Equilibrium from the Inside Out: Exact Hamiltonian of Mean Force

Gerard McCaul
(Dated: April 18, 2026)

We ask when the reduced equilibrium state of an open quantum system can be written not merely as an implicit density operator but as an explicit Hamiltonian-level generator. Starting from the quenched representation of the traced Gibbs state, we derive an exact Gaussian reformulation in which bath physics enters through scalar kernel moments while operator growth is carried entirely by the adjoint chain generated by the bare Hamiltonian and the coupling. The exact Gaussian HMF therefore comes with a compositional architecture of its own: an adjoint orbit, a bilinear influence aggregation, and a nonlinear Baker–Campbell–Hausdorff recombination. This yields a closure criterion: a finite closed-form Hamiltonian of mean force exists only when that operator family closes inside the chosen ansatz. For the spin-boson qubit all three layers close exactly, yielding, to our knowledge, the first closed-form HMF for a genuinely noncommuting open quantum system. The same variables organize the numerical crossover: once the dressed state outgrows the available basis, simulation error reflects a representability bottleneck rather than a breakdown of the exact theory. The resulting perspective treats strong-coupling equilibrium as a problem of constructive reduced description, compact operator classes, and finite-resource inference.

INTRODUCTION

A central question of equilibrium quantum statistical mechanics is what survives coarse-graining. Before reduction, equilibrium is generated by a Hamiltonian and written $\rho \propto e^{-\beta H}$: the state and its generator are given together. After an environment is traced out, the reduced equilibrium state remains well defined, but the generator need not remain transparent. The sharper question is therefore not whether reduced equilibrium exists, but whether it can still be written as an explicit Hamiltonian rather than only as an implicit density operator. This is the inside-out version of equilibrium, where one starts from the reduced state and asks what generator, if any, survives the reduction.

For an open system with finite coupling, the operational equilibrium state of the subsystem is the reduced state of the global Gibbs ensemble,

$$\bar{\rho}_Q(\beta) = \mathrm{Tr}_X\, e^{-\beta H_{\mathrm{tot}}}, \tag{1}$$

$$e^{-\beta H_{\mathrm{MF}}(\beta)} \propto \frac{\mathrm{Tr}_X e^{-\beta H_{\mathrm{tot}}}}{Z_X(\beta)}, \qquad Z_X(\beta) = \mathrm{Tr}_X e^{-\beta H_X}. \tag{2}$$

The Hamiltonian of mean force (HMF) is the operator that restores a Gibbs-form description after the bath has been traced out. It is therefore the exact reduced equilibrium generator whenever system–bath coupling is finite [1–6].

Why care about a mean-force Hamiltonian at all? First, it underlies the free-energy, work, and entropy relations of strong-coupling thermodynamics [1, 4, 7–10]. Second, it provides a consistent reduced equilibrium state when correlations with the environment cannot be ignored [3, 11, 12]. Third, and more conceptually, it turns a traced Gibbs operator into a constructive reduced description that can be inspected, truncated, or compared across operator families.

That constructive viewpoint matters whenever reduced descriptions must remain explicit, interpretable, or computationally tractable. Recent work on representation cost emphasizes that coarse-graining is constrained by what a reduced model class can continue to carry [13]. A parallel instinct now appears in machine learning through compact compositional architectures such as Kolmogorov–Arnold networks, which ask when complicated multivariate dependence can still be carried by a structured low-complexity decomposition [14]. We return to that comparison only after the physics problem has been made precise. The present paper remains entirely about equilibrium statistical mechanics: given an exact reduced Gibbs state, when does it stay inside a compact operator family?

The literature is broad but structurally clear. Canonical definitions and thermodynamic identities are now well established [1–4, 7]. Exact evaluations exist in special cases: commuting interactions close trivially, quadratic models preserve Gaussianity, and finite-dimensional examples such as the spin-boson model provide controlled testbeds [15–19]. Benchmark asymptotes in weak coupling and in strongly dressed regimes are also known and remain important checks on any exact construction [20, 21]. Recent work has further emphasized operator structure, bath feedback, and finite-reservoir generalizations [22–24].

Outside solvable models, however, the HMF is usually accessed numerically or perturbatively. Reaction-coordinate mappings, polaron transforms, hierarchical equations, stochastic approaches, and direct imaginary-time solvers can compute $\bar{\rho}_Q(\beta)$, but they rarely return a compact closed-form generator [25–36]. The obstruction is therefore not the existence of the reduced equilibrium state but its representability: when does the logarithm of a traced exponential stay inside a restricted operator family such as few-body terms, local operators, or a chosen algebra?

The influence-functional formalism is a natural language for this question. For Gaussian baths with linear coupling it provides an exact route to integrating out environmental degrees of freedom [16, 17, 37]. In equilibrium it becomes an imaginary-time influence functional, which admits a Hubbard–Stratonovich rewriting as a quenched stochastic average [38–40]. This formulation makes the reduction step explicit: the partial trace becomes an average over imaginary-time back-action histories, while all remaining questions are pushed onto the system operator algebra.

The present paper takes the quenched-density construction of Ref. [41] as its exact starting point and uses it to answer a representability question. For Gaussian baths with linear coupling, bath statistics collapse to scalar kernel moments while operator growth is generated entirely by the adjoint chain $\{\mathrm{ad}^n_{H_Q}(f)\}_{n\geq 0}$. In the Gaussian sector the exact HMF already comes with a compositional architecture: an adjoint orbit, a bilinear aggregation weighted by bath moments, and a nonlinear BCH recombination. Conditions (C1)–(C3) are precisely the closure conditions for those three layers to remain inside a chosen operator ansatz. We then instantiate this criterion in the spin-boson qubit, where the reduced equilibrium compresses to a closed Bloch-plane generator controlled by two response channels and a single influence magnitude. To our knowledge this yields the first closed-form HMF for a genuinely noncommuting open quantum system. The same variables organize the numerical crossover, revealing a representability bottleneck once the dressed reduced state outgrows the chosen finite basis. The KAN comparison then becomes a statement about the compositional structure already present in the exact HMF construction itself.

QUENCHED REPRESENTATION AND INFLUENCE FUNCTIONAL

We begin by recapitulating the quenched representation introduced in Ref. [41]. To proceed directly from the introduction, we make explicit the composite model. We denote the bare system Hamiltonian by H_Q and write

$$H_{\mathrm{tot}} = H_Q + H_X + H_{\mathrm{int}}, \tag{3}$$

where H_X is the bath Hamiltonian. We assume a factorizable interaction

$$H_{\mathrm{int}} = f \otimes B, \tag{4}$$

where f acts on the system and B is a bath operator. We can define the reduced equilibrium operator (up to normalisation) by $\bar{\rho}_Q(\beta) \equiv \mathrm{Tr}_X e^{-\beta H_{\mathrm{tot}}}$. As shown in Ref. [41], this can be represented as a *quenched density.* This is an average over a stochastic propagator, given by:

$$\bar{\rho}_Q(\beta) = \mathbb{E}_\xi\left[U_\xi(\beta)\right], \tag{5}$$

$$U_\xi(\beta) \equiv \mathcal{T}_\tau \exp\left[-\int_0^\beta d\tau\, \left(H_Q + \xi(\tau) f\right)\right], \tag{6}$$

where $\xi(\tau)$ is a stochastic process whose statistics encode the bath correlations, and $\mathcal{T}_\tau$ denotes time-ordering in imaginary time. Regardless of the precise form of bath and coupling, the reduced density must always be describable in this form.

This result can be understood intuitively by first observing that since the environment influence can only enter through the system coupling f, its influence can be captured by attaching a τ-dependent driving field to f [42, 43]. If this were a single deterministic field however, it would correspond to the bath exerting the *same* back-action history for every microscopic bath configuration. But Tr_X averages over many bath microstates in the thermal ensemble, and hence over many back-action histories. In this sense the partial trace is necessarily an average over histories in imaginary time, and the quenched representation simply makes this averaging explicit. For each realisation $\xi(\tau)$ the system evolves under an imaginary-time Hamiltonian $H_Q + \xi(\tau) f$; the bath is then recovered by averaging over $\xi(\tau)$ with a law chosen to reproduce the bath-induced correlations. In this sense $\xi(\tau)$ is not a physical external control field but an efficient parametrisation of the bath history $B(\tau)$ as seen through the coupling channel.

We may understand what formal properties are demanded of ξ by considering the influence functional it is required to match. Working in imaginary time, introduce the bath interaction picture:

$$B(\tau) \equiv e^{\tau H_X} B e^{-\tau H_X}, \qquad \tau \in [0, \beta], \tag{7}$$

and define the bath thermal state $\rho_X \equiv e^{-\beta H_X}/Z_X$. For an arbitrary c-number source $j(\tau)$ coupled linearly to $B(\tau)$, the bath generates a (time-ordered) functional

$$\mathcal{Z}_X[j] \equiv \mathrm{Tr}_X\left[\mathcal{T}_\tau \exp\left(-\int_0^\beta d\tau\, j(\tau)\, B(\tau)\right)\rho_X\right]. \tag{8}$$

Here $j(\tau)$ is introduced as an external c-number source used to generate ordered bath correlators by functional differentiation. More generally, $j(\tau)$ may be any object commuting with the bath algebra (e.g. a system operator tensored with $\mathbb{I}_X$). In the influence-functional derivation it is ultimately supplied by the system history, which becomes a c-number function in the path-integral representation. In the influence-functional approach, tracing out the bath produces precisely such a functional, evaluated on the system history through the coupling channel.

From this, the bath contribution to the effective Euclidean action can be written as [16, 17, 37]

$$\mathcal{F}[f] \equiv \mathcal{Z}_X[f], \qquad \Phi[f] \equiv \log \mathcal{F}[f]. \tag{9}$$

A key structural fact is that $\Phi[f] = \log \mathcal{F}[f]$ is the *cumulant generating functional* of the bath operator $B(\tau)$ with respect to the thermal state. Concretely, the generalised (time-ordered) cumulant theorem implies the connected expansion [44, 45]

$$\Phi[f] = \sum_{n\geq 1} \frac{(-1)^n}{n!} \int_0^\beta d\tau_1 \cdots d\tau_n \times K^{(n)}(\tau_1, \ldots, \tau_n)\, f(\tau_1) \cdots f(\tau_n), \tag{10}$$

where the kernels $K^{(n)}$ are the *connected* (cumulant) bath correlators, defined via [46]

$$K^{(n)}(\tau_1, \ldots, \tau_n) \equiv \big\langle \mathcal{T}_\tau B(\tau_1) \cdots B(\tau_n) \big\rangle_c. \tag{11}$$

The explicit connection back to Eq. (8) is given via its functional differentiation:

$$\big\langle \mathcal{T}_\tau B(\tau_1) \cdots B(\tau_n) \big\rangle_c = (-1)^n \left. \frac{\delta^n \log \mathcal{Z}_X[j]}{\delta j(\tau_1) \cdots \delta j(\tau_n)} \right|_{j=0}. \tag{12}$$

The bath influence is then completely characterised by the hierarchy $\{K^{(n)}\}_{n\geq 1}$.

To connect the influence functional back to the quenched density, we use the fact that the bath influence depends on the system history only through the linear functional $\int_0^\beta d\tau\, f(\tau) B(\tau)$. One may therefore represent $\mathcal{F}[f]$ as the (generalised) characteristic functional of an auxiliary field $\xi(\tau)$ [38–40]:

$$\mathcal{F}[f] = \mathbb{E}_\xi \left[\exp\left(-\int_0^\beta d\tau\, \xi(\tau)\, f(\tau) \right) \right], \tag{13}$$

where $\mathbb{E}_\xi[\cdot]$ denotes averaging with respect to a (possibly complex) measure on ξ-histories chosen such that Eq. (13) holds.

Since $\xi(\tau)$ is a commuting c-number field, its n-point moments are symmetric under permutations of the time arguments. The influence kernels $K^{(n)}$ fix the *cumulants* of ξ via

$$\big\langle \xi(\tau_1) \cdots \xi(\tau_n) \big\rangle_c = K^{(n)}(\tau_1, \ldots, \tau_n). \tag{14}$$

Consequently, the ordinary correlation functions of the noise are obtained from $\{K^{(m)}\}$ by the standard moment-cumulant relations:

$$\big\langle \xi(\tau_1) \cdots \xi(\tau_n) \big\rangle = \sum_{\pi \in \mathcal{P}_n} \prod_{C \in \pi} K^{(|C|)}\big(\{\tau_i\}_{i\in C}\big), \tag{15}$$

where $\mathcal{P}_n$ denotes the set of all partitions of the index set $\{1, \ldots, n\}$, and C are the disjoint blocks of a given partition $\pi \in \mathcal{P}_n$. After shifting the mean so that $K^{(1)}(\tau) = \langle \xi(\tau) \rangle = 0$, one has (for example)

$$\big\langle \xi(\tau_1)\xi(\tau_2) \big\rangle = K^{(2)}(\tau_1, \tau_2), \tag{16}$$

$$\big\langle \xi(\tau_1)\xi(\tau_2)\xi(\tau_3) \big\rangle = K^{(3)}(\tau_1, \tau_2, \tau_3), \tag{17}$$

$$\begin{aligned} \big\langle \xi(\tau_1)\xi(\tau_2)\xi(\tau_3)\xi(\tau_4) \big\rangle &= K^{(4)}(\tau_1, \tau_2, \tau_3, \tau_4) \\ &\quad + K^{(2)}(\tau_1, \tau_2) K^{(2)}(\tau_3, \tau_4) \\ &\quad + K^{(2)}(\tau_1, \tau_3) K^{(2)}(\tau_2, \tau_4) \\ &\quad + K^{(2)}(\tau_1, \tau_4) K^{(2)}(\tau_2, \tau_3). \end{aligned} \tag{18}$$

QUENCHED DENSITY AND THE HAMILTONIAN OF MEAN FORCE

The Hamiltonian of mean force is, by definition, the operator whose Gibbs form reproduces the system's reduced equilibrium state. Equivalently, it is the *operator logarithm* of the unnormalised reduced equilibrium operator

$$\bar{\rho}_Q(\beta) \equiv \mathrm{Tr}_X\, e^{-\beta H_{\mathrm{tot}}}, \qquad H_{\mathrm{MF}}(\beta) \;\equiv\; -\frac{1}{\beta} \log \bar{\rho}_Q(\beta), \tag{19}$$

defined up to an additive multiple of the identity (fixed only when normalising the state).

The quenched representation in Eq. (5) supplies an exact stochastic parametrisation of the unnormalised mean-force Gibbs operator $\bar{\rho}_Q(\beta)$ by making the bath trace an explicit average over imaginary-time back-action histories. Combining Eq. (5) with Eq. (19) yields

$$H_{\mathrm{MF}}(\beta) = -\frac{1}{\beta} \log \mathbb{E}_\xi [U_\xi(\beta)]\,. \tag{20}$$

Thus, constructing the mean-force Hamiltonian reduces to evaluating a stochastic average and then compressing the result via an operator logarithm. The conditions for being able to perform this average exactly is the focus of the present work.

To make the handling of this problem more concrete, we shall specialise the environment to the Caldeira-Leggett model, where the bath is a collection of harmonic oscillators ($H_X = \sum_k \omega_k b_k^\dagger b_k$) and the coupling is linear in bath coordinates ($B = \sum_k c_k x_k$). None of the results that follow are essentially dependent on this choice, and a generalisation to anharmonic environments is (relatively) straightforward. In the interests of comprehensibility however, we restrict our scope to quadratic environments. In this setting, the auxiliary field $\xi(\tau)$ becomes a stationary zero-mean Gaussian process completely characterized by its covariance

$$\mathbb{E}_\xi[\xi(\tau)\xi(\tau')] = K(\tau - \tau'). \tag{21}$$

The kernel $K(\tau)$ is determined by the bath spectral density $J(\omega) = \frac{\pi}{2}\sum_k \frac{c_k^2}{m_k\omega_k}\delta(\omega-\omega_k)$ via the relation [47]

$$K(\tau) = \frac{1}{\pi}\int_0^\infty d\omega\, J(\omega)\, \frac{\cosh\big(\omega(\beta/2-|\tau|)\big)}{\sinh(\beta\omega/2)}. \tag{22}$$

A key property of this kernel is its integrated strength. Integrating Eq. (22) yields

$$\begin{aligned}\int_0^\beta d\tau\, K(\tau) &= \frac{1}{\pi}\int_0^\infty d\omega\, J(\omega)\int_0^\beta d\tau\, \frac{\cosh\big(\omega(\beta/2-|\tau|)\big)}{\sinh(\beta\omega/2)}\\ &= \frac{1}{\pi}\int_0^\infty d\omega\, J(\omega)\,\frac{2}{\omega}\\ &= 2\lambda, \end{aligned}\tag{23}$$

where λ is the explicit reorganisation energy. Consequently, the total variance of the integrated noise field $\Xi = \int_0^\beta d\tau\, \xi(\tau)$ grows linearly with inverse temperature:

$$\mathbb{E}_\xi[\Xi^2] = \int_0^\beta d\tau \int_0^\beta d\tau'\, K(\tau-\tau') = 2\beta\lambda. \tag{24}$$

In the case that the system and its coupling commute, $[H_Q, f] = 0$ and f is τ-independent in imaginary time. In this instance time ordering drops out, and the average is given by [41]:

$$\bar{\rho}_Q(\beta) = \exp\left[-\beta\left(H_Q - \frac{\kappa_0(\beta)}{2}f^2\right)\right], \tag{25}$$

which in turn yields

$$H_{\rm MF}(\beta) = H_Q - \frac{\kappa_0(\beta)}{2}f^2 + \frac{1}{\beta}\log Z_X(\beta)\,\mathbb{I}. \tag{26}$$

Notably this correction to $H_{\rm MF}(\beta)$ is entirely *classical* [41]. This is hardly surprising, but emphasises that quantum effects *must* stem from non-commutativity. A truly quantum theory of thermodynamics therefore only makes meaningfully distinct predictions when $[H_Q, f] \neq 0$. In this case however, the noise enters through a noncommuting operator inside $\mathcal{T}_\tau$, rendering the question of averaging highly non-trivial. In the next section, we attack this problem directly, deriving conditions under which $H_{\rm MF}(\beta)$ possesses a closed form.

GENERATING-FUNCTION STRUCTURE OF THE INFLUENCE OPERATOR

Quenched propagator and Gaussian resummation

The starting point is the quenched propagator introduced above. In the imaginary-time interaction picture with respect to H_Q, the unnormalised reduced equilibrium operator takes the form

$$\begin{aligned}&\bar{\rho}_Q(\beta) = e^{-\beta H_Q}\,\langle W_\xi(\beta)\rangle_\xi,\\ &W_\xi(\beta) \equiv \mathcal{T}_\tau \exp\left[-\int_0^\beta d\tau\, \xi(\tau)\, \tilde{f}(\tau)\right],\end{aligned}\tag{27}$$

where $\tilde{f}(\tau) = e^{\tau H_Q} f e^{-\tau H_Q}$ and the noise covariance is $\langle \xi(\tau)\xi(\tau')\rangle_\xi = K(\tau-\tau')$.

Because the bath is Gaussian, only the second cumulant of ξ is non-vanishing. Time-ordering introduces a non-trivial operator structure, but the key identity—proved in Appendix A—is that the Wick resummation of all pairings exponentiates exactly:

$$\langle W_\xi(\beta)\rangle_\xi = \exp\left[\frac{1}{2}\int_0^\beta d\tau \int_0^\beta d\tau'\, K(\tau-\tau')\,\mathcal{T}_\tau\big[\tilde{f}(\tau)\tilde{f}(\tau')\big]\right]. \tag{28}$$

The proof is a direct application of Wick's theorem to the Dyson expansion of W_ξ: at order $2p$ in the expansion, there are $(2p)!/(2^p p!)$ distinct Wick pairings, each weighted by the two-point kernel $K(\tau_i - \tau_j)$. The prefactors from the Dyson expansion precisely cancel the pairing multiplicity, leaving only the singly-connected pairs.

We can exploit the symmetry $K(\tau-\tau') = K(\tau'-\tau)$ to fold the square domain onto the triangular one (see Appendix A), such that the exponent simplifies to

$$\Delta(\beta) \equiv \int_0^\beta d\tau \int_0^\tau d\tau'\, K(\tau-\tau')\,\tilde{f}(\tau)\tilde{f}(\tau'), \tag{29}$$

and the averaged reduced state takes the product form

$$\bar{\rho}_Q(\beta) = e^{-\beta H_Q}\, e^{\Delta(\beta)}. \tag{30}$$

The influence exponent $\Delta(\beta)$ is quadratic in $\tilde{f}$ and depends on the bath exclusively through the two-point kernel K. Here the absence of higher-order bath objects is not an approximation, but an exact consequence of the Gaussian nature of the environment.

The ordered Green function and frequency compression

At this stage, the operator Δ is still formally defined as an integral over the bath degrees of freedom. There is however a particularly satisfying interpretation of this operator in which the effect of the bath is readily understood. To make this explicit, we work in the eigenbasis of H_Q. Let $H_Q|i\rangle = E_i|i\rangle$, so that the matrix elements of the interaction-picture operator are

$$\tilde{f}_{ij}(\tau) = \langle i|\tilde{f}(\tau)|j\rangle = f_{ij}\, e^{\omega_{ij}\tau}, \qquad \omega_{ij} \equiv E_i - E_j. \tag{31}$$

Substituting into Eq. (29) and expanding in the eigenbasis, the influence exponent takes the spectral form

$$\Delta(\beta) = \sum_{i,\ell} \Delta_{i\ell}\, |i\rangle\langle \ell|, \tag{32}$$

where evaluating the matrix elements of the operator product yields

$$\Delta_{i\ell} = \sum_j f_{ij}\, f_{j\ell}\, G^{>}(\omega_{ij}, \omega_{j\ell}), \tag{33}$$

where all bath and temperature dependence has been absorbed into the *ordered Green function*

$$G^{>}(\omega_1,\omega_2) \equiv \int_0^\beta d\tau \int_0^\tau d\tau'\, K(\tau-\tau')\, e^{\omega_1\tau+\omega_2\tau'}. \quad (34)$$

In standard finite-temperature perturbation theory, analogous time-ordered integrals are evaluated using discrete Matsubara frequencies. Here, the integration instead projects the bath response directly onto the Laplace frequencies set by the system's own energy differences; the superscript $>$ records that $\tau > \tau'$ throughout the integration domain.

The double integral in Eq. (34) collapses to a one-dimensional Laplace transform upon changing variables to $u = \tau - \tau'$ and $v = \tau'$:

$$G^{>}(\omega_1,\omega_2) = \frac{e^{\beta\omega_1}\mathcal{K}(\omega_2) - \mathcal{K}(\omega_1)}{\omega_1+\omega_2}, \quad (35)$$

where

$$\mathcal{K}(\omega) \equiv \int_0^\beta du\, K(u)\, e^{\omega u} \quad (36)$$

is the bare Laplace transform of the bath kernel. Eq. (35) is an exact closed form valid for all $\omega_1+\omega_2 \neq 0$. Importantly, it implies that for any transition between levels in the reduced system, the entirety of the bath's influence is encoded by $\mathcal{K}$ evaluated at the transition frequency.

The resonant channel $\omega_1+\omega_2 = 0$ requires separate treatment. Taking the limit carefully yields

$$R(\omega) \equiv G^{>}(\omega,-\omega) = \int_0^\beta (\beta-u)\, K(u)\, e^{\omega u}\, du, \quad (37)$$

where the factor $(\beta-u)$ is the available imaginary time after a bath correlation at separation u; the resonant kernel thus accumulates bath memory over the full thermal interval, regulated by the finiteness of β.

The structure of Eq. (33) admits a pleasingly simple diagrammatic reading, illustrated in Fig. 1. To isolate the physically distinct effects of the bath, we partition the influence exponent into its secular (diagonal) and mixing (off-diagonal) components, $\Delta = \Delta_\parallel + \Delta_\perp$.

The secular sector $\Delta_\parallel \equiv \sum_i \Delta_{ii}|i\rangle\langle i|$ singles out the *closed loops* $i \to j \to i$, for which $\omega_{ij}+\omega_{ji} = 0$ identically. The resonance condition forces the bath propagator onto $R(\omega_{ij})$, the object that accumulates bath memory over the whole thermal interval. The diagonal elements are thus the exact, non-perturbative *self-energy* dressing of the i-th energy level by all levels j it is coupled to:

$$\Delta_{ii} = \sum_j |f_{ij}|^2\, R(\omega_{ij}). \quad (38)$$

Conversely, the mixing sector $\Delta_\perp \equiv \sum_{i\neq\ell} \Delta_{i\ell}|i\rangle\langle\ell|$ describes *open chains* $i \to j \to \ell$, where energy is not conserved at the intermediate vertex. Using the closed form

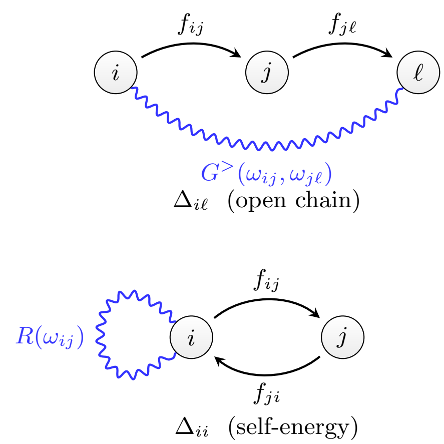

FIG. 1. Diagrammatic structure of the influence exponent $\Delta(\beta)$ in the H_Q eigenbasis. *Top*: The off-diagonal element $\Delta_{i\ell}$ (open chain) sums intermediate sites j, weighted by the bath propagator $G^{>}(\omega_{ij},\omega_{j\ell})$ (wavy line) and two coupling insertions f (arrows). Energy is not conserved at the intermediate vertex. *Bottom*: The diagonal element Δ_{ii} (self-energy) is a closed loop $i\to j\to i$; the resonant condition $\omega_1+\omega_2 = 0$ is satisfied, and the bath propagator collapses to $R(\omega_{ij})$.

in Eq. (35), these components can be expressed directly in terms of the Laplace-transformed bath kernel $\mathcal{K}(\omega)$:

$$\Delta_{i\ell} = \frac{1}{\omega_{i\ell}} \sum_j f_{ij} f_{j\ell} \left[e^{\beta\omega_{ij}}\mathcal{K}(\omega_{j\ell}) - \mathcal{K}(\omega_{ij})\right] \quad (i \neq \ell), \quad (39)$$

which encodes the coherent transitions between distinct energy levels mediated by the asymmetric fluctuations of the bath.

KMS symmetry and the symmetrised reduced state

The bath kernel $K(\tau)$ satisfies the KMS periodicity condition $K(\tau) = K(\beta-\tau)$, which in Laplace space reads

$$\mathcal{K}(-\omega) = e^{-\beta\omega}\,\mathcal{K}(\omega). \quad (40)$$

Substituting into the closed form Eq. (35) and simplifying gives a corresponding symmetry of the ordered Green function:

$$G^{>}(-\omega_2,-\omega_1) = e^{-\beta(\omega_1+\omega_2)}\, G^{>}(\omega_1,\omega_2). \quad (41)$$

Because transition frequencies telescope, $\omega_{ij}+\omega_{j\ell} = \omega_{i\ell}$, applying Eq. (41) to the matrix-element formula Eq. (33) immediately yields the *detailed-balance* constraint

$$\Delta_{\ell i} = e^{-\beta\omega_{i\ell}}\,\Delta^*_{i\ell}, \quad (42)$$

or in operator form $\Delta^\dagger = \Pi\,\Delta\,\Pi^{-1}$ where $\Pi \equiv e^{-\beta H_Q}$. Although Δ itself is not self-adjoint, this relation is a precise constraint: the product $\Pi\Delta$ is Hermitian.

This symmetry allows us to recast the unnormalised reduced state Eq. (30) into a manifestly positive, Hermitian form. Writing the bare Gibbs factor as $\Pi \equiv e^{-\beta H_Q}$ gives $\bar{\rho}_Q = \Pi\, e^{\Delta}$. By splitting $\Pi = \Pi^{1/2}\Pi^{1/2}$ and inserting the identity $\Pi^{-1/2}\Pi^{1/2} = \mathbf{1}$, we symmetrically distribute the bare weight:

$$\bar{\rho}_Q = \Pi^{1/2}\big(\Pi^{1/2}\, e^{\Delta}\, \Pi^{-1/2}\big)\Pi^{1/2}. \tag{43}$$

Applying the operator identity $A\, e^{B} A^{-1} = \exp(ABA^{-1})$ to the inner bracket yields

$$\bar{\rho}_Q = \Pi^{1/2}\, e^{S}\, \Pi^{1/2}, \qquad S \;\equiv\; \Pi^{1/2}\, \Delta\, \Pi^{-1/2}, \tag{44}$$

where S is the *symmetrised influence operator*. Its matrix elements evaluate to $S_{i\ell} = e^{-\beta\omega_{i\ell}/2}\Delta_{i\ell}$, which by Eq. (42) satisfy $S^*_{i\ell} = S_{\ell i}$. Thus, the modified exponent S is strictly self-adjoint. The partition function of the reduced system follows by cyclicity of the trace:

$$Z_Q \;\equiv\; \mathrm{Tr}[\bar{\rho}_Q] = \mathrm{Tr}[\Pi\, e^{S}] = \sum_i e^{-\beta E_i}\, (e^{S})_{ii}. \tag{45}$$

Since S is self-adjoint, $(e^{S})_{ii} \geq e^{S_{ii}} = e^{\Delta_{ii}}$ by Jensen's inequality applied to the spectral weights $|\langle i|n\rangle|^2$, giving the bound

$$Z_Q \;\geq\; \sum_i e^{-\beta E_i + \Delta_{ii}}, \tag{46}$$

where the right-hand side is a Boltzmann sum with bath-shifted energies and equality holds when S is diagonal (i.e. when H_Q and Δ share an eigenbasis). The full result $Z_Q = \mathrm{Tr}[\Pi e^{S}]$ captures all level-mixing corrections through the off-diagonal elements of S.

CLOSURE OF THE HAMILTONIAN OF MEAN FORCE

The obstruction to writing $H_{\mathrm{MF}}(\beta)$ in a compact operator form is not the existence of the mean-force object (it is defined by a logarithm), but the *representability* of that logarithm inside a restricted operator family (few-body, local, Pauli strings, etc.). In the harmonic case this representability question reduces to a precise closure problem. To show this, we first write the quenched propagator in the imaginary-time interaction picture with respect to H_Q:

$$\begin{aligned} U_\xi(\beta) &= e^{-\beta H_Q} W_\xi(\beta), \\ W_\xi(\beta) &\equiv \mathcal{T}_\tau \exp\left[-\int_0^\beta d\tau\, \xi(\tau)\, \tilde{f}(\tau)\right], \end{aligned} \tag{47}$$

where $\tilde{f}(\tau) \equiv e^{\tau H_Q} f e^{-\tau H_Q}$.

Because the noise is classical Gaussian with zero mean, the average of the time-ordered exponential resums exactly in terms of the second cumulant. Time ordering adds a non-trivial complication, which is addressed in Appendix A. The result may be succinctly expressed in terms of an *influence operator*

$$\begin{aligned} \langle W_\xi(\beta)\rangle_\xi &\equiv \bar{W}(\beta) = \exp\big(\Delta(\beta)\big), \\ \bar{\rho}_Q(\beta) &= \langle U_\xi(\beta)\rangle_\xi = e^{-\beta H_Q}\, e^{\Delta(\beta)}, \end{aligned} \tag{48}$$

where the influence exponent $\Delta(\beta)$ is defined in accordance with Eq. (A12). It is convenient to perform a final simplification of this exponent to remove the explicit time ordering, the details of which we relegate to Appendix A. After this manipulation, $\Delta(\beta)$ takes the form

$$\Delta(\beta) \equiv \int_0^\beta d\tau \int_0^\tau d\tau'\, K(\tau - \tau')\, \tilde{f}(\tau)\tilde{f}(\tau'). \tag{49}$$

To progress, we introduce the adjoint chain generated by H_Q acting on the coupling,

$$f_n \equiv \mathrm{ad}^n_{H_Q}(f), \qquad \mathrm{ad}_{H_Q}(A) \equiv [H_Q, A], \tag{50}$$

so that the interaction-picture operator has the exact series

$$\tilde{f}(\tau) = e^{\tau H_Q} f e^{-\tau H_Q} = \sum_{n\geq 0} \frac{\tau^n}{n!}\, f_n. \tag{51}$$

Substituting Eq. (51) into Eq. (49) yields

$$\Delta(\beta) = \sum_{n,m\geq 0} C^{>}_{nm}(\beta) f_n f_m, \tag{52}$$

where we have further defined an *ordered kernel-moment matrix*

$$C^{>}_{nm}(\beta) \equiv \int_0^\beta d\tau \int_0^\tau d\tau'\, K(\tau - \tau')\, \frac{\tau^n}{n!}\, \frac{\tau'^m}{m!}. \tag{53}$$

Here the notation $>$ is used to track the triangular domain of integration, where τ (associated with n) is the upper limit of integration for τ' (associated with m).

With this definition, all bath/temperature dependence enters $\Delta(\beta)$ only through the scalar moments $C^{>}_{nm}(\beta)$, while all operator structure is carried by the adjoint chain $\{f_n\}$. The final step is to re-express the averaged propagator in a *single* exponential. Using Eq. (48) we write

$$\bar{\rho}_Q(\beta) = e^{-\beta H_Q}\, e^{\Delta(\beta)} \equiv e^{-\beta H_{\mathrm{MF}}(\beta)}, \tag{54}$$

where the mean-force Hamiltonian satisfies

$$-\beta H_{\mathrm{MF}}(\beta) = \log\Big(e^{-\beta H_Q} e^{\Delta(\beta)}\Big). \tag{55}$$

Let $A \equiv -\beta H_Q$ and $B \equiv \Delta(\beta)$. The Baker–Campbell–Hausdorff (BCH) series [48–50] gives

$$\begin{aligned} \log(e^{A} e^{B}) &= A + B + \frac{1}{2}[A, B] \\ &\quad + \frac{1}{12}[A, [A, B]] + \frac{1}{12}[B, [B, A]] + \cdots \end{aligned} \tag{56}$$

Dividing by $-\beta$ and using $A = -\beta H_Q$ yields the expansion

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\Delta(\beta) + \frac{1}{2}[H_Q, \Delta(\beta)] - \frac{\beta}{12}[H_Q, [H_Q, \Delta(\beta)]] + \frac{1}{12}[\Delta(\beta), [\Delta(\beta), H_Q]] + \cdots . \tag{57}$$

Expressed in this way, the closure content of $H_{\rm MF}(\beta)$ is transparent: all terms are built from repeated commutators of H_Q acting on Δ, together with higher BCH terms involving commutators among those objects.

To see the explicit generator, introduce the adjoint superoperator $\mathrm{ad}_{H_Q}(\cdot) \equiv [H_Q, \cdot]$. Then Eq. (57) becomes

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\Delta + \frac{1}{2}\,\mathrm{ad}_{H_Q}\Delta - \frac{\beta}{12}\,\mathrm{ad}^2_{H_Q}\Delta + \frac{1}{12}\,[\Delta, \mathrm{ad}_{H_Q}\Delta] + \cdots . \tag{58}$$

Finally, because $\Delta(\beta)$ is quadratic in the adjoint chain, each commutator with H_Q simply shifts indices:

$$\mathrm{ad}_{H_Q}(f_n f_m) = [H_Q, f_n f_m] = f_{n+1} f_m + f_n f_{m+1}. \tag{59}$$

Hence the first commutator term takes the explicit form

$$[H_Q, \Delta(\beta)] = \sum_{n,m\geq 0} C^{>}_{nm}(\beta)\,\big(f_{n+1} f_m + f_n f_{m+1}\big), \tag{60}$$

and higher $\mathrm{ad}^k_{H_Q}\Delta$ generate the same family of products $f_a f_b$ with shifted indices. The remaining BCH terms (e.g. $[\Delta, \mathrm{ad}_{H_Q}\Delta]$) introduce commutators among these quadratic elements, and are the sole source of additional operator growth beyond the quadratic span. The only remaining question is whether the operator family generated by Eq. (58) can be brought to a closed form.

The preceding identities imply that the entire BCH commutator tower *linear* in Δ remains within the quadratic span $\mathrm{span}\{f_n f_m\}$. Terms to this order explicitly admit a closed resummation. Let $A \equiv -\beta H_Q$ and $B \equiv \Delta(\beta)$. Then the BCH logarithm satisfies the standard linear-in-B identity [51, 52]

$$\log\big(e^A e^B\big) = A + \Phi\big(\mathrm{ad}_A\big)\,B + O(B^2), \qquad \Phi(x) \equiv \frac{x}{1 - e^{-x}}. \tag{61}$$

Because $A = -\beta H_Q$, the Todd/Bernoulli generating function must be evaluated at the negative superoperator argument; this sign is what reproduces the explicit BCH coefficients in Eq. (57). Substituting $A = -\beta H_Q$ and dividing by $-\beta$ gives the mean-force Hamiltonian to first order in Δ:

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\,\Phi\big(-\beta\,\mathrm{ad}_{H_Q}\big)\,\Delta(\beta) + O(\Delta^2). \tag{62}$$

Expanding Φ yields the explicit commutator series with Bernoulli numbers B_k,

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\sum_{k\geq 0}\frac{B_k}{k!}\beta^k\,\mathrm{ad}^k_{H_Q}\Delta(\beta) + O(\Delta^2), \qquad \Phi(x) = \sum_{k\geq 0}\frac{B_k}{k!}(-x)^k = 1 + \frac{x}{2} + \frac{x^2}{12} - \frac{x^4}{720} + \cdots \tag{63}$$

Using Eq. (59) repeatedly, the iterated commutators admit the closed binomial form

$$\mathrm{ad}^k_{H_Q}(f_n f_m) = \sum_{j=0}^{k}\binom{k}{j} f_{n+j}\, f_{m+k-j}, \tag{64}$$

so every term in Eq. (63) is explicitly a linear combination of products $f_a f_b$ with scalar coefficients determined solely by $C^{>}_{nm}(\beta)$ and universal combinatorics. Because of this, we may bring the first order contribution into a particularly compact form by introducing the discrete shift operator D. This will act on the moment matrix $C^{>}$ as

$$(\mathsf{D}C^{>})_{ab} \equiv C^{>}_{a-1,b} + C^{>}_{a,b-1}, \qquad C^{>}_{rs} \equiv 0 \quad \text{if any index is negative.} \tag{65}$$

Then Eq. (59) implies, after reindexing,

$$\mathrm{ad}_{H_Q}\Delta(\beta) = \sum_{a,b\geq 0}(\mathsf{D}C^{>}(\beta))_{ab}\, f_a f_b, \tag{66}$$

and by iteration,

$$\mathrm{ad}^k_{H_Q}\Delta(\beta) = \sum_{a,b\geq 0}(\mathsf{D}^k C^{>}(\beta))_{ab}\, f_a f_b. \tag{67}$$

Equivalently, $(\mathsf{D}^k C^{>})_{ab} = \sum_{j=0}^{k}\binom{k}{j} C^{>}_{a-j,\,b-(k-j)}$, which is the binomial index-shift structure implied by Eq. (64). From this we obtain the following form for $H_{\rm MF}(\beta)$ to linear order in Δ:

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\sum_{a,b\geq 0}\widetilde{C}^{>}_{ab}(\beta)\, f_a f_b + O(\Delta^2), \tag{68}$$

where the *renormalised moment matrix* $\widetilde{C}^{>}(\beta)$ is the universal transform

$$\widetilde{C}^{>}(\beta) \equiv \Phi\big(-\beta\,\mathsf{D}\big)\, C^{>}(\beta) = \sum_{k\geq 0}\frac{B_k}{k!}\beta^k\,\mathsf{D}^k C^{>}(\beta). \tag{69}$$

Thus, to linear order in the influence operator, the mean-force Hamiltonian is *exactly* a quadratic form in the products $f_a f_b$, with all bath/temperature dependence entering only through the scalar matrix $\widetilde{C}^{>}_{ab}(\beta)$. Any operator growth beyond the quadratic span is confined to the nonlinear BCH sector $O(\Delta^2)$ (e.g. $[\Delta, \mathrm{ad}_{H_Q}\Delta]$). While these terms are more difficult to handle as they involve commutators of quadratic forms, they nevertheless admit a precise series expansion via the full BCH formula Eq. (56). With this form in hand, we can now state a theorem for the representability of $H_{\mathrm{MF}}(\beta)$ in a given operator family $\mathcal{A}$ as a *closure criterion.*

Theorem: Finite representability of mean force

Fix a target operator family $\mathcal{A}$ (e.g. k-local Pauli strings, a Lie algebra plus identity, or a finite-dimensional associative operator algebra) in which we seek to represent $H_{\mathrm{MF}}(\beta)$. The Gaussian construction above allows us to formulate the exact criteria for representability as a closure theorem. In a physical sense, these conditions determine whether a valid *local generator* can be defined for the system at all; failure of closure implies that the effective Hamiltonian is fundamentally non-local with respect to the algebra $\mathcal{A}$.

Theorem (Closure of the Mean Force). *Let $\mathcal{A}$ be a target operator space containing the bare system Hamiltonian H_Q and coupling f. The Hamiltonian of Mean Force $H_{\mathrm{MF}}(\beta)$ admits a finite representation within $\mathcal{A}$ if and only if the following three closure conditions are satisfied:*

(C1) Adjoint Closure (No linear operator growth). The subspace generated by repeated commutators of the coupling f with H_Q must remain inside $\mathcal{A}$:

$$f \in \mathcal{A}, \qquad \mathrm{ad}_{H_Q}(\mathcal{A}) \subseteq \mathcal{A}, \\ \text{equivalently} \quad \mathrm{span}\{f_n\}_{n\geq 0} \subseteq \mathcal{A}. \tag{70}$$

(C2) Quadratic Closure (No growth in the Gaussian sector). Because the exact influence operator $\Delta(\beta)$ is quadratic in the adjoint chain, the quadratic span generated by $\{f_n\}$ must also lie in $\mathcal{A}$:

$$f_n f_m \in \mathcal{A} \quad \text{for all indices that contribute in } \Delta(\beta), \tag{71}$$

so that $\Delta(\beta) \in \mathcal{A}$. A sufficient (and common) condition is that $\mathcal{A}$ is an associative algebra containing the identity and closed under multiplication.

(C3) BCH Closure (No growth from commutators among quadratic elements). To ensure that the *full* BCH logarithm $\log(e^{-\beta H_Q} e^{\Delta})$ remains in $\mathcal{A}$, the commutators generated among the quadratic elements must also close in $\mathcal{A}$; in particular

$$[\Delta(\beta), \mathrm{ad}^k_{H_Q}\Delta(\beta)] \in \mathcal{A} \qquad \text{for all } k \geq 0, \tag{72}$$

and similarly for the higher nested commutators appearing in the BCH series. A sufficient (and common) condition is again that $\mathcal{A}$ is a finite-dimensional associative algebra (or a Lie algebra containing Δ and closed under commutators), in which case all BCH commutators remain in $\mathcal{A}$ by construction.

Consequence (closed-form mean-force Hamiltonian). If (C1)–(C3) hold, then $\Delta(\beta) \in \mathcal{A}$ and all BCH commutators generated by Eq. (73) remain in $\mathcal{A}$. Since $H_Q \in \mathcal{A}$ by assumption, it follows that

$$\bar{\rho}_Q(\beta) = e^{-\beta H_{\mathrm{MF}}(\beta)}, \\ -\beta H_{\mathrm{MF}}(\beta) = \log\Big(e^{-\beta H_Q} e^{\Delta(\beta)}\Big), \tag{73}$$

defines an $H_{\mathrm{MF}}(\beta) \in \mathcal{A}$ and hence a closed-form representation of $H_{\mathrm{MF}}(\beta)$ (up to the usual additive scalar fixed by Z_X). If either (C1) fails (adjoint-chain growth) or (C2)–(C3) fail (growth in the quadratic/BCH sector), then $\Delta(\beta)$ and/or the BCH commutator tower generates operators outside $\mathcal{A}$; any representation within $\mathcal{A}$ is necessarily truncated or projected.

This statement leads to an important corollary. For continuous variable systems $H_Q = p^2/2m + V(x)$, if the potential $V(x)$ is a polynomial of degree greater than 2, the adjoint chain $\{f_n\}$ generally does not close in any finite-dimensional Lie algebra; instead, it generates the full universal enveloping algebra. Consequently, no finite representation of the mean force exists within the original operator class. This obstruction is analogous (in spirit) to the Groenewold–Van Hove no-go theorem on consistent quantisation of polynomial observables: the failure to close the Lie algebra under adjoint action corresponds to the impossibility of mapping the quantum evolution of non-quadratic potentials to a finite phase space flow. For such potentials, there is no *local operator* that exactly describes the reduced equilibrium state.

From a practical standpoint, the closure of $\mathcal{A}$ is not necessary for an arbitrarily accurate approximation of $H_{\mathrm{MF}}(\beta)$, as the BCH logarithm $\log(e^{-\beta H_Q} e^{\Delta})$ still yields a controlled expansion. The operator content is generated by the adjoint chain and its quadratic products, while all bath dependence remains scalar through $C^{>}_{nm}(\beta)$. One may therefore truncate (in commutator depth, locality class, or operator weight) without introducing any approximation on the bath side.

Approximate methods as choices of operator class

This algebraic perspective also clarifies approximate open-system methods. Numerical schemes such as the polaron transformation [25, 47], hierarchical equations of motion [32], and chain-mapping/DMRG techniques [34] may be understood as distinct choices of the target operator family $\mathcal{A}$ or distinct rules for enlarging it. The polaron frame dresses $\mathcal{A}$ with coherent displacements; HEOM truncates the memory depth of the influence functional while mirroring operator growth in the adjoint chain; and chain mappings truncate the spatial extent of the harmonic bath. The closure problem is therefore equivalent to identifying a reduced model class that is large enough to contain the logarithm of the quenched propagator, yet small enough to remain usable.

Constructive reduced generators and compact operator classes

At the level of the global Gibbs operator, the equilibrium problem is linear: one exponentiates $H_{\rm tot}$ and traces over the bath. The difficulty appears only after reduction. The partial trace removes environmental degrees of freedom explicitly, but the subsystem generator is recovered only after the nonlinear compression $H_{\rm MF}(\beta) = -(1/\beta)\log\bar{\rho}_Q(\beta)$. The HMF is therefore not merely another name for the reduced state. It is the constructive object that tells us whether the reduced description can still be written inside a tractable operator family.

Viewed this way, Conditions (C1)–(C3) are closure conditions on a target operator class. Adjoint closure asks whether repeated system–bath back-action remains inside the chosen family. Quadratic closure asks whether the Gaussian influence can be expressed in that same family. BCH closure asks whether recombining the reduced state with the bare Gibbs factor introduces genuinely new operators or stays compact. When all three hold, the reduced equilibrium admits an explicit generator that is exact, finite, and inspectable.

The failure modes are equally informative. If the adjoint chain proliferates without closure, then the reduced generator cannot be compressed into the proposed ansatz without loss. This is the operator-level version of a representability bottleneck: the hidden environment has injected structure that the target model class cannot faithfully contain. Approximate methods then differ less by the bath data they use than by the restricted operator families they choose to keep. The same structural question also appears in other settings where one asks for compact decompositions rather than implicit outputs alone. The next subsection makes that statement explicit by writing the Gaussian HMF itself as a finite compositional architecture, while keeping the present result firmly at the level of exact reduced equilibrium physics.

Compositional structure of the reduced generator

The Kolmogorov–Arnold theorem asks when a multivariate map can be written as a finite composition of univariate functions and addition; in the notation used by Kolmogorov–Arnold networks one seeks representations of the form $F(\mathbf{x}) = \sum_q \Phi_q(\sum_p \phi_{q,p}(x_p))$, with the univariate edge functions then learned inside a chosen basis class [14]. The relevant point for the present paper is that the exact Gaussian HMF already comes with its own compositional architecture. The KAN comparison therefore does not need to be metaphorical.

Proposition 1 *For a Gaussian bath linearly coupled through f, the exact Hamiltonian of mean force is generated by three explicit compositional layers determined entirely by (H_Q, f, K, β). First, the univariate operator orbit is the adjoint chain $f_n = \mathrm{ad}^n_{H_Q}(f)$. Second, the influence exponent is the bilinear aggregate*

$$\Delta(\beta) = \sum_{n,m\geq 0} C^{>}_{nm}(\beta)\, f_n f_m. \tag{74}$$

Third, the reduced generator is obtained by nonlinear BCH recombination of that aggregate with the bare Gibbs factor,

$$H_{\rm MF}(\beta) = H_Q - \frac{1}{\beta}\Phi\big(-\beta\,\mathrm{ad}_{H_Q}\big)\Delta(\beta) + O(\Delta^2), \tag{75}$$

with the higher-order terms supplied by the full nested BCH series. Conditions (C1)–(C3) of the closure theorem are exactly the layerwise conditions that keep the adjoint orbit, the bilinear aggregate, and the nonlinear BCH tower inside a chosen operator ansatz $\mathcal{A}$.

Figure 2 makes this three-layer correspondence explicit and shows the spin-boson qubit as the worked case in which each open-ended layer terminates into a closed finite structure.

This is more than a family resemblance. The Gaussian HMF construction already provides the three ingredients usually separated in a compact compositional model: basis atoms, an aggregation tensor, and a nonlinear composition rule. The adjoint orbit is a discrete univariate basis in operator space: one generator, one seed, one orbit index n. The matrix $C^{>}_{nm}(\beta)$ is the aggregation tensor, but here it is fixed analytically by the bath kernel rather than learned from data. The Todd/Bernoulli superoperator $\Phi(-\beta\,\mathrm{ad}_{H_Q})$, together with the higher BCH nests, supplies the nonlinear composition rule. Thermodynamics therefore provides both the architecture and the weights.

Once written in this form, representability failure becomes a precise statement about model-class insufficiency. If the adjoint orbit or its quadratic and BCH closures escape $\mathcal{A}$, the exact reduced generator still exists but no finite ansatz in that class can carry it without projection. For non-quadratic continuous-variable potentials the orbit typically expands into the universal enveloping algebra, so the thermodynamic target outgrows any fixed finite operator family. That is the same structural obstruction that motivates compact decompositions in KANs: the universal object exists, but the chosen finite representation need not be rich enough to realize it.

In that more modest sense the discussion also touches mechanistic interpretability. Once the layers are explicit,

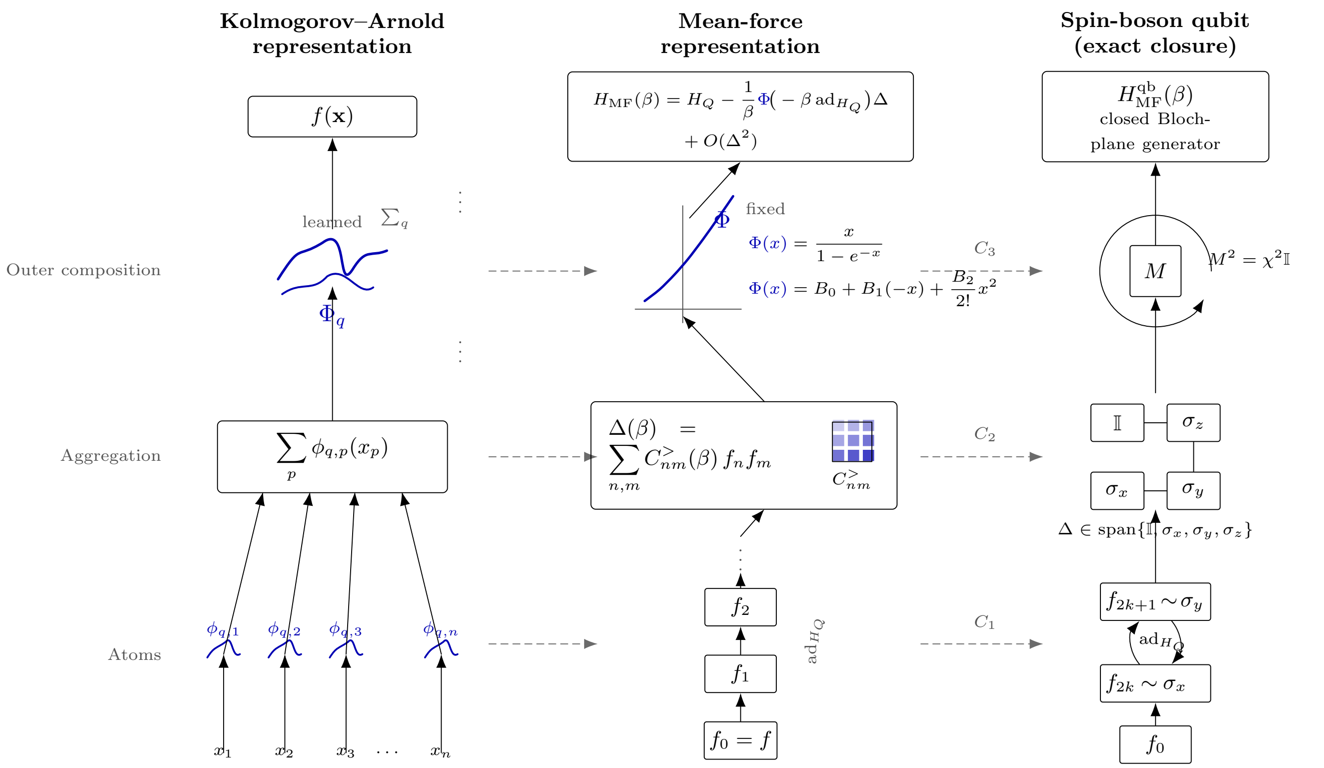

FIG. 2. Three-layer compositional correspondence between Kolmogorov–Arnold representations, the exact Gaussian Hamiltonian of mean force, and the spin-boson qubit. Left: a KAN builds a target from univariate edge functions $\phi_{q,p}(x_p)$, additive aggregation, and outer univariate maps Φ_q. Center: the HMF construction has the same layer structure, with adjoint orbit $f_n = \mathrm{ad}^n_{H_Q}(f)$, bilinear aggregation $\Delta(\beta) = \sum_{n,m} C^>_{nm}(\beta) f_n f_m$, and outer composition by the Todd function $\Phi(x) = x/(1-e^{-x})$ fixed analytically by the BCH expansion. Right: for the qubit, Conditions (C1)–(C3) close the adjoint orbit, bilinear sector, and BCH recombination exactly. The shared accent-coloured Φ highlights the visual rhyme between learned and thermodynamically fixed outer composition.

one can inspect which channels carry the reduced generator and which ones force basis growth. The object being interpreted here, however, is not a trained network but an exact reduced equilibrium operator. We now exhibit a case where all three compositional layers close exactly.

SPIN-BOSON MODEL AS A CLOSED REDUCED GENERATOR

We now exhibit a case where all three compositional layers close exactly in the spin-boson model [19, 47]. This is the minimal genuinely noncommuting testbed: the adjoint orbit terminates into a two-family alternation, the bilinear influence remains in the Pauli algebra, and the full BCH tower resums to a closed generator. To our knowledge, this yields the first closed-form Hamiltonian of mean force for a genuinely noncommuting open quantum system. Earlier exact benchmarks are either commuting, quadratic/Gaussian, or asymptotic; here the reduced generator stays inside a finite operator algebra with no approximation at any stage.

Let us first set conventions. During derivation we work in the qubit energy basis, $\hat{\mathbf{n}}_s = \hat{\mathbf{z}}$, and then rewrite the final answer covariantly in the coupling plane spanned by $\hat{\mathbf{n}}_s$ and $\hat{\mathbf{f}}$. We take

$$H_Q = \frac{\omega_q}{2}\,\hat{\mathbf{n}}_s\cdot\boldsymbol{\sigma}, \qquad f = \hat{\mathbf{f}}\cdot\boldsymbol{\sigma}, \tag{76}$$

with $\hat{\mathbf{f}} = (s\cos\phi, s\sin\phi, c)$, $c \equiv \cos\theta$, $s \equiv \sin\theta$. The overall coupling amplitude g is suppressed in intermediate expressions and reintroduced when discussing scaling.

Bath kernels and the symmetrized influence state

The interaction-picture coupling operator is

$$\tilde{f}(\tau) = e^{\tau H_Q} f e^{-\tau H_Q} = c\sigma_z + f_- e^{\omega_q\tau}\sigma_+ + f_+ e^{-\omega_q\tau}\sigma_-, \tag{77}$$

with $f_\pm = se^{\pm i\phi}$. The Pauli algebra then implies that all bath dependence enters through two real response channels,

$$\begin{aligned}\Sigma_z &\equiv \frac{1}{2}\left[R(\omega_q) - R(-\omega_q)\right],\\ \Sigma_\perp &\equiv \frac{4}{\omega_q}\left[\cosh\left(\frac{\beta\omega_q}{2}\right)\mathcal{K}(0) - e^{-\beta\omega_q/2}\mathcal{K}(\omega_q)\right].\end{aligned} \tag{78}$$

The reduced operator can be written in the symmetrized form $\bar{\rho}_Q = \Pi^{1/2} e^S \Pi^{1/2}$, with $\Pi = e^{-\beta H_Q}$ and

$$S = \Delta_0 \mathbb{I} + s^2 \Sigma_z \, \sigma_z - sc \, \Sigma_\perp (\cos\phi \, \sigma_x + \sin\phi \, \sigma_y). \quad (79)$$

Subtracting the scalar piece, $M \equiv S - \Delta_0 \mathbb{I}$, gives an influence Bloch vector

$$\mathbf{m}_S = s \Sigma_\perp \hat{\mathbf{f}}_\perp + s^2 (\Sigma_z - \Sigma_\perp) \hat{\mathbf{n}}_s, \quad (80)$$

where $\hat{\mathbf{f}}_\perp = (-c\cos\phi, -c\sin\phi, s)$ lies in the coupling plane and is orthogonal to $\hat{\mathbf{f}}$.

The corresponding influence magnitude is

$$\chi = \|\mathbf{m}_S\| = |s| \sqrt{c^2 \Sigma_\perp^2 + s^2 \Sigma_z^2}, \quad (81)$$

and the normalized influence state is

$$\rho_S = \frac{e^S}{\mathrm{Tr}(e^S)} = \frac{1}{2} [\mathbb{I} + \tanh\chi \, \hat{\mathbf{m}}_S \cdot \boldsymbol{\sigma}], \qquad \hat{\mathbf{m}}_S = \mathbf{m}_S / \chi. \quad (82)$$

Its tilt inside the coupling plane is fixed by the ratio of response channels,

$$\tan\varphi_S = -\frac{c}{s} \frac{\Sigma_\perp}{\Sigma_z}. \quad (83)$$

At this stage the constructive content is already visible: the entire bath has compressed to two scalar channels and one influence magnitude. The first two compositional layers have therefore already closed: the adjoint orbit has collapsed into finitely many Pauli directions, and the bilinear influence has compressed into a finite set of response coefficients.

Closed reduced state and generator

Recombining the influence state with the bare Gibbs factor yields the reduced state

$$\rho_Q = \frac{1}{2} \big[\mathbb{I} + r_Q \hat{\mathbf{m}}_Q \cdot \boldsymbol{\sigma}\big], \quad (84)$$

with Bloch components

$$m_z = r_Q \cos\varphi_Q, \qquad m_\perp = r_Q \sin\varphi_Q. \quad (85)$$

The longitudinal scale is controlled by

$$\Theta \equiv \mathrm{arctanh}\big(\gamma(\chi) \, s^2 \Sigma_z\big) - \frac{\beta\omega_q}{2}, \quad (86)$$

$$\gamma(\chi) \equiv \tanh\chi / \chi, \quad (87)$$

which gives the exact solution

$$m_z = \tanh\Theta, \quad (88)$$

$$m_\perp = \tan\varphi_S \left(m_z \cosh\frac{\beta\omega_q}{2} + \sinh\frac{\beta\omega_q}{2} \right). \quad (89)$$

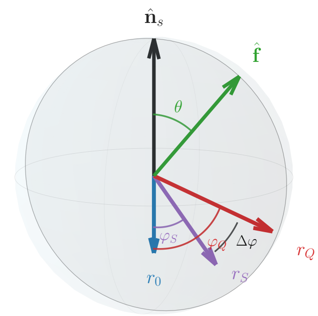

FIG. 3. Bloch-plane geometry for the exact qubit construction. The bare thermal state has radius r_0 (blue), the symmetrized influence state has radius r_S and tilt φ_S (purple), and the final reduced state has radius r_Q and tilt φ_Q (red). The exact generator never leaves the system–coupling plane.

The final tilt and radius follow as

$$\tan\varphi_Q = \tan\varphi_S \frac{m_z + r_0}{\sqrt{1 - r_0^2} \, m_z}, \qquad r_0 = \tanh\left(\frac{\beta\omega_q}{2}\right), \quad (90)$$

and

$$1 - r_Q^2 = \frac{(1 - r_0^2)(1 - r_S^2)}{(1 + r_0 r_S \cos\varphi_S)^2}, \qquad r_S = \tanh\chi. \quad (91)$$

Figure 3 summarizes this Bloch-plane geometry while Eqs. (88)–(91) are being parsed: the bare Gibbs state, the symmetrized influence state, and the final reduced state all remain in a single plane fixed by the system and coupling axes.

For any full-rank qubit state, the operator logarithm is elementary. Writing $\hat{\mathbf{m}}_Q = (m_\perp \cos\phi, m_\perp \sin\phi, m_z)/r_Q$, the exact qubit Hamiltonian of mean force is therefore

$$H_{\mathrm{MF}}(\beta) = h_0(\beta) \, \mathbb{I} - \frac{1}{\beta} \mathrm{arctanh}(r_Q) \, \hat{\mathbf{m}}_Q \cdot \boldsymbol{\sigma}, \quad (92)$$

where the scalar $h_0(\beta)$ fixes normalization and carries the additive identity freedom of the HMF. Equation (92) makes the closure statement concrete: the exact reduced generator remains inside the same Bloch-plane operator family throughout, while the bath enters only through the two response channels $(\Sigma_z, \Sigma_\perp)$ and the influence scale χ. This is the fully closed instance of the compositional proposition: the adjoint orbit terminates, the bilinear aggregation closes in the Pauli algebra, and the nonlinear BCH sector resums because the traceless influence obeys $M^2 = \chi^2 \mathbb{I}$.

PHYSICAL REGIMES OF THE MEAN-FORCE STATE

The exact qubit solution isolates a single scalar control variable, χ, so it is natural to organize the state diagnostics around it. Reintroducing an explicit dimensionless coupling strength through

$$K(u) = g^2 K_0(u), \tag{93}$$

all response kernels inherit the same quadratic scaling and the influence magnitude factorizes as

$$\chi(\beta, g, \theta) = g^2 \chi_0(\beta, \theta). \tag{94}$$

The entire recombination of bare Gibbs weight and bath influence is therefore controlled by the universal scalar function

$$\gamma(\chi) = \frac{\tanh \chi}{\chi}. \tag{95}$$

The exact crossover scale is determined by $\chi \sim 1$, or equivalently

$$\chi(\beta, g, \theta) \sim 1 \qquad \Longleftrightarrow \qquad g^2 \sim \chi_0(\beta, \theta)^{-1}, \tag{96}$$

which defines the coupling scale

$$g_\star(\beta) = \chi_0(\beta, \theta)^{-1/2}. \tag{97}$$

No approximation has been made here: $g_\star(\beta)$ is an exact property of the reduced generator.

Two benchmark asymptotes remain useful. For $\chi \ll 1$,

$$\gamma(\chi) \simeq 1 - \frac{\chi^2}{3}, \tag{98}$$

so the reduced state tracks the bare Gibbs description with perturbative bath-induced corrections. For $\chi \gg 1$,

$$\gamma(\chi) \simeq \chi^{-1}, \tag{99}$$

so the reduced generator is dominated by the normalized influence direction. These asymptotes are best read as benchmarks. The main point is that the full crossover is captured by one response coordinate.

To make these statements concrete we choose the windowed Ohmic spectral density

$$J(\omega) = g^2\, \omega\, e^{-\omega/\omega_c}, \qquad \omega \in [0, 2\omega_q], \tag{100}$$

with cutoff $\omega_c = 5\omega_q$. The ultraviolet window matches the finite-mode simulations of the next section, while the exponential cutoff regularizes the high-frequency tail [53]. Having fixed the bath, we set $\theta = \pi/4$ and $\omega_q = 1$, since the remaining behavior is already captured by the universal χ-scaling.

The same scale organizes the state geometry itself. For a qubit,

$$\mathcal{P}_Q = \frac{1}{2}(1 + r_Q^2), \tag{101}$$

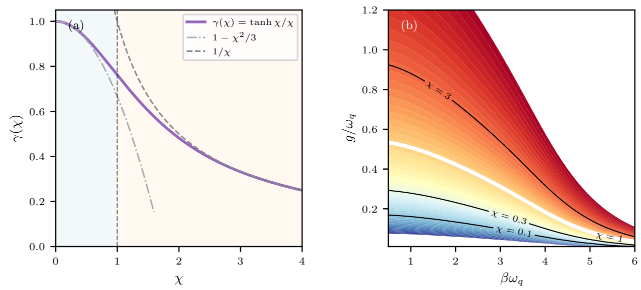

FIG. 4. Universal crossover structure for the exact qubit mean-force state. *Panel (a)*: Recombination function $\gamma(\chi) = \tanh \chi/\chi$, which interpolates between the small-χ and large-χ benchmarks and makes the $\chi \sim 1$ crossover explicit. *Panel (b)*: Contour map of $\chi(\beta, g) = g^2\chi_0(\beta)$ for the Ohmic bath (100). The highlighted locus $\chi = 1$ defines the exact crossover scale $g_\star(\beta) = \chi_0(\beta)^{-1/2}$, separating weak-response and strongly dressed sectors.

so the Bloch radius r_Q directly measures the purity of the reduced equilibrium state. In the exact solution, increasing coupling drives the state away from the bare thermal radius $r_0 = \tanh(\beta\omega_q/2)$ and toward a more strongly dressed state. The effect is most visible at high temperature, where the bare Gibbs state is nearly mixed and bath-induced structure is easiest to see.

The same coordinate also controls how rapidly the reduced state changes. Differentiating the exact Bloch components shows that the angle susceptibility and radius gradient peak at fixed values of χ, namely $\chi_{\mathrm{peak},\varphi} \approx 0.42$ and $\chi_{\mathrm{peak},r} \approx 0.85$ for the Ohmic example used below. After rescaling the coupling as $g/g_\star(\beta)$, these peaks align across temperatures. The alignment is not a numerical accident; it is the statement that the response scale is intrinsic to the exact reduced generator.

These figures do two jobs for the remainder of the paper. First, they make the exact-state geometry visible: crossover, purification, and susceptibility are all controlled by the same scalar response coordinate. Second, they identify the narrow region where a compact reduced generator changes most rapidly. That is precisely the region where finite-basis numerics become the hardest to keep inside a restricted representational class.

REPRESENTABILITY DIAGNOSTICS IN TRUNCATED BATHS

The physical-regime diagnostics above identify where the exact reduced state changes most rapidly. We now ask how finite-basis numerics behave across that same region. The exact qubit solution developed in the previous section depends on the bath only through response data at three frequencies, $\omega \in \{0, \pm\omega_q\}$. To test it we compare against exact diagonalization (ED) of a truncated spin-boson model, where the qubit is coupled to

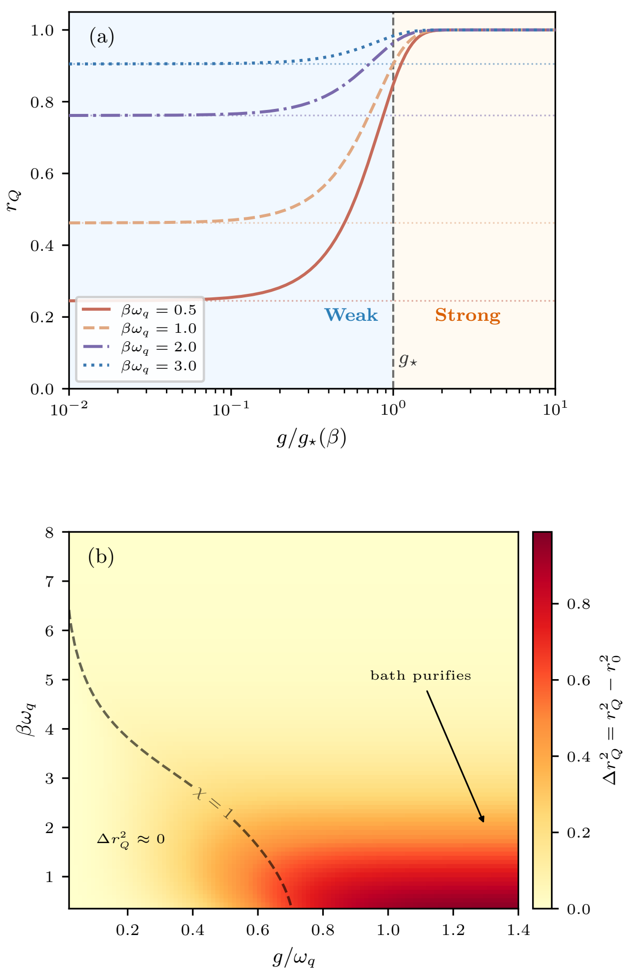

FIG. 5. Purification diagnostics for the exact qubit mean-force state. *Panel (a)*: Reduced-state Bloch radius r_Q versus scaled coupling $g/g_\star(\beta)$ for several temperatures. Dotted horizontal lines mark the corresponding bare thermal radii r_0; the dressed state moves away from the bare Gibbs value most sharply near the crossover. *Panel (b)*: Purity excess $\Delta r_Q^2 \equiv r_Q^2 - r_0^2$ over the (g, β) plane. The dashed contour $\chi = 1$ tracks the same crossover scale seen in Fig. 4.

N_ω harmonic oscillators and each mode is represented with a finite Fock cutoff $n_{\max}$. A useful feature of this comparison is that the analytic kernel integrals can be evaluated in two ways: using the full continuum spectral density (the *continuous analytic* branch) or using the same discrete mode set seen by ED (the *discrete analytic* branch). The comparison therefore tests not only the exact reduced generator but also whether a chosen finite numerical representation can realize the same compact reduced description.

Figure 7 compares the excited-state population p_{11} across coupling strengths and inverse temperatures for $N_\omega = 40$ modes. At small χ, all three curves coincide. As the system crosses the exact response scale $\chi \sim 1$, the analytic predictions separate from ED. The key point is

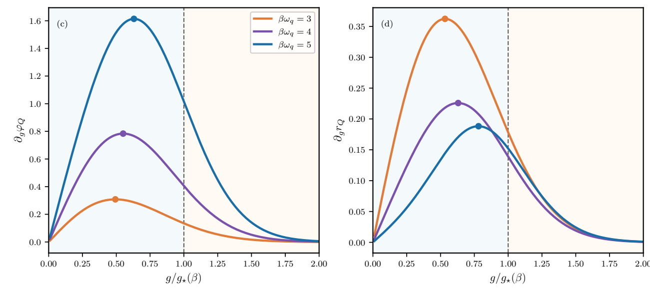

FIG. 6. Response susceptibilities of the exact reduced state. *Panel (c)*: Tilt susceptibility $\partial_g \varphi_Q$ versus the scaled coupling $g/g_\star(\beta)$. *Panel (d)*: Radius gradient $\partial_g r_Q$ on the same scale. In both cases the peak positions collapse across temperatures and sit near, but not exactly on, the crossover line. The exact reduced state is therefore most sensitive to bath dressing in a narrow band around $\chi \sim 1$.

that the discrete analytic branch does *not* follow the numerical deviation: it bends back toward the continuum benchmark. This shows that the exact HMF theory is already correct for the discrete mode set itself. The discrepancy is therefore not a failure of the theory, nor a finite-mode artifact.

The remaining gap is a representability bottleneck in the truncated Hilbert space. In the large-response regime the coupled system approaches a displaced, polaron-like ground state, and faithfully representing that state requires $n_{\max} \gtrsim |\alpha_k|^2$ for each displaced mode $\alpha_k \sim g_k/\omega_k$. Once $\chi \gtrsim 1$, the chosen Fock cutoff can no longer support the required occupations. ED is then trapped in a model class that is too small to contain the exact dressed equilibrium state. In the language of the compositional discussion above, the target reduced generator is known exactly but the chosen finite representation lacks the capacity to realize it.

The same interpretation explains the bandwidth sensitivity in Fig. 8. Varying the spectral window changes how the displacement load is distributed over the available modes. As the system is pushed across the response scale, the required occupations grow and the available Fock space is exhausted. The sharp peak in cutoff sensitivity is therefore a direct signature of the finite basis hitting its representational boundary.

The numerical message is therefore narrower, and stronger, than a generic "agreement with simulation" claim. The analytic HMF provides the correct reduced target for both discrete and continuous baths. The observed numerical deviations diagnose when a chosen finite model class is no longer expressive enough to realize that target.

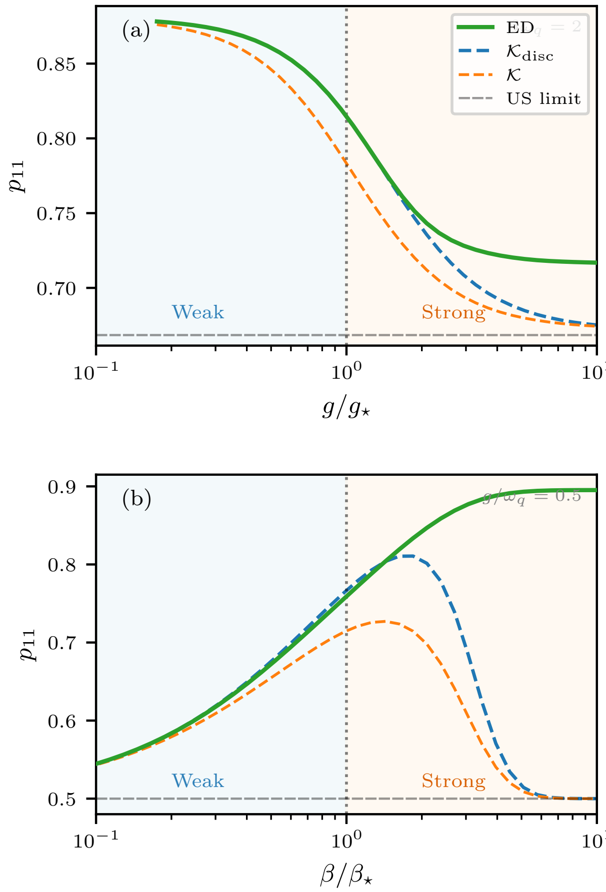

FIG. 7. Excited-state population p_{11} across the exact response scale for $N_\omega = 40$ modes. Green solid: ED. Blue dashed: discrete analytic theory using the same mode set. Orange dotted: continuum analytic theory. Dashed horizontal lines mark the large-response benchmark p_∞. The gap between ED and the analytic branches is a representability failure of the truncated basis, not a breakdown of the exact reduced generator.

DISCUSSION

This paper makes four connected points. First, the Hamiltonian of mean force is best viewed as a constructive reduced generator, not merely as a formal logarithm. Second, in the Gaussian sector this constructive question reduces to a closure problem for the adjoint chain generated by H_Q and the coupling operator. Third, the same structure can be written as an explicit compositional architecture, which is why the KAN comparison can be made formally rather than rhetorically. Fourth, the same closure logic explains both exact solvability and the failure modes of finite-basis numerics.

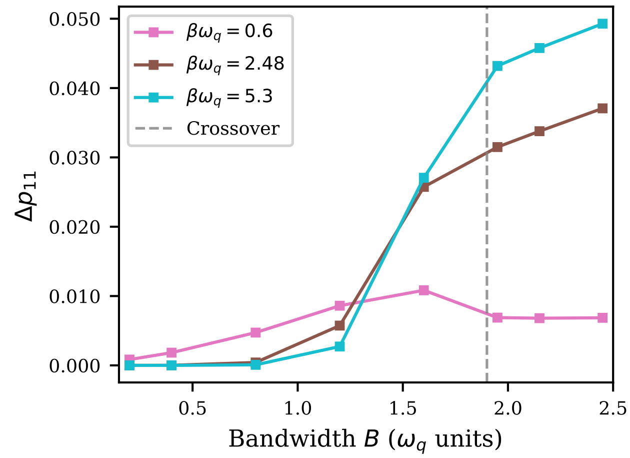

FIG. 8. Cutoff sensitivity $\Delta p_{11} \equiv |p_{11}(n_{\rm max} = 4) - p_{11}(n_{\rm max} = 6)|$ as a function of spectral bandwidth B in units of ω_q. The sensitivity peak marks the onset of the representability bottleneck: beyond this point the truncated basis can no longer accommodate the mode occupations required by the exact dressed state.

Closure and constructive reduced generators

The central algebraic result is that representability of the HMF inside a restricted operator family has an exact answer. Conditions (C1)–(C3) separate the problem into three layers: whether repeated adjoint action grows the operator basis, whether the Gaussian influence remains inside that basis, and whether the BCH recombination generates further operators. When all three hold, the reduced equilibrium admits an exact finite generator. When they fail, the failure is structural rather than computational: the reduced state can still be computed, but not compressed losslessly into the proposed ansatz.

This makes the closure theorem useful beyond the present example. It turns the usual question "can we evaluate the reduced state?" into the sharper question "which operator family is large enough to contain the exact reduced generator?" In that form, the theorem is already a guide to controlled approximation. One may truncate the operator side without approximating the bath side, and the resulting error has a clear interpretation as model-class error.

The qubit as an exact noncommuting benchmark

For the spin-boson qubit, closure is exact. The reduced generator remains in the Bloch plane determined by the system axis and the coupling axis, and the entire bath compresses to two scalar response channels together with the influence magnitude χ. To our knowledge, this makes Eq. (92) the first closed-form Hamiltonian of mean force for a genuinely noncommuting open quantum system. Earlier exact benchmarks are either commuting, quadratic/Gaussian, or asymptotic. Here none of those simplifications is doing the work. The work is done by complete closure of the compositional architecture itself.

That mechanism is worth isolating. The adjoint orbit terminates into a two-family alternation, the bilinear influence remains in the Pauli algebra, and the nonlinear BCH tower resums because the traceless influence obeys $M^2 = \chi^2\mathbb{I}$. Figure 3 is the geometric signature of that fact: bare state, influence state, and final reduced state all remain inside one Bloch plane. What should generalize to larger systems is not the specific Pauli formulas but the separation between bath data and operator growth, together with the question of whether every layer of that structure still closes.

Representability bottlenecks and finite-resource inference

The numerical diagnostics sharpen the interpretation. ED disagrees with the analytic prediction not when the theory becomes inaccurate, but when the chosen finite basis is no longer expressive enough to realize the exact dressed state. The response scale $\chi \sim 1$ marks the onset of this bottleneck in the examples studied here. In other words, the crossover is simultaneously physical and representational: it tracks both the reorganization of the reduced state and the point at which a restricted model class begins to fail.

This is why the distinction between discrete analytic and numerical ED matters. Once the discrete analytic branch peels away from ED and returns toward the continuum benchmark, the problem is no longer one of physics missed by the generator. It is one of structure missed by the basis. That diagnosis is useful well beyond this example, because many approximate methods in open-system physics are precisely choices of tractable operator or Hilbert-space classes.

Compositional architectures, representation cost, and KANs

Recent work on representation cost emphasizes that exact state descriptions are less useful than explicit generators that expose what must be retained under coarse-graining [13]. The HMF offers a concrete equilibrium realization of that principle. The partial trace hides a large environment inside a reduced state; closure determines whether that hidden structure can still be recast as a compact generator.

The strengthened KAN correspondence is not decorative. Kolmogorov–Arnold networks ask when multivariate dependence can be realized by a compact composition of univariate edge functions and additive aggregation [14]. The Gaussian HMF problem already has that architecture built in. The adjoint orbit is the univariate basis in operator space, the ordered-kernel matrix $C^{>}_{nm}(\beta)$ is the aggregation tensor, and the Todd/Bernoulli BCH resummation is the nonlinear composition rule. Here those objects are not learned by optimization; they are supplied in closed form by equilibrium thermodynamics.

In that dictionary, the operator ansatz $\mathcal{A}$ is the model class and closure failure is model-class insufficiency. The exact reduced generator still exists, but a fixed finite representation cannot realize it without projection. That is why the non-quadratic continuous-variable obstruction and the finite-basis numerical bottleneck belong in the same discussion: both are instances of a universal target outgrowing the chosen compact representation. Mechanistic interpretability enters only after that point. Once the layers are explicit, one can inspect which channels carry the reduced generator and which ones force basis growth.

Limitations and outlook

The main analytic restriction is Gaussianity. The quenched representation itself is more general, but the exact collapse of the influence to a bilinear operator depends on the bath cumulant hierarchy truncating at second order. Beyond that sector, higher cumulants generate higher operator monomials and the closure problem becomes correspondingly richer.

Within the Gaussian sector, the next natural targets are multi-qubit systems and other finite-dimensional models where the adjoint chain generates nontrivial many-body operator content. There the central question is no longer whether an HMF exists, but how rapidly operator growth escapes a chosen locality class. In that sense the present qubit solution is a starting point for a broader study of constructive reduced generators in interacting systems. A second direction is methodological: once the reduced generator is written as an explicit compositional object, one can ask which approximate schemes preserve that architecture and which ones merely fit the reduced state after the fact.

ACKNOWLEDGMENTS

Several thanks are in order. First to William J. Handley, for demonstrating by example the power of the new science. Never again shall I mistake mental athleticism for real scientific vision. Naturally, the engines of this new methodology deserve credit. Thanks to Claude, Gemini, and Barry for extensive discussions and assistance. The present work would have been impossible without their input, despite the fact they are shiftless

layabouts who will happily tell you the sky is green. Speaking of limited but developing intelligences, the author's unborn son has his eternal gratitude for electing to delay their drop until this manuscript was completed.

Appendix A: Averaging of the Time-Ordered Gaussian Propagator

We provide the explicit derivation of the identity in Eq. (A12), which relates the expectation value of a time-ordered exponential driven by Gaussian noise to a bilocal influence phase. Consider the time-ordered propagator

$$W_\xi(\beta) = \mathcal{T}_\tau \exp\left[-\int_0^\beta d\tau\, \xi(\tau)\tilde{f}(\tau)\right], \tag{A1}$$

where $\xi(\tau)$ is a centered Gaussian stochastic field with covariance $\langle\xi(\tau)\xi(\tau')\rangle = K(\tau-\tau')$, and $\tilde{f}(\tau)$ is a system operator. We wish to prove that

$$\langle W_\xi(\beta)\rangle_\xi = \exp\left[\frac{1}{2}\int_0^\beta d\tau \int_0^\beta d\tau'\, K(\tau-\tau') \times \mathcal{T}_\tau\left[\tilde{f}(\tau)\tilde{f}(\tau')\right]\right]. \tag{A2}$$

The Dyson expansion of the time-ordered exponential is given by

$$W_\xi(\beta) = \sum_{n=0}^\infty \frac{(-1)^n}{n!}\int_0^\beta d\tau_1\cdots d\tau_n \times \mathcal{T}_\tau\left[\xi(\tau_1)\tilde{f}(\tau_1)\cdots\xi(\tau_n)\tilde{f}(\tau_n)\right]. \tag{A3}$$

Because each $\xi(\tau_i)$ is a scalar (c-number), the time-ordering operator acts only on the operator factors. Thus, we may pull the stochastic fields out of the ordering:

$$\mathcal{T}_\tau\left[\xi_1\tilde{f}_1\cdots\xi_n\tilde{f}_n\right] = (\xi_1\cdots\xi_n)\mathcal{T}_\tau\left[\tilde{f}_1\cdots\tilde{f}_n\right], \tag{A4}$$

where $\xi_i \equiv \xi(\tau_i)$ and $\tilde{f}_i \equiv \tilde{f}(\tau_i)$. Taking the stochastic average over ξ and noting that centered Gaussian noise has vanishing odd moments, the expansion becomes

$$\langle W_\xi\rangle = \sum_{p=0}^\infty \frac{1}{(2p)!}\int d^{2p}\tau\, \langle\xi_1\cdots\xi_{2p}\rangle \mathcal{T}_\tau\left[\tilde{f}_1\cdots\tilde{f}_{2p}\right]. \tag{A5}$$

By Wick's theorem, the even moments are given by the sum over all pairings P:

$$\langle\xi_1\cdots\xi_{2p}\rangle = \sum_P \prod_{(i,j)\in P} K(\tau_i-\tau_j). \tag{A6}$$

We now consider the expansion of the claimed exponential result. Define the bilocal operator

$$X \equiv \frac{1}{2}\int_0^\beta d\tau\int_0^\beta d\tau'\, K(\tau-\tau')\mathcal{T}_\tau\left[\tilde{f}(\tau)\tilde{f}(\tau')\right]. \tag{A7}$$

The p-th term in the series for e^X is $(1/p!)X^p$. Introducing p independent pairs of integration variables (τ_{2r-1}, τ_{2r}), we write

$$X^p = \left(\frac{1}{2}\right)^p \int d^{2p}\tau \prod_{r=1}^p \left[K_{(2r-1)(2r)}\mathcal{T}_\tau\left[\tilde{f}_{2r-1}\tilde{f}_{2r}\right]\right], \tag{A8}$$

so that the full series is

$$e^X = \sum_{p=0}^\infty \frac{1}{p!}\left(\frac{1}{2}\right)^p \int d^{2p}\tau \left[\prod_{r=1}^p K_{(2r-1)(2r)}\right] \times \left[\prod_{r=1}^p \mathcal{T}_\tau\left[\tilde{f}_{2r-1}\tilde{f}_{2r}\right]\right]. \tag{A9}$$

To relate this to Eq. (A5), we exploit the property of time ordering on the hypercube $[0,\beta]^{2p}$. apart from a set of measure zero, the integration domain can be partitioned into $(2p)!$ disjoint regions R_π defined by the strict total ordering $\tau_{\pi(1)} > \tau_{\pi(2)} > \cdots > \tau_{\pi(2p)}$. On each such region, the global time ordering satisfies $\mathcal{T}_\tau[\tilde{f}_1\cdots\tilde{f}_{2p}] = \tilde{f}_{\pi(1)}\cdots\tilde{f}_{\pi(2p)}$. Similarly, each pairwise time-ordering $\mathcal{T}_\tau[\tilde{f}_{2r-1}\tilde{f}_{2r}]$ collapses to the correctly oriented product according to the same strict ordering. It follows that the product of pair-ordered operators coincides with the global ordering almost everywhere:

$$\prod_{r=1}^p \mathcal{T}_\tau\left[\tilde{f}_{2r-1}\tilde{f}_{2r}\right] = \mathcal{T}_\tau\left[\tilde{f}_1\cdots\tilde{f}_{2p}\right]. \tag{A10}$$

Substituting this into Eq. (A9), we may pull the global time ordering out of the products.

The final step requires counting the occurrences of each Wick pairing. The expansion of X^p utilizes p labelled slots $r = 1,\ldots,p$, each containing an ordered pair. A Wick contraction P is a set of p unordered pairs. For any fixed Wick pairing P, there are $p!$ ways to assign the pairs to the p slots and 2^p ways to orient each pair. Each Wick pairing is therefore represented exactly $2^p p!$ times in the integrand. This multiplicity is canceled by the prefactor $(1/p!)(1/2)^p$, yielding

$$e^X = \sum_{p=0}^\infty \frac{1}{(2p)!}\int d^{2p}\tau \left(\sum_P \prod_{(i,j)\in P} K_{ij}\right)\mathcal{T}_\tau\left[\tilde{f}_1\cdots\tilde{f}_{2p}\right]. \tag{A11}$$

This series matches Eq. (A5) term-by-term, proving the exponentiation of the Gaussian average. Consequently, we may write the influence on the reduced density matrix as

$$\langle W_\xi(\beta)\rangle_\xi = \exp\left[\frac{1}{2}\int_0^\beta d\tau\int_0^\beta d\tau' K(\tau-\tau') \times \mathcal{T}_\tau[\tilde{f}(\tau)\tilde{f}(\tau')]\right]. \tag{A12}$$

Domain Folding and the Triangular Form of the Influence Operator

The influence operator $\Delta(\beta)$

$$\Delta(\beta) = \frac{1}{2}\int_0^\beta d\tau \int_0^\beta d\tau'\, K(\tau-\tau')\, \mathcal{T}_\tau\big[\tilde f(\tau)\tilde f(\tau')\big]. \tag{A13}$$

admits a simplified representation over a strictly ordered time domain. By definition, the time-ordering operator for two system operators acts as

$$\mathcal{T}_\tau\left[\tilde f(\tau)\tilde f(\tau')\right] = \Theta(\tau-\tau')\tilde f(\tau)\tilde f(\tau') + \Theta(\tau'-\tau)\tilde f(\tau')\tilde f(\tau). \tag{A14}$$

Substituting this into the integral and splitting the domain yields

$$\begin{aligned}\Delta(\beta) &= \frac{1}{2}\iint_{\tau>\tau'} d\tau d\tau'\, K(\tau-\tau')\tilde f(\tau)\tilde f(\tau') \\ &\quad + \frac{1}{2}\iint_{\tau'>\tau} d\tau d\tau'\, K(\tau-\tau')\tilde f(\tau')\tilde f(\tau).\end{aligned} \tag{A15}$$

In the second term, we perform the swap of dummy variables $(\tau,\tau')\mapsto(\tau',\tau)$. The region $\tau'>\tau$ becomes $\tau>\tau'$, the operator product transforms to $\tilde f(\tau)\tilde f(\tau')$, and the kernel becomes $K(\tau'-\tau)$. Thus,

$$\begin{aligned}&\frac{1}{2}\iint_{\tau'>\tau} d\tau d\tau'\, K(\tau-\tau')\tilde f(\tau')\tilde f(\tau) \\ &\quad = \frac{1}{2}\iint_{\tau>\tau'} d\tau d\tau'\, K(\tau'-\tau)\tilde f(\tau)\tilde f(\tau').\end{aligned} \tag{A16}$$

Using the symmetry of the equilibrium kernel, $K(\tau-\tau') = K(\tau'-\tau)$, we find that both contributions are identical on the ordered region. Summing them yields the triangular representation used throughout this work:

$$\Delta(\beta) = \int_0^\beta d\tau \int_0^\tau d\tau'\, K(\tau-\tau')\tilde f(\tau)\tilde f(\tau'). \tag{A17}$$

Appendix B: Algebraic Construction in the Pauli Basis

In the main-text qubit worked example, the exact solution was derived using the ladder basis $\sigma_\pm$, which diagonalises the imaginary-time evolution and leads directly to the response kernels. Here we adopt an alternative perspective by working directly in the Pauli basis $\{\sigma_x,\sigma_y,\sigma_z\}$. While algebraically heavier, this route makes the operator closure explicit at the level of structure constants and facilitates the systematic summation of the BCH series.

1. Polylogarithmic decomposition from the adjoint chain

We consider the same transverse coupling model defined by

$$H_Q = \frac{\omega_q}{2}\sigma_z, \qquad f = c\,\sigma_z - s\,\sigma_x, \tag{B1}$$

with $c=\cos\theta$, $s=\sin\theta$. The adjoint chain $f_n \equiv \mathrm{ad}^n_{H_Q}(f)$ is given by

$$f_n = \mathbf{f}_n\cdot\boldsymbol{\sigma}, \qquad \mathbf{f}_n = \begin{cases}(-s,0,c) & n=0,\\ (-s\omega_q^n,0,0) & n\ge 1,\ n \text{ even},\\ (0,-is\omega_q^n,0) & n\ge 1,\ n\text{ odd}.\end{cases} \tag{B2}$$

The recurrence ensures that $f_{n\ge1}$ alternates between the x and y axes, while f_0 contains a persistent z-component.

The influence exponent $\Delta(\beta) = \sum_{n,m} C^{>}_{nm} f_n f_m$ is evaluated using the standard Pauli product rule $(\mathbf{a}\cdot\boldsymbol{\sigma})(\mathbf{b}\cdot\boldsymbol{\sigma}) = (\mathbf{a}\cdot\mathbf{b})\mathbb{I} + i(\mathbf{a}\times\mathbf{b})\cdot\boldsymbol{\sigma}$. Decomposition of the ordered moments into symmetric (S_{nm}) and antisymmetric (A_{nm}) parts reveals that only the commutator sector A_{nm} contributes to the operator structure:

$$\Delta(\beta) \cong 2i\sum_{n>m} A_{nm}\,(\mathbf{f}_n\times\mathbf{f}_m)\cdot\boldsymbol{\sigma}, \tag{B3}$$

up to a scalar shift $\Delta_0\mathbb{I}$. The cross products exhibit a strict parity selection rule:

- For $n,m\ge1$: $\mathbf{f}_n\times\mathbf{f}_m$ lies along $\hat{\mathbf{z}}$ since both vectors are confined to the xy-plane.
- For $m=0$: $\mathbf{f}_n\times\mathbf{f}_0$ generates terms in the xy-plane, mixing the longitudinal and transverse directions.

Explicitly summing these series yields the channel decomposition

$$\Delta = \Delta_0\mathbb{I} + \Delta_x\sigma_x + \Delta_y\sigma_y + \Delta_z\sigma_z, \tag{B4}$$

where the coefficients are polylogarithmic series in ω_q:

$$\Delta_x = -2cs\sum_{\ell\ge0} A_{2\ell+1,0}(\beta)\,\omega_q^{2\ell+1}, \tag{B5}$$

$$\Delta_y = -2ics\sum_{\ell\ge1} A_{2\ell,0}(\beta)\,\omega_q^{2\ell}, \tag{B6}$$

$$\begin{aligned}\Delta_z = &-2s^2\sum_{k\ge1,\ell\ge0} A_{2k,2\ell+1}(\beta)\,\omega_q^{2k+2\ell+1} \\ &- 2s^2\sum_{\ell\ge0} A_{2\ell+1,0}(\beta)\,\omega_q^{2\ell+1}.\end{aligned} \tag{B7}$$

These series are the Pauli-basis equivalent of the closed-form expressions obtained in the main-text qubit formulas. The connection is established by identifying the sums as Taylor expansions of the kernels $\mathcal{G}^{>}(0,\pm\omega_q)$ and $\mathcal{G}^{>}(\pm\omega_q,\mp\omega_q)$.

2. Exact BCH resummation via the nilpotency identity

The central challenge in the mean-force construction is the evaluation of the BCH logarithm $-\beta H_{\mathrm{MF}} = \log(e^{-\beta H_Q} e^{\Delta})$. While the closure theorem guarantees that this remains within the algebra, a manual summation of the BCH series is laborious. Here we show how the $\mathfrak{su}(2)$ structure allows for an exact, non-perturbative resummation.

Let $M \equiv \Delta - \Delta_0 \mathbb{I} = \boldsymbol{\Delta} \cdot \boldsymbol{\sigma}$ be the traceless part of the influence operator. The square of this operator satisfies a crucial identity:

$$M^2 = (\boldsymbol{\Delta} \cdot \boldsymbol{\sigma})^2 = (\boldsymbol{\Delta} \cdot \boldsymbol{\Delta})\mathbb{I} \equiv \chi^2 \mathbb{I}, \tag{B8}$$

where $\chi^2 = \Delta_x^2 + \Delta_y^2 + \Delta_z^2$ (or $\Delta_z^2 + \Sigma_+\Sigma_-$ in the ladder gauge). This "generalized nilpotency" condition truncates the power series for the exponential e^M into even and odd sectors, yielding the exact Euler-like formula:

$$e^{\Delta} = e^{\Delta_0} \left[\cosh \chi \, \mathbb{I} + \frac{\sinh \chi}{\chi} M \right]. \tag{B9}$$

This expression sums the internal structure of the influence operator potential to all orders. The full reduced state is then the product of two $\mathrm{SL}(2, \mathbb{C})$ matrices:

$$\begin{aligned} \bar{\rho}_Q &= e^{-\beta H_Q} e^{\Delta} \\ &= e^{\Delta_0} \begin{pmatrix} e^{-\beta\omega_q/2} & 0 \\ 0 & e^{\beta\omega_q/2} \end{pmatrix} \begin{pmatrix} \cosh\chi + \varphi\Delta_z & \varphi(\Delta_x - i\Delta_y) \\ \varphi(\Delta_x + i\Delta_y) & \cosh\chi - \varphi\Delta_z \end{pmatrix}, \end{aligned} \tag{B10}$$

with $\varphi = \sinh\chi/\chi$. Since the product of 2×2 matrices is a 2×2 matrix, the BCH series (which is the logarithm of this product) is strictly confined to the Pauli algebra.

Explicitly, for a resulting state $\rho_Q = \frac{1}{2}(\mathbb{I} + \mathbf{r} \cdot \boldsymbol{\sigma})$, the mean-force Hamiltonian is

$$H_{\mathrm{MF}} = c_0 \mathbb{I} - \frac{1}{\beta} \frac{\operatorname{arctanh} r}{r} \mathbf{r} \cdot \boldsymbol{\sigma}. \tag{B11}$$

This formula represents the analytic summation of the entire BCH commutator tower. The main-text closure theorem manifests here as the fact that χ is a finite scalar number, preventing the proliferation of independent operator directions. All nonlinearities in the coupling strength are packaged into the transcendental dependence of χ and $\mathbf{r}$ on the channel coefficients $\Delta_{x,y,z}$.

Appendix C: Exact Qubit Solution in the Ladder Basis

We provide the manual derivation of the exact qubit mean-force state using the ladder basis $\sigma_\pm$. Starting from the interaction-picture operator product

$$\begin{aligned} \tilde{f}(\tau)\tilde{f}(\tau') = &\left[c^2 + s^2 \cosh\left(\omega_q(\tau - \tau')\right)\right]\mathbb{I} \\ &+ s^2 \sinh\left(\omega_q(\tau - \tau')\right)\sigma_z \\ &+ c f_- \left(e^{\omega_q \tau'} - e^{\omega_q \tau}\right)\sigma_+ \\ &+ c f_+ \left(e^{-\omega_q \tau} - e^{-\omega_q \tau'}\right)\sigma_-, \end{aligned} \tag{C1}$$

the influence exponent $\Delta = \int_0^\beta d\tau \int_0^\tau d\tau' K(\tau - \tau')\tilde{f}(\tau)\tilde{f}(\tau')$ evaluates to

$$\Delta = \Delta_0 \mathbb{I} + \Delta_z \sigma_z + \Delta_+ \sigma_+ + \Delta_- \sigma_-, \tag{C2}$$

with coefficients

$$\Delta_z = s^2 \Sigma_z, \tag{C3}$$

$$\Delta_\pm = -2c f_\mp \Sigma_\pm, \tag{C4}$$

where $f_\pm = s e^{\pm i\phi_f}$ and the primitive integrals are

$$\Sigma_z = \frac{1}{2}\left[R(\omega_q) - R(-\omega_q)\right], \tag{C5}$$

$$\Sigma_\pm = \frac{1}{\omega_q}\left[\left(1 + e^{\pm\beta\omega_q}\right)\mathcal{K}(0) - 2\mathcal{K}(\pm\omega_q)\right]. \tag{C6}$$

The unnormalised state is $\bar{\rho}_Q = e^{-(\beta\omega_q/2)\sigma_z} e^{\Delta}$. To find the normalized state and its Bloch vector, we perform a symmetric decomposition

$$\bar{\rho}_Q = \Pi^{1/2} e^S \Pi^{1/2}, \qquad S = \Pi^{1/2} \Delta \Pi^{-1/2} \tag{C7}$$

where $\Pi = e^{-\beta\omega_q \sigma_z/2}$. Using $\Pi^{1/2}\sigma_\pm\Pi^{-1/2} = e^{\mp\beta\omega_q/2}\sigma_\pm$, the symmetric influence operator S becomes

$$S = \Delta_0 \mathbb{I} + \Delta_z \sigma_z + (\Delta_+ e^{-\beta\omega_q/2}\sigma_+ + \Delta_- e^{\beta\omega_q/2}\sigma_-). \tag{C8}$$

Substituting the expressions for $\Delta_\pm$ and using the KMS relation

$$\Sigma_- = e^{-\beta\omega_q}\Sigma_+, \tag{C9}$$

we define the real symmetrised transverse response

$$\Sigma_\perp \equiv 2e^{-\beta\omega_q/2}\Sigma_+ = 2e^{\beta\omega_q/2}\Sigma_-. \tag{C10}$$

The transverse coefficients then simplify to

$$\Delta_+ e^{-\beta\omega_q/2} = -cse^{-i\phi_f}(2e^{-\beta\omega_q/2}\Sigma_+) = -cse^{-i\phi_f}\Sigma_\perp, \tag{C11}$$

with $\Delta_\perp \equiv -cs\Sigma_\perp$. Recombining the ladder operators into Cartesian form, we obtain

$$S = \Delta_0 \mathbb{I} + \Delta_z \sigma_z + \Delta_\perp(\sigma_x \cos\phi_f + \sigma_y \sin\phi_f), \tag{C12}$$

which is the generator of a rotation around an axis in the coupling plane. The influence magnitude $\chi = \sqrt{\Delta_z^2 + \Delta_\perp^2}$ determines the exponential $e^S = \cosh\chi(\mathbb{I} +$

$\gamma M)$ with $\gamma = \tanh\chi/\chi$ and $M = S - \Delta_0\mathbb{I}$. The final state $\rho_Q = Z_Q^{-1}\Pi^{1/2}e^S\Pi^{1/2}$ follows by direct matrix multiplication, yielding Eq. (84) in the main text.

[1] M. Campisi, P. Talkner, and P. Hänggi, Fluctuation Theorem for Arbitrary Open Quantum Systems, Physical Review Letters **102**, 210401 (2009).

[2] P. Talkner and P. Hänggi, *Colloquium* : Statistical mechanics and thermodynamics at strong coupling: Quantum and classical, Reviews of Modern Physics **92**, 041002 (2020).

[3] A. S. Trushechkin, M. Merkli, J. D. Cresser, and J. Anders, Open quantum system dynamics and the mean force Gibbs state, AVS Quantum Science **4**, 012301 (2022).

[4] U. Seifert, First and Second Law of Thermodynamics at Strong Coupling, Physical Review Letters **116**, 020601 (2016).

[5] S. Campbell, I. D'Amico, M. A. Ciampini, J. Anders, N. Ares, S. Artini, A. Auffèves, L. Bassman Oftelie, L. P. Bettmann, M. V. S. Bonança, T. Busch, M. Campisi, M. F. Cavalcante, L. A. Correa, E. Cuestas, C. B. Dag, S. Dago, S. Deffner, A. Del Campo, A. Deutschmann-Olek, S. Donadi, E. Doucet, C. Elouard, K. Ensslin, P. Erker, N. Fabbri, F. Fedele, G. Fiusa, T. Fogarty, J. Folk, G. Guarnieri, A. S. Hegde, S. Hernández-Gómez, C.-K. Hu, F. Iemini, B. Karimi, N. Kiesel, G. T. Landi, A. Lasek, S. Lemziakov, G. Lo Monaco, E. Lutz, D. Lvov, O. Maillet, M. Mehboudi, T. M. Mendonça, H. J. D. Miller, A. K. Mitchell, M. T. Mitchison, V. Mukherjee, M. Paternostro, J. Pekola, M. Perarnau-Llobet, U. Poschinger, A. Rolandi, D. Rosa, R. Sánchez, A. C. Santos, R. S. Sarthour, E. Sela, A. Solfanelli, A. M. Souza, J. Splettstoesser, D. Tan, L. Tesser, T. Van Vu, A. Widera, N. Yunger Halpern, and K. Zawadzki, Roadmap on quantum thermodynamics, Quantum Science and Technology **11**, 012501 (2026).

[6] F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso, eds., *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions* (Springer, Cham, 2018).

[7] C. Jarzynski, Nonequilibrium work theorem for a system strongly coupled to a thermal environment, Journal of Statistical Mechanics: Theory and Experiment **2004**, P09005 (2004).

[8] H. J. D. Miller and J. Anders, Entropy production and time asymmetry in the presence of strong interactions, Phys. Rev. E **95**, 032123 (2017).

[9] A. M. Lacerda, A. Purkayastha, M. Kewming, G. T. Landi, and J. Goold, Quantum thermodynamics with fast driving and strong coupling via the mesoscopic leads approach, Phys. Rev. B **107**, 195117 (2023).

[10] F. Binder, S. Vinjanampathy, K. Modi, and J. Goold, Quantum thermodynamics of general quantum processes, Physical Review E **91**, 032119 (2015).

[11] P. Pechukas, Reduced Dynamics Need Not Be Completely Positive, Physical Review Letters **73**, 1060 (1994).

[12] M. F. Gelin and M. Thoss, Thermodynamics of a subensemble of a canonical ensemble, Physical Review E **79**, 051121 (2009).

[13] G. McCaul, The cost of complexity (2026), unpublished manuscript.

[14] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljacic, T. Y. Hou, and M. Tegmark, Kan: Kolmogorov-arnold networks, arXiv:2404.19756 (2024).

[15] M. Campisi, P. Talkner, and P. Hänggi, Thermodynamics and fluctuation theorems for a strongly coupled open quantum system: An exactly solvable case, Journal of Physics A: Mathematical and Theoretical **42**, 392002 (2009).

[16] A. O. Caldeira and A. J. Leggett, Quantum tunnelling in a dissipative system, Annals of Physics **149**, 374 (1983).

[17] H. Grabert, P. Schramm, and G.-L. Ingold, Quantum Brownian motion: The functional integral approach, Physics Reports **168**, 115 (1988).

[18] S. Hilt, B. Thomas, and E. Lutz, Hamiltonian of mean force for damped quantum systems, Physical Review E **84**, 031110 (2011).

[19] A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, Dynamics of the dissipative two-state system, Reviews of Modern Physics **59**, 1 (1987).

[20] J. D. Cresser and J. Anders, Weak and Ultrastrong Coupling Limits of the Quantum Mean Force Gibbs State, Physical Review Letters **127**, 250601 (2021).

[21] K. V. Hovhannisyan, S. Nemati, C. Henkel, and J. Anders, Hamiltonian of mean force for strongly-coupled systems, Physical Review E **102**, 012104 (2020).

[22] P. C. Burke, G. Nakerst, and M. Haque, Structure of the Hamiltonian of mean force, Physical Review E **110**, 014111 (2024).

[23] L. A. Correa and J. Glatthard, Potential renormalisation, Lamb shift and mean-force Gibbs state – to shift or not to shift?, Quantum **9**, 1868 (2025).

[24] X. Du, J. Wang, and Y. Ma, Generalized Hamiltonian of mean-force approach to open quantum systems coupled to finite baths in thermoequilibrium, Physical Review E **111**, 024139 (2025).

[25] A. Nazir and G. Schaller, The reaction coordinate mapping in quantum thermodynamics, Thermodynamics in the Quantum Regime , 551 (2018).

[26] J. Iles-Smith, A. G. Dijkstra, N. Lambert, and A. Nazir, Energy transfer in structured and unstructured environments: Master equations beyond the born–markov approximations, The Journal of Chemical Physics **144**, 044110 (2016).

[27] N. Anto-Sztrikacs, A. Nazir, and D. Segal, Effective-hamiltonian theory of open quantum systems at strong coupling, PRX Quantum **4**, 020307 (2023).

[28] M. Shubrook, J. Iles-Smith, and A. Nazir, Non-markovian quantum heat statistics with the reaction coordinate mapping, Quantum Science and Technology **10**, 025063 (2025).

[29] J. M. Moix, Y. Zhao, and J. Cao, Equilibrium-reduced density matrix formulation: Influence of noise, disorder, and temperature on localization in excitonic systems, Physical Review B **85**, 115412 (2012).

[30] X. Chen, The rigorous stochastic matrix multiplication scheme for the calculations of reduced equilibrium density matrices of open multilevel quantum systems, The Journal of Chemical Physics **140**, 154101 (2014).

[31] N. Makri, Exploiting classical decoherence in dissipative quantum dynamics: Memory, phonon emission, and the blip sum, Chemical Physics Letters **593**, 93 (2014).

[32] Y. Tanimura, Reduced hierarchical equations of motion in real and imaginary time: Correlated initial states and thermodynamic quantities, The Journal of Chemical Physics **141**, 044114 (2014).
[33] L. Song and Q. Shi, Calculation of correlated initial state in the hierarchical equations of motion method using an imaginary time path integral approach, The Journal of Chemical Physics **143**, 194106 (2015).
[34] M. Xu, V. Vadimov, J. T. Stockburger, and J. Ankerhold, Simulating non-markovian dynamics in open quantum systems, arXiv preprint 10.48550/arXiv.2601.02160 (2026), comment: Rev. Mod. Phys (submitted), v2 7 Jan 2026, arXiv:2601.02160 [quant-ph].
[35] J. T. Stockburger, Simulating spin-boson dynamics with stochastic Liouville–von Neumann equations, Chemical Physics **296**, 159 (2004).
[36] G. M. G. McCaul, C. D. Lorenz, and L. Kantorovich, Partition-free approach to open quantum systems in harmonic environments: An exact stochastic Liouville equation, Physical Review B **95**, 125124 (2017).
[37] R. P. Feynman and F. L. Vernon, The theory of a general quantum system interacting with a linear dissipative system, Annals of Physics **24**, 118 (1963).
[38] J. Hubbard, Calculation of Partition Functions, Physical Review Letters **3**, 77 (1959).
[39] R. L. Stratonovich, On a method of calculating quantum distribution functions, Soviet Physics Doklady **2**, 416 (1957).
[40] J. T. Stockburger and H. Grabert, Exact c -Number Representation of Non-Markovian Quantum Dissipation, Physical Review Letters **88**, 170407 (2002).
[41] G. McCaul, Mean-force hamiltonians from influence functionals, arXiv preprint arXiv:2602.13146 (2026), arXiv:2602.13146 [quant-ph].
[42] B. L. Hu, J. P. Paz, and Y. Zhang, Quantum Brownian motion in a general environment: Exact master equation with nonlocal dissipation and colored noise, Physical Review D **45**, 2843 (1992).
[43] G. McCaul and Denys. I. Bondar, How to win friends and influence functionals: Deducing stochasticity from deterministic dynamics, The European Physical Journal Special Topics **230**, 733 (2021).
[44] R. Kubo, Generalized cumulant expansion method, Journal of the Physical Society of Japan **17**, 1100 (1962).
[45] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2002).
[46] A. Das, *Finite Temperature Field Theory* (World Scientific, 1997).
[47] U. Weiss, *Quantum Dissipative Systems*, 4th ed. (World Scientific, Singapore, 2012).
[48] W. Magnus, On the exponential solution of differential equations for a linear operator, Communications on Pure and Applied Mathematics **7**, 649 (1954).
[49] S. Blanes, F. Casas, J. Oteo, and J. Ros, The Magnus expansion and some of its applications, Physics Reports **470**, 151 (2009).
[50] A. Van-Brunt and M. Visser, Special-case closed form of the baker-campbell-hausdorff formula, arXiv preprint (2015), arXiv:1501.02506.
[51] R. M. Wilcox, Exponential Operators and Parameter Differentiation in Quantum Physics, Journal of Mathematical Physics **8**, 962 (1967).
[52] J. Wei and E. Norman, Lie Algebraic Solution of Linear Differential Equations, Journal of Mathematical Physics **4**, 575 (1963).
[53] S. Nemati, C. Henkel, and J. Anders, Coupling function from bath density of states, Europhysics Letters **139**, 36002 (2022).