WHEN NOISE HELPS: PROBABILISTIC KINETIC SOLVING FOR REACTION NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting the outcome of a chemical reaction from first principles would let us screen molecules, materials and drugs on the computer instead of in the lab. Reaction networks, which map out the possible pathways a reaction can take, are a standard tool for this. A kinetic solver turns a network into a reaction outcome, but its reaction rates depend exponentially on the energy barriers computed by quantum chemistry. Even small barrier errors, within reach of the most expensive methods, can flip the predicted product. We introduce a kinetic solver that treats barriers as distributions instead of fixed numbers, and predicts the expected outcome over that uncertainty. This gives more accurate and more reproducible predictions, at a cost that shrinks as uncertainty grows. We validate these gains on synthetic networks spanning a wide range of topologies, temperatures, and reaction times, and qualitatively on real reaction networks from the literature. As machine learning pushes reaction network exploration towards full automation, manual error curation becomes infeasible, and uncertainty aware solvers offer a path to reliable chemistry prediction at scale.

1 INTRODUCTION

Chemical reactions are the physical world's engine: the transformation of matter underpins everything from the synthesis of molecules and materials to biochemical energy transfer and life itself. Quantum-mechanical simulations offer a path to design chemistry from first principles, shifting from observing nature to programming it, but face a fundamental tension between accuracy and scale. The Schrödinger equation is exact but intractable beyond the hydrogen atom. Approximate methods, from electronic structure approaches like coupled cluster and Density Functional Theory (DFT) to classical force fields, trade fidelity for feasibility, where each step towards greater speed or scale tends to sacrifice accuracy.

Reaction networks organize this landscape as a directed graph: nodes are stable molecules (reactants, intermediates, products), edges are elementary reaction steps that transform one molecule into another, and each edge carries an energy barrier that sets how fast that step occurs. A kinetic solver consumes these barriers to compute the system's dynamics: how the population of each node evolves over time and which pathways dominate the flux between them. Solving the dynamics on a reaction network provides a proxy to the experimentally observable outcomes from first principles.

However, exploring a reaction network requires thousands of elementary steps, each demanding hundreds of electronic structure queries, which pushes practitioners toward cheaper methods. Meanwhile, the kinetic solver is exponentially sensitive to barrier errors: even a 2 kcal/mol error, within reach of the most accurate methods, can reroute the flux and flip the reaction outcome. Additional errors from vibrational analysis and implicit solvation accumulate on top. As a result, even the most accurate methods, at a cost practically impossible over a full network, could still misrepresent the experimental outcome.

Every computational method carries uncertainty, but standard kinetic solvers discard it by treating barriers as exact. In this work, we introduce a kinetic solver that leverages barrier uncertainty, predicting the expected population rather than the outcome of a single best guess evaluation. This yields predictions that are both more accurate with respect to the ground truth populations and more reliable, with lower variance under perturbations of the electronic structure estimates. While this

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

expectation is more expensive to estimate, it also smooths the signal, so the larger the uncertainty the cheaper the approximation we can afford. Our contributions are as follows:

- We formulate a probabilistic kinetic solver that leverages uncertainty in free-energy barriers, yielding more accurate and reliable predictions than standard deterministic solvers.
- We derive an approximation of the solver whose cost shrinks as barrier uncertainty grows.
- We validate the method on synthetic networks across topologies, times, and temperatures, quantifying gains in population RMSE and top-k, and qualitatively on real networks.

2 Background

2.1 Potential Energy Surface

For a system of N atoms, a potential energy surface (PES) maps each nuclear geometry $\mathbf{x} \in \mathbb{R}^{3N}$ to its energy $E(\mathbf{x})$. Local minima represent stable configurations, while first-order saddle points represent transition states between them. A minimum energy path connecting two minima through a transition state defines an elementary reaction step. The energy difference between the transition state and reactant minimum defines its activation barrier, with higher barriers corresponding to slower reactions.

Because the PES is not known exactly, computational methods trade accuracy for cost. These range from classical (Cornell et al., 1995; Jorgensen et al., 1996) and machine-learning force fields (Batzner et al., 2022; Batatia et al., 2022; Wood et al., 2026) to semiempirical methods (Bannwarth et al., 2019), DFT (Kohn & Sham, 1965; Perdew et al., 1996; Zhao & Truhlar, 2008; Furness et al., 2020; Luise et al., 2025), and coupled cluster (Purvis III & Bartlett, 1982).

2.2 Reaction Network

A reaction network (Unsleber & Reiher, 2020) represents a PES as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The node set $\mathcal{V} = \{v_i\}_{i=1}^{M}$ contains one node for each minimum $\mathbf{x}_i$, and an edge $(i, j) \in \mathcal{E}$ exists when a transition state $\mathbf{x}_{ij}^{\ddagger}$ connects minima i and j. Each elementary reaction step contributes two directed edges, $i \to j$ and $j \to i$, with distinct activation free energies and rates. The Eyring equation (Eyring, 1935) relates the barrier $\Delta G_{ij}^{\ddagger}$ to the rate constant:

$$k_{ij} = \frac{k_B T}{h} \exp\left(-\frac{\Delta G_{ij}^{\ddagger}}{k_B T}\right) \tag{1}$$

where h is Planck's constant, k_B the Boltzmann constant, and T the temperature. The barrier combines the electronic energy difference, and free-energy corrections: $\Delta G_{ij}^{\ddagger} = G_{ij}^{\ddagger} - G_i = E(\mathbf{x}_{ij}^{\ddagger}) - E(\mathbf{x}_i) + \Delta G_{ij}^{\mathrm{corr}} = \Delta E_{ij}^{\ddagger} + \Delta G_{ij}^{\mathrm{corr}}$. Both terms carry errors from the electronic-structure method and the approximations used for zero-point, vibrational, rotational, and entropic corrections (Cramer, 2013). Because the Eyring equation amplifies these errors exponentially, small barrier errors can produce large rate errors, as analyzed in Section 3.

2.3 Kinetic Simulation

Given a reaction network, a kinetic solver computes the population vector $\boldsymbol{p}(t) \in \mathbb{R}^M$, where $p_i(t)$ is the fraction of material at node v_i. The rate constants from Eq. 1 form the rate matrix $\boldsymbol{K} \in \mathbb{R}^{M \times M}$, which defines the population dynamics:

$$K_{ij} = \begin{cases} k_{ji}, & i \neq j, \\ -\sum_{j \neq i} k_{ij}, & i = j, \end{cases} \qquad \frac{d\boldsymbol{p}}{dt} = \boldsymbol{K}\boldsymbol{p}(t). \tag{2}$$

The off-diagonal entry K_{ij} describes flow from node j to node i, while the diagonal ensures that each column sums to zero and the total population is conserved. Given an initial population $\boldsymbol{p}^{\mathrm{init}} = \boldsymbol{p}(0)$, the solution is

$$\boldsymbol{p}(t) = \exp(t\boldsymbol{K})\, \boldsymbol{p}^{\mathrm{init}}. \tag{3}$$

Forming the matrix exponential densely costs $O(M^3)$. Sparse ODE integrators cost $O(|\mathcal{E}|)$ per step, but reaction rate matrices are often stiff because their rates span many orders of magnitude, forcing explicit solvers to take prohibitively many small steps (Wanner & Hairer, 1996). Krylov methods instead approximate $\exp(t\boldsymbol{K})\boldsymbol{p}^{\text{init}}$ directly, require $O(|\mathcal{E}|)$ work per iteration, and handle stiffness effectively (Saad, 1992). Our method builds on these Krylov methods and, by leveraging uncertainty, enables smaller, more efficient subspaces, as described in Section 4.

2.4 Barrier Accuracy Across Levels of Theory

Table 1 summarizes reported barrier errors and estimated costs across levels of theory, establishing the error range studied throughout the paper.

Table 1: Electronic-barrier errors on BH76 (Goerigk et al., 2017) and estimated wall clock costs for an organic molecule ($C_{16}H_{21}Cl_2F_3N_2O_7$, 30 heavy atoms) on 16 CPU cores, with single points run embarrassingly in parallel (one per core). RMSE is measured against CCSD(T)/CBS references. Costs assume $\sim 10^3$ single point calculations per elementary step (Unsleber et al., 2022) and $\sim 10^3$ steps per network. *GFN2-xTB reports MAE (Bannwarth et al., 2019), which lower-bounds RMSE.

	BH76	Wall-clock cost		
Method	RMSE (kcal/mol)	Single point	Elem. step	Network
ωB97X-V (Mardirossian, 2014)	2.16	$\sim$5 min	$\sim$3 days	$\sim$9 yr
B3LYP-D3(BJ) (Grimme et al., 2011)	6.45	$\sim$2 min	$\sim$1.5 days	$\sim$4 yr
PBE (Perdew et al., 1996)	10.39	$\sim$30 s	$\sim$8 h	$\sim$11 mo
GFN2-xTB (Bannwarth et al., 2019)	$\gtrsim 13.9^*$	$\sim$10 ms	$\sim$10 s	$\sim$3 h

These errors are optimistic lower bounds. BH76 systems have at most five heavy atoms, while relevant small molecules have a median of $\sim$32 (Leeson et al., 2021). It also measures only electronic barriers, excluding free energy correction errors of up to $\sim$5 kcal/mol (Besora et al., 2018) and solvation. Even a 3 kcal/mol barrier error changes a room-temperature rate by $\sim \times 150$.

Higher accuracy remains costly and using traditional methods like ωB97X-V would require years per network. Machine learning force fields can approach similar high level accuracy at substantially lower costs (Kovács et al., 2025; Fu et al., 2025; Wood et al., 2026), with reported costs near 23 min per elementary step on one L4 GPU (Gugler et al., 2026). These advances enable automated network exploration (Kendall & Zádor, 2026; Gugler et al., 2026), where kinetic solvers must predict outcomes without hand curation.

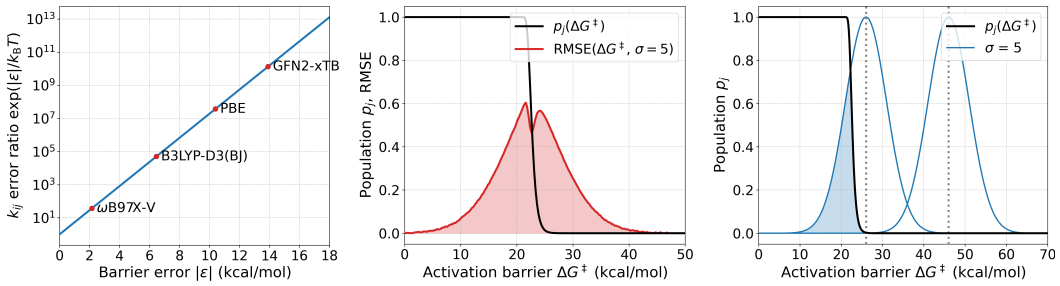
3 Barrier Sensitivity in Reaction Kinetics

In this section, we show that small perturbations in activation free energies can produce qualitatively different reaction outcomes. Because the Eyring rate depends exponentially on the activation free energy, a barrier error ϵ causes the rate to be multiplied or divided by $\exp(|\epsilon|/k_{\mathrm{B}}T)$. This ratio is plotted in Figure 1A. Barrier errors of 4 kcal/mol produce rate errors of approximately $\times 10^3$, while errors of 14 kcal/mol, typical of semiempirical methods such as GFN2-xTB (see Table 1), produce errors of approximately $\times 10^{10}$.

To connect rate errors to reaction outcomes, we consider the simplest network, a two state network with product j more stable than reactant i at room temperature, $T = 300$ K, and $t = 1$ h. Let $p_j(\Delta G^\ddagger)$ denote the product population. Given a noisy barrier $\Delta\hat{G}^\ddagger = \Delta G^\ddagger + \epsilon$, with $\epsilon \sim \mathcal{N}(0, \sigma^2)$, we define the population MSE as:

$$\mathrm{MSE}_{p_j}(\Delta G^\ddagger, \sigma) = \mathbb{E}_{\epsilon\sim\mathcal{N}(0,\sigma^2)}\left[\left(p_j(\Delta\hat{G}^\ddagger) - p_j(\Delta G^\ddagger)\right)^2\right], \tag{4}$$

Figure 1B reports the corresponding RMSE for $\sigma = 5$ kcal/mol as a function of the barrier. Near the narrow transition between reaction and no reaction for the ground truth $p_j(\Delta G^\ddagger)$, a standard barrier uncertainty can produce large population errors that even reverse the predicted outcome.



(a) (b) (c)

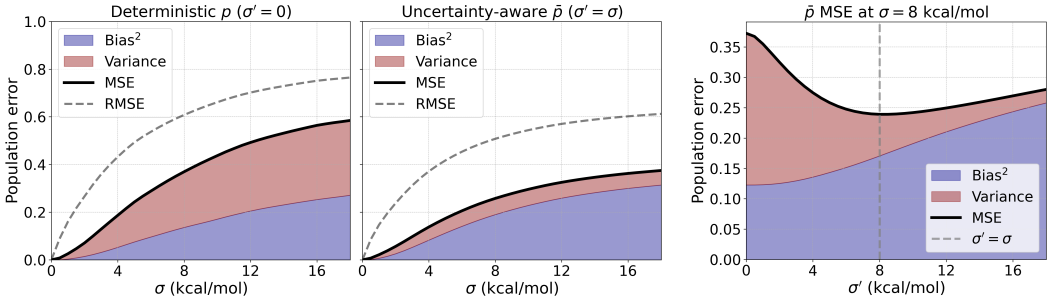
Figure 1: Barrier sensitivity in a two-state network. (a) Rate-error factor $\exp(|\epsilon|/k_{\mathrm{B}}T)$ against barrier error $|\epsilon|$. Markers show the factors at each method's BH76 electronic-barrier RMSE. (b) Ground-truth product population and deterministic prediction RMSE against the free-energy barrier under Gaussian noise with $\sigma = 5$ kcal/mol. (c) Product population with Gaussian barrier distributions illustrating different uncertainty levels. All panels use $T = 300$ K and $t = 1$ h.

.

Figure 1C provides geometric intuition for this sensitivity. The population $p_j(\Delta G^\ddagger)$ transitions from 1 to 0 over a narrow range. We overlay barrier distributions with $\sigma = 5$ kcal/mol centered at 26 and 46 kcal/mol. At 46 kcal/mol, the distribution remains in the non-reactive region, while at 26 kcal/mol, it overlaps the transition region despite the deterministic prediction of zero product. Simply integrating the population over this Gaussian barrier distribution yields an expected product population of 26.28%, which the deterministic prediction discards.

4 Probabilistic Kinetic Solver

4.1 Accounting for Uncertainty Reduces Error Variance



(a) Sweeping the true error σ. (b) Sweeping the assumed error σ'.

Figure 2: Bias-variance decomposition of the population MSE on networks composed of $M = 100$ nodes at $T = 400$ K and $t = 1$ h. (a) Error against the true barrier error σ, for the deterministic $\boldsymbol{p}$ (left) and uncertainty-aware $\bar{\boldsymbol{p}}$ (right) predictors. $\bar{\boldsymbol{p}}$ reduces the total error, it lowers variance significantly and it slightly increases bias. (b) Error against the assumed error σ' at fixed $\sigma = 8$ kcal mol^{-1}; the deterministic predictor is the left endpoint $\sigma' = 0$, and the minimum falls at $\sim\sigma' = \sigma$ (dashed line).

We show that predicting the expected population under barrier uncertainty yields lower population RMSE than a deterministic point estimate. It also yields more robust predictions by reducing their sensitivity to perturbations in the estimated barriers.

Consider a reaction network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and let $\boldsymbol{\Delta G}^\ddagger \in \mathbb{R}^d$ collect its $d = |\mathcal{E}|$ activation free energies. An electronic-structure method (Section 2.4) with barrier error σ returns the noisy estimates:

$$\boldsymbol{\Delta \hat{G}}^\ddagger = \boldsymbol{\Delta G}^\ddagger + \boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\sigma), \tag{5}$$

where $\boldsymbol{\Sigma}_\sigma$ is the barrier-error covariance induced by the network topology and the error scale σ. Let $p_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger)$ be the deterministic population of node i at fixed t, T, $\boldsymbol{p}^{\text{init}}$. The uncertainty-aware predictor is the expected population over plausible barriers under an assumed error scale σ':

$$\bar{p}_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger,\, \sigma') = \mathbb{E}_{\boldsymbol{\varepsilon}'\sim\mathcal{N}(\mathbf{0},\boldsymbol{\Sigma}_{\sigma'})}\left[p_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger + \boldsymbol{\varepsilon}')\right]. \tag{6}$$

approximated in practice by L Monte Carlo samples, and reducing to $\bar{p}_i = p_i$ at $\sigma' = 0$. Notice $\boldsymbol{\Sigma}_\sigma$ represents the true barrier error covariance, while $\boldsymbol{\Sigma}_{\sigma'}$ is the covariance assumed by the predictor for error scale σ'. We draw correlated barrier noise without forming $\boldsymbol{\Sigma}_{\sigma'}$ by sampling $\eta_i \sim \mathcal{N}(0, (1-\alpha)\sigma'^2)$ per minimum and $\eta^\ddagger_{ij} = \eta^\ddagger_{ji} \sim \mathcal{N}(0, \alpha\sigma'^2)$ per transition state, then setting $\varepsilon'_{ij} = \eta^\ddagger_{ij} - \eta_i$. This preserves detailed balance in every sample and induces $\boldsymbol{\Sigma}_{\sigma'}$ from the network topology. The parameter $\alpha \in [0,1]$ splits variance between minima and transition states; we use $\alpha = 1/2$. We extend the population MSE metric from Section 3, to the full population vector of an M-species network and to $\bar{\boldsymbol{p}}$:

$$\text{MSE}_{\bar{\boldsymbol{p}}}(\boldsymbol{\Delta}\boldsymbol{G}^\ddagger, \sigma, \sigma') = \frac{1}{2}\sum_{i=1}^{M}\mathbb{E}_{\boldsymbol{\varepsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{\Sigma}_\sigma)}\left[\left(\bar{p}_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger, \sigma') - p_i(\boldsymbol{\Delta}\boldsymbol{G}^\ddagger)\right)^2\right]. \tag{7}$$

The factor $1/2$ avoids double-counting displaced population, and the metric reduces to Eq. 4 at $M = 2$ (Appendix A). We report its square root as population RMSE.

Figure 2a compares the two strategies across different network topologies with $M = 100$ nodes and $\sigma = \sigma'$. The uncertainty aware predictor achieves lower population error than the deterministic predictor across the entire range, with biggest improvements at higher uncertainties.

Figure 2b tests robustness to misspecification of the assumed σ' at fixed $\sigma = 8$ kcal/mol. The minimum lies near $\sigma' = \sigma$ and is shallow: underestimating σ' never exceeds the deterministic RMSE, and moderate overestimation still helps. We use $\sigma' = \sigma$ in the remaining experiments.

The bias-variance decomposition explains the mechanism: averaging over plausible barriers introduces a modest bias but removes a larger variance component, making $\bar{\boldsymbol{p}}$ less sensitive to noisy barrier estimates than $\boldsymbol{p}$. We quantify this sensitivity as $\mathcal{V}(\sigma') = \frac{1}{2}\sum_i \text{Var}_{\boldsymbol{\varepsilon}}\left[\bar{p}_i(\boldsymbol{\Delta}\boldsymbol{G}^\ddagger + \boldsymbol{\varepsilon}, \sigma')\right]$, which equals the variance component of Eq. 7 (Appendix A) and recovers $\boldsymbol{p}$ at $\sigma' = 0$. Its decrease with σ' in Figures 2a and 2b confirms this reduced sensitivity.

4.2 Leveraging noise to simplify $\bar{\boldsymbol{p}}$

We approximate the noise-averaged predictor $\bar{p}_i$ with a lower-cost surrogate $\bar{q}_i$:

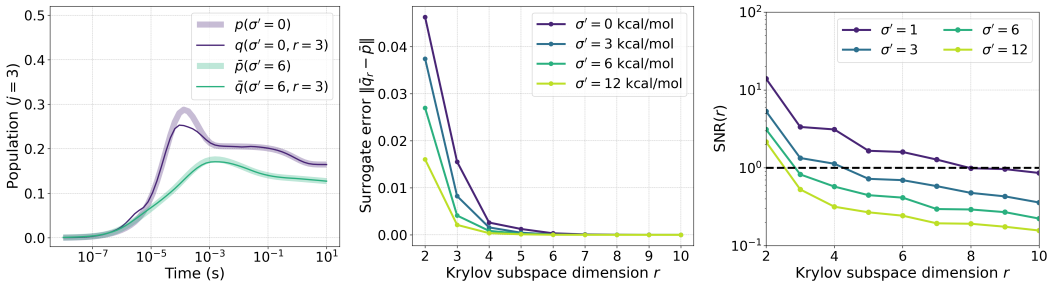
$$\bar{q}_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger,\, \sigma') = \mathbb{E}_{\boldsymbol{\varepsilon}'\sim\mathcal{N}(\mathbf{0},\boldsymbol{\Sigma}_{\sigma'})}\left[q_i(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger + \boldsymbol{\varepsilon}')\right]. \tag{8}$$

Here q_i is the rank-r Krylov approximation of p_i for each matrix exponential. Computing $\bar{q}_i$ requires one kinetic solve for each of the L Monte Carlo samples. Its cost grows with the number of Krylov vectors r. Our key observation is that barrier uncertainty limits the spectral resolution the ensemble can resolve, so a smaller r suffices as long as the averaged surrogate matches, $\bar{q}_i \approx \bar{p}_i$, even if each individual q_i is a coarser approximation of p_i. Under detailed balance, $\boldsymbol{K}$ has real eigenvalues $\lambda_i \le 0$ and real eigenvectors $\boldsymbol{u}_i$, so the propagated population diagonalizes as:

$$\boldsymbol{p}(t) = e^{t\boldsymbol{K}}\boldsymbol{p}^{\text{init}} = \sum_{i=1}^{M} c_i\boldsymbol{u}_i e^{\lambda_i t} \approx \boldsymbol{q}^{(r)}(t) = \sum_{i=1}^{r}\hat{c}_i\hat{\boldsymbol{u}}_i e^{\hat{\lambda}_i t}. \tag{9}$$

where $\boldsymbol{U} = [\boldsymbol{u}_1, \ldots, \boldsymbol{u}_M]$, $c_i = (\boldsymbol{U}^{-1}\boldsymbol{p}^{\text{init}})_i$, and $\hat{\lambda}_i$, $\hat{\boldsymbol{u}}_i$, and $\hat{c}_i$ are the corresponding Ritz approximations obtained in the Krylov subspace, so q_i is the i-th entry of the rank-r truncation $\boldsymbol{q}^{(r)}(t)$. Each mode contributes with weight $|c_i|e^{\lambda_i t}$, so faster-decaying modes (large $|\lambda_i|$) contribute least. A first-order perturbation analysis over barrier realizations (Appendix B) shows that their absolute uncertainty grows with $|\lambda_i|$, namely $\sqrt{\text{Var}_{\boldsymbol{\varepsilon}'}(\delta\lambda_i)} \propto (\sigma'/k_BT)\,|\lambda_i|$. The modes that contribute least are therefore also the ones most swamped by noise. We identify and discard these modes directly using the SNR criterion below.

Shift-and-invert Arnoldi (Saad, 1992) initialized with $\boldsymbol{p}^{\text{init}}$ builds the Krylov subspace so dominant contributions appear first. We stop as soon as the added contribution falls below its variation across



(a) (b) (c)

Figure 3: Barrier uncertainty enables more aggressive Krylov truncation. a Exact and rank-3 surrogate populations for σ'=0 and σ'=6. b Surrogate error as a function of Krylov rank and barrier uncertainty. c SNR as a function of rank, with the first crossing below 1 at ranks $8, 5, 3, 3$ for σ'=$1, 3, 6, 12$ kcal/mol, respectively. All panels use M=100, T=1000 K, t=0.01 s, and 128 Monte Carlo samples.

barrier samples. Let $\boldsymbol{q}_\ell^{(r)}$ be the sample-ℓ population at iteration r, $\bar{\boldsymbol{q}}^{(r)}$ its mean, $\boldsymbol{s}_r = \bar{\boldsymbol{q}}^{(r)} - \bar{\boldsymbol{q}}^{(r-1)}$ the added contribution, and $\hat{\boldsymbol{s}}_r = \boldsymbol{s}_r / \|\boldsymbol{s}_r\|$:

$$\mathrm{SNR}(r) = \frac{\|\boldsymbol{s}_r\|^2}{\hat{\boldsymbol{s}}_r^\top \boldsymbol{\Sigma}_r \, \hat{\boldsymbol{s}}_r}, \qquad \boldsymbol{\Sigma}_r = \mathrm{Cov}_\ell \left[\boldsymbol{q}_\ell^{(r)} - \boldsymbol{q}_\ell^{(r-1)} \right]. \tag{10}$$

We stop at the first r with $\mathrm{SNR}(r) < 1$: further modes fall below the uncertainty floor.

Figure 3 illustrates the mechanism on networks with $M = 100$ nodes at $T = 1000$ K; we pick this high temperature to activate more pathways and stress the truncation rule. In (a), at rank $r = 3$ the gap between $\bar{\boldsymbol{q}}$ and $\bar{\boldsymbol{p}}$ shrinks as σ' grows. In (b), at fixed rank the surrogate error decreases with σ' across 100 networks. Finally, (c) adds the growing noise floor to give the SNR criterion, which stops even earlier: the crossing moves to smaller rank as σ' increases ($8, 5, 3, 3$ at $\sigma' = 1, 3, 6, 12$ kcal/mol).

4.3 Algorithm Overview

Algorithm 1 returns the mean uncertainty-aware population $\bar{\boldsymbol{q}}$. Its inputs are those of a deterministic kinetic solver, augmented by the barrier uncertainty σ' and number of Monte Carlo samples L. All computations are embarrassingly parallel across Monte Carlo samples except the ensemble average and SNR check, which aggregate all sample populations after each Krylov iteration.

Algorithm 1 Probabilistic kinetic solver

Input: Network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, estimated barriers $\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger$, temperature T, time t, initial population $\boldsymbol{p}^{\mathrm{init}}$, barrier uncertainty σ', and number of Monte Carlo samples L
Output: Mean uncertainty-aware population $\bar{\boldsymbol{q}}$

Sample $\boldsymbol{\Delta}\tilde{\boldsymbol{G}}_\ell^\ddagger \sim \mathcal{N}(\boldsymbol{\Delta}\hat{\boldsymbol{G}}^\ddagger, \boldsymbol{\Sigma}_{\sigma'})$ for $\ell = 1, \ldots, L$ ▷ Section 4.1
$\tilde{\boldsymbol{K}}_\ell \leftarrow \textsc{RateMatrix}(\mathcal{G}, \boldsymbol{\Delta}\tilde{\boldsymbol{G}}_\ell^\ddagger, T)$ ▷ Eqs. 1 and 2
Initialize $\boldsymbol{V}_\ell^{(1)} \leftarrow [\boldsymbol{p}^{\mathrm{init}}]$ and $\boldsymbol{H}_\ell^{(0)}$ as empty
for $r = 1, 2, \ldots, M$ **do**
  $(\boldsymbol{V}_\ell^{(r+1)}, \boldsymbol{H}_\ell^{(r)}) \leftarrow \textsc{ArnoldiStep}(t\tilde{\boldsymbol{K}}_\ell, \boldsymbol{V}_\ell^{(r)}, \boldsymbol{H}_\ell^{(r-1)})$ ▷ (Katrutsa et al., 2019)
  $\boldsymbol{q}_\ell^{(r)} \leftarrow \textsc{PopulationFromSubspace}(\boldsymbol{V}_\ell^{(r)}, \boldsymbol{H}_\ell^{(r)})$ ▷ Appendix C
  $\bar{\boldsymbol{q}}^{(r)} \leftarrow \frac{1}{L} \sum_\ell \boldsymbol{q}_\ell^{(r)}$
  if $r > 1$ and $\left[\mathrm{SNR}\big(\{\boldsymbol{q}_\ell^{(r)} - \boldsymbol{q}_\ell^{(r-1)}\}_\ell\big) < 1$ **or** $\|\bar{\boldsymbol{q}}^{(r)} - \bar{\boldsymbol{q}}^{(r-1)}\| < 10^{-6}\right]$ **then** ▷ Eq. 10
    break
return $\bar{\boldsymbol{q}}^{(r)}$

5 EXPERIMENTS

We report four studies: synthetic networks with known ground truth (§ 5.1); time and temperature sweeps (§ 5.2); runtime and sampling convergence (§ 5.3); and two first principles networks as case studies (§ 5.4).

Predictors. We compare three predictors. The *deterministic* predictor $\boldsymbol{p}$ propagates the point estimate barriers through a high-rank Krylov propagator (baseline). The *Monte Carlo* predictor $\bar{\boldsymbol{p}}$ (Eq. 6) averages the same propagator over L barrier samples. The *surrogate Monte Carlo* predictor $\bar{\boldsymbol{q}}$ (Eq. 8) adds the SNR cutoff of Section 4.2 to truncate its rank per sample, making each sample cheaper. Defaults: $L = 128$, $\sigma' = \sigma$. Since $\bar{\boldsymbol{q}}$ matches $\bar{\boldsymbol{p}}$ in accuracy everywhere we tested, we report $\bar{\boldsymbol{p}}$ only for timings (§ 5.3).

Synthetic networks. We grow each network breadth-first from one reactant. Each expanded state proposes four elementary steps. Each step closes a cycle to an existing state with probability 0.35; otherwise it creates a new state. We sample state energies uniformly from $[-23, 23]$ kcal/mol and barriers uniformly up to 46 kcal/mol above the higher-energy endpoint. Branching and ring-closure probabilities are set so the degree and cycle distributions are close to those of real networks (Appendix D).

5.1 UNCERTAINTY PROPAGATION IMPROVES PREDICTIONS

We compare $\boldsymbol{p}$ and $\bar{\boldsymbol{q}}$ (with 128 Monte Carlo samples) on synthetic networks with $M \in \{10, 100, 1000\}$ states, corrupting the barriers with $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\sigma)$ (Section 4.1) and taking the uncorrupted populations as ground truth. We report population RMSE (Eq. 7) and top-K accuracy, which measures how often the true dominant state is among the K largest predicted populations. Each (M, σ) point averages 10 000 networks at $T = 400$ K and $t = 1$ h. Figure 4 shows $\bar{\boldsymbol{q}}$ achieves lower RMSE and higher top-K accuracy across network sizes.

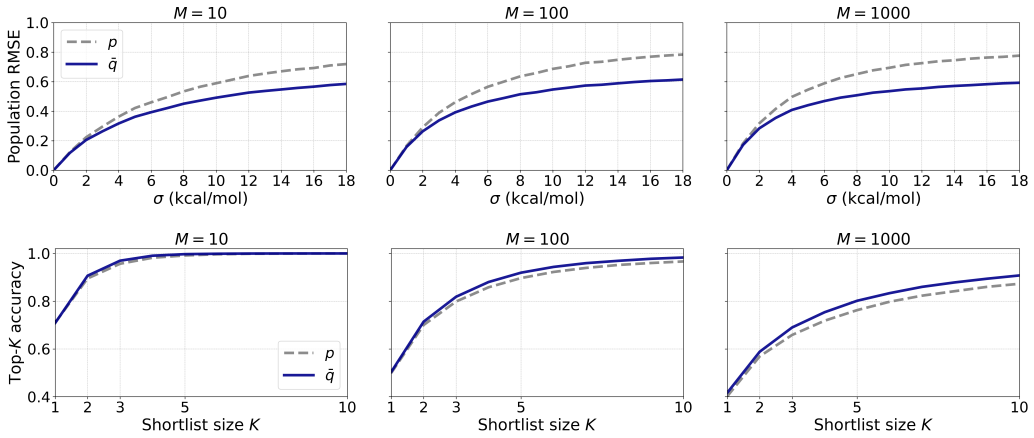


Figure 4: Synthetic networks with $M \in \{10, 100, 1000\}$ at $T = 400$ K and $t = 1$ h. Top: population RMSE vs true barrier error σ. Bottom: top-K accuracy at $\sigma = 6$ kcal mol^{-1}.

5.2 TIME AND TEMPERATURE DEPENDENCE

We test whether the advantage of $\bar{\boldsymbol{q}}$ persists across kinetic conditions. We sweep reaction time on a log-spaced grid $t \in \{10^{-4}, 10^{-3}, \ldots, 10^{4}\}$ s and temperature $T \in \{300, 400, \ldots, 1000\}$ K, averaging 1000 networks with M=100 states per cell. We set $\sigma' = \sigma \in \{3, 6, 12\}$ kcal mol^{-1} and use 128 Monte Carlo samples. We report the RMSE gain $\mathrm{RMSE}(\boldsymbol{p}) - \mathrm{RMSE}(\bar{\boldsymbol{q}})$; positive values favor $\bar{\boldsymbol{q}}$. Figure 5 shows that $\bar{\boldsymbol{q}}$ improves accuracy throughout the tested grid. The gain increases with barrier uncertainty, longer reaction times, and higher temperatures.

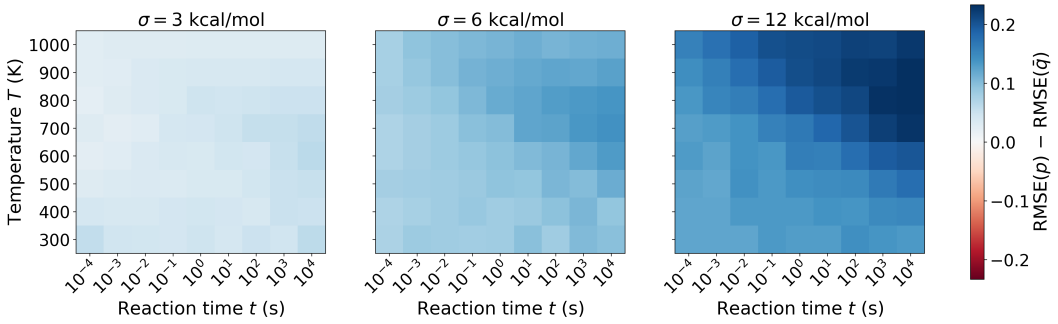
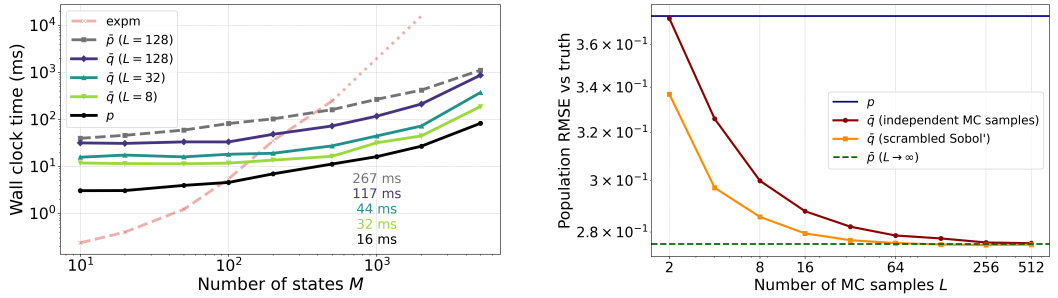


Figure 5: RMSE gain of $\bar{\boldsymbol{q}}$ over $\boldsymbol{p}$ across reaction time and temperature at three barrier uncertainties. Each cell averages 1000 synthetic networks with $M = 100$ states. Positive values (blue) favor $\bar{\boldsymbol{q}}$.

5.3 Compute cost and sampling convergence

We test whether uncertainty propagation remains practical as network size M and MC sample count L grow. Both studies share $T = 1000$ K, $t = 0.01$ s, and $\sigma' = \sigma = 6$ kcal mol^{-1}, a harder regime: short times and high temperatures activate more modes and demand the higher Krylov ranks. For cost, on 16 cores we time $\boldsymbol{p}$, $\bar{\boldsymbol{p}}$, $\bar{\boldsymbol{q}}$, and dense `expm` for $M = 10$ to 5000, using $L = 128$ for $\bar{\boldsymbol{p}}$ and $L \in \{8, 32, 128\}$ for $\bar{\boldsymbol{q}}$; each point is the median over seven networks. For sampling convergence, we vary L from 2 to 512 with independent MC and scrambled Sobol′ draws (Sobol, 1967; Owen, 1998), averaged over 1000 networks with $M = 100$.

Figure 6a shows that SNR truncation makes $\bar{\boldsymbol{q}}$ from 1.3 to 2.3× cheaper than $\bar{\boldsymbol{p}}$ across the tested M. At $M = 1000$, $\bar{\boldsymbol{q}}$ with $L = 128$ costs 117 ms, or ∼7 times the 16 ms deterministic solve rather than 128 times. Figure 6b shows that independent sampling converges at the expected $1/L$ rate in RMSE, while scrambled Sobol′ sampling converges faster. About 64 samples capture most of the accuracy gain; we use $L = 128$ Sobol′ elsewhere.



(a) Cost.

(b) Sampling convergence.

Figure 6: Compute cost and sampling convergence. a Wall-clock time against network size; the dotted line extrapolates dense `expm` beyond $M = 500$. b Population RMSE against the number of barrier samples for independent MC and scrambled Sobol′ sampling.

5.4 Qualitative case study: real quantum-chemical networks

Finally, we use the aqueous ozonolysis networks of Petrus et al. (2025), both generated with Chemoton (Unsleber et al., 2022): tetramethylethene (CRN-T, 117 states, 157 steps) and ethene (CRN-E, 246 states, 473 steps). Geometries were optimized at LC-PBE (Iikura et al., 2001)/def2-TZVP (Weigend & Ahlrichs, 2005) and energies refined with DLPNO-CCSD(T) (Riplinger & Neese, 2013)/cc-pVTZ (Dunning Jr, 1989) and CPCM (Barone & Cossi, 1998) water. The released data provide raw electronic energies without free-energy corrections, so our populations may differ slightly. Experimental times at 298 K are $\sim 10^{-3}$ s (CRN-T) and $\sim 10^{5}$ s (CRN-E) (Dowideit & von

Sonntag, 1998). We set $\sigma' = 4$ kcal,mol^{-1}, since single points on LC-PBE geometries, missing free-energy corrections, and implicit solvation leave a few kcal/mol of barrier uncertainty (§ 2.4).

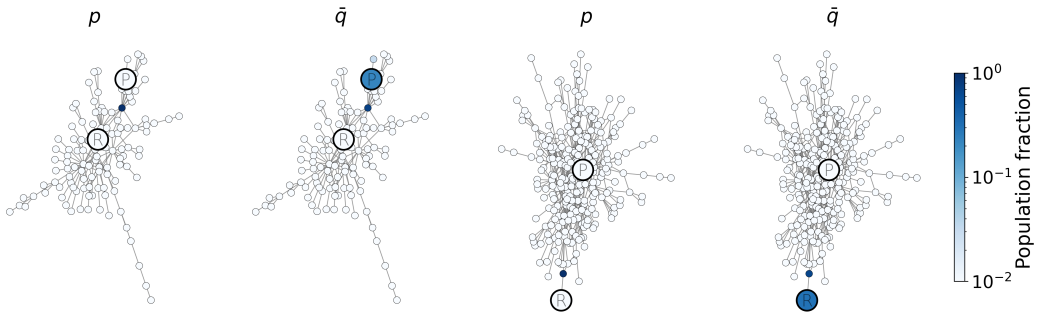


Figure 7: Predicted population fractions on CRN-T (left) and CRN-E (right) under $\boldsymbol{p}$ and $\bar{\boldsymbol{q}}$ at experimental times. Circled R and P mark the entrance flask and product.

CRN-T. In Figure 7 (left), $\bar{\boldsymbol{q}}$ recovers the experimental product at the reported time with $\sim$23% of the mass, while $\boldsymbol{p}$ leaves it empty. Figure 8 explains why: $\boldsymbol{p}$ rises two orders of magnitude too late, whereas $\bar{\boldsymbol{q}}$ reaches appreciable conversion within the experimental window, though it still underestimates the reported yield.

CRN-E. The entrance flask has no low barrier exit toward the reported product. The O_3 + C_2H_4 cycloaddition step is missing from the network, so both predictors drain into an off-pathway trap. Uncertainty propagation reweights existing pathways, it cannot create a missing one.

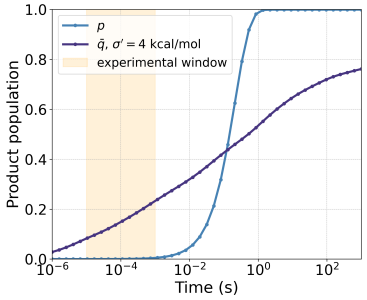


Figure 8: Product population on CRN-T as a function of time.

6 Related Work

Automated network builders (Unsleber & Reiher, 2020; Unsleber et al., 2022; Gugler et al., 2026; Kendall & Zádor, 2026; Petrus et al., 2025) now generate thousands of computed barriers with non-negligible errors. Prior responses to the uncertainty concern act upstream of the prediction. Proppe et al. (2016) propagate barrier uncertainty as a diagnostic, producing error bars on concentrations and ranking which barriers to recompute at higher level. KiNetX and IRES-KIEA turn that ranking into a refinement loop that iterates exploration and higher-level recomputation (Proppe & Reiher, 2019; Bensberg & Reiher, 2024). YAKS uses the propagated spread as an exploration policy, choosing which reaction to add next (Woulfe & Savoie, 2025). We instead leverage the barrier uncertainty in the prediction itself, computing the expected concentration given the barrier distribution.

Beyond the Krylov backbone introduced in Section 2, scalable kinetic solvers include the canonical expm implementation (Sidje, 1998), adaptive schemes with data-driven spectral shifts (Katrutsa et al., 2019), and pre-solver network reductions via graph topology (Hirono et al., 2021) or algebraic lumping (Eilertsen et al., 2026). These methods truncate against a fixed numerical or structural target; our threshold is derived from the barrier uncertainty itself.

7 Conclusion

We introduced a probabilistic kinetic solver that treats energy barriers as distributions and predicts the expected population under barrier uncertainty. It improves RMSE and top-K accuracy across different settings and recovers product formation on a real ozonolysis network where the deterministic solver fails. A limitation of our method is that it assumes unbiased Gaussian barrier errors, while real errors are often systematic and correlated. Coupling this solver with automated, machine-learning-driven reaction network exploration, where manual error curation is infeasible, is a natural next step toward reliable chemistry prediction at scale.

AI USE STATEMENT

We used LLMs in this work to assist with the writing of the paper, for retrieval and discovery of related work, and to assist with writing code for our experiments. All AI-assisted content was reviewed and tested by the authors, and all cited references were verified. We take full responsibility for the final content of this paper.

REFERENCES

Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. Gfn2-xtb—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of chemical theory and computation*, 15 (3):1652–1671, 2019.

Vincenzo Barone and Maurizio Cossi. Quantum calculation of molecular energies and energy gradients in solution by a conductor solvent model. *The Journal of Physical Chemistry A*, 102(11): 1995–2001, 1998.

Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in neural information processing systems*, 35:11423–11436, 2022.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1): 2453, 2022.

Moritz Bensberg and Markus Reiher. Uncertainty-aware first-principles exploration of chemical reaction networks. *The Journal of Physical Chemistry. a*, 128(22):4532, 2024.

Maria Besora, Pietro Vidossich, Agusti Lledos, Gregori Ujaque, and Feliu Maseras. Calculation of reaction free energies in solution: a comparison of current approaches. *The Journal of Physical Chemistry A*, 122(5):1392–1399, 2018.

Wendy D Cornell, Piotr Cieplak, Christopher I Bayly, Ian R Gould, Kenneth M Merz, David M Ferguson, David C Spellmeyer, Thomas Fox, James W Caldwell, and Peter A Kollman. A second generation force field for the simulation of proteins, nucleic acids, and organic molecules. *Journal of the American Chemical Society*, 117(19):5179–5197, 1995.

Christopher J Cramer. *Essentials of computational chemistry: theories and models*. John Wiley & Sons, 2013.

Peter Dowideit and Clemens von Sonntag. Reaction of ozone with ethene and its methyl-and chlorine-substituted derivatives in aqueous solution. *Environmental Science & Technology*, 32 (8):1112–1119, 1998.

Thom H Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen. *The Journal of chemical physics*, 90(2):1007–1023, 1989.

Justin Eilertsen, Valery G Romanovski, Santiago Schnell, and Sebastian Walcher. Lumping of reaction networks: Generic and critical parameters. *arXiv preprint arXiv:2606.28895*, 2026.

Henry Eyring. The activated complex in chemical reactions. *The Journal of chemical physics*, 3(2): 107–115, 1935.

Xiang Fu, Brandon M Wood, Luis Barroso-Luque, Daniel S Levine, Meng Gao, Misko Dzamba, and C Lawrence Zitnick. Learning smooth and expressive interatomic potentials for physical property prediction. *arXiv preprint arXiv:2502.12147*, 2025.

James W Furness, Aaron D Kaplan, Jinliang Ning, John P Perdew, and Jianwei Sun. Accurate and numerically efficient r2scan meta-generalized gradient approximation. *The journal of physical chemistry letters*, 11(19):8208–8215, 2020.

Lars Goerigk, Andreas Hansen, Christoph Bauer, Stephan Ehrlich, Asim Najibi, and Stefan Grimme. A look at the density functional theory zoo with the advanced gmtkn55 database for general main group thermochemistry, kinetics and noncovalent interactions. *Physical Chemistry Chemical Physics*, 19(48):32184–32215, 2017.

Stefan Grimme, Stephan Ehrlich, and Lars Goerigk. Effect of the damping function in dispersion corrected density functional theory. *Journal of computational chemistry*, 32(7):1456–1465, 2011.

Stefan Gugler, Max Eissler, Khaled Kahouli, and Klaus-Robert Müller. Reactionatlas: Ab origine exploration of chemical reaction networks with machine learning. *arXiv preprint arXiv:2606.30778*, 2026.

Yuji Hirono, Takashi Okada, Hiroyasu Miyazaki, and Yoshimasa Hidaka. Structural reduction of chemical reaction networks based on topology. *Physical Review Research*, 3(4):043123, 2021.

Hisayoshi Iikura, Takao Tsuneda, Takeshi Yanai, and Kimihiko Hirao. A long-range correction scheme for generalized-gradient-approximation exchange functionals. *The Journal of Chemical Physics*, 115(8):3540–3544, 2001.

William L Jorgensen, David S Maxwell, and Julian Tirado-Rives. Development and testing of the opls all-atom force field on conformational energetics and properties of organic liquids. *Journal of the american chemical society*, 118(45):11225–11236, 1996.

Alexandr Katrutsa, Mike Botchev, and Ivan Oseledets. Practical shift choice in the shift-and-invert krylov subspace evaluations of the matrix exponential. *arXiv preprint arXiv:1909.13059*, 2019.

Daniel T Kendall and Judit Zádor. Benchmarking the uma foundation interatomic potential for gas-phase chemical kinetics. *The Journal of Physical Chemistry A*, 2026.

Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133, 1965.

Dávid Péter Kovács, J Harry Moore, Nicholas J Browning, Ilyes Batatia, Joshua T Horton, Yixuan Pu, Venkat Kapil, William C Witt, Ioan-Bogdan Magdau, Daniel J Cole, et al. Mace-off: Short-range transferable machine learning force fields for organic molecules. *Journal of the American Chemical Society*, 147(21):17598–17611, 2025.

Paul D Leeson, A Patricia Bento, Anna Gaulton, Anne Hersey, Emma J Manners, Chris J Radoux, and Andrew R Leach. Target-based evaluation of 'drug-like'properties and ligand efficiencies. *Journal of medicinal chemistry*, 64(11):7210, 2021.

Giulia Luise, Chin-Wei Huang, Thijs Vogels, Derk P Kooi, Sebastian Ehlert, Stephanie Lanius, Klaas JH Giesbertz, Amir Karton, Deniz Gunceler, Stefano Battaglia, et al. Accurate and scalable exchange-correlation with deep learning. *arXiv preprint arXiv:2506.14665*, 2025.

Narbe Mardirossian and Martin Head-Gordon. ωb97x-v: A 10-parameter, range-separated hybrid, generalized gradient approximation density functional with nonlocal correlation, designed by a survival-of-the-fittest strategy. *Physical Chemistry Chemical Physics*, 16(21):9904–9924, 2014.

Art B Owen. Scrambling sobol'and niederreiter–xing points. *Journal of complexity*, 14(4):466–489, 1998.

John P Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical review letters*, 77(18):3865, 1996.

Enric Petrus, Livia A Hunkeler, Markus Reiher, Urs Von Gunten, and Thomas B Hofstetter. Automated reaction exploration of ozonation processes for model olefins in water. *Environmental Science & Technology*, 59(49):26806–26818, 2025.

Jonny Proppe and Markus Reiher. Mechanism deduction from noisy chemical reaction networks. *Journal of chemical theory and computation*, 15(1):357–370, 2019.

Jonny Proppe, Tamara Husch, Gregor N Simm, and Markus Reiher. Uncertainty quantification for quantum chemical models of complex reaction networks. *Faraday discussions*, 195:497–520, 2016.

George D Purvis III and Rodney J Bartlett. A full coupled-cluster singles and doubles model: The inclusion of disconnected triples. *The Journal of chemical physics*, 76(4):1910–1918, 1982.

Christoph Riplinger and Frank Neese. An efficient and near linear scaling pair natural orbital based local coupled cluster method. *The Journal of chemical physics*, 138(3), 2013.

Yousef Saad. Analysis of some krylov subspace approximations to the matrix exponential operator. *SIAM Journal on Numerical Analysis*, 29(1):209–228, 1992.

Roger B Sidje. Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software (TOMS)*, 24(1):130–156, 1998.

Ilya M Sobol. Distribution of points in a cube and approximate evaluation of integrals. *USSR Computational mathematics and mathematical physics*, 7:86–112, 1967.

Jan P Unsleber and Markus Reiher. The exploration of chemical reaction networks. *Annual review of physical chemistry*, 71(1):121–142, 2020.

Jan P Unsleber, Stephanie A Grimmel, and Markus Reiher. Chemoton 2.0: autonomous exploration of chemical reaction networks. *Journal of Chemical Theory and Computation*, 18(9):5393, 2022.

Gerhard Wanner and Ernst Hairer. *Solving ordinary differential equations II*, volume 375. Springer Berlin Heidelberg New York, 1996.

Florian Weigend and Reinhart Ahlrichs. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for h to rn: Design and assessment of accuracy. *Physical chemistry chemical physics*, 7(18):3297–3305, 2005.

JH Wilkinson. The algebraic eigenvalue problem. *Handbook for automatic computation*, 2:191, 1971.

Brandon Wood, Misko Dzamba, Xiang Fu, Meng Gao, Muhammed Shuaibi, Luis Barroso-Luque, Kareem Abdelmaqsoud, Vahe Gharakhanyan, John Kitchin, Daniel Levine, et al. Uma: A family of universal models for atoms. *Advances in Neural Information Processing Systems*, 38:129391–129427, 2026.

Michael Woulfe and Brett M Savoie. Chemical reaction networks from scratch with reaction prediction and kinetics-guided exploration. *Journal of Chemical Theory and Computation*, 21(3): 1276–1291, 2025.

Yan Zhao and Donald G Truhlar. The m06 suite of density functionals for main group thermochemistry, thermochemical kinetics, noncovalent interactions, excited states, and transition elements: two new functionals and systematic testing of four m06-class functionals and 12 other functionals. *Theoretical chemistry accounts*, 120(1):215–241, 2008.

A Bias-Variance Decomposition of Population Error

The population MSE of Eq. 7 decomposes as

$$\mathrm{MSE}_{\boldsymbol{p}} = \frac{1}{2}\sum_{i=1}^{M}\left(\mathbb{E}_{\boldsymbol{\varepsilon}}[\bar{p}_i] - p_i\right)^2 + \frac{1}{2}\sum_{i=1}^{M}\mathrm{Var}_{\boldsymbol{\varepsilon}}[\bar{p}_i]. \tag{11}$$

The first term measures systematic error, while the second measures sensitivity to perturbations in the estimated barriers. The variance term is exactly $\mathcal{V}(\sigma')$ defined in the main text. For $M = 2$, population conservation makes the two component errors equal and opposite, so the factor $1/2$ recovers the two-state metric of Eq. 4.

B Eigenvalue Perturbation Under Log-Normal Rate Noise

We show that at first order in the barrier noise, the absolute uncertainty of every eigenvalue is proportional to σ'/k_BT times its own magnitude.

Setup. Let $\boldsymbol{K}$ have off-diagonal entries $K_{ab} = k_{ab}$ (rate $b \to a$) and diagonal $K_{bb} = -\sum_{a\neq b} k_{ab}$, and let $\boldsymbol{v}_i, \boldsymbol{w}_i$ be the right and left eigenvectors of a simple eigenvalue λ_i with $\boldsymbol{w}_i^\top \boldsymbol{v}_i = 1$. As in Section 4.2, t is absorbed into $\boldsymbol{K}$. Barrier noise enters the Eyring rates multiplicatively,

$$\tilde{k}_{ab} = k_{ab}\, e^{\varepsilon_{ab}}, \qquad \varepsilon_{ab} \sim \mathcal{N}(0, s^2) \text{ independent}, \qquad s \equiv \frac{\sigma'}{k_BT}. \tag{12}$$

To first order, $\delta K_{ab} = k_{ab}\varepsilon_{ab}$ off diagonal and $\delta K_{bb} = -\sum_{a\neq b} k_{ab}\varepsilon_{ab}$ on it, by conservation of probability.

Derivation. Wilkinson's first-order shift $\delta\lambda_i = \boldsymbol{w}_i^\top \delta\boldsymbol{K}\boldsymbol{v}_i$ (Wilkinson, 1971) together with $\lambda_i = \boldsymbol{w}_i^\top \boldsymbol{K}\boldsymbol{v}_i$ yields

$$\lambda_i = \sum_{a\neq b} g_{ab}^{(i)}, \qquad \delta\lambda_i = \sum_{a\neq b} g_{ab}^{(i)}\,\varepsilon_{ab}, \qquad g_{ab}^{(i)} \equiv k_{ab}\, v_i(b)\big(w_i(a) - w_i(b)\big), \tag{13}$$

so the eigenvalue and its perturbation share the same edge coefficients. Independence of the ε_{ab} then gives

$$\sqrt{\mathrm{Var}(\delta\lambda_i)} = s\,\|\boldsymbol{g}^{(i)}\|_2 = r_i\,\frac{\sigma'}{k_BT}\,|\lambda_i|, \qquad r_i \equiv \frac{\|\boldsymbol{g}^{(i)}\|_2}{\left|\sum_{a\neq b} g_{ab}^{(i)}\right|} \geq \frac{1}{\sqrt{M(M-1)}}. \tag{14}$$

Since r_i is bounded away from zero, $\sqrt{\mathrm{Var}(\delta\lambda_i)} \propto (\sigma'/k_BT)|\lambda_i|$: the absolute uncertainty of every mode is proportional to σ'/k_BT times its own magnitude, up to the bounded factor r_i.

C Shift-and-Invert Krylov Propagation

This appendix records the propagation scheme summarized in Section 4.2.

Projection. Arnoldi iteration on $\boldsymbol{A} = (\boldsymbol{I} - t\boldsymbol{K})^{-1}$, started from $\boldsymbol{p}^{\mathrm{init}}$, returns an orthonormal basis $\boldsymbol{V} \in \mathbb{R}^{M\times r}$ of the Krylov subspace $\mathrm{span}\{\boldsymbol{p}^{\mathrm{init}}, \boldsymbol{A}\boldsymbol{p}^{\mathrm{init}}, \ldots, \boldsymbol{A}^{r-1}\boldsymbol{p}^{\mathrm{init}}\}$ together with the projected operator $\boldsymbol{H} = \boldsymbol{V}^\top \boldsymbol{A}\boldsymbol{V} \in \mathbb{R}^{r\times r}$. The populations follow as

$$\begin{aligned}
\boldsymbol{p}(t) &= \boldsymbol{U}\,\mathrm{diag}(e^{t\lambda_i})\,\boldsymbol{U}^{-1}\boldsymbol{p}^{\mathrm{init}} && \text{eigendecomposition } (M\times M)\\
&\approx \boldsymbol{V}\,f(\boldsymbol{H})\,\boldsymbol{V}^\top\boldsymbol{p}^{\mathrm{init}} && \text{Krylov approximation } (r\times r)\\
&= \boldsymbol{V}\,\boldsymbol{Q}\,\mathrm{diag}(e^{t\hat{\lambda}_i})\,\boldsymbol{Q}^{-1}\,\boldsymbol{V}^\top\boldsymbol{p}^{\mathrm{init}} && \text{expand } f(\boldsymbol{H}),\ \hat{\lambda}_i = (1-1/h_i)/t \approx \lambda_i\\
&= \sum_{i=1}^{r} \hat{c}_i\,\hat{\boldsymbol{u}}_i\,e^{t\hat{\lambda}_i} && \hat{c}_i = \big(\boldsymbol{Q}^{-1}\boldsymbol{V}^\top\boldsymbol{p}^{\mathrm{init}}\big)_i,\ \hat{\boldsymbol{u}}_i = (\boldsymbol{V}\boldsymbol{Q})_{:,i}
\end{aligned} \tag{15}$$

where h_i and $\boldsymbol{Q} \in \mathbb{R}^{r\times r}$ are the eigenvalues and eigenvectors of $\boldsymbol{H}$. The transfer function is $f(h) = \exp(1 - 1/h)$, which composes the shift-and-invert map $t\hat{\lambda} = 1 - 1/h$ with the exponential in a single evaluation. Entries with $\mathrm{Re}(h) \leq 0$ correspond to modes outside the physical spectrum and are set to zero.

Algorithm 2 Shift-and-invert Arnoldi propagation for one rate matrix

Input: Rate matrix $\boldsymbol{K}$, time t, initial population $\boldsymbol{p}^{\text{init}} = \boldsymbol{e}_1$, maximum rank $r_{\max}$, tolerance τ
Output: Population approximation $\boldsymbol{q}(t) \approx e^{t\boldsymbol{K}}\boldsymbol{p}^{\text{init}}$

$\mathcal{F} \leftarrow \textsc{SparseLU}(\boldsymbol{I} - t\boldsymbol{K})$ ▷ shift-and-invert factorization, computed once
$f(h) \leftarrow \exp(1 - 1/h)$ ▷ maps $\text{spec}(\mathcal{F}^{-1}) \to \text{spec}(e^{t\boldsymbol{K}})$
$\boldsymbol{V}_1 \leftarrow [\boldsymbol{p}^{\text{init}}]$; $\boldsymbol{q}^{(0)} \leftarrow \boldsymbol{0}$
for $r = 1$ **to** $r_{\max}$ **do**
  $(\boldsymbol{V}_{r+1}, \boldsymbol{H}_r) \leftarrow \textsc{ArnoldiStep}(\mathcal{F}, \boldsymbol{V}_r, \boldsymbol{H}_{r-1})$
  $\boldsymbol{q}^{(r)}(t) \leftarrow \textsc{PopulationFromSubspace}(\boldsymbol{V}_r, \boldsymbol{H}_r)$
  if $\|\boldsymbol{q}^{(r)}(t) - \boldsymbol{q}^{(r-1)}(t)\|_1 < \tau$ **then**
    return $\boldsymbol{q}^{(r)}(t)$ ▷ converged
return $\boldsymbol{q}^{(r_{\max})}(t)$

function $\textsc{ArnoldiStep}(\mathcal{F}, \boldsymbol{V}_r, \boldsymbol{H}_{r-1})$
  $\boldsymbol{w} \leftarrow \textsc{Solve}(\mathcal{F}, \boldsymbol{v}_r)$
  Orthogonalize $\boldsymbol{w}$ twice against $\boldsymbol{V}_r$; store coefficients in column r of $\boldsymbol{H}_r$ ▷ double Gram–Schmidt
  $\beta_r \leftarrow \|\boldsymbol{w}\|_2$; $\boldsymbol{H}_r[r{+}1, r] \leftarrow \beta_r$; $\boldsymbol{v}_{r+1} \leftarrow \boldsymbol{w}/\beta_r$
  return $(\boldsymbol{V}_{r+1}, \boldsymbol{H}_r)$

function $\textsc{PopulationFromSubspace}(\boldsymbol{V}_r, \boldsymbol{H}_r)$
  $(\{h_i\}, \boldsymbol{Q}) \leftarrow \textsc{EigenDecompose}(\boldsymbol{H}_r)$
  $\boldsymbol{c} \leftarrow \textsc{Solve}(\boldsymbol{Q}, \boldsymbol{e}_1)$
  $\boldsymbol{q} \leftarrow \boldsymbol{V}_r\ \text{Re}\big[\boldsymbol{Q}\,(f(h_i) \odot \boldsymbol{c})\big]$
  Clip negative entries and normalize $\boldsymbol{q}$
  return $\boldsymbol{q}$

D Synthetic Network Examples

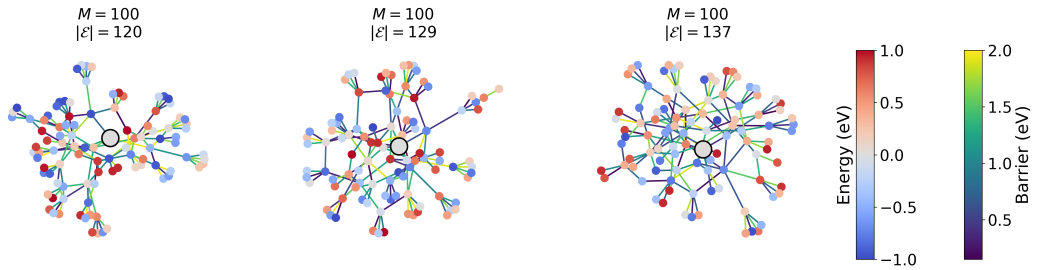


Figure 9: Three representative synthetic network realizations with $M = 100$ states and ring-closure probability $p = 0.35$. The outlined node is the initial reactant. Node and edge colors indicate state energies and activation barriers, respectively. $|\mathcal{E}|$ is the number of elementary reaction steps.