

Symmetry Adapted Hierarchical Equations of Motion for Exact Simulations of Large Polariton Systems

M. Elious Mondal*

Department of Chemistry, University of Rochester, Rochester, New York 14627, USA

Pengfei Huo†

*Department of Chemistry, University of Rochester, Rochester, New York 14627, USA and
The Institute of Optics, Hajim School of Engineering,
University of Rochester, Rochester, New York 14627, USA*

Hierarchical equations of motion (HEOM) provide exact dynamics of open quantum systems coupled to harmonic baths, but their computational cost becomes prohibitive for systems with many independent local environments. In this work, we develop a symmetry-adapted HEOM formalism to significantly reduce the computational cost for the permutationally invariant Holstein-Tavis-Cummings (HTC) model. The method removes redundant information in two stages. First, all auxiliary density operators (ADOs) related only by relabeling identical molecules and their bath channels are replaced by a single canonical occupation-pattern representative. Second, molecules with the same local hierarchy occupation produce repeated matrix elements within each representative, allowing only the distinct complex variables to be propagated instead of the full $(N + 1) \times (N + 1)$ ADO matrices. The resulting matrix-free equations are evaluated using precomputed connections and molecular multiplicities. At fixed hierarchy depth L and number of bath correlation exponentials m , the number of canonical representatives becomes independent of the ensemble size for $N \geq L$ and the number of unique variables saturates for $N \geq L + 2$. The formulation easily extends to multiple-exponential bath decompositions, arbitrary initial density operators, static disorders, and cavity loss. Our benchmarks reproduce conventional HEOM dynamics while requiring far fewer propagated variables and substantially less memory.

I. INTRODUCTION

Collective strong light-matter coupling hybridizes molecular excitations with a confined photon mode to form polaritonic states distributed across a molecular ensemble. [1, 2] The Holstein-Tavis-Cummings (HTC) model provides a widely used microscopic description of this setting: the cavity couples collectively to the molecular transitions, whereas vibrational and environmental fluctuations generally remain local to individual molecules. These local environments drive relaxation, dephasing, and population exchange between bright, dark, and photonic sectors, and their memory can be important on the timescales of polariton dynamics [3–5]. Accurately treating such non-Markovian dynamics therefore requires a method that retains molecularly resolved environmental fluctuations while remaining computationally practical as the number of molecules increases.

Hierarchical equations of motion (HEOM) [6–9] provide a nonperturbative and systematically convergent description of harmonic environments once their correlation functions are represented by finite sums of exponential terms [7, 10, 11]. Their direct application to a molecular ensemble, however, encounters a severe dimensional bottleneck. For N molecules, m exponential terms per local bath, and total hierarchy depth L , conventional HEOM

contains

$$\mathcal{N}_{\text{ADO}}(N, m, L) = \binom{Nm + L}{L}$$

Auxiliary density operators (ADO) and a total of

$$\mathcal{N}_{\text{var}}(N, m, L) = \mathcal{N}_{\text{ADO}}(N, m, L)(N + 1)^2$$

complex variables when every ADO is stored as a dense matrix in the single excitation basis. For identical molecules, much of this growth represents redundant information. Conventional HEOM separately retains every molecular assignment of the same local hierarchy occupations and stores many matrix elements that are repeated because the associated molecules are physically interchangeable.

Several complementary developments have improved the efficiency of HEOM, including scaled ADOs [12, 13], filtering of negligible hierarchy elements [14–16], compact bath correlation decompositions [17, 18], parallel propagation [19–22], and low-rank or tensor representations [10–13, 17, 23–31]. Permutationally invariant formulations have also reduced the computational cost of quantum trajectories [32], while collective structure and Hamiltonian sparsity have been exploited in polariton simulations [33–37]. Mean-field dynamical approaches have been employed to exactly evaluate observables such as photon populations and linear spectra in the thermodynamic limit ($N \rightarrow \infty$) [38–48]. However, these methods are not applicable for determining polariton or localized populations, nor can they be reliably used for finite,

* mmondal@ur.rochester.edu

† pengfei.huo@rochester.edu

small values of N . These approaches do not directly remove the sources of redundancy present in molecule-resolved HEOM. Reducing only the system Hilbert space is insufficient because independent local environments preserve molecule-specific bath histories and mediate dynamics between collective bright states and the molecular dark state sector. An efficient HEOM treatment must therefore compress the hierarchy labels and the internal ADO representation together.

In this work, we develop a two-stage symmetry-adapted HEOM construction for identical HTC ensembles. First, ADOs related by simultaneous relabeling of molecules and their associated local hierarchy occupations are replaced by one canonical occupation-pattern representative. Second, molecules carrying the same local hierarchy occupation are grouped into occupation categories, and only the resulting unique matrix elements are propagated. Canonical raising and lowering connections, molecular multiplicities, grouped Hamiltonian sums, and aligned target accessors are precomputed and assembled into a matrix-free right-hand side, so neither full ADO matrices nor a dense HEOM generator is constructed during propagation. This procedure is an exact algebraic reduction of the same finite HEOM and introduces no approximation beyond the chosen bath correlation decomposition, hierarchy closure, system Hamiltonian, and numerical integrator. At fixed m and L , the number of canonical ADO representatives saturates for $N \geq L$, and the fully compressed unique variable dimension saturates for $N \geq L + 2$. This saturation concerns computational dimension, not physical convergence with ensemble size. The equations can retain an explicit dependence on N through molecular multiplicities and through the single-molecule coupling g , including calculations in which the collective Rabi splitting is held fixed. The formalism can be easily extended to systems beyond the HTC Hamiltonian as long as they have the permutation relabelling symmetry of the different system subcomponents. The method can also be easily adapted to account for static inhomogeneous diagonal and off-diagonal coupling disorders.

We validate the reduced propagation against conventional QuTiP HEOM [11, 49] calculations for the accessible small ensembles and use the existing calculations to examine three classes of initial conditions: a permutation symmetric upper-polariton state, a localized molecular excitation, and a coherent superposition of a cavity photon and a localized molecular excitation. The extension to a multi-exponential bath correlation decomposition and its numerical illustration are presented in the Supporting Information. The present formulation assumes identical molecular transition energies, identical molecule-cavity couplings within each calculation, identical independent local baths, local diagonal system-bath coupling, one cavity mode, two-level molecules, and the single-excitation manifold. The inhomogeneous static disorder and cavity loss are demonstrated and discussed in the supplementary Information. Within these assump-

tions, the method provides a direct route to systematically convergent open system dynamics for ensemble sizes that are inaccessible to dense conventional HEOM. Optical response functions [33, 50, 51], higher excitation manifolds [5, 34], and ensembles of identical multichromophoric units [50] are natural extensions of the same compression strategy.

II. MODEL AND CONVENTIONAL HEOM

A. Identical Holstein-Tavis-Cummings Hamiltonian

We consider N identical two-level molecules coupled to one cavity mode. We set $\hbar = 1$, use the rotating-wave approximation, and restrict the Hilbert space to the single excitation manifold. The ordered basis is

$$\mathcal{B}_N = \{|c\rangle, |e_1\rangle, |e_2\rangle, \dots, |e_N\rangle\}, \quad (2)$$

where $|c\rangle$ contains one cavity photon and electronically unexcited molecules, while $|e_i\rangle$ contains an electronic excitation on the i^{th} molecule and no cavity photon. The Tavis-Cummings (TC) system Hamiltonian is

$$\hat{H}_S = \omega_{\text{cav}}|c\rangle\langle c| + \epsilon_x \sum_{i=1}^N |e_i\rangle\langle e_i| + g \sum_{i=1}^N (|c\rangle\langle e_i| + |e_i\rangle\langle c|). \quad (3)$$

Here, ω_{cav} is the cavity frequency, ϵ_x is the molecular transition energy entering the propagated system Hamiltonian, and g is the single-molecule light-matter coupling. The normalized bright state is

$$|B\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |e_i\rangle. \quad (4)$$

The collective Rabi splitting on resonance is

$$\Omega_R = 2g\sqrt{N}. \quad (5)$$

With $\Delta = \omega_{\text{cav}} - \epsilon_x$ and $\tan(2\theta) = \Omega_R/\Delta$, the TC upper and lower polariton eigenstates are,

$$|+\rangle = \sin\theta|B\rangle + \cos\theta|c\rangle, \quad (6a)$$

$$|-\rangle = \cos\theta|B\rangle - \sin\theta|c\rangle, \quad (6b)$$

with $\theta \in [0, \pi/2]$ being controlled by the cavity detuning. The remaining polaritonic subspace is described by the dark projector

$$\hat{P}_D = \sum_{i=1}^N |e_i\rangle\langle e_i| - |B\rangle\langle B|. \quad (6c)$$

B. Local environments and bath decomposition

Each molecule is coupled diagonally to an independent harmonic environment. The total Hamiltonian is

$$\hat{H} = \hat{H}_S + \hat{H}_B + \hat{H}_{SB}, \quad (7a)$$




FIG. 1. A schematic representation of the symmetry adaptation of HEOM for the HTC system. (a) is the ADO structure for the conventional HEOM hierarchy. (b) is the reduction of ADOs within different layers to representatives due to permutation symmetry. (c) is the further reduction of the ADO representatives to their unique variables.

with the bath ($\hat{H}_B$) and the system-bath ($\hat{H}_{SB}$) interactions defined via

$$\hat{H}_B = \sum_{i=1}^N \sum_{\xi} \omega_{\xi} \hat{b}_{i\xi}^{\dagger} \hat{b}_{i\xi}, \quad (7b)$$

$$\hat{H}_{SB} = \sum_{i=1}^N \hat{Q}_i \otimes \hat{B}_i, \quad \hat{Q}_i = |e_i\rangle\langle e_i|, \quad (7c)$$

$$\hat{B}_i = \sum_{\xi} g_{i\xi} (\hat{b}_{i\xi}^{\dagger} + \hat{b}_{i\xi}). \quad (7d)$$

For the harmonic environments introduced in Eq. 7b, the local spectral density on the i^{th} molecule is

$$J_i(\omega) = \pi \sum_{\xi} |g_{i\xi}|^2 \delta(\omega - \omega_{i\xi}). \quad (8)$$

The local baths are mutually independent and statistically identical, *i.e.*,

$$J_i(\omega) = J(\omega) = \pi \sum_{\xi} |g_{\xi}|^2 \delta(\omega - \omega_{\xi}). \quad (9)$$

The corresponding equilibrium bath correlation functions satisfy

$$\langle \hat{B}_i(t) \hat{B}_j(0) \rangle_B = \delta_{ij} C(t). \quad (10)$$

With the spectral density convention in Eq. 8, the local bath correlation function is

$$C(t) = \frac{1}{\pi} \int_0^{\infty} d\omega J(\omega) \left[\coth\left(\frac{\beta\omega}{2}\right) \cos(\omega t) - i \sin(\omega t) \right], \quad (11)$$

where $\beta = (k_B T)^{-1}$. The correlation function entering the finite HEOM is represented as

$$C_m(t) = \sum_{k=0}^{m-1} c_k e^{-\nu_k t}, \quad (12a)$$

$$C_m^*(t) = \sum_{k=0}^{m-1} \check{c}_k e^{-\nu_k^* t}. \quad (12b)$$

If the channel k' satisfies $\nu_{k'} = \nu_k^*$, then $\check{c}_k = c_{k'}^*$. For a real, self-conjugate rate, $\check{c}_k = c_k^*$. This distinction is required when the correlation function contains conjugate complex decay rates. The coefficients c_k may be complex. Eq. 12 can be obtained from a Matsubara or Padé decomposition or from a controlled numerical fit [4, 17, 18, 51].

C. Conventional scaled HEOM

In the conventional HEOM, each auxiliary density operator (ADO) is labeled by a collection of nonnegative integers,

$$\mathbf{n} = (n_{10}, \dots, n_{1,m-1}; \dots; n_{N0}, \dots, n_{N,m-1}), \quad (13)$$

where $n_{ik} \in \mathbb{N}_0$ is the hierarchy occupation associated with molecule $i \in \{1, \dots, N\}$ and exponential component $k \in \{0, \dots, m-1\}$. Its tier and damping rate are

$$|\mathbf{n}| = \sum_{i=1}^N \sum_{k=0}^{m-1} n_{ik}, \quad (14a)$$

$$\Gamma_{\mathbf{n}} = \sum_{i=1}^N \sum_{k=0}^{m-1} n_{ik} \nu_k. \quad (14b)$$

The physical reduced density operator is the zeroth-tier ADO, $\hat{\rho}_S(t) = \tilde{\hat{\rho}}_0(t)$. We use the scaled convention [10, 12, 13]

$$\tilde{\hat{\rho}}_{\mathbf{n}} = \left[\prod_{i=1}^N \prod_{k=0}^{m-1} n_{ik}! |c_k|^{n_{ik}} \right]^{-1/2} \hat{\rho}_{\mathbf{n}}. \quad (15)$$

Let $\mathbf{e}_{ik}$ denote the hierarchy index that increases n_{ik} by one while leaving all other occupations unchanged. The conventional scaled HEOM is [12],

$$\begin{aligned} \frac{\partial}{\partial t} \tilde{\hat{\rho}}_{\mathbf{n}} = & -i \left[\hat{H}_S, \tilde{\hat{\rho}}_{\mathbf{n}} \right] - \Gamma_{\mathbf{n}} \tilde{\hat{\rho}}_{\mathbf{n}} \\ & - i \sum_{i=1}^N \sum_{k=0}^{m-1} \sqrt{(n_{ik}+1)|c_k|} \left[\hat{Q}_i, \tilde{\hat{\rho}}_{\mathbf{n}+\mathbf{e}_{ik}} \right] \\ & - i \sum_{i=1}^N \sum_{k=0}^{m-1} \sqrt{\frac{n_{ik}}{|c_k|}} \left(c_k \hat{Q}_i \tilde{\hat{\rho}}_{\mathbf{n}-\mathbf{e}_{ik}} - \check{c}_k \tilde{\hat{\rho}}_{\mathbf{n}-\mathbf{e}_{ik}} \hat{Q}_i \right). \end{aligned} \quad (16)$$

We retain labels with $|\mathbf{n}| \leq L$. A hard cutoff omits upward connections from tier L . There are Nm hierarchy directions, so the number of ADOs through depth L is

$$\mathcal{N}_{\text{ADO}}(N, m, L) = \binom{Nm+L}{L}. \quad (17)$$

Since each ADO is a dense $(N+1) \times (N+1)$ matrix in $\mathcal{B}_N$, the number of propagated complex variables in the conventional dense representation is

$$\mathcal{N}_{\text{var}}(N, m, L) = \mathcal{N}_{\text{ADO}}(N, m, L)(N+1)^2. \quad (18)$$

III. MOLECULAR-RELABELING COVARIANCE

Consider a relabeling π that assigns the original molecular label i to the new label $\pi(i)$. The corresponding operator satisfies

$$\hat{U}_{\pi} |c\rangle = |c\rangle, \quad (19a)$$

$$\hat{U}_{\pi} |e_i\rangle = |e_{\pi(i)}\rangle. \quad (19b)$$

Uniform molecular energies and couplings imply

$$\hat{U}_{\pi} \hat{H}_S \hat{U}_{\pi}^{\dagger} = \hat{H}_S, \quad (20a)$$

$$\hat{U}_{\pi} \hat{Q}_i \hat{U}_{\pi}^{\dagger} = \hat{Q}_{\pi(i)}. \quad (20b)$$

Each molecule i is associated with a complete local hierarchy vector

$$\mathbf{n}_i = (n_{i0}, n_{i1}, \dots, n_{i,m-1}), \quad (21)$$

where n_{ik} is the occupation number of the k^{th} exponential term in the bath correlation decomposition for molecule i . A molecular relabeling must therefore move the entire vector $\mathbf{n}_i$, rather than permuting its individual components. If the permutation π sends molecule j to the new label $i = \pi(j)$, then the hierarchy vector assigned to position i after relabeling is the vector that previously belonged to molecule $j = \pi^{-1}(i)$. Accordingly, we define

$$(\pi \mathbf{n})_{ik} = n_{\pi^{-1}(i),k}, \quad (22)$$

for every exponential index k . Thus, the relabeling changes only the molecular index and leaves the ordering of the bath exponential components within each local hierarchy vector unchanged. This rearrangement preserves the tier and damping rate,

$$|\pi \mathbf{n}| = |\mathbf{n}|, \quad \Gamma_{\pi \mathbf{n}} = \Gamma_{\mathbf{n}}, \quad (23)$$

and maps neighboring labels consistently according to $\pi(\mathbf{n} \pm \mathbf{e}_{ik}) = \pi \mathbf{n} \pm \mathbf{e}_{\pi(i)k}$. Every term in Eq. 16 transforms in the same manner. For example, for an arbitrary operator $\hat{A}$,

$$-i \left[\hat{H}_S, \hat{U}_{\pi} \hat{A} \hat{U}_{\pi}^{\dagger} \right] = \hat{U}_{\pi} \left(-i \left[\hat{H}_S, \hat{A} \right] \right) \hat{U}_{\pi}^{\dagger}, \quad (24a)$$

$$\left[\hat{Q}_{\pi(i)}, \hat{U}_{\pi} \hat{A} \hat{U}_{\pi}^{\dagger} \right] = \hat{U}_{\pi} \left[\hat{Q}_i, \hat{A} \right] \hat{U}_{\pi}^{\dagger}, \quad (24b)$$

$$\begin{aligned} c_k \hat{Q}_{\pi(i)} \hat{U}_{\pi} \hat{A} \hat{U}_{\pi}^{\dagger} - \check{c}_k \hat{U}_{\pi} \hat{A} \hat{U}_{\pi}^{\dagger} \hat{Q}_{\pi(i)} \\ = \hat{U}_{\pi} \left(c_k \hat{Q}_i \hat{A} - \check{c}_k \hat{A} \hat{Q}_i \right) \hat{U}_{\pi}^{\dagger}. \end{aligned} \quad (24c)$$

Given any solution $\{\tilde{\hat{\rho}}_{\mathbf{n}}(t)\}$, define the relabelled hierarchy

$$\tilde{\hat{\rho}}'_{\pi \mathbf{n}}(t) = \hat{U}_{\pi} \tilde{\hat{\rho}}_{\mathbf{n}}(t) \hat{U}_{\pi}^{\dagger}. \quad (25)$$

The identities above show that the primed hierarchy satisfies the same HEOM generator as the original hierarchy, but with correspondingly relabelled initial data. This is molecular-relabeling covariance and generally relates two different initial-value problems.

A relation between ADOs within one propagated hierarchy requires the initial hierarchy itself to satisfy

$$\tilde{\hat{\rho}}_{\pi \mathbf{n}}(0) = \hat{U}_{\pi} \tilde{\hat{\rho}}_{\mathbf{n}}(0) \hat{U}_{\pi}^{\dagger}. \quad (26)$$

Uniqueness of the finite linear initial-value problem then gives

$$\tilde{\hat{\rho}}_{\pi \mathbf{n}}(t) = \hat{U}_{\pi} \tilde{\hat{\rho}}_{\mathbf{n}}(t) \hat{U}_{\pi}^{\dagger} \quad (27)$$

at every later time.

For a factorized preparation, all higher-tier ADOs initially ($t = 0$) vanish. Full compression therefore requires the initial system operator to be invariant under every retained molecular relabeling. If d molecules are distinguished by the initial operator, only relabelings that leave those labels fixed are used, while the remaining $N - d$ molecules are compressed. A general noninvariant initial operator may alternatively be reconstructed from relabelled representative propagations using linearity as discussed later in Sec. VI.

Generator covariance requires identical mean molecular energies, molecule-cavity couplings, local coupling operators, bath decompositions, and a relabeling-compatible hierarchy closure. A generic fixed realization of energetic disorder is not invariant, whereas the ensemble average over independent and identically distributed energetic shifts remains permutation covariant, as shown in the Supporting Information. Molecular transition dipoles do not occur in the HEOM generator written here and thus unequal dipoles may instead break preparation or readout symmetry.

IV. FIRST REDUCTION: CANONICAL AUXILIARY DENSITY OPERATORS

A. One exponential per local bath

We begin with the simplest bath decomposition, containing one exponential term per molecule, $m = 1$. In this case, each molecule i carries a single nonnegative hierarchy occupation n_i , and a molecule-resolved ADO is labeled by

$$\mathbf{n} = (n_1, n_2, \dots, n_N). \quad (28)$$

The value of n_i specifies the hierarchy occupation associated with the local bath of molecule i , while the total hierarchy tier is

$$|\mathbf{n}| = \sum_{i=1}^N n_i. \quad (29)$$

Because the molecules and their local baths are identical, simultaneously relabeling the molecules and their hierarchy occupations does not produce an independent dynamical object. Thus, two hierarchy labels that differ only in the ordering of their entries belong to the same relabeling class. For example, the labels $(2, 1, 0, 0)$ and $(0, 2, 1, 0)$ assign the same set of occupations to the molecular ensemble, but attach them to different molecular labels. Their corresponding ADOs are related by the molecular permutation described in Sec. III and therefore need not be propagated independently.

We select one representative from each such class by arranging the nonzero occupations in descending order and suppressing the trailing zeros. The resulting canonical occupation pattern is

$$\Lambda = [\Lambda_1, \Lambda_2, \dots, \Lambda_{r_\Lambda}], \quad \Lambda_1 \geq \Lambda_2 \geq \dots \geq \Lambda_{r_\Lambda} > 0, \quad (30)$$

where r_Λ is the number of molecules with nonzero hierarchy occupation. The tier of the pattern is

$$|\Lambda| = \sum_{a=1}^{r_\Lambda} \Lambda_a. \quad (31)$$

Thus, Λ records how the total hierarchy occupation is distributed among the occupied molecular baths, without retaining the physically redundant assignment of those occupations to particular molecular labels. Mathematically, the patterns Λ are integer partitions of the hierarchy tier, subject to the condition $r_\Lambda \leq N$.

For a system containing N molecules, the canonical pattern is embedded into a molecule resolved hierarchy label by appending the required number of zeros:

$$\mathbf{n}^\Lambda = (\Lambda_1, \dots, \Lambda_{r_\Lambda}, 0, \dots, 0). \quad (32)$$

We retain the ADO associated with this representative and denote it by

$$\hat{R}_\Lambda(t) = \tilde{\hat{\rho}}_{\mathbf{n}^\Lambda}(t). \quad (33)$$

All molecule-resolved ADOs whose hierarchy labels are permutations of $\mathbf{n}^\Lambda$ can be recovered from $\hat{R}_\Lambda(t)$ by applying the corresponding molecular relabeling (Eq. 19b).

For example, the hierarchy labels

$$(2, 1, 0, 0), \quad (0, 2, 1, 0), \quad (1, 0, 0, 2) \longrightarrow \Lambda = [2, 1] \quad (34)$$

all reduce to the same canonical pattern. This pattern states that one molecule carries hierarchy occupation 2, another carries occupation 1, and all remaining molecules carry occupation 0. It does not specify which molecular labels carry the two nonzero occupations. The empty pattern $\emptyset$, for which every molecule has zero hierarchy occupation, represents the physical reduced density operator:

$$\hat{R}_\emptyset(t) = \hat{\rho}_S(t). \quad (35)$$

To describe a canonical pattern more conveniently, we introduce the multiplicity M_q^Λ as the number of molecules carrying hierarchy occupation q . For a nonzero occupation q , this number is

$$M_q^\Lambda = \sum_{a=1}^{r_\Lambda} \delta_{\Lambda_a, q}, \quad q > 0, \quad (36)$$

where $\delta_{\Lambda_a, q}$ is the Kronecker delta. Thus, each occurrence of q in the canonical pattern contributes one to M_q^Λ . Since r_Λ molecules have nonzero occupations, the remaining $N - r_\Lambda$ molecules have zero occupation:

$$M_0^\Lambda = N - r_\Lambda. \quad (37)$$

The multiplicities satisfy two useful relations. First, every molecule belongs to exactly one occupation cate-

gory. Second, the sum of all occupations gives the hierarchy tier. Therefore,

$$\sum_{q=0}^L M_q^\Lambda = N, \quad (38a)$$

$$\sum_{q=1}^L q M_q^\Lambda = |\Lambda|, \quad (38b)$$

where L is the maximum retained hierarchy tier.

Although a canonical pattern does not specify which molecule carries each occupation, it represents every molecule-resolved hierarchy label obtained by assigning these occupations to the N molecular labels. The number of distinct assignments is

$$\mathcal{N}_{\text{ord}}(\Lambda; N) = \frac{N!}{\prod_{q=0}^L M_q^\Lambda!}. \quad (39)$$

The numerator counts all possible orderings of the N occupations. However, exchanging two molecules that carry the same occupation does not produce a new hierarchy label. The factorial $M_q^\Lambda!$ removes these repeated orderings for occupation q . This also applies to the zero occupations, so the factor $M_0^\Lambda!$ must be included in the denominator.

For example, consider $N = 4$ and the canonical pattern $\Lambda = [2, 1]$. One molecule carries occupation 2, one carries occupation 1, and the remaining two molecules carry occupation 0. The corresponding multiplicities are

$$M_2^\Lambda = 1, \quad M_1^\Lambda = 1, \quad M_0^\Lambda = 2. \quad (40)$$

The number of distinct molecule resolved hierarchy labels represented by this pattern is therefore

$$\mathcal{N}_{\text{ord}}([2, 1]; 4) = \frac{4!}{1! 1! 2!} = 12. \quad (41)$$

Thus, instead of independently propagating these twelve relabeling-related ADOs, we retain only the single representative $\hat{R}_{[2,1]}(t)$.

B. Canonical targets and aligned matrices

The equation of motion for an ADO at pattern Λ couples it to ADOs in the neighboring hierarchy tiers. These neighboring ADOs are obtained by increasing or decreasing the hierarchy occupation of one molecule. Let $\mathbf{e}_i$ denote the N -component vector whose i th entry is one and whose remaining entries are zero. Starting from the representative label $\mathbf{n}^\Lambda$, raising the occupation of molecule i gives

$$\mathbf{m}_{\Lambda i}^+ = \mathbf{n}^\Lambda + \mathbf{e}_i, \quad (42)$$

whereas lowering its occupation gives

$$\mathbf{m}_{\Lambda i}^- = \mathbf{n}^\Lambda - \mathbf{e}_i, \quad n_i^\Lambda > 0. \quad (43)$$

We refer to $\mathbf{m}_{\Lambda i}^+$ and $\mathbf{m}_{\Lambda i}^-$ as raw target labels because their entries remain attached to the original molecular labels.

A raw target label is not necessarily in the descending order used to store canonical representatives. We therefore apply the canonicalization operation $\mathcal{C}$, which sorts the positive occupations in descending order and removes the trailing zeros. The canonical target patterns are

$$\Lambda_i^+ = \mathcal{C}(\mathbf{m}_{\Lambda i}^+), \quad (44a)$$

$$\Lambda_i^- = \mathcal{C}(\mathbf{m}_{\Lambda i}^-). \quad (44b)$$

Only the representatives $\hat{R}_{\Lambda_i^+}$ and $\hat{R}_{\Lambda_i^-}$ associated with these canonical patterns are stored and propagated.

Canonicalization may change the assignment of occupations to molecular labels. This distinction matters because the corresponding HEOM term contains the local system operator $\hat{Q}_i$ of the molecule whose occupation was raised or lowered. Before using a stored target ADO in that term, its molecular indices must therefore be returned to the ordering of the raw target label.

For example, consider $N = 4$ and $\Lambda = [2, 2]$, with representative label

$$\mathbf{n}^\Lambda = (2, 2, 0, 0). \quad (45)$$

Lowering the occupation of molecule 1 produces the raw target

$$\mathbf{m}_{\Lambda 1}^- = (1, 2, 0, 0). \quad (46)$$

After sorting, the corresponding canonical target is

$$\Lambda_1^- = [2, 1], \quad \mathbf{n}^{\Lambda_1^-} = (2, 1, 0, 0). \quad (47)$$

The stored target places occupation 2 on molecule 1 and occupation 1 on molecule 2, whereas the raw target requires the opposite assignment. Consequently, molecules 1 and 2 must be exchanged before this stored representative can be used in the lowering term associated with molecule 1.

More generally, let $\pi_{\Lambda i}^\pm$ denote a molecular relabeling that converts the stored canonical ordering of $\Lambda_i^\pm$ into the molecular ordering of the raw target $\mathbf{m}_{\Lambda i}^\pm$. The target ADO expressed in the required raw ordering is then

$$\hat{R}_{\Lambda_i^\pm}^{[\mathbf{m}_{\Lambda i}^\pm]} = \hat{U}_{\pi_{\Lambda i}^\pm} \hat{R}_{\Lambda_i^\pm} \hat{U}_{\pi_{\Lambda i}^\pm}^\dagger. \quad (48)$$

We call this the aligned target matrix. It is obtained by relabeling a stored representative and is not a separate ADO that must be propagated.

Using these aligned targets, the equation of motion can be written entirely in terms of canonical representatives. For a hierarchy truncated at tier L , we define

$$\chi_L(\Lambda) = \begin{cases} 1, & |\Lambda| < L, \\ 0, & |\Lambda| = L. \end{cases} \quad (49)$$

This factor prevents raising operations from generating ADOs above the maximum retained tier. The one-exponential canonical-ADO equation is

$$\begin{aligned} \frac{\partial}{\partial t} \hat{R}_\Lambda = & -i[\hat{H}_S, \hat{R}_\Lambda] - \nu|\Lambda|\hat{R}_\Lambda \\ & -i\chi_L(\Lambda) \sum_{i=1}^N \sqrt{(n_i^\Lambda + 1)|c|} \left[\hat{Q}_i, \hat{R}_{\Lambda_i^+}^{[\mathbf{m}_{\Lambda_i^+}^+]} \right] \\ & -i \sum_{\substack{i=1 \\ n_i^\Lambda > 0}}^N \sqrt{\frac{n_i^\Lambda}{|c|}} \left(c\hat{Q}_i \hat{R}_{\Lambda_i^-}^{[\mathbf{m}_{\Lambda_i^-}^-]} - c^* \hat{R}_{\Lambda_i^-}^{[\mathbf{m}_{\Lambda_i^-}^-]} \hat{Q}_i \right). \end{aligned} \quad (50)$$

The first line contains the system evolution and the hierarchy damping. The second line couples $\hat{R}_\Lambda$ to the next higher tier by raising one molecular occupation. The final two lines couple it to the next lower tier by lowering a nonzero occupation. In every coupling term, the target representative is first aligned with the molecular ordering required by the corresponding local operator $\hat{Q}_i$.

Eq. 50 is obtained directly from the conventional HEOM using the permutation covariance relation in Eq. 27. Canonicalization and alignment only identify and reconstruct relabeling related ADOs. They introduce no additional approximation relative to the conventional HEOM with the same bath decomposition, hierarchy depth, and closure.

We next count the number of ADOs retained after this first symmetry reduction. At tier ℓ , every canonical pattern is a descending list of positive integers whose sum is ℓ and whose length does not exceed N . The number of such patterns is the restricted integer partition number $p_N(\ell)$. Including all tiers from zero through L , the total number of canonical representatives is

$$\mathcal{N}_{\text{ADO}}^{\text{can}}(N, m=1, L) = \sum_{\ell=0}^L p_N(\ell). \quad (51)$$

A pattern at tier ℓ can contain at most ℓ nonzero entries because every nonzero occupation is at least one. Since $\ell \leq L$, no retained pattern can involve more than L molecules with nonzero occupations. Therefore, once $N \geq L$, increasing the number of molecules does not generate any additional canonical patterns. In this regime, the restriction on the pattern length becomes irrelevant, and

$$\mathcal{N}_{\text{ADO}}^{\text{can}}(N, m=1, L) = \sum_{\ell=0}^L p(\ell), \quad N \geq L, \quad (52)$$

where $p(\ell)$ is the unrestricted integer partition number.

Fig. 2 illustrates the first source of computational savings. The molecule resolved hierarchy of conventional HEOM grows combinatorially with N , whereas canonicalization removes hierarchy labels that differ only by relabeling identical molecules. At fixed L , no retained pattern can contain more than L nonzero molecular occupations, so the number of canonical representatives becomes independent of N for $N \geq L$. This saturation




FIG. 2. Number of ADOs retained through hierarchy depth L for a one-exponential decomposition per local bath. The curves labeled “normal” correspond to the molecule-resolved hierarchy used by conventional HEOM, whereas those labeled “sym” show the canonical ADO representatives obtained after molecular relabeling. The conventional hierarchy grows rapidly with the ensemble size N , while the number of canonical representatives saturates once $N \geq L$.

applies only to the number of ADO representatives. If every canonical ADO is stored as a full $(N+1) \times (N+1)$ matrix, its memory and propagation costs continue to increase quadratically with N . The unique-variable reduction introduced below removes this remaining matrix redundancy.

V. SECOND REDUCTION: UNIQUE VARIABLES AND MATRIX-FREE PROPAGATION

The first symmetry reduction replaces all relabeling-related ADOs by a single canonical representative $\hat{R}_\Lambda$. However, each representative is still an $(N+1) \times (N+1)$ matrix in the single excitation basis. Storing the full matrix would therefore leave a computational cost that grows quadratically with the number of molecules.

A second reduction is possible because many matrix elements within a canonical representative are identical. We first identify these repeated entries and count the number of unique variables. We then derive equations that propagate the unique variables directly, without reconstructing the full ADO matrices.

A. Repeated matrix elements within a representative

Consider two molecules i and j that carry the same hierarchy occupation in the representative label $\mathbf{n}^\Lambda$:

$$n_i^\Lambda = n_j^\Lambda. \quad (53)$$

Exchanging these two molecules leaves the representative hierarchy label unchanged. The permutation covariance

relation in Eq. 27 therefore gives

$$\hat{R}_\Lambda = \hat{U}_{ij} \hat{R}_\Lambda \hat{U}_{ij}^\dagger, \quad n_i^\Lambda = n_j^\Lambda. \quad (54)$$

Consequently, simultaneously exchanging the corresponding molecular rows and columns cannot change the ADO matrix.

This means that a matrix element does not depend on the individual labels of the molecules involved. It depends only on the hierarchy occupation categories carried by its row and column. Molecules with the same occupation are therefore indistinguishable within the representative ADO.

This symmetry does not imply that $\hat{R}_\Lambda$ is Hermitian. Although the physical reduced density operator is Hermitian, higher-tier ADOs are generally non-Hermitian. The matrix elements associated with the two opposite directions of a coherence must therefore be retained as independent complex variables.

Let $\mathcal{G}_\Lambda$ denote the list of occupation categories present in $\mathbf{n}^\Lambda$, including the zero-occupation category, and let

$$G_\Lambda = |\mathcal{G}_\Lambda| \quad (55)$$

be the number of such categories. For each $q \in \mathcal{G}_\Lambda$, we define the cavity population variable

$$Z^\Lambda = \langle c | \hat{R}_\Lambda | c \rangle, \quad (56)$$

the two cavity-molecule coherence variables

$$X_q^\Lambda = \langle e_i | \hat{R}_\Lambda | c \rangle, \quad n_i^\Lambda = q, \quad (57a)$$

$$Y_q^\Lambda = \langle c | \hat{R}_\Lambda | e_i \rangle, \quad n_i^\Lambda = q, \quad (57b)$$

and the molecular population variable

$$P_q^\Lambda = \langle e_i | \hat{R}_\Lambda | e_i \rangle, \quad n_i^\Lambda = q. \quad (58)$$

Because every molecule in category q is equivalent within the representative, the definitions do not depend on which particular molecule i is selected from that category.

For two different occupation categories q and r , we define the ordered intermolecular coherence

$$S_{q,r}^\Lambda = \langle e_i | \hat{R}_\Lambda | e_j \rangle, \quad n_i^\Lambda = q, \quad n_j^\Lambda = r, \quad q \neq r. \quad (59)$$

The ordering of the category indices matters because higher-tier ADOs need not be Hermitian. In general,

$$S_{q,r}^\Lambda \neq S_{r,q}^\Lambda. \quad (60)$$

If category q contains at least two molecules, we must also retain the off-diagonal coherence between two different molecules in that same category:

$$T_q^\Lambda = \langle e_i | \hat{R}_\Lambda | e_j \rangle, \quad i \neq j, \quad n_i^\Lambda = n_j^\Lambda = q. \quad (61)$$

The variable T_q^Λ exists only when the multiplicity satisfies

$$M_q^\Lambda \geq 2. \quad (62)$$

As an explicit example, consider $N = 4$ and the canonical pattern $\Lambda = [2, 1]$. Its representative molecule-resolved label is

$$\mathbf{n}^{[2,1]} = (2, 1, 0, 0). \quad (63)$$

The three occupation categories are therefore 2, 1, and 0. In the ordered basis

$$\mathcal{B}_4 = \{|c\rangle, |e_1\rangle, |e_2\rangle, |e_3\rangle, |e_4\rangle\}, \quad (64)$$

the representative has the structure

$$\hat{R}_{[2,1]} \longrightarrow \begin{pmatrix} Z & Y_2 & Y_1 & Y_0 & Y_0 \\ X_2 & P_2 & S_{2,1} & S_{2,0} & S_{2,0} \\ X_1 & S_{1,2} & P_1 & S_{1,0} & S_{1,0} \\ X_0 & S_{0,2} & S_{0,1} & P_0 & T_0 \\ X_0 & S_{0,2} & S_{0,1} & T_0 & P_0 \end{pmatrix}_{\mathcal{B}_4}^{[2,1]}. \quad (65)$$

The superscript $[2, 1]$ is understood to apply to every variable on the right-hand side. The full matrix contains 25 entries, but only 17 of them are distinct. For example, molecules 3 and 4 both belong to the zero-occupation category. Their cavity coherences are therefore both described by X_0 and Y_0 , their populations are both described by P_0 , and their mutual coherences are described by T_0 .

The advantage becomes more pronounced as N increases. For the same pattern $\Lambda = [2, 1]$ with $N = 10$, the full matrix contains

$$(N + 1)^2 = 11^2 = 121 \quad (66)$$

entries. Nevertheless, the occupation categories remain 2, 1, and 0, so the number of unique variables remains 17.

B. Number of unique variables and saturation

We now count the unique variables associated with a general canonical pattern. Let h_Λ be the number of occupation categories containing at least two molecules:

$$h_\Lambda = \sum_{q \in \mathcal{G}_\Lambda} \Theta(M_q^\Lambda - 2), \quad (67)$$

where $\Theta(x)$ equals one for $x \geq 0$ and zero otherwise. Each such category contributes one variable of type T_q^Λ . For a pattern containing G_Λ occupation categories, the unique variables consist of:

1. one cavity population variable Z^Λ ;
2. G_Λ variables of each type X_q^Λ , Y_q^Λ , and P_q^Λ ;
3. $G_\Lambda(G_\Lambda - 1)$ ordered variables $S_{q,r}^\Lambda$ with $q \neq r$;
4. h_Λ same-category coherences T_q^Λ .

The number of distinct complex variables in $\hat{R}_\Lambda$ is therefore

$$\begin{aligned} u_{\text{var}}(\Lambda; N, 1) &= 1 + 3G_\Lambda + G_\Lambda(G_\Lambda - 1) + h_\Lambda \\ &= (G_\Lambda + 1)^2 + h_\Lambda. \end{aligned} \quad (68)$$

Summing over all canonical patterns through hierarchy depth L gives the total number of unique variables

$$\mathcal{N}_{\text{var}}^{\text{unq}}(N, 1, L) = \sum_{|\Lambda| \leq L} [(G_\Lambda + 1)^2 + h_\Lambda]. \quad (69)$$




FIG. 3. Number of propagated complex variables through hierarchy depth L for a one-exponential decomposition per local bath. Curves labeled “sym” correspond to storing every entry of the full $(N + 1) \times (N + 1)$ matrix associated with each canonical ADO; curves labeled “unq” show the dimension after retaining only the unique variables. The full-matrix dimension continues to grow with N , whereas the unique-variable dimension saturates for $N \geq L + 2$.

Fig. 3 demonstrates the additional reduction obtained by propagating unique variables. Although the canonical-ADO count has already saturated for $N \geq L$, retaining full ADO matrices preserves an $(N + 1)^2$ dependence. The unique variable representation instead depends only on the occupation categories present in each canonical pattern and on whether a category contains one or at least two molecules. For a fixed pattern, the dependence on N enters only through the number of molecules in the zero-occupation category,

$$M_0^\Lambda = N - r_\Lambda. \quad (70)$$

Once $M_0^\Lambda \geq 2$, adding another zero-occupation molecule changes only the value of this multiplicity. It does not introduce a new occupation category or a new type of unique variable. Every pattern retained through depth L satisfies

$$r_\Lambda \leq L. \quad (71)$$

Therefore, if $N \geq L + 2$, even a pattern containing the maximum possible number L of nonzero occupations has

at least two molecules in the zero category. All canonical patterns have then reached their final list of unique variable types. Consequently,

$$\mathcal{N}_{\text{var}}^{\text{unq}}(N, 1, L) = \mathcal{N}_{\text{var}}^{\text{unq}}(L + 2, 1, L), \quad N \geq L + 2. \quad (72)$$

The first symmetry reduction makes the number of canonical ADO representatives independent of N for $N \geq L$. The second reduction goes further and it makes the total number of propagated complex variables independent of N for $N \geq L + 2$. The saturated dimension can still grow rapidly with the hierarchy depth L and with the number of exponential terms in the bath correlation decomposition.

C. Direct propagation of the unique variables

The unique variable representation becomes computationally useful only if its variables can be propagated without reconstructing the full ADO matrices. We therefore rewrite the HEOM directly in terms of Z , X , Y , P , S , and T .

There are two types of contributions to the derivative of a unique variable. The system Hamiltonian couples matrix elements within the same canonical representative. The hierarchy terms connect the representative to canonical patterns in the neighboring hierarchy tiers. We consider these two contributions separately.

1. Global vector of propagated variables

For every canonical pattern Λ , we collect its unique variables into a local vector $|R_\Lambda(t)\rangle$. This is a bookkeeping vector rather than a quantum state. Its entries are the unique variables belonging to Λ , placed in a fixed order. The complete propagated vector is obtained by stacking these local vectors according to a fixed ordering of the canonical patterns:

$$|R(t)\rangle = \begin{pmatrix} |R_\emptyset(t)\rangle \\ |R_{[1]}(t)\rangle \\ |R_{[2]}(t)\rangle \\ |R_{[1,1]}(t)\rangle \\ \vdots \end{pmatrix}. \quad (73)$$

Thus, every entry of $|R(t)\rangle$ is identified by its canonical pattern and its variable type within that pattern. Because the HEOM are linear, the evolution can be written formally as

$$\frac{\partial}{\partial t}|R(t)\rangle = \hat{\mathcal{M}}_u|R(t)\rangle. \quad (74)$$

The generator $\hat{\mathcal{M}}_u$ is not constructed as a dense matrix. Instead, its action is evaluated directly from pre-computed connections between the unique variables. The

following derivation explains how these connections are obtained.

2. Hamiltonian terms from grouped molecular sums

The cavity-molecule coupling in $\hat{H}_S$ generates sums over molecular rows and columns. In a full matrix calculation, these sums would contain one term for every molecule. In the unique variable representation, all molecules in the same occupation category contribute identical entries. Their contributions can therefore be combined using the multiplicities M_q^Λ .

Consider any molecule i belonging to category q . The sum over the molecular columns of its row is

$$\mathcal{R}_q^\Lambda = \sum_{j=1}^N \langle e_i | \hat{R}_\Lambda | e_j \rangle, \quad n_i^\Lambda = q. \quad (75)$$

This sum contains the diagonal entry P_q^Λ , the off-diagonal entries connecting molecule i to other molecules in category q , and the entries connecting it to molecules in every different category. Grouping identical terms gives

$$\mathcal{R}_q^\Lambda = P_q^\Lambda + (M_q^\Lambda - 1) T_q^\Lambda + \sum_{\substack{r \in \mathcal{G}_\Lambda \\ r \neq q}} M_r^\Lambda S_{q,r}^\Lambda. \quad (76)$$

Similarly, the sum over the molecular rows of column i is

$$\mathcal{C}_q^\Lambda = \sum_{j=1}^N \langle e_j | \hat{R}_\Lambda | e_i \rangle, \quad n_i^\Lambda = q, \quad (77)$$

which becomes

$$\mathcal{C}_q^\Lambda = P_q^\Lambda + (M_q^\Lambda - 1) T_q^\Lambda + \sum_{\substack{r \in \mathcal{G}_\Lambda \\ r \neq q}} M_r^\Lambda S_{r,q}^\Lambda. \quad (78)$$

If $M_q^\Lambda = 1$, no second molecule exists in category q , and the term containing T_q^Λ is absent. For the example $\Lambda = [2, 1]$ with $N = 4$, the zero-category contains two molecules. The row sum associated with a zero-occupation molecule is

$$\mathcal{R}_0^{[2,1]} = P_0^{[2,1]} + T_0^{[2,1]} + S_{0,2}^{[2,1]} + S_{0,1}^{[2,1]}. \quad (79)$$

The term $T_0^{[2,1]}$ accounts for the other molecule in the zero-category. The two S variables account for the molecules in categories 2 and 1.

For a larger system, many identical terms are absorbed into the multiplicity factors. For example, if category r contains one hundred molecules, its entire contribution to $\mathcal{R}_q^\Lambda$ is evaluated as $100 S_{q,r}^\Lambda$ rather than as a sum over one hundred separately stored entries.

For the one-exponential hierarchy, the damping rate of pattern Λ is

$$\Gamma_\Lambda = \nu |\Lambda|. \quad (80)$$

Using the grouped sums, the Hamiltonian and damping contributions to two representative variables are

$$\left. \frac{\partial Z^\Lambda}{\partial t} \right|_{H+\Gamma} = -ig \sum_{q \in \mathcal{G}_\Lambda} M_q^\Lambda (X_q^\Lambda - Y_q^\Lambda) - \Gamma_\Lambda Z^\Lambda, \quad (81a)$$

$$\left. \frac{\partial X_q^\Lambda}{\partial t} \right|_{H+\Gamma} = -i(\epsilon_x - \omega_{\text{cav}}) X_q^\Lambda - ig(Z^\Lambda - \mathcal{R}_q^\Lambda) - \Gamma_\Lambda X_q^\Lambda. \quad (81b)$$

The first equation combines the identical contributions of the M_q^Λ molecules belonging to category q . The second replaces a sum over all molecular columns by the grouped row sum $\mathcal{R}_q^\Lambda$. The remaining component equations are obtained in the same manner and are given in the Supporting Information.

3. Hierarchy terms and aligned target access

The upward and downward hierarchy terms connect variables belonging to different canonical patterns. Changing the occupation of molecule i produces the raw targets $\mathbf{m}_{\Lambda i}^+$ and $\mathbf{m}_{\Lambda i}^-$ introduced in Eqs. 42 and (43). Canonicalizing these raw labels identifies the stored target patterns Λ_i^+ and Λ_i^- .

The ordering of the molecules in a stored target representative may differ from the ordering required by the raw hierarchy term. Formally, the required target matrix is obtained by aligning the stored representative:

$$\hat{R}_{\Lambda'}^{[\mathbf{m}]} = \hat{U}_\pi \hat{R}_\Lambda \hat{U}_\pi^\dagger, \quad \Lambda' = \mathcal{C}(\mathbf{m}). \quad (82)$$

Explicitly constructing this matrix would restore the $(N+1)^2$ storage that the unique variable reduction is intended to remove. Instead, we access only the individual matrix entry required by a particular component equation.

For a raw target $\mathbf{m}$ with canonical pattern Λ' , define the aligned target accessor

$$\mathcal{A}_{\Lambda'}^{[\mathbf{m}]}(a, b) = \langle a | \hat{R}_{\Lambda'}^{[\mathbf{m}]} | b \rangle, \quad |a\rangle, |b\rangle \in \mathcal{B}_N. \quad (83)$$

The accessor first maps the row and column labels into the canonical ordering of Λ' . It then determines whether the requested entry is represented by a Z , X , Y , P , S , or T variable and reads that variable directly from the propagated vector. It does not construct the full target matrix.

For targets generated from pattern Λ by changing the occupation of molecule i , we use the shorter notation

$$\mathcal{A}_{\Lambda i}^\pm(a, b) \equiv \mathcal{A}_{\Lambda_i^\pm}^{[\mathbf{m}_{\Lambda i}^\pm]}(a, b). \quad (84)$$

To derive the hierarchy contribution to a general unique variable, write

$$V_\alpha^\Lambda = \langle a_\alpha^\Lambda | \hat{R}_\Lambda | b_\alpha^\Lambda \rangle. \quad (85)$$

The index α identifies one of the retained Z , X , Y , P , S , or T variables, while a_α^Λ and b_α^Λ specify its row and column basis states.

The local system-bath coupling operator is $\hat{Q}_i = |e_i\rangle\langle e_i|$. For any matrix $\hat{A}$, its commutator with $\hat{Q}_i$ satisfies

$$\langle a | [\hat{Q}_i, \hat{A}] | b \rangle = \delta_{a,e_i} \langle e_i | \hat{A} | b \rangle - \delta_{b,e_i} \langle a | \hat{A} | e_i \rangle. \quad (86)$$

This identity has a simple interpretation. The first term is present only when the row of the propagated variable corresponds to molecule i . The second term is present only when its column corresponds to molecule i . Using Eq. 86, the upward hierarchy contribution is

$$\begin{aligned} \left. \frac{\partial V_\alpha^\Lambda}{\partial t} \right|_{\text{up}} &= -i\chi_L(\Lambda) \sum_{i=1}^N \sqrt{(n_i^\Lambda + 1) |c|} \\ &\times \left[\delta_{a_\alpha^\Lambda, e_i} \mathcal{A}_{\Lambda i}^+(e_i, b_\alpha^\Lambda) - \delta_{b_\alpha^\Lambda, e_i} \mathcal{A}_{\Lambda i}^+(a_\alpha^\Lambda, e_i) \right]. \end{aligned} \quad (87)$$

The factor $\chi_L(\Lambda)$ prevents an upward connection when $|\Lambda| = L$. For the downward contribution, we define

$$w_i^\Lambda = \sqrt{\frac{n_i^\Lambda}{|c|}}. \quad (88)$$

Only molecules with nonzero hierarchy occupation can be lowered. The downward contribution is

$$\begin{aligned} \left. \frac{\partial V_\alpha^\Lambda}{\partial t} \right|_{\text{down}} &= -i \sum_{\substack{i=1 \\ n_i^\Lambda > 0}}^N w_i^\Lambda c \delta_{a_\alpha^\Lambda, e_i} \mathcal{A}_{\Lambda i}^-(e_i, b_\alpha^\Lambda) \\ &+ i \sum_{\substack{i=1 \\ n_i^\Lambda > 0}}^N w_i^\Lambda c^* \delta_{b_\alpha^\Lambda, e_i} \mathcal{A}_{\Lambda i}^-(a_\alpha^\Lambda, e_i). \end{aligned} \quad (89)$$

The apparently large molecular sums in Eqs. 87 and (89) contain very few nonzero terms. Their action can be understood by considering the type of the propagated variable:

- For $Z^\Lambda = \langle c | \hat{R}_\Lambda | c \rangle$, neither the row nor the column is a molecular excitation. Both Kronecker deltas vanish, so there is no direct hierarchy contribution.
- For $X_q^\Lambda = \langle e_i | \hat{R}_\Lambda | c \rangle$, only the row delta is nonzero. The hierarchy term therefore requires one aligned target entry.
- For $Y_q^\Lambda = \langle c | \hat{R}_\Lambda | e_i \rangle$, only the column delta is nonzero. Again, only one aligned target entry is required.
- For P_q^Λ , $S_{q,r}^\Lambda$, or T_q^Λ , both basis labels are molecular excitations. At most one row contribution and one column contribution can survive.

Thus, every unique variable requires at most two aligned target accesses in each hierarchy direction. The sums over N molecules are useful for writing the equations compactly, but they are not evaluated as molecule-sized loops during time propagation.

4. Precomputed matrix-free right-hand side

Before propagation begins, every nonzero contribution to the derivative of every unique variable is identified. For each contribution, the precomputation records the source variable, the target variable, the alignment required to access it, its multiplicity, and its numerical coefficient.

The right-hand side evaluation then requires only three operations: read the appropriate entries of the global vector, multiply them by their precomputed coefficients, and add the results to the derivatives. The full ADO matrices, molecular permutation matrices, and dense generator $\hat{M}_u$ are never constructed.

After the saturation condition $N \geq L + 2$ is reached, increasing N changes only multiplicity factors such as M_q^Λ . It does not enlarge the canonical-pattern list, the unique variable vector, or the number of runtime connections. The storage requirement and the right-hand-side (RHS) operation count therefore become independent of the molecular ensemble size.

D. Physical observables

The physical reduced density operator is the zeroth-tier representative,

$$\hat{\rho}_S(t) = \hat{R}_\emptyset(t). \quad (90)$$

For $N \geq 2$, the empty pattern contains only the zero-occupation category, with

$$M_0^\emptyset = N. \quad (91)$$

The complete physical density operator is therefore specified by only five unique variables:

$$Z^\emptyset, \quad X_0^\emptyset, \quad Y_0^\emptyset, \quad P_0^\emptyset, \quad T_0^\emptyset. \quad (92)$$

The trace, cavity population, and total molecular-excitation population are

$$\text{Tr}[\hat{\rho}_S] = Z^\emptyset + N P_0^\emptyset, \quad (93a)$$

$$P_{\text{cav}} = \langle c | \hat{\rho}_S | c \rangle = Z^\emptyset, \quad (93b)$$

$$P_{\text{exc}} = \sum_{i=1}^N \langle e_i | \hat{\rho}_S | e_i \rangle = N P_0^\emptyset. \quad (93c)$$

The bright state population is

$$P_B = \langle B | \hat{\rho}_S | B \rangle = P_0^\emptyset + (N - 1) T_0^\emptyset. \quad (93d)$$

The total population in the dark molecular subspace is

$$P_D = \text{Tr}[\hat{\rho}_S \hat{P}_D] = (N - 1) (P_0^\emptyset - T_0^\emptyset). \quad (93e)$$

Using the mixing convention in Eq. 6, the upper- and lower-polariton populations are

$$\begin{aligned} P_+ &= \langle +|\hat{\rho}_S|+ \rangle \\ &= Z^\varnothing \cos^2 \theta + P_B \sin^2 \theta \\ &\quad + \sqrt{N} \sin \theta \cos \theta (X_0^\varnothing + Y_0^\varnothing), \end{aligned} \quad (93f)$$

$$\begin{aligned} P_- &= \langle -|\hat{\rho}_S|- \rangle \\ &= Z^\varnothing \sin^2 \theta + P_B \cos^2 \theta \\ &\quad - \sqrt{N} \sin \theta \cos \theta (X_0^\varnothing + Y_0^\varnothing). \end{aligned} \quad (93g)$$

All of these observables are evaluated directly from the five unique variables of the physical representative. The dense $(N+1) \times (N+1)$ reduced density matrix is never reconstructed.

When the finite bath correlation decomposition contains $m > 1$ exponential terms, each molecule carries a local hierarchy vector $\alpha_i = (n_{i0}, \dots, n_{i,m-1})$ rather than a scalar occupation. A molecular relabeling moves this complete vector without changing the order of its exponential channels, and canonical representatives are formed from repeated local hierarchy vectors. The same unique-variable construction, grouped Hamiltonian sums, aligned target accessors, and matrix-free propagation therefore apply. At fixed m and L , the canonical-pattern count still saturates for $N \geq L$, and the unique-variable dimension saturates for $N \geq L+2$. The number of retained patterns can nevertheless grow rapidly with both m and L . The complete derivation, worked example, and numerical illustration for a multi-exponential decomposition are provided in the Supporting Information.

VI. INITIAL CONDITIONS AND DISTINGUISHED MOLECULES

The symmetry reduction used during propagation must also be consistent with the initial hierarchy. We first consider factorized system-bath initial conditions, then discuss localized initial states that distinguish particular molecules, and finally show how arbitrary initial density operators can be treated using linearity.

A. Permutation-symmetric initial states

For a factorized initial system-bath state, the physical zeroth-tier ADO contains the chosen system density operator, while all higher-tier ADOs initially vanish:

$$\hat{R}_\varnothing(0) = \hat{\rho}_S(0), \quad (94a)$$

$$\hat{R}_\Lambda(0) = 0, \quad |\Lambda| > 0. \quad (94b)$$

The vanishing higher-tier ADOs automatically satisfy every molecular relabeling symmetry. The only remaining

question is whether the initial physical density operator is unchanged by molecular relabeling:

$$\hat{U}_\pi \hat{\rho}_S(0) \hat{U}_\pi^\dagger = \hat{\rho}_S(0). \quad (95)$$

When this condition holds, the initial state can be represented directly using the fully compressed variables of the empty pattern.

The cavity population, bright state population, and polariton populations all satisfy Eq. 95. For example, the initial cavity population

$$\hat{\rho}_S(0) = |c\rangle\langle c| \quad (96)$$

is represented by

$$Z^\varnothing(0) = 1, \quad X_0^\varnothing(0) = Y_0^\varnothing(0) = P_0^\varnothing(0) = T_0^\varnothing(0) = 0. \quad (97)$$

For an initial bright state population,

$$\hat{\rho}_S(0) = |B\rangle\langle B|, \quad (98)$$

every molecular matrix element is equal. The corresponding compressed variables are

$$P_0^\varnothing(0) = T_0^\varnothing(0) = \frac{1}{N}, \quad Z^\varnothing(0) = X_0^\varnothing(0) = Y_0^\varnothing(0) = 0. \quad (99)$$

Using the polariton convention in Eq. 6, an initial upper-polariton population is represented by

$$Z^\varnothing(0) = \cos^2 \theta, \quad (100a)$$

$$X_0^\varnothing(0) = Y_0^\varnothing(0) = \frac{\sin \theta \cos \theta}{\sqrt{N}}, \quad (100b)$$

$$P_0^\varnothing(0) = T_0^\varnothing(0) = \frac{\sin^2 \theta}{N}. \quad (100c)$$

Similarly, an initial lower-polariton population is represented by

$$Z^\varnothing(0) = \sin^2 \theta, \quad (101a)$$

$$X_0^\varnothing(0) = Y_0^\varnothing(0) = -\frac{\sin \theta \cos \theta}{\sqrt{N}}, \quad (101b)$$

$$P_0^\varnothing(0) = T_0^\varnothing(0) = \frac{\cos^2 \theta}{N}. \quad (101c)$$

Thus, cavity, bright, and polariton initial populations require only the five unique variables of the empty pattern.

B. Localized initial states and distinguished molecules

A localized molecular population,

$$\hat{\rho}_S(0) = |e_1\rangle\langle e_1|, \quad (102)$$

is not unchanged by a permutation that exchanges molecule 1 with another molecule. Molecule 1 must

therefore remain explicitly identifiable during the propagation.

The state is nevertheless unchanged by any relabeling that leaves molecule 1 fixed and permutes only molecules $2, \dots, N$:

$$\hat{U}_\pi |e_1\rangle \langle e_1| \hat{U}_\pi^\dagger = |e_1\rangle \langle e_1|, \quad \pi(1) = 1. \quad (103)$$

Consequently, only molecule 1 must be treated as distinguished. Its system row, system column, and local hierarchy label remain explicit. The hierarchy labels of molecules $2, \dots, N$ can still be sorted and compressed exactly as before.

More generally, suppose an initial operator distinguishes d molecular labels. These d labels are kept explicit, while the remaining $N - d$ molecules are treated as equivalent and their local hierarchy labels are sorted canonically. The effective size of the compressible molecular ensemble is therefore

$$N_{\text{sym}} = N - d. \quad (104)$$

The canonical-pattern count saturates when the compressible part contains at least L molecules:

$$N - d \geq L. \quad (105)$$

The unique variable dimension saturates when it contains at least $L + 2$ molecules:

$$N - d \geq L + 2. \quad (106)$$

Equivalently, the saturation conditions are

$$N \geq L + d \quad \rightarrow \quad \text{canonical pattern}, \quad (107a)$$

$$N \geq L + d + 2 \quad \rightarrow \quad \text{unique variables}. \quad (107b)$$

For a single localized excitation, $d = 1$, and these conditions become

$$N \geq L + 1 \quad \rightarrow \quad \text{canonical pattern}, \quad (108a)$$

$$N \geq L + 3 \quad \rightarrow \quad \text{unique variables}. \quad (108b)$$

A localized initial state therefore enlarges the compressed representation relative to a fully permutation-symmetric initial state. However, it does not restore the combinatorial growth of the molecule-resolved hierarchy because the remaining $N - 1$ molecules are still treated collectively.

C. General initial density operators

The HEOM evolution is linear in the initial density operator. We denote the reduced dynamical map by $\hat{\mathcal{G}}_t$, such that

$$\hat{\rho}_S(t) = \hat{\mathcal{G}}_t [\hat{\rho}_S(0)]. \quad (109)$$

For two initial operators $\hat{A}$ and $\hat{B}$ and complex coefficients a and b , linearity gives

$$\hat{\mathcal{G}}_t [a\hat{A} + b\hat{B}] = a\hat{\mathcal{G}}_t[\hat{A}] + b\hat{\mathcal{G}}_t[\hat{B}]. \quad (110)$$

An arbitrary initial operator in the single excitation manifold can be expanded as

$$\begin{aligned} \hat{\rho}_S(0) = & \rho_{cc}(0)|c\rangle\langle c| \\ & + \sum_{i=1}^N \rho_{ic}(0)|e_i\rangle\langle c| + \sum_{i=1}^N \rho_{ci}(0)|c\rangle\langle e_i| \\ & + \sum_{i=1}^N \rho_{ii}(0)|e_i\rangle\langle e_i| \\ & + \sum_{\substack{i,j=1 \\ i \neq j}}^N \rho_{ij}(0)|e_i\rangle\langle e_j|. \end{aligned} \quad (111)$$

The five terms of Eq. 111 contain five distinct ordered operator types: a cavity population, a molecule-cavity coherence, a cavity-molecule coherence, a molecular population, and an ordered intermolecular coherence.

Molecular covariance of the dynamical map gives

$$\hat{\mathcal{G}}_t [\hat{U}_\pi \hat{A} \hat{U}_\pi^\dagger] = \hat{U}_\pi \hat{\mathcal{G}}_t [\hat{A}] \hat{U}_\pi^\dagger. \quad (112)$$

Therefore, the response to any molecularly relabeled operator can be obtained from the response to one representative operator of the same type. For example, if $\pi(1) = i$, then

$$\hat{\mathcal{G}}_t [|e_i\rangle\langle c|] = \hat{U}_\pi \hat{\mathcal{G}}_t [|e_1\rangle\langle c|] \hat{U}_\pi^\dagger. \quad (113)$$

Similarly, if $\pi(1) = i$ and $\pi(2) = j$, then

$$\hat{\mathcal{G}}_t [|e_i\rangle\langle e_j|] = \hat{U}_\pi \hat{\mathcal{G}}_t [|e_1\rangle\langle e_2|] \hat{U}_\pi^\dagger, \quad i \neq j. \quad (114)$$

It follows that the evolution of an arbitrary initial operator can be reconstructed from the representative propagations summarized in Table I.

If the reduced map preserves the adjoint operation,

$$\hat{\mathcal{G}}_t [\hat{A}^\dagger] = \hat{\mathcal{G}}_t [\hat{A}]^\dagger, \quad (115)$$

the response to $|c\rangle\langle e_1|$ can be obtained by taking the adjoint of the response to $|e_1\rangle\langle c|$. In that case, only four independent representative propagations are required. For bath decompositions containing complex decay rates ν_k , Eq. 115 must not be assumed automatically. It holds only when the complex exponential channels occur in the required conjugate pairs and the hierarchy truncation and closure preserve the corresponding conjugation relation. Without these conditions, the two ordered cavity-molecule coherences should be propagated separately.

The number of representative propagations needed to describe an arbitrary initial operator is therefore independent of N . However, if one explicitly requests every

Representative	Interpretation	Distinguished labels	Possible adjoint reduction
$ c\rangle\langle c $	Cavity population	0	Not applicable
$ e_1\rangle\langle c $	Molecule-cavity coherence	1	Paired with $ c\rangle\langle e_1 $
$ c\rangle\langle e_1 $	Cavity-molecule coherence	1	Paired with $ e_1\rangle\langle c $
$ e_1\rangle\langle e_1 $	Localized molecular population	1	Not applicable
$ e_1\rangle\langle e_2 $	Ordered intermolecular coherence	2	Its adjoint follows same relabeling

TABLE I. Representative initial system operators required for linear reconstruction. The number of distinguished labels specifies how many molecular identities remain explicit during the corresponding propagation. The two ordered cavity-molecule coherences can be related by an adjoint operation only when the bath decomposition, hierarchy truncation, and closure preserve the required conjugation relation.

entry of the final dense density matrix, the reconstruction must produce $(N + 1)^2$ matrix elements and consequently requires $O(N^2)$ output storage. This quadratic cost belongs to the requested dense output itself, not to the compressed HEOM propagation.

VII. RESULTS AND DISCUSSION

We use the numerical examples below to verify the symmetry-adapted propagation and to illustrate how the population dynamics change with molecular ensemble size. For all the simulations presented below, the molecule and cavity excitation energies are taken to be, $\omega_{\text{cav}} = \epsilon_x = 10000 \text{ cm}^{-1}$ and the the Rabi splitting is taken to be $\Omega_R = 500 \text{ cm}^{-1}$. The single molecule light-matter coupling changes with N according to Eq. 5,

$$g(N) = \frac{\Omega_R}{2\sqrt{N}}. \quad (116)$$

Consequently, the resonant UP and LP energies remain unchanged as N varies, while the coupling of any individual molecule to the cavity decreases as $N^{-1/2}$. For the one-exponential decomposition per local bath, the retained Drude-Lorentz contribution is generated from

$$J_{\text{DL}}(\omega) = \frac{2\lambda_{\text{DL}}\gamma_{\text{DL}}\omega}{\omega^2 + \gamma_{\text{DL}}^2}, \quad (117)$$

where λ_{DL} is the reorganization energy and γ_{DL} is the bath decay rate. We use $\lambda_{\text{DL}} = 50 \text{ cm}^{-1}$, $\gamma_{\text{DL}} = 18 \text{ cm}^{-1}$, and $T = 300 \text{ K}$, for which $k_B T \approx 208 \text{ cm}^{-1}$. The residual contribution from the omitted Matsubara terms is treated with the Ishizaki-Tanimura terminator. [7] Unless stated otherwise, the symmetry adapted calculations use hierarchy depth $L = 25$ and time step $\Delta t = 0.5 \text{ fs}$. The $N = 2$ comparisons with the QuTiP [11, 49] implementation of conventional HEOM use $L = 15$.

A. Single exponential bath

We now demonstrate the ability of the unique variable reduced HEOM method to simulate large polaritonic systems. We choose the upper polariton (UP) state, defined




FIG. 4. Population dynamics for an initial upper-polariton population with a one-exponential decomposition per local bath. (a) Comparison between symmetry-adapted HEOM and the QuTiP implementation of conventional HEOM for $N = 2$. Colored curves show the symmetry-adapted calculation, and the overlaid black traces show the corresponding conventional HEOM results for the UP, LP, and total dark populations. (b) UP, (c) LP, (d) total dark, (e) cavity, and (f) total exciton populations for $N = 5, 10, 25, 10^2, 10^3$, and 10^{12} , as identified by the legend.

in Eq. 6a, as the initial state of the system. On resonance, the initially populated upper polariton contains equal cavity and bright molecular contributions,

$$|+\rangle = \frac{1}{\sqrt{2}}(|c\rangle + |B\rangle). \quad (118)$$

In Fig. 4(a), we compare the symmetry-adapted unique variable calculation with conventional HEOM for $N = 2$ system. The UP, LP, and total dark state populations, which are represented by the red, blue, and yellow curves respectively, fully agree with the conventional HEOM simulations of QuTiP, which are represented by the corresponding black curves in the figure. This agreement

confirms that the canonical-ADO and unique variable reductions preserve the dynamics of the corresponding finite HEOM calculation.

In Fig. 4(b)-(d) we present the UP, LP and the total dark-state population dynamics by varying the N at the same Rabi splitting of $\Omega_R = 500 \text{ cm}^{-1}$. For $N = 5$, represented by the red curves, there is an initial fast decay of the UP state upto $\approx 100 \text{ fs}$, which is accompanied by the fast increase of the Dark state population. The LP populations does not significantly increase at this stage. The UP population then follows a slower decay, and we see a slow increment in the dark state as well as the LP populations. Around 700 fs , the Dark state starts decaying slowly and transferring more population to the LP. For $N = 10$, represented by the orange curves in the figure, we see a similar behaviour of the UP, LP, and dark state populations. The second decay of the UP population is faster than in the $N = 5$ case and there is also an overall decrease in the amount of population transfer from the dark to LP state. With increasing N value, the second decay behaviour of the UP population slowly disappears and for $N = 100$ (represented by the green curve), there is no longer any visible second decay timescale. At $N = 100$, there is still a small finite population transfer to the LP state from the dark state. The dynamics for $N = 10^3$ and $N = 10^{12}$, represented by light and dark blue curves respective, have no visible difference in the timescale considered and we can safely assume that for these parameters, the system has reached the thermodynamic limit, with the LP state only acquiring populations transiently and coherently during the very early stages of the dynamics (upto roughly 250 fs) and all the population transfers to dark states. At this limit, the UP and dark state populations represent a simple kinetic rate model kind of population transfer. [52, 53].

In Fig. 4(e)-(f), we present the population dynamics for the cavity and total exciton states, respectively, when the initial condition is the UP state and thus both the cavity and the exciton have 0.5 population at $t = 0 \text{ fs}$. For $N = 5$ case, we see an initial fast decay of the cavity populations to the excitons up to 100 fs . After 100 fs , the rate of transfer slows down, and around 700 fs , there is a reversal of net population flow, causing the cavity population to increase. For $N = 10$ and $N = 25$, we see a similar behavior, but the overall magnitude of the cavity population is reduced. At $N = 100$, we no longer observe the revival behavior of the cavity population. At $N = 10^3$, the system reached the thermodynamic limit, and the behavior converges, with all population getting transferred to the excitons. It is to be noted that even at $N = 10^{12}$, the initial transient population dynamics still shows coherent oscillations.

B. Dynamics from a localized molecular excitation




FIG. 5. Population dynamics with unique variable HEOM for the localized initial condition of Eq. 119. Panels (a)-(c) present the populations of the first site, cavity and the total population of other sites, respectively, for different number of molecules (N) coupled to cavity. Panels (d)-(f) present the total dark populations, and the populations of the upper and lower polariton states respectively, for different number of molecules (N) coupled to cavity. In all the panels, the different $N = 5, 10, 25, 10^2, 10^3, 10^{12}$ cases are represented by solid line curves with red, orange, yellow, green, blue and dark blue respectively.

We next consider the nonsymmetric initial condition

$$\hat{\rho}_S(0) = |e_1\rangle\langle e_1|. \quad (119)$$

This propagation retains molecule 1 as a distinguished label while compressing the remaining $N - 1$ molecules.

In Fig. 5(a)-(c), we present the site-1 (initially localized excitation), cavity and total site (excluding site-1) population dynamics for varying N . For $N = 5$ system (the red curves in all panels), there is an initial fast decay of the localized site-1 population which is accompanied by initial coherent increase of populations in cavity and other sites as can be seen in panels (b) and (c). After the initial transient period, the populations reach a steady state. The steady cavity population is approximately $0.2 \simeq 1/5$. The $N = 10$ and $N = 25$ systems (represented by the orange and yellow curves respectively in all panels) show a similar saturation behaviour where the steady state cavity populations show a $1/N$ behaviour. Unlike $N = 5$ system, the $N = 10$ and $N = 25$ systems do not show two separate timescale behaviour in the decay and with increase in N , the rate of decay of the site-1 populations decrease. For $N = 100$, there is almost

no population in the cavity mode and almost all of the population is redistributed to the other sites. For 10^3 molecules the rate of redistribution is much lower and for $N = 10^{12}$, there is no visible transfer to other states within the 2 ps time window and a localized excitations stays localized without any redistribution.

On resonance, the localized state contains bright state weight $1/N$ and dark state weight $1 - 1/N$. The initial TC-basis (Eq. 6) populations are therefore

$$P_+(0) = P_-(0) = \frac{1}{2N}, \quad (120a)$$

$$P_D(0) = 1 - \frac{1}{N}. \quad (120b)$$

The nonzero initial polariton populations arise from the projection of $|e_1\rangle$ onto the collective bright state. In Fig. 5(d)-(f) we present the TC-eigenbasis populations which provide a complementary interpretation. For $N = 5$, Eq. 120a gives an initial dark population of 0.8 and initial UP and LP populations of 0.1 each. There is a fast decay of the dark state population in the first 100 fs to the LP and UP states which is followed by a slower decay. The initial spike in the UP population is followed by a decay back to the dark states. Most of the dark state decay is to the LP state at later stage of dynamics. With increasing N , the initial polariton participation decreases as $1/N$, while the dark contribution approaches unity. The subsequent changes in the UP and LP populations are therefore progressively reduced with increasing N . At higher N values, the population gets trapped in the dark subspace with no transfers to the bright polariton manifold. Together, the site and TC-eigenbasis representations demonstrate that distinguished molecule propagation retains both the localized component and its collective bright state contribution without restoring the full molecule-resolved hierarchy.

C. Dynamics from a coherent cavity-site superposition

The final single exponential example uses the coherent nonsymmetric initial state

$$|\Psi_S(0)\rangle = \sqrt{0.3}|e_1\rangle + \sqrt{0.7}|c\rangle, \quad (121a)$$

$$\hat{\rho}_S(0) = |\Psi_S(0)\rangle\langle\Psi_S(0)|. \quad (121b)$$

Its nonzero populations and coherences are

$$\langle e_1|\hat{\rho}_S(0)|e_1\rangle = 0.3, \quad (122a)$$

$$\langle c|\hat{\rho}_S(0)|c\rangle = 0.7, \quad (122b)$$

$$\langle e_1|\hat{\rho}_S(0)|c\rangle = \langle c|\hat{\rho}_S(0)|e_1\rangle = \sqrt{0.21}. \quad (122c)$$

In Fig. 6(a)-(c) we present the site-1, cavity and total site (excluding site-1) population dynamics for different N values for the initial condition of Eq. 121a. For $N = 5$, represented by the red curve in all panels, there is an initial fast decay of the site-1 and cavity populations to




FIG. 6. Population dynamics for the cavity-site superposition initial condition in Eq. 121a. Panels (a), (b), and (c) show the population of molecule 1, the cavity population, and the total population of the remaining molecular sites, respectively. Panels (d), (e), and (f) show the total dark state manifold, UP, and LP populations. Curves correspond to $N = 5, 10, 25, 10^2, 10^3$, and 10^{12} , as identified by the legend.

the other sites. After this initial transfer to the sites, there is coherent exchange of populations between the cavity and the sites, following which there is no further transfer and the cavity population reaches a steady state of approximately 0.2 ($\simeq 1/5$). For increasing N , we see a similar initially coherent transfer from the cavity to the other sites with the cavity steady state population showing a $1/N$ behaviour. For the site-1, with increasing N , the rate of decay of population slows down and at $N = 10^{12}$ we see that there is no transfer of population from site-1 and that it remains at its initial population of 0.3. The increase of the total population of the other sites shows a similar behaviour as in the $N = 5$ case and with increase in N , the increase of the other site populations comes mostly from the cavity decay. At $N = 10^{12}$, all of the cavity population is transferred to the other sites and there is no exchange with the site-1 population.

Using the UP and LP convention in Eq. 6, the resonant initial TC-eigenbasis populations are

$$P_D(0) = 0.3 \left(1 - \frac{1}{N}\right), \quad (123a)$$

$$P_{\pm}(0) = \frac{1}{2} \left(\sqrt{0.7} \pm \sqrt{\frac{0.3}{N}} \right)^2, \quad (123b)$$

The different initial UP, LP, and dark populations are therefore fixed by the coherent cavity-bright projection of the prepared state and vary explicitly with N .




FIG. 7. Disorder-averaged population dynamics following initialization in the upper polariton of the mean resonant HTC Hamiltonian in the presence of independent Gaussian static energetic disorder and identical local Drude-Lorentz environments. Panels show the populations of the reference (a) upper polariton, (b) lower polariton, and (c) total dark manifold for $N = 5, 10, 25, 10^2, 10^3$, and 10^{12} . The reference states are defined by the disorder-free TC Hamiltonian.

In Fig. 6(d)-(f) we present the total dark state, UP and LP population dynamics for different N values for the initial condition of Eq. 121a. The initial values in Eq. 123 account for the ensemble-size dependent intercepts in Fig. 6(d)-(f). For smaller ensembles, the localized molecular component has a larger bright state projection, producing a stronger imbalance between the initial UP and LP populations. As N increases, this molecular bright component decreases, $P_D(0)$ approaches 0.3, and both polariton populations approach 0.35.

The polariton states in Fig. 6(d)-(f) show very similar behaviours as observed in Fig. 4. For small N values there are two timescales of decay of the UP population: an initial fast decay followed by a slower decay and with increase in N , the second decay mechanism disappears and we have an overall fast decay of UP state. The LP state on the other hand, at small N values, shows an increase of population after its initial decrease. With increase of N value, this behaviour disappears and at $N = 10^3$ the population of LP only decays. The decay of UP and LP states is accompanied by the increase of dark state populations and the amount of population transferred to dark states increases with N with the total dark state population becoming 1 at long times for $N = 10^3$. It is to be noted that even at $N = 10^{12}$, unlike Fig. 4, the initial population dynamics of polariton states show coherent oscillations.

D. Static Energetic Disorder

We now present the polariton dynamics in presence of diagonal inhomogeneous static disorder [29, 33, 38]. The TC-system Hamiltonian in Eq. 3 is modified as,

$$\hat{H}_S(\delta) = \hat{H}_S + \sum_{i=1}^N \delta \epsilon_i \hat{Q}_i, \quad \hat{Q}_i = |e_i\rangle\langle e_i|. \quad (124)$$

The exciton diagonal energies is sampled from an inhomogeneous gaussian distribution,

$$P(\delta) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\delta \epsilon_i^2}{2\sigma^2}\right). \quad (125)$$

Since the distribution is static and follows a gaussian statistics, we can represent the diagonal disorder as a diagonal bath with a non-decaying correlation function 11,

$$C_{\text{static}}(t) = \sigma^2 \cdot e^{-0 \cdot t} = \sigma^2 \quad (126)$$

This adds an HEOM bath channel to every molecule with the expansion coefficient σ^2 and decay rate of 0.

Fig. 7 shows the disorder-averaged dynamics following initialization in the upper polariton (Eq. 6a) of the mean resonant TC Hamiltonian (Eq. 3). We use $\omega_{\text{cav}} = \epsilon_x = 10000 \text{ cm}^{-1}$, $\Omega_R = 100 \text{ cm}^{-1}$, $\sigma = 25 \text{ cm}^{-1}$. The collective Rabi splitting is held fixed as N is varied according to Eq. 5. In addition to the static channel, each molecule is coupled to an identical Drude-Lorentz environment with $\lambda_{\text{DL}} = 50 \text{ cm}^{-1}$, $\gamma_{\text{DL}} = 18 \text{ cm}^{-1}$, $T = 300 \text{ K}$ for which $k_B T \simeq 208 \text{ cm}^{-1}$. We used the multi-exponential formalism to represent the two channels (Drude-Lorentz and static disorder) where the total hierarchy depth converged at $L = 14$ for the above parameters. The details of multi-exponential formalism of symmetry-adapted HEOM and the static disorder is discussed in the Supplementary Information.

The static energy fluctuations mix the permutation-symmetric bright state with the molecular dark manifold in each realization. After disorder averaging, this appears as rapid loss of bright-state coherence and growth of the total dark population. The dynamical Drude-Lorentz environments act simultaneously and allow energy relaxation in addition to the inhomogeneous dephasing generated by the static channel.

For $N = 5$, the UP population decreases rapidly, while the LP population develops a pronounced early-time maximum and remains appreciable over the displayed interval. The LP curve also exhibits a weak late-time upturn after its initial decrease. As N increases,

N	$\mathcal{N}_{\text{ADO}}^{\text{naive}}$	$\mathcal{M}_{\text{ADO}}^{\text{naive}}$ (GB)	$\mathcal{N}_{\text{ADO}}^{\text{can}}$	$\mathcal{N}_{\text{var}}^{\text{unq}}$	$\mathcal{M}_{\text{var}}^{\text{unq}}$ (GB)	$\mathcal{R}$ -factor
5	1.43×10^5	8.21×10^{-2}	2.60×10^3	6.21×10^4	9.94×10^{-4}	8.26×10^1
10	1.84×10^8	3.55×10^2	7.53×10^3	2.44×10^5	3.91×10^{-3}	9.10×10^4
25	1.26×10^{14}	1.37×10^9	9.30×10^3	3.07×10^5	4.91×10^{-3}	2.79×10^{11}
10^2	1.30×10^{26}	2.12×10^{22}	9.30×10^3	3.07×10^5	4.91×10^{-3}	4.32×10^{24}
10^3	8.90×10^{49}	1.43×10^{48}	9.30×10^3	3.07×10^5	4.91×10^{-3}	2.91×10^{50}
10^{12}	6.45×10^{274}	1.03×10^{291}	9.30×10^3	3.07×10^5	4.91×10^{-3}	2.10×10^{293}

TABLE II. Analytical hierarchy sizes and raw state-vector storage for one exponential term per molecular bath and total hierarchy depth $L = 25$. The two storage columns count one copy of the corresponding complex-double state and are not measurements of peak resident memory. The canonical-ADO and unique-variable columns refer to a fully permutation-symmetric propagation. The reduction factor $\mathcal{R}$ is the ratio of naive ADO storage to unique-variable storage.

the remaining UP and LP populations are reduced and a progressively larger fraction of the excitation accumulates in the dark manifold. For the largest ensembles, the dark population approaches unity within the plotted time window, and the curves for $N = 10^3$ and $N = 10^{12}$ are nearly coincident. This agreement is a physical convergence with ensemble size for the chosen parameters and it is distinct from the earlier saturation of the computational representation with N .

E. Comment on the computational resources utilized

Table II summarizes the analytical storage requirements for the one-exponential bath calculations at hierarchy depth $L = 25$. The molecular ensemble sizes are the same as those used in the population dynamics figures. In the single excitation basis $\mathcal{B}_N = \{|c\rangle, |e_1\rangle, \dots, |e_N\rangle\}$, every conventional ADO is a dense $(N+1) \times (N+1)$ complex matrix. For one exponential term per molecule, the conventional hierarchy and its raw storage are

$$\mathcal{N}_{\text{ADO}}^{\text{naive}}(N, 1, 25) = \binom{N+25}{25}, \quad (127a)$$

$$\mathcal{M}_{\text{ADO}}^{\text{naive}} = 16(N+1)^2 \mathcal{N}_{\text{ADO}}^{\text{naive}} \text{ bytes}, \quad (127b)$$

where each complex double-precision number occupies 16 bytes. The unique-variable storage is

$$\mathcal{M}_{\text{var}}^{\text{unq}} = 16 \mathcal{N}_{\text{var}}^{\text{unq}} \text{ bytes}. \quad (128)$$

The memory values in Table II are reported in decimal gigabytes, with 1 GB = 10^9 bytes. The reduction factor reported below compares the complete number of complex entries in the molecule-resolved hierarchy with the number of propagated unique variables,

$$\mathcal{R}(N) = \frac{\mathcal{N}_{\text{ADO}}^{\text{naive}}(N, 1, 25)(N+1)^2}{\mathcal{N}_{\text{var}}^{\text{unq}}(N, 1, 25)} = \frac{\mathcal{M}_{\text{ADO}}^{\text{naive}}}{\mathcal{M}_{\text{var}}^{\text{unq}}}. \quad (129)$$

The conventional hierarchy becomes impractical at remarkably small ensemble sizes. At $N = 10$, storing only one copy of all ADO matrices would already require 3.55×10^2 GB, whereas the corresponding unique-variable

vector requires only 3.91×10^{-3} GB, or 3.91 MB. The storage reduction at this ensemble size is 9.10×10^4 . At $N = 25$, the conventional one-vector lower bound reaches 1.37×10^9 GB, while the unique-variable vector remains only 4.91 MB, corresponding to a reduction by a factor of 2.79×10^{11} . The conventional estimates for the larger ensembles are astronomical and demonstrate that a direct molecule-resolved HEOM calculation is not merely expensive, but impossible on existing state-of-art computational hardware.

The two symmetry reductions remove different parts of this growth. At $L = 25$, the canonical hierarchy contains 9296 ADO representatives for every $N \geq 25$, instead of the rapidly increasing number of molecule-labeled ADOs. If these representatives were retained as full matrices, however, their storage would still grow as $(N+1)^2$. The second reduction removes this remaining matrix redundancy. The fully symmetric hierarchy contains exactly 306743 unique variables at $N = 25$, 306750 at $N = 26$, and 306751 for every $N \geq 27$. Its raw state-vector storage therefore saturates at only 4.91 MB, independent of any further increase in the number of molecules. The enormous calculation at $N = 10^{12}$ consequently uses the same number of propagated variables as the calculation at $N = 27$; only the multiplicity coefficients and the N -dependent system parameters change.

The memory values in Table II deliberately do not include the work arrays required by the time integrator. A conventional 4th Runge-Kutta implementation commonly retains the current state, four stage vectors, and one or more temporary or output vectors. Its state-sized working storage is therefore typically about six to seven times the one-vector values reported in the table, although the precise factor depends on the implementation and can be reduced with a low-storage integrator.

Localized molecular excitation and cavity-site superposition dynamics each distinguish one molecule and therefore use the same partially symmetric hierarchy structure. At $L = 25$, the corresponding canonical-pattern count saturates at 41391 for $N \geq 26$, and the unique-variable count saturates at 1715600 for $N \geq 28$. One copy of this state vector occupies 27.45 MB. This is approximately 5.59 times the storage of the fully symmetric vector, but it remains independent of N beyond the sat-

uration threshold and is still negligible compared with the conventional hierarchy. The localized and cavity-site superposition cases require identical structural memory because both retain one distinguished molecular label; their different populations and coherences alter only the initialized variable values, not the propagated dimension.

VIII. CONCLUSIONS

We have developed a symmetry adapted HEOM formulation for identical Holstein-Tavis-Cummings ensembles with independent, identical local environments. The reduction proceeds in two stages. Molecularly relabeled hierarchy elements are first represented by one canonical ADO occupation pattern. Repeated matrix elements within each representative are then replaced by a compact set of unique complex variables. These variables are propagated directly using precomputed hierarchy connections and molecular multiplicities, without constructing full ADO matrices or a dense HEOM generator.

The reduction is algebraically equivalent to the corresponding finite conventional HEOM defined by the same bath correlation decomposition, hierarchy depth, closure, and system-bath model. Canonicalization and unique-variable propagation introduce no additional dynamical approximation. The numerical comparison presented here shows that the reduced and conventional HEOM population dynamics are indistinguishable on the plotted scale.

For fixed hierarchy depth L and number of exponential terms m , the number of canonical ADO representatives becomes independent of ensemble size for $N \geq L$, and the unique-variable dimension saturates for $N \geq L + 2$. Fully permutation-symmetric initial operators use the maximally compressed representation. Localized populations and coherent nonsymmetric states are treated by retaining the required molecular labels explicitly while continuing to compress the remaining equivalent molecules. Linearity and molecular relabeling further permit general initial system operators, under the assumed initial bath preparation, to be reconstructed from a finite set of

representative propagations.

The complete reduction requires identical molecular transition energies and cavity couplings, identical independent local environments, and local diagonal system-bath interactions. The present formulation is also restricted to two level molecules, one cavity mode, and the single excitation manifold. Moreover, although the propagated dimension saturates with N , its value can still grow rapidly with the number of bath correlation exponentials and hierarchy depth. Extensions to optical response functions, higher excitation manifolds, and ensembles of identical multichromophoric units will be considered separately in the forthcoming papers.

DATA AVAILABILITY

The data that support the plots within this paper and other findings of this study are available from the corresponding authors upon a reasonable request.

CODE AVAILABILITY

The source code for this project are available from the corresponding authors upon a reasonable request.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Award under Grant No. CHE-2244683. M.E.M. appreciates the support from the Agnes M. and George Messersmith Fellowship by the University of Rochester. Computing resources were provided by the Center for Integrated Research Computing (CIRC) at the University of Rochester. We appreciate valuable suggestions from Sebastian Montillo Vega, Eric Koessler, and Santanu Poddar.

-
- [1] Arkajit Mandal, Michael AD Taylor, Braden M Weight, Eric R Koessler, Xinyang Li, and Pengfei Huo. Theoretical advances in polariton chemistry and molecular cavity quantum electrodynamics. *Chem. Rev.*, 123(16): 9786–9879, 2023.
 - [2] Wenxiang Ying, M. Elious Mondal, Eric R. Koessler, Sebastian Montillo Vega, and Pengfei Huo. Collective effects in polariton chemistry and photophysics. *Ann. Rev. Phys. Chem.*, 77(Volume 77, 2026):201–224, 2026. ISSN 1545-1593.
 - [3] Benjamin XK Chng, Wenxiang Ying, Yifan Lai, A Nickolas Vamivakas, Steven T Cundiff, Todd D Krauss, and Pengfei Huo. Mechanism of molecular polariton decoherence in the collective light-matter couplings regime. *J. Phys. Chem. Lett.*, 15(47):11773–11783, 2024.
 - [4] Deping Hu, Benjamin XK Chng, Wenxiang Ying, and Pengfei Huo. Trajectory-based non-adiabatic simulations of the polariton relaxation dynamics. *J. Chem. Phys.*, 162(12), 2025.
 - [5] M. Elious Mondal, Seongje Park, Minjung Son, A Nickolas Vamivakas, Steven T Cundiff, Todd D Krauss, and Pengfei Huo. Long-lived polaritonic coherence and polaron decoupling effects in 2D electronic spectra. *Proc. Natl. Acad. Sci. USA*, 123(17):e2530454123, 2026.
 - [6] Yoshitaka Tanimura and Ryogo Kubo. Time evolution of a quantum system in contact with a nearly gaussian-

- markoffian noise bath. *J. Physical. Soc. Japan*, 58(1): 101–114, 1989.
- [7] Akihito Ishizaki and Yoshitaka Tanimura. Quantum dynamics of system strongly coupled to low-temperature colored noise bath: Reduced hierarchy equations approach. *J. Physical. Soc. Japan*, 74(12):3131–3134, 2005.
 - [8] Yoshitaka Tanimura. Stochastic liouville, langevin, fokker-planck, and master equation approaches to quantum dissipative systems. *J. Physical. Soc. Japan*, 75(8): 082001, 2006.
 - [9] Yoshitaka Tanimura. Numerically “exact” approach to open quantum dynamics: The hierarchical equations of motion (HEOM). *J. Chem. Phys.*, 153(2), 2020.
 - [10] Shuming Bai, Shuocang Zhang, Chenghong Huang, and Qiang Shi. Hierarchical equations of motion for quantum chemical dynamics: Recent methodology developments and applications. *Acc. Chem. Res.*, 57(21):3151–3160, 2024.
 - [11] Neill Lambert, Tarun Raheja, Simon Cross, Paul Mencil, Shah Nawaz Ahmed, Alexander Pitchford, Daniel Burgarth, and Franco Nori. QuTiP-BoFiN: A bosonic and fermionic numerical hierarchical-equations-of-motion library with applications in light-harvesting, quantum control, and single-molecule electronics. *Phys. Rev. Res.*, 5(1):013181, 2023.
 - [12] Qiang Shi, Liping Chen, Guangjun Nan, Rui-Xue Xu, and YiJing Yan. Efficient hierarchical liouville space propagator to quantum dissipative dynamics. *J. Chem. Phys.*, 130(8), 2009.
 - [13] Jing Zhu, Sabre Kais, Patrick Rebentrost, and Alán Aspuru-Guzik. Modified scaled hierarchical equation of motion approach for the study of quantum coherence in photosynthetic complexes. *J. Phys. Chem. B*, 115(6): 1531–1537, 2011.
 - [14] Hou-Dao Zhang, Qin Qiao, Rui-Xue Xu, Xiao Zheng, and YiJing Yan. Efficient steady-state solver for hierarchical quantum master equations. *J. Chem. Phys.*, 147(4), 2017.
 - [15] Story Temen, Amber Jain, and Alexey V Akimov. Hierarchical equations of motion in the libra software package. *Int. J. Quantum Chem.*, 120(22):e26373, 2020.
 - [16] Yi-Te Huang, Po-Chen Kuo, Neill Lambert, Mauro Cirio, Simon Cross, Shen-Liang Yang, Franco Nori, and Yueh-Nan Chen. An efficient julia framework for hierarchical equations of motion in open quantum systems. *Commun. Phys.*, 6(1):313, 2023.
 - [17] Hao Liu, Lili Zhu, Shuming Bai, and Qiang Shi. Reduced quantum dynamics with arbitrary bath spectral densities: Hierarchical equations of motion based on several different bath decomposition schemes. *J. Chem. Phys.*, 140(13), 2014.
 - [18] Zi-Hao Chen, Yao Wang, Xiao Zheng, Rui-Xue Xu, and YiJing Yan. Universal time-domain prony fitting decomposition for optimized hierarchical quantum master equations. *J. Chem. Phys.*, 156(22), 2022.
 - [19] Christoph Kreisbeck, Tobias Kramer, Mirta Rodriguez, and Birgit Hein. High-performance solution of hierarchical equations of motion for studying energy transfer in light-harvesting complexes. *J. Chem. Theory. Comput.*, 7(7):2166–2174, 2011.
 - [20] Masashi Tsuchimoto and Yoshitaka Tanimura. Spins dynamics in a dissipative environment: Hierarchical equations of motion approach using a graphics processing unit (GPU). *J. Chem. Theory. Comput.*, 11(8):3859–3865, 2015.
 - [21] Tobias Kramer, Matthias Noack, Alexander Reinefeld, Mirta Rodriguez, and Yaroslav Zelinskyy. Efficient calculation of open quantum system dynamics and time-resolved spectroscopy with distributed memory HEOM (DM-HEOM). *J. Comput. Chem.*, 39(22):1779–1794, 2018.
 - [22] Daochi Zhang, Lyuzhou Ye, Jiaan Cao, Yao Wang, Rui-Xue Xu, Xiao Zheng, and YiJing Yan. HEOM-QUICK2: A general-purpose simulator for fermionic many-body open quantum systems—an update. *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, 14(4):e1727, 2024.
 - [23] Lachlan Lindoy. *New developments in open quantum system dynamics*. PhD thesis, University of Oxford, 2019.
 - [24] Yaming Yan, Meng Xu, Tianchu Li, and Qiang Shi. Efficient propagation of the hierarchical equations of motion using the tucker and hierarchical tucker tensors. *J. Chem. Phys.*, 154(19), 2021.
 - [25] Yaling Ke. Tree tensor network state approach for solving hierarchical equations of motion. *J. Chem. Phys.*, 158(21), 2023.
 - [26] Lachlan P Lindoy, Daniel Rodrigo-Albert, Yannic Rath, and Ivan Rungger. pyTTN: An open-source toolbox for open and closed system quantum dynamics simulations using tree tensor networks. *J. Chem. Phys.*, 163(20), 2025.
 - [27] Juan C Rodriguez-Betancourt, Michelle C Anderson, Luchang Niu, Xinxian Chen, and Ignacio Franco. TENSO: Software package for numerically exact open quantum dynamics based on efficient tree tensor network decomposition of the hierarchical equations of motion. *J. Chem. Theory. Comput.*, 2026.
 - [28] Xinxian Chen and Ignacio Franco. Tree tensor network hierarchical equations of motion based on time-dependent variational principle for efficient open quantum dynamics in structured thermal environments. *J. Chem. Phys.*, 163(10), 2025.
 - [29] Tianchu Li, Pranay Venkatesh, Qiang Shi, and Andrés Montoya-Castillo. For molecular polaritons, disorder and phonon timescales control the activation of dark states in the thermodynamic limit. *arXiv preprint arXiv:2603.06868*, 2026.
 - [30] Etienne Mangaud, Amine Jaouadi, Alex Chin, and Michèle Desouter-Lecomte. Survey of the hierarchical equations of motion in tensor-train format for non-markovian quantum dynamics. *Eur. Phys. J. Spec. Top.*, 232(12):1847–1869, 2023.
 - [31] Weizhong Guan, Peng Bao, Jiawei Peng, Zhenggang Lan, and Qiang Shi. mpsqd: A matrix product state based python package to simulate closed and open system quantum dynamics. *J. Chem. Phys.*, 161(12), 2024.
 - [32] Elliot W Lloyd, Aleksandra A Ziolkowska, and Jonathan Keeling. Permutation-symmetric quantum trajectories. *arXiv preprint arXiv:2605.11103*, 2026.
 - [33] M. Elious Mondal, A. Nickolas Vamivakas, Steven T. Cundiff, Todd D. Krauss, and Pengfei Huo. Polariton spectra under the collective coupling regime. I. Efficient simulation of linear spectra and quantum dynamics. *J. Chem. Phys.*, 162(1):014114, 01 2025.
 - [34] M. Elious Mondal, A. Nickolas Vamivakas, Steven T Cundiff, Todd D Krauss, and Pengfei Huo. Polariton spectra under the collective coupling regime. II. 2D non-linear spectra. *J. Chem. Phys.*, 162(7):074110, 2025.

- [35] Benjamin XK Chng, M. Elious Mondal, Wenxiang Ying, and Pengfei Huo. Quantum dynamics simulations of exciton polariton transport. *Nano Lett.*, 25(4):1617–1622, 2025.
- [36] Benjamin XK Chng, Braden M Weight, M. Elious Mondal, and Pengfei Huo. Ab Initio Polariton Transport Dynamics with the Classical Path Approximation. *Nano Lett.*, 26(18):6055–6062, 2026.
- [37] Saeed Rahmanian Koshkaki, Arshath Manjalingal, Logan Blackham, and Arkajit Mandal. Exciton-polariton dynamics in multilayered materials. *Nat. Commun.*, 17(1):6156, 2026.
- [38] Tianchu Li, Chenghong Huang, Shuming Bai, and Qiang Shi. Theoretical methods based on linear response theory to simulate dynamics and absorption spectra of molecular polaritons. *J. Chem. Phys.*, 162(9), 2025.
- [39] Anqi Mu, Nathan Ng, Andrew J Millis, and David R Reichman. Beyond mean-field dynamics of the dicke model with non-markovian dephasing. *Phys. Rev. Lett.*, 136(9):093601, 2026.
- [40] Piper Fowler-Wright, Brendon W Lovett, and Jonathan Keeling. Efficient many-body non-markovian dynamics of organic polaritons. *Phys. Rev. Lett.*, 129(17):173001, 2022.
- [41] Piper Fowler-Wright, Kristín B Arnardóttir, Peter Kirton, Brendon W Lovett, and Jonathan Keeling. Determining the validity of cumulant expansions for central spin models. *Phys. Rev. Res.*, 5(3):033148, 2023.
- [42] Piper Fowler-Wright. Mean-field and cumulant approaches to modelling organic polariton physics. *arXiv preprint arXiv:2405.09812*, 2024.
- [43] Gerald E Fux, Piper Fowler-Wright, Joel Beckles, Eoin P Butler, Paul R Eastham, Dominic Gribben, Jonathan Keeling, Dainius Kilda, Peter Kirton, Ewen DC Lawrence, et al. OQuPy: A python package to efficiently simulate non-markovian open quantum systems with process tensors. *J. Chem. Phys.*, 161(12), 2024.
- [44] Wenxiang Ying and Abraham Nitzan. Excitation-density-controlled regimes of collective light-matter dynamics. *J. Phys. Chem. Lett.*, 2026.
- [45] Juan B. Pérez-Sánchez, Arghadip Koner, Nathaniel P. Stern, and Joel Yuen-Zhou. Simulating molecular polaritons in the collective regime using few-molecule models. *Proc. Natl. Acad. Sci. USA*, 120(15):e2219223120, 2023.
- [46] Juan B Pérez-Sánchez, Arghadip Koner, Sricharan Raghavan-Chitra, and Joel Yuen-Zhou. CUT-E as a 1/N expansion for multiscale molecular polariton dynamics. *J. Chem. Phys.*, 162(6), 2025.
- [47] Joel Yuen-Zhou and Arghadip Koner. Linear response of molecular polaritons. *J. Chem. Phys.*, 160(15), 2024.
- [48] Lachlan P. Lindoy, Arkajit Mandal, and David R. Reichman. Investigating the collective nature of cavity-modified chemical kinetics under vibrational strong coupling. *Nanophotonics*, 13(14):2617–2633, 2024.
- [49] Neill Lambert, Eric Giguère, Paul Menczel, Boxi Li, Patrick Hopf, Gerardo Suárez, Marc Gali, Jake Lishman, Rushiraj Gadhvi, Rochisha Agarwal, et al. QuTiP 5: The quantum toolbox in Python. *Phys. Rep.*, 1153:1–62, 2026.
- [50] M. Elious Mondal, Eric R Koessler, Justin Provazza, A Nickolas Vamivakas, Steven T Cundiff, Todd D Krauss, and Pengfei Huo. Quantum dynamics simulations of the 2D spectroscopy for exciton polaritons. *J. Chem. Phys.*, 159(9):094102, 2023.
- [51] Wenxiang Ying, M Elious Mondal, and Pengfei Huo. Theory and quantum dynamics simulations of exciton-polariton motional narrowing. *J. Chem. Phys.*, 161(6), 2024.
- [52] Yifan Lai, Wenxiang Ying, and Pengfei Huo. Non-equilibrium rate theory for polariton relaxation dynamics. *J. Chem. Phys.*, 161(10), 2024.
- [53] Yifan Lai, Wenxiang Ying, Todd D Krauss, and Pengfei Huo. Analytic rate theory of polariton relaxation that explains long polariton lifetime. *J. Chem. Phys.*, 164(2), 2026.