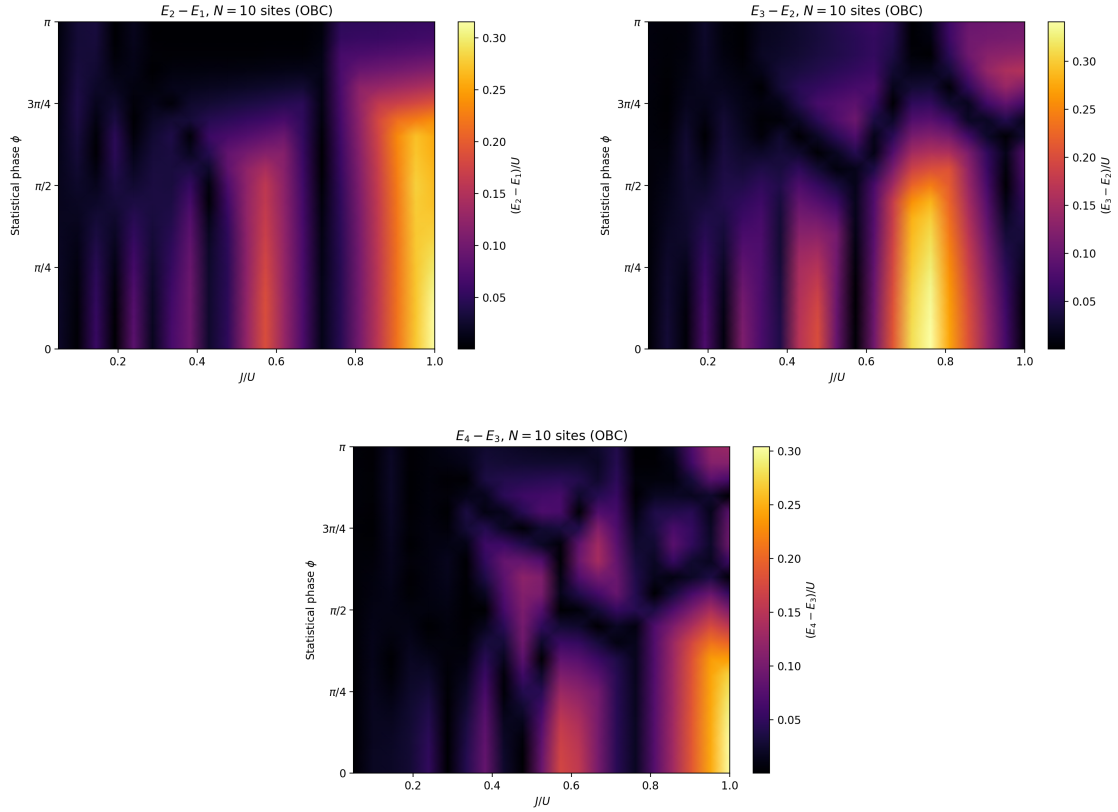


Figure 9: Consecutive level spacings $E_{k+1} - E_k$ for $k = 1, 2, 3$ (top left, top right, bottom) over the same $(J/U, \phi)$ plane and for the same chain as Fig. 8. The sharp dark ridges are near-degeneracies between adjacent excited levels, which for a spectrum computed over all particle-number sectors correspond to crossings between states of different total occupation.

D The four second-order transitions

This appendix carries out, for each of the four processes of the bond, the second-order sum (15)–(16) over the two virtual paths, and collects the resulting amplitudes quoted in Sec. 2.3. Throughout, the resonance condition $2C_{11} = C_{02}$ is used and the abbreviations K_1, K_2, K_3 of (17) apply.

D.1 Transition $|0,1\rangle \to |1,0\rangle$

Here $E_i = E_f = C_{01}$ (degenerate), with intermediates $|1,1\rangle$ and $|0,0\rangle$. *Path A* (via $|1,1\rangle$, $E_m = 2C_{11}$): $\langle 1,1|\, H_\Omega\, |0,1\rangle = \Omega_\alpha e^{i\alpha}$ and $\langle 1,0|\, H_\Omega\, |1,1\rangle = \Omega_\alpha e^{-i2\alpha}$, giving $\Omega_\alpha^2 e^{-i\alpha}/(C_{01} - 2C_{11})$. *Path B* (via $|0,0\rangle$, $E_m = 2C_{00}$): $\langle 0,0|\, H_\Omega\, |0,1\rangle = \Omega_\alpha e^{-i2\alpha}$ and $\langle 1,0|\, H_\Omega\, |0,0\rangle = \Omega_\alpha e^{i\alpha}$, giving $\Omega_\alpha^2 e^{-i\alpha}/(C_{01} - 2C_{00})$. Both carry the net phase $e^{i(\alpha - 2\alpha)} = e^{-i\alpha}$, so

$$A_{01\to 10} = \Omega_\alpha^2\, e^{-i\alpha} \left(\frac{1}{C_{01} - 2C_{11}} + \frac{1}{C_{01} - 2C_{00}} \right) \xrightarrow{2C_{11} = C_{02}} A^R_{01\to 10} = \Omega_\alpha^2\, e^{-i\alpha}\, K_1. \quad (101)$$

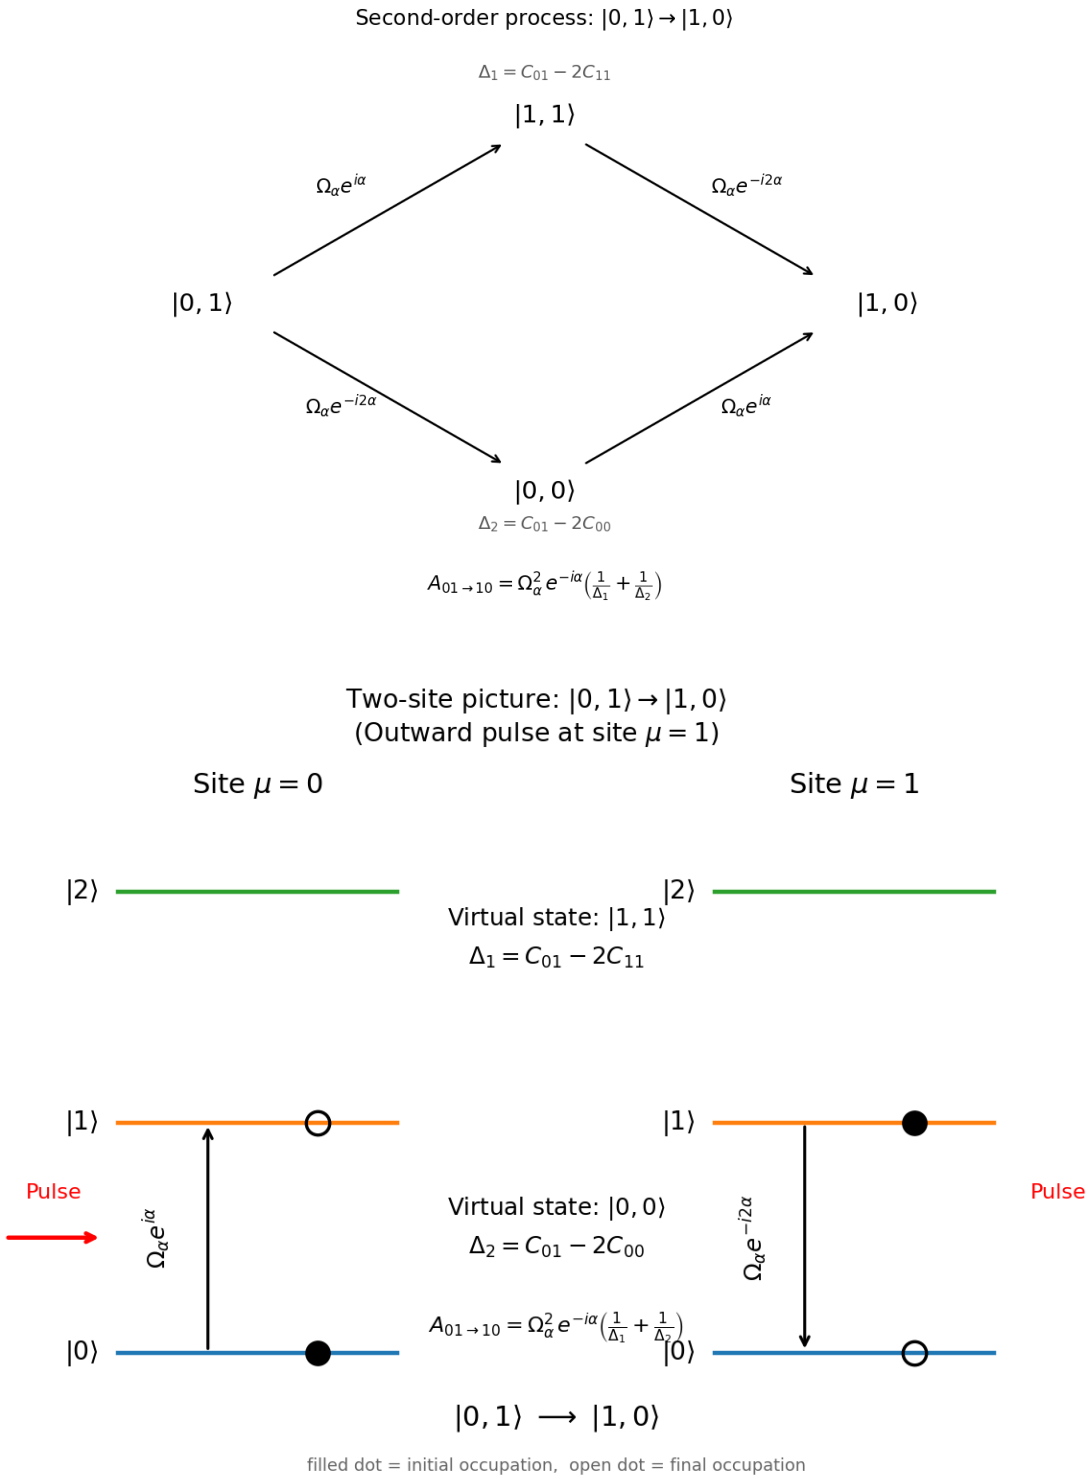


Figure 10: Transition $|0,1\rangle \to |1,0\rangle$: second-order path diagram (top) and two-site level picture (bottom). Here and in Figs. 12–17 the two atoms of the bond are labelled $\mu = 0, 1$ in the figures, corresponding to the sites j and $j+1$ of the text.

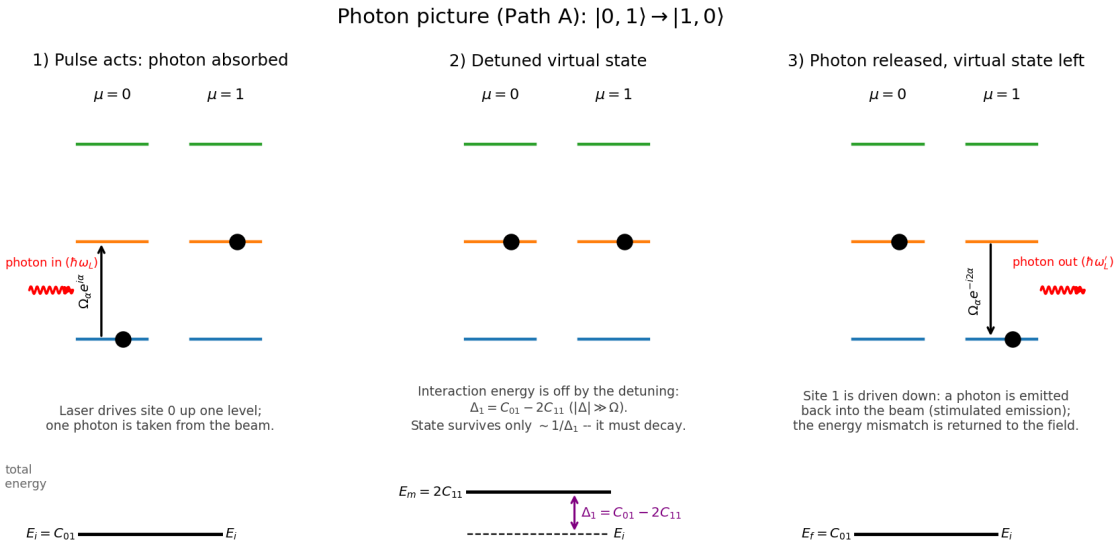


Figure 11: Photon picture of $|0,1\rangle \to |1,0\rangle$ (Path A): a photon is absorbed on site 0, the system dwells in a virtual state detuned by $\Delta \gg \Omega$, and a photon is re-emitted on site 1; Path B reverses the order.

D.2 Transition $|0,2\rangle \to |1,1\rangle$

Here $E_i = C_{02}$, $E_f = 2C_{11}$ (non-degenerate; use (16)), with intermediates $|1,2\rangle$ ($E_m = C_{12}$) and $|0,1\rangle$ ($E_m = C_{01}$). Both paths have numerator $\Omega_\alpha \Omega_\beta\, e^{i(\alpha-2\beta)}$, e.g. Path A: $\langle 1,2|\, H_\Omega\, |0,2\rangle = \Omega_\alpha e^{i\alpha}$ then $\langle 1,1|\, H_\Omega\, |1,2\rangle = \Omega_\beta e^{-i2\beta}$. Thus

$$\begin{aligned} A_{02\to 11} &= \tfrac{1}{2}\Omega_\alpha\Omega_\beta\, e^{i(\alpha-2\beta)}\Big(\tfrac{1}{C_{02}-C_{12}} + \tfrac{1}{2C_{11}-C_{12}} + \tfrac{1}{C_{02}-C_{01}} + \tfrac{1}{2C_{11}-C_{01}}\Big) \\ &\xrightarrow{2C_{11}=C_{02}} A^{R}_{02\to 11} = \Omega_\alpha\Omega_\beta\, e^{i(\alpha-2\beta)}\, K_2, \end{aligned} \tag{102}$$

where the paired denominators coincide at resonance and the factor $\frac{1}{2}$ cancels.

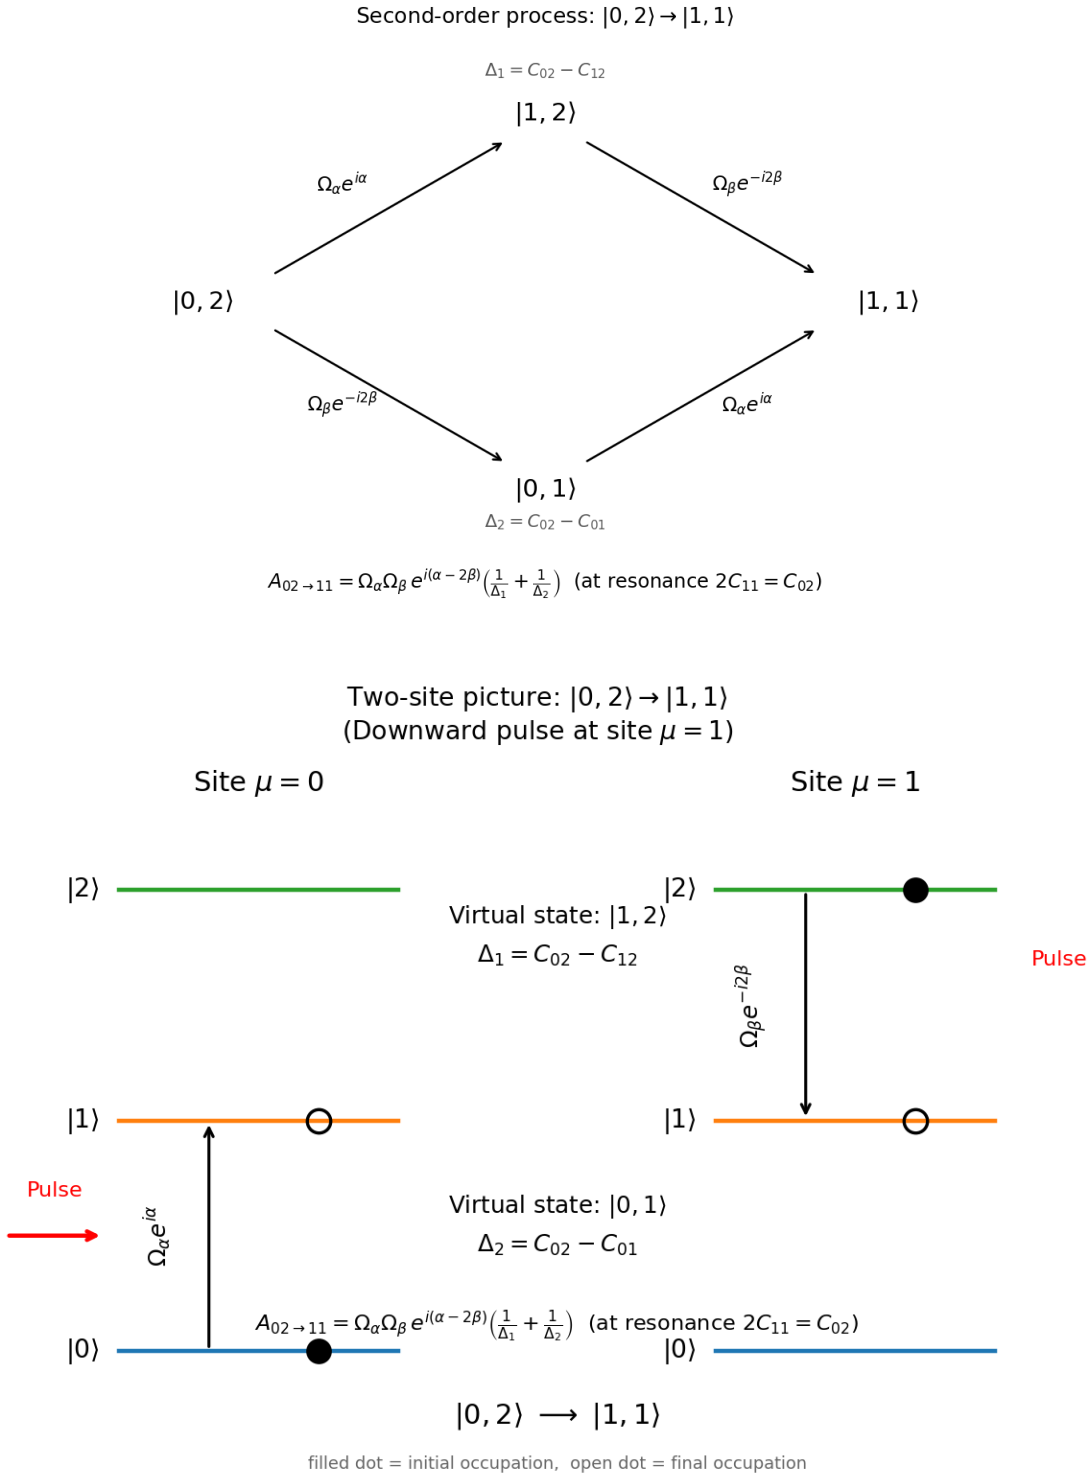


Figure 12: Transition $|0,2\rangle \to |1,1\rangle$: second-order path diagram (top) and two-site level picture (bottom).

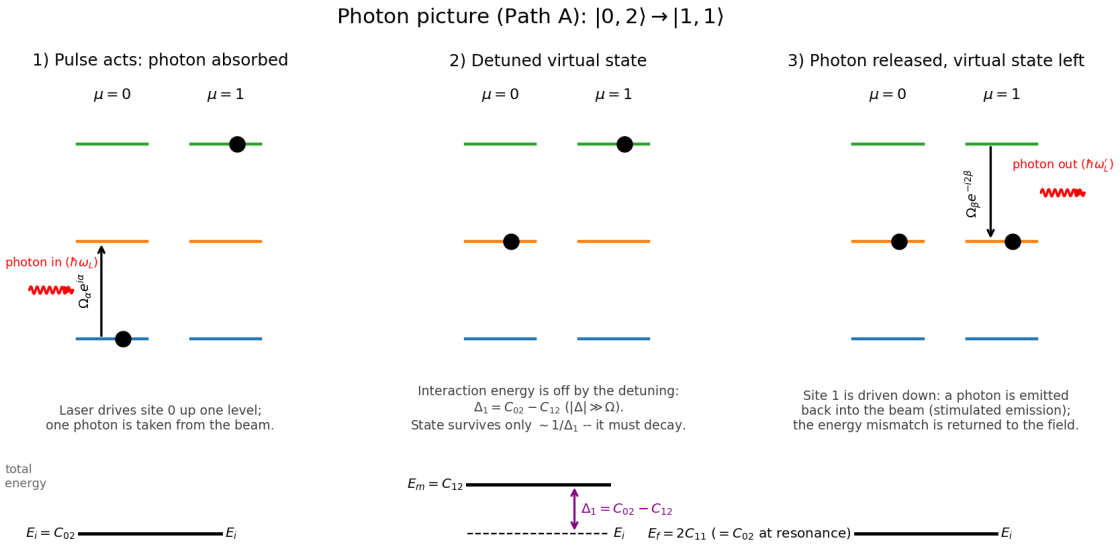


Figure 13: Photon picture of $|0,2\rangle \to |1,1\rangle$ (Path A via the detuned virtual state; Path B reverses the order).

D.3 Transition $|1,1\rangle \to |2,0\rangle$

Here $E_i = 2C_{11}$, $E_f = C_{02}$ (non-degenerate), with intermediates $|2,1\rangle$ ($E_m = C_{12}$) and $|1,0\rangle$ ($E_m = C_{01}$). Both paths have numerator $\Omega_\alpha \Omega_\beta\, e^{i(\beta-2\alpha)}$, e.g. Path A: $\langle 2,1|\, H_\Omega\, |1,1\rangle = \Omega_\beta e^{i\beta}$ then $\langle 2,0|\, H_\Omega\, |2,1\rangle = \Omega_\alpha e^{-i2\alpha}$. Thus

$$\begin{aligned} A_{11\to 20} &= \tfrac{1}{2}\Omega_\alpha \Omega_\beta\, e^{i(\beta-2\alpha)} \Big(\tfrac{1}{2C_{11}-C_{12}} + \tfrac{1}{C_{02}-C_{12}} + \tfrac{1}{2C_{11}-C_{01}} + \tfrac{1}{C_{02}-C_{01}} \Big) \\ &\xrightarrow{2C_{11}=C_{02}} \; A^R_{11\to 20} = \Omega_\alpha \Omega_\beta\, e^{i(\beta-2\alpha)}\, K_2. \end{aligned} \tag{103}$$

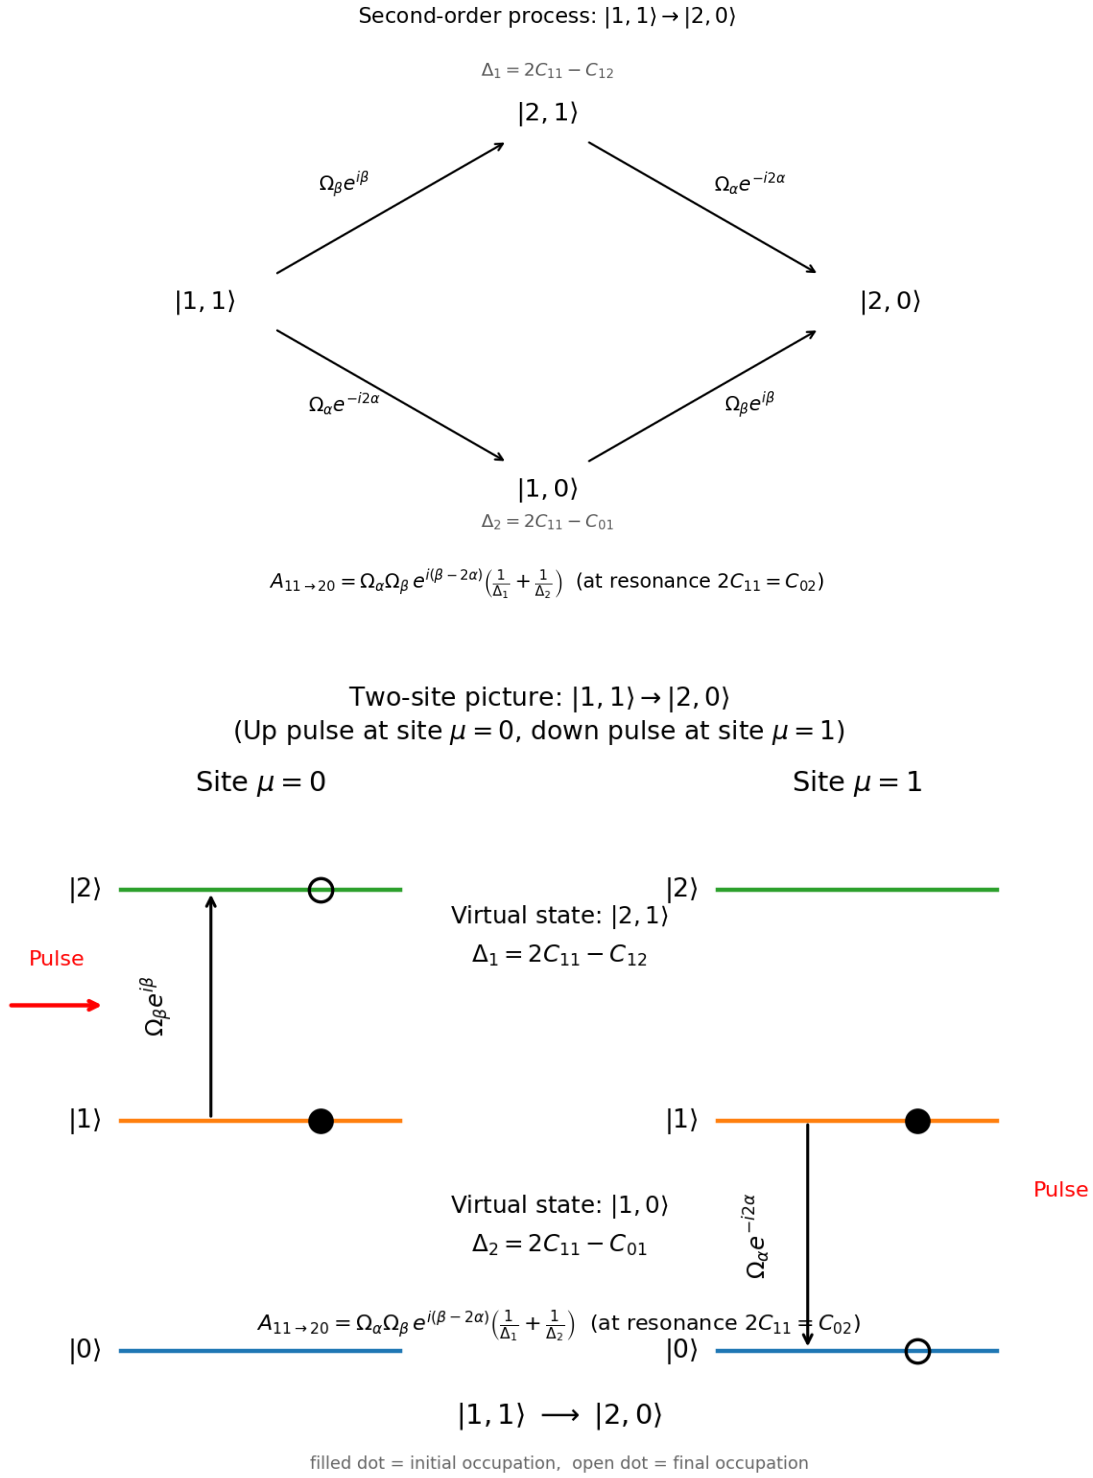


Figure 14: Transition $|1,1\rangle \to |2,0\rangle$: second-order path diagram (top) and two-site level picture (bottom).

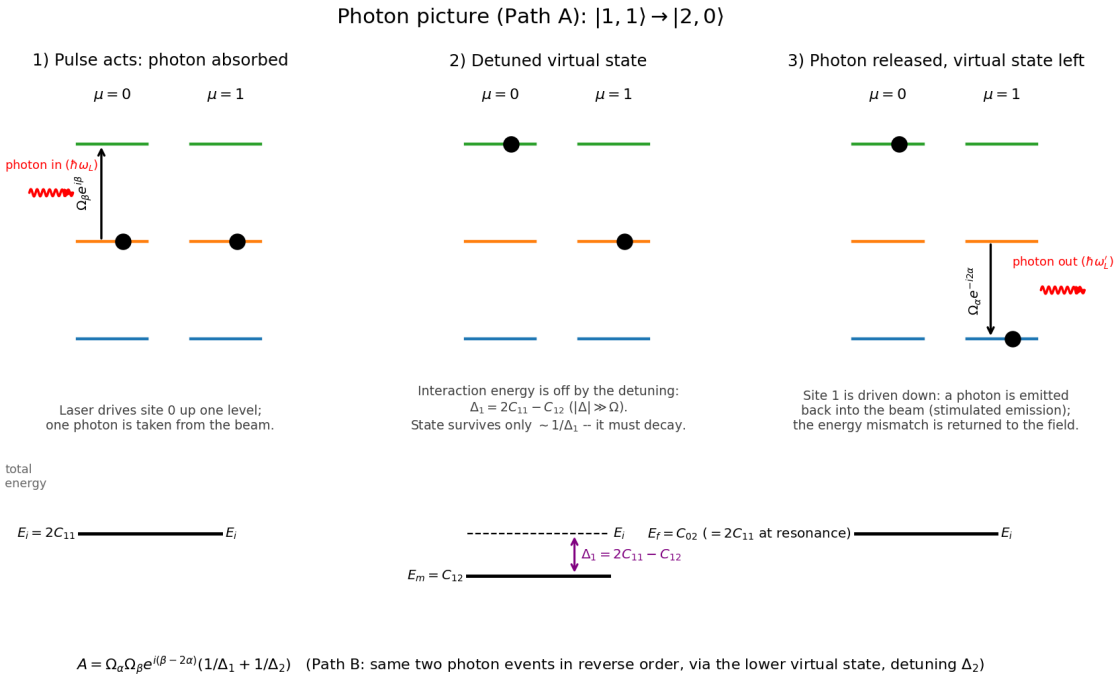


Figure 15: Photon picture of $|1,1\rangle \to |2,0\rangle$ (Path A via the detuned virtual state; Path B reverses the order).

D.4 Transition $|1,2\rangle \to |2,1\rangle$

Here $E_i = E_f = C_{12}$ (degenerate; use (15)), with intermediates $|2,2\rangle$ ($E_m = 2C_{22}$) and $|1,1\rangle$ ($E_m = 2C_{11}$). Both paths carry the phase $e^{i\beta}e^{-i2\beta} = e^{-i\beta}$ (e.g. Path A: $\langle 2,2|\, H_\Omega\, |1,2\rangle = \Omega_\beta e^{i\beta}$ then $\langle 2,1|\, H_\Omega\, |2,2\rangle = \Omega_\beta e^{-i2\beta}$), so

$$A_{12\to 21} = \Omega_\beta^2\, e^{-i\beta}\left(\frac{1}{C_{12}-2C_{22}} + \frac{1}{C_{12}-2C_{11}}\right) \xrightarrow{2C_{11}=C_{02}} A^R_{12\to 21} = \Omega_\beta^2\, e^{-i\beta}\, K_3. \quad (104)$$

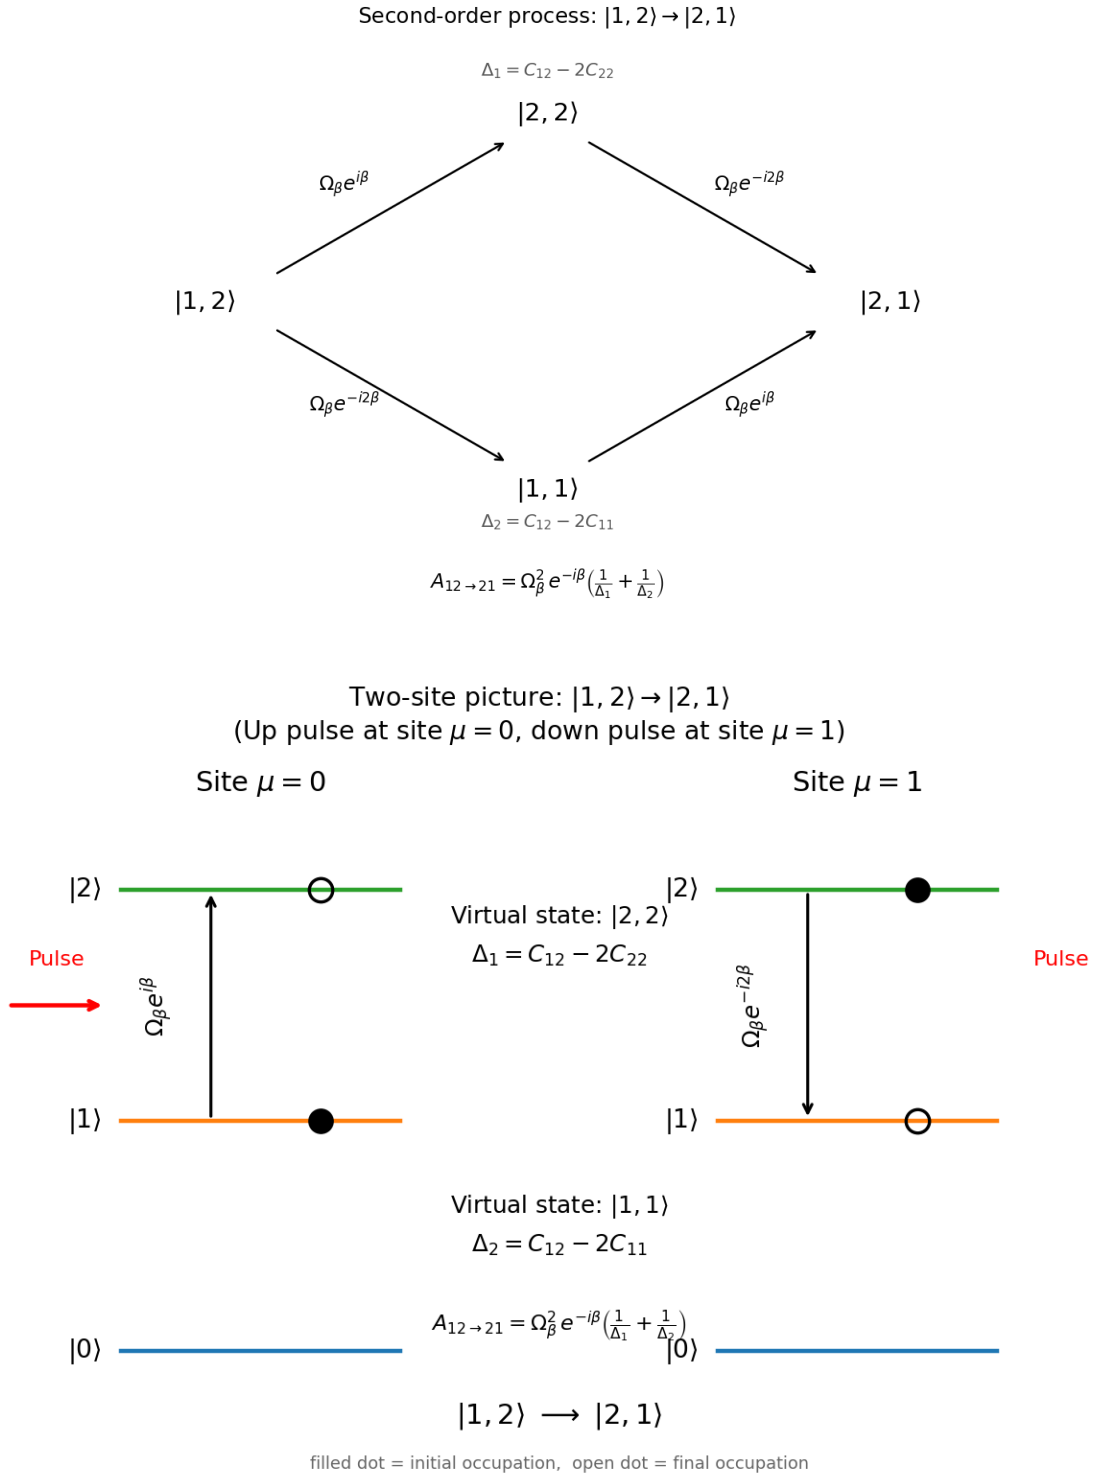


Figure 16: Transition $|1,2\rangle \to |2,1\rangle$: second-order path diagram (top) and two-site level picture (bottom).

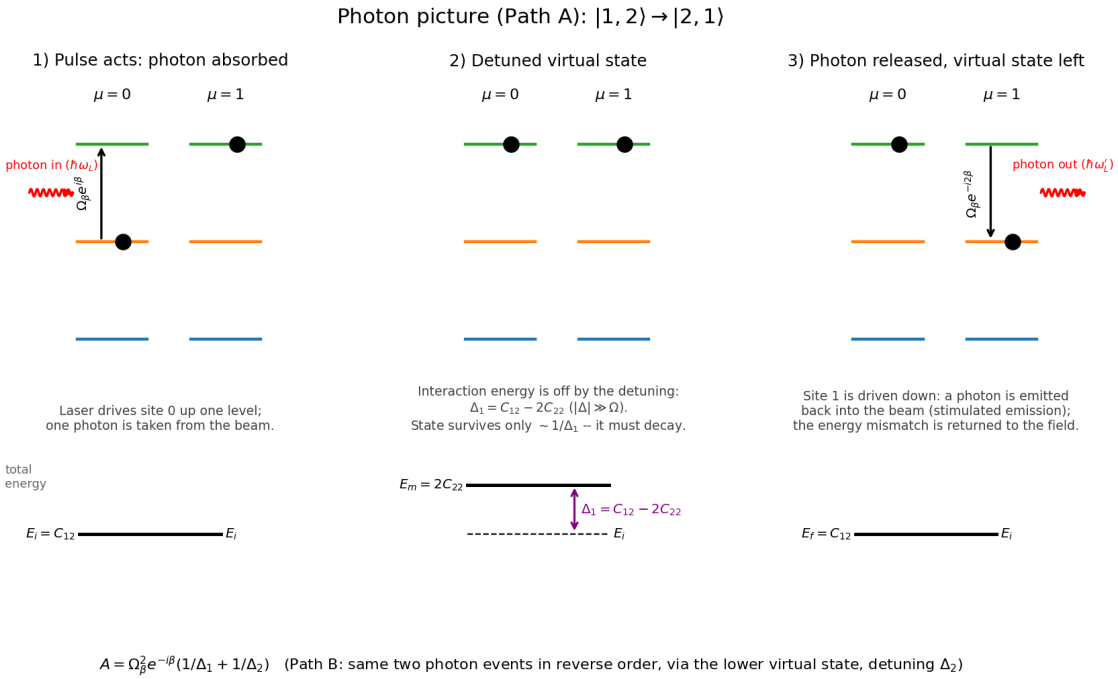


Figure 17: Photon picture of $|1,2\rangle \to |2,1\rangle$ (Path A via the detuned virtual state; Path B reverses the order).

E Floquet averaging and Bessel functions

This appendix derives the Floquet factor (31) quoted in Sec. 3.3, starting from the general pulsed drive (26) of finite duty cycle.

Integrating δ_{2j} over one evolution gives the accumulated phase

$$\Phi_j(t) = \nu_j \hbar \omega_j t + \chi_j \left(N + \Theta(s - T_j')\right) + \kappa\, \Theta(T_j' - s) \sin(\omega' s), \tag{105}$$

where we introduced the modulation index and the per-pulse phase increment

$$\kappa \equiv \frac{A}{\omega'}, \qquad \chi_j \equiv \kappa \sin\!\left(\omega' T_j'\right), \qquad N = \left\lfloor \frac{t}{T_0} \right\rfloor, \qquad s = t - N T_0, \tag{106}$$

so that N counts completed drive periods and s is the time elapsed within the current one. Exponentiating (105),

$$e^{-i\Phi_j(t)} = e^{-i\nu_j \hbar \omega_j t}\; e^{-iN\chi_j}\; e^{-i\chi_j \Theta(s - T_j')}\; e^{-i\kappa\, \Theta(T_j' - s) \sin(\omega' s)}, \tag{107}$$

and applying the Jacobi–Anger identity

$$e^{-iz\sin\theta} = \sum_{\ell=-\infty}^{\infty} J_\ell(z)\, e^{-i\ell\theta} \tag{108}$$

to the oscillatory factor yields the piecewise continuous representation

$$e^{-i\Phi_j(t)} = e^{-i\nu_j \hbar \omega_j t}\, e^{-iN\chi_j} \begin{cases} \displaystyle\sum_{\ell=-\infty}^{\infty} J_\ell(\kappa)\, e^{-i\ell\omega' s}, & 0 \le s \le T_j', \\ e^{-i\chi_j}, & T_j' < s < T_0. \end{cases} \tag{109}$$

Two successive simplifications reduce this to the form used in the main text. In the *stroboscopic case* $\omega' T_j' = 2\pi m_j$ with $m_j \in \mathbb{Z}$ one has $\chi_j = 0$, so the discrete phase staircase accumulated over completed periods drops out entirely and

$$e^{-i\Phi_j(t)} = e^{-i\nu_j \hbar\omega_j t} \left[\Theta(T_j' - s) \sum_{\ell=-\infty}^{\infty} J_\ell(\kappa)\, e^{-i\ell\omega' s} + \Theta(s - T_j') \right]. \tag{110}$$

If in addition $T_j' = T_0$, so that there is no gap between successive pulses, then $s = t \pmod{T_0}$ and hence $e^{-i\ell\omega' s} = e^{-i\ell\omega' t}$. The piecewise structure disappears and the phase factor collapses to the pure Jacobi–Anger form

$$e^{-i\Phi_j(t)} = e^{-i\tilde{\nu}_j t} \sum_{\ell=-\infty}^{\infty} J_\ell(\kappa)\, e^{-i\ell\omega' t}, \qquad \tilde{\nu}_j \equiv \nu_j \hbar\omega_j, \tag{111}$$

whose conjugate is the Floquet factor $\mathcal{F}_j(t) = e^{i\Phi_j(t)}$ of (31).