university of groningen

Faculty of Science and Engineering

Master Thesis

Modelling Exciton Dissipation via the Numerical Integration of the Schrödinger Equation

From Linear Shifts to Non-Markovian Memory Kernels

Author:
Senan Bird
S5407729

Supervisor:
Prof. dr. T.L.C. Jansen

MSc Physics (Quantum Universe)

February 2026

Abstract

Numerical Integration of the Schrödinger Equation is a methodology used to simulate exciton transfer in molecular aggregates. By treating the electronic system and the nuclear bath within a mixed quantum-classical framework, Numerical Integration of the Schrödinger Equation simplifies the complexity of these interactions by assuming a high-temperature limit. The system Hamiltonian is propagated via a wavefunction that evolves according to an Ornstein–Uhlenbeck process, which simulates environmental noise and system-bath coupling. However, this classical approximation does not fully capture the thermalisation of molecular interactions or the Stokes shift. Unlike numerically exact methods, such as Hierarchy of Pure States or Hierarchical Equations of Motion, standard Numerical Integration of the Schrödinger Equation fails to reach a proper Boltzmann distribution.

This thesis implements two computationally efficient schemes to incorporate the Stokes shift into the Numerical Integration of the Schrödinger Equation framework. The first method utilises a Linear Reorganisation Shift within the Hamiltonian, updated at each propagation step. The second method employs a memory kernel, allowing the Hamiltonian to retain a temporary "memory" of the dissipative effects. Both modifications demonstrate improved accuracy relative to the standard Numerical Integration of the Schrödinger Equation model. The memory kernel, in particular, yields superior long-term results as the system reaches the steady state, while the Linear Reorganisation Shift offers occasional improvements under specific parameter sets in the short/medium term before the system/bath has stabilised. These methods provide a scalable solution for simulating large-scale molecular networks, significantly enhancing accuracy with limited additional computational overhead.

Contents

Acknowledgements **1**

1 Introduction **2**

2 Theory **6**
2.1 Frenkel Exciton Model 6
2.2 NISE 8
2.3 Stokes Shift 12
2.4 HOPS 20
2.5 HEOM 24

3 Methods **26**
3.1 NISE 26
3.2 Linear Reorganisation Shift 27
3.3 Memory Kernel Shift 29

4 Results **32**
4.1 Average Site Energy 32
4.2 Exciton Population Dynamics 36

5 Discussion **45**

6 Conclusion **47**

A Appendix **49**
A.1 Computational Scaling of the HOPS Method 49
A.2 Reversing Energy 50

Acknowledgements

I would like to thank my supervisor, Prof. dr. T.L.C. Jansen, for his guidance and patience throughout this thesis. He provided invaluable insights into the physics of exciton population transfer, while his mentorship in computational techniques—particularly navigating high-performance computing clusters and developing foundational skills in C—has been really valuable and hopefully I can transfer these skills with me in the future.

I'd like to acknowledge the work that Hans Douwe Wobben has done on his thesis for creating the HOPS program. I used this heavily throughout my thesis as a benchmark and for important sanity checks. I am also grateful to the Theory of Condensed Matter research group for providing a welcoming atmosphere that made this year a memorable experience. Additionally, I would like to thank my parents and my friend Alisha for their time in proofreading this manuscript and providing helpful feedback that improved the clarity and readability of my thesis.

All simulations presented in this paper were conducted in the High Performance Computing cluster, Hábrók. All HOPS simulations were done using Hans's code that was run in Hábrók.

AI was used in the debugging of the Python and C/C++ programs and in the creation of the plots to analyse the data. Grammarly was used when writing to help with grammar, spell checking and improving readability. Gemini was used to improve the LaTeX code when formatting this document in Overleaf.

Chapter 1

Introduction

In the initial steps of natural photosynthesis, absorbed sunlight is converted into excitation energy with near-unity quantum efficiency [1]. This efficiency relies on the strictly directional flow of energy through dense networks of pigment-protein complexes (PPCs). A well-studied model system for understanding these ultrafast exciton dynamics is the Light-Harvesting Complex 2 (LH2) found in photosynthetic purple bacteria, as shown in Figure 1.1.

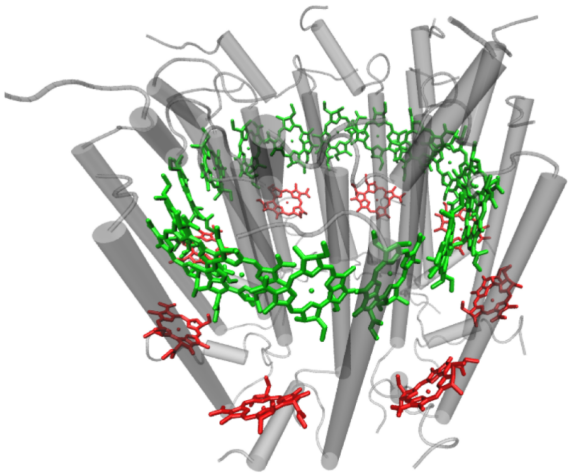

Figure 1.1. Schematic overview of the LH2 complex of Rhodospirillum molischianum. The BChla chromophores form two concentric rings mounted on the protein scaffold. The B850 ring coloured in green contains twice as many chromophores as the B800 ring shown in red.[2]

Accounting for the directional energy flow is computationally demanding as it requires an explicit description of the feedback from the bath. At room temperature, the surrounding protein matrices undergo continuous thermal fluctuations [1]. Because the pigments within LH2 are densely packed, their electrodynamic couplings are acutely sensitive to these fluctuating intermolecular distances,

leading to significant energetic disorder [1].

The core principle of natural light-harvesting is its ability to maintain a "downhill" energy flow despite the constant thermal noise created by the vibrating protein environment [1]. To properly model these dynamics in complexes like LH2, or to design better artificial light-harvesting technologies [3], we need theoretical methods that can accurately simulate how energy relaxes.

Numerically exact methods for calculating exciton population transfer exist, such as Hierarchy of Pure States (HOPS) [4] and Hierarchical Equations of Motion (HEOM) [5]. These methods calculate the full Hamiltonian that captures the system and the system bath coupling:

$$H = H_{\mathrm{S}} + H_{\mathrm{B}} + H_{\mathrm{SB}} \tag{1.1}$$

While these methods are numerically exact, they are computationally expensive. For instance, HEOM is calculated using a reduced density matrix, which requires $N \times N$ memory space for a system with N eigenstates. HOPS does improve on this by taking a wavefunction approach using complex time contour integrals. This allows the memory requirements for an N-level system to scale linearly with N [5].

However, despite this spatial memory advantage, the computational time required by HOPS scales aggressively when applied to extended physical networks like LH2. To accurately capture the persistent memory effects-or non-Markovian dynamics-of the sluggish protein scaffold, HOPS introduces a vast mathematical hierarchy of auxiliary wavefunctions to account for the interaction of the bath.

The size of this hierarchy depends entirely on the number of environmental modes coupled to the system. In photosynthetic complexes, each pigment typically interacts with its own independent local thermal bath[6]. Consequently, as the number of system sites N increases, the number of distinct environmental bath modes grows proportionally.

This linear increase in bath modes triggers a combinatorial explosion in the number of auxiliary equations that must be propagated simultaneously. Theoretically, the total computational cost per time step scales as a severe polynomial, roughly $\mathcal{O}\left(N^{k_{\max}}\right)$ (see A.1 for more information), where $k_{\max}$ is the truncation depth of the hierarchy. Therefore, while HOPS elegantly sidesteps the density matrix memory bottleneck, simulating large, densely packed biological networks with independent local environments remains heavily constrained by computational time.

By making the classical high-temperature assumption $k_B T \gg \hbar\omega$ [7], we can approximate the bath and the system-bath Hamiltonian $H_B + H_{SB}$ via a fluctuating stochastic field. This cuts down on the computational time as we are no longer computing this quantum mechanically. This is the basis for the Numerically Integrated Schrödinger Equation (NISE)[8]. A primary advantage of NISE is its

superior scalability for large systems compared to exact methods like HOPS or HEOM [9]. While NISE typically exhibits $\mathcal{O}\left(N^3\right)$ scaling, the implementation of a 'sparse' propagation method, which assumes that most off-diagonal coupling elements are zero—can further reduce the scaling to $\mathcal{O}\left(N^2\right)$ [9].

When making the assumption of $k_B T \gg \hbar\omega$, this neglects the "back reaction" of the system onto the bath[10]. This means that the Stokes shift gets neglected as the reorganisation energy tends to zero as the temperature approaches infinity. This causes NISE to improperly simulate the thermodynamic effects of the system/bath. TNISE is a solution to imitate the thermalised effect and bring the population down to the Boltzmann distribution, but it doesn't explicitly capture the Stokes shift[8]. Another example in a recent paper, a machine-learned correction term has been developed for the NISE approach that can be trained on smaller, numerically feasible systems to accurately capture dissipative quantum dynamics and the Stokes shift in larger systems[11].

While this thesis primarily focuses on exciton transport within light-harvesting antennas like LH2, accurately modelling the Stokes shift has implications for the immediate next step in energy conversion: charge transfer. In natural photosynthesis, electronic excitations must be delivered to specialised reaction centers to initiate charge separation [12]. Similarly, in artificial technologies like organic photovoltaics (OPV), excitons reach an interface where the excited electron and hole separate to form a charge transfer (CT) state. The Stokes shift—driven by the reorganisation of the surrounding environment—is absolutely essential for this process. As the environment reorganises around the newly separated charges, it lowers the energy of the CT state, effectively creating an energetic "trap." This thermodynamic relaxation prevents the electron from jumping back and recombining with the hole, ensuring that the flow of energy remains strictly directional. Because standard NISE lacks a mechanism for this dissipation, it cannot accurately model the efficiency of charge separation. Therefore, introducing the Stokes shift into the NISE framework is a necessary step toward simulating the full energy conversion process from initial sunlight capture to stable charge separation [12].

The fundamental difference in how these methods handle system-bath interactions is illustrated in Figure 1.2. Standard NISE (Panel A) employs a one-way stochastic driving mechanism; the classical bath perturbs the system, but the system cannot exert a back-action on the bath. Consequently, the lack of dissipation drives the system toward an infinite-temperature steady state, resulting in equal site populations. In contrast, numerically exact methods like HEOM and HOPS (Panel B) incorporate bidirectional feedback. This allows the system to dissipate energy back into the quantum hierarchy, enabling it to properly thermalise to a steady-state Boltzmann distribution.

To address these limitations, this thesis proposes a middle ground between the exactness of HEOM/HOPS and the efficiency of NISE. By implementing two distinct correction schemes—the Linear Reorganisation Shift and the Memory Kernel Shift—we attempt to artificially re-introduce the 'back-

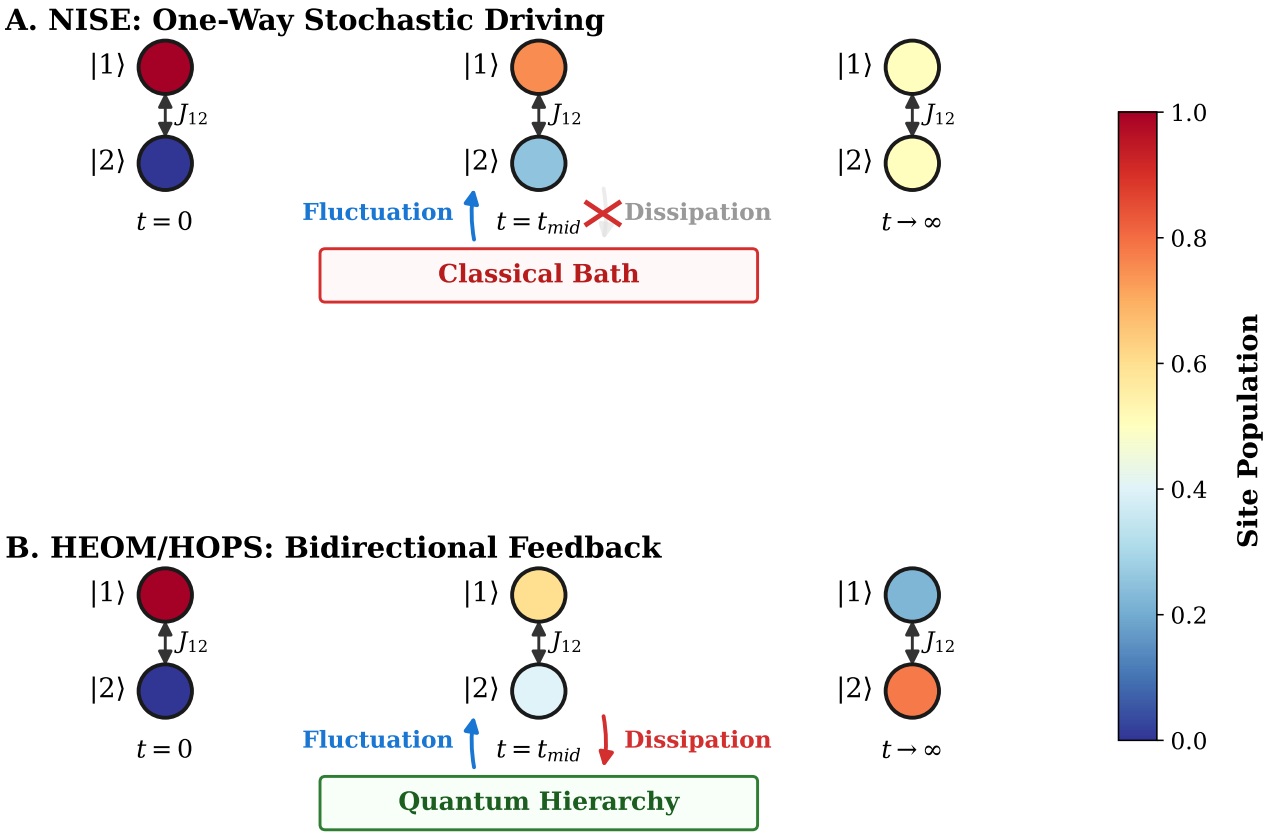

Figure 1.2. Schematic comparison of exciton dynamics and thermalisation in a dimer model. **(A)** The NISE approach lacks system-to-bath back-action (dissipation), driving the system to an equal population distribution at $t \to \infty$. **(B)** HEOM/HOPS includes bidirectional feedback, allowing the system to relax to a proper thermal equilibrium.

action' of the system onto the bath. These modifications aim to recover the Stokes shift within a mixed quantum-classical framework, thereby enabling the simulation of true thermodynamic relaxation without the combinatorial cost of auxiliary hierarchies like HOPS or HEOM.

The remainder of this thesis is structured as follows: Section 2 establishes the theoretical groundwork, building from the foundations of the Frenkel Exciton Model [13]. This section provides a detailed explanation of the NISE method and describes how HOPS and HEOM simulate exciton transfer with numerical accuracy. We specifically justify the selection of HOPS as the "gold standard" against, which our two distinct NISE correction schemes are benchmarked.

The implementation of the Linear Reorganisation Shift and Memory Kernel Shift methods is covered in Section 3, with the resulting data presented in Section 4. Finally, the thesis concludes with a discussion of these results and their implications for future research.

Chapter 2

Theory

2.1 Frenkel Exciton Model

In solid state physics, an exciton is a bound pair of a negatively charged electron and a positively charged hole that is bound together by their mutual coulomb attraction to form a bound state [14]. In the Frenkel Exciton model, the exciton is tightly bound. This means the excited electron and the hole it leaves behind reside on the same specific atom or molecule within a crystal lattice. Unlike the Wannier-Mott model (where the electron and hole are loosely bound and orbit each other over a distance of many unit cells) [15], the Frenkel Exciton has a radius comparable to the size of the unit cell itself. Because it is so tightly bound, it acts as a highly localised excitation that can hop from one site (molecule/atom) to another.

The Frenkel Exciton is typically described as a Frenkel Exciton Hamiltonian. The Hamiltonian considers a system of N sites that can be excited from the ground state. For instance a site could be a molecule and system of 10 sites could be 10 molecules all bound together. Each site can host an exciton localised on this site. If the system has a single localised excitation at site m while all other sites remain in the ground state, then it can be expressed as the creation operator $\left|m\right\rangle = e_m^\dagger \left|0\right\rangle$ where $|0\rangle$ indicates the aggregate ground state (the vacuum state with zero excitons). The $e_m^\dagger$ generates an exciton that is localised on the ground state on the mth site. In a system of N sites, one site will be excited in the excited state $|1\rangle$ while the remaining sites will be in their ground state $|0\rangle$.In a system of N sites the Hamiltonian describing the system can be written as:

$$H_S = \sum_{m=1}^{N} \epsilon_m |m\rangle\langle m| + \sum_{m=1}^{N} \sum_{n \neq m}^{N} J_{mn} |m\rangle\langle n| \tag{2.1}$$

where H_S denotes the system Hamiltonian, ϵ_m represents the energy of site m and J represents the

resonance coupling between states m and n. This off-diagonal term is responsible for the transfer of the excitation between molecules, allowing the localised excited states to mix and form delocalised exciton eigenstates. [16] In a dynamic system, ϵ is time-dependent and fluctuates over time. In theory the J coupling may change over time as molecules move relative to each other, but for simplicity, these have been left as a constant in our simulations so our system Hamiltonian that is time dependent becomes:

$$H_s(t) = \sum_{m=1}^{N} \epsilon(t)_m |m\rangle\langle m| + \sum_{m=1}^{N} \sum_{n \neq m}^{N} J_{mn} |m\rangle\langle n| \tag{2.2}$$

Showing this in Matrix form for a two site dimer system this is written:

$$H = \begin{pmatrix} \epsilon_1 & J_{12} \\ J_{12}^* & \epsilon_2 \end{pmatrix} \tag{2.3}$$

To find the exciton eigenstates of the system, we must diagonalise the Hamiltonian matrix. Physically, this process transforms our perspective from a basis of localised sites (where the excitation sits on a single molecule) to an exciton basis (where the excitation is a delocalised wave shared across the aggregate).

Mathematically, this involves solving the eigenvalue problem:

$$H_S|\alpha\rangle = E_\alpha|\alpha\rangle \tag{2.4}$$

where $|\alpha\rangle$ represents the delocalised exciton eigenstate and E_α is its corresponding energy eigenvalue. The eigenstate $|\alpha\rangle$ is formed by a linear combination of the localised site states, $|\alpha\rangle = \sum_{m=1}^{N} c_{\alpha m}|m\rangle$, where the expansion coefficients $c_{\alpha m}$ define the probability amplitude of the excitation residing on site m.

This gives a snapshot of the exciton eigenstates at a moment in time. To understand how this system evolves we have to understand the Time-Dependent Schrödinger Equation (TDSE). This allows us to track how the exciton moves, decoheres and transfers energy to the site aggregate. This requires the tracking of the systems overall wavefunction $|\psi(t)\rangle$. The dynamics of this wavefunction is given by Equation 2.5:

$$i\hbar \frac{\partial}{\partial t}|\psi(t)\rangle = H_S(t)|\psi(t)\rangle \tag{2.5}$$

In an idealised static system, this is a straightforward analytical solution. If the thermal bath is completely frozen, the Hamiltonian does not change over time ($H_S(t) = H_S$). In this idealised scenario,

the TDSE is:

$$i\hbar\frac{\partial}{\partial t}|\psi(t)\rangle = H_S|\psi(t)\rangle \tag{2.6}$$

Because H_S is a constant operator with respect to time, this is a standard first-order linear differential equation. We can solve it by integrating both sides, which yields the formal analytical solution:

$$|\psi(t)\rangle = \exp\left(-\frac{i}{\hbar}H_S t\right)|\psi(0)\rangle \tag{2.7}$$

Here, $|\psi(0)\rangle$ is the initial state of the exciton at time $t = 0$, and the exponential term is known as the time-evolution operator, $U(t, 0)$ [17].

However, as established in Equation 2.2, our system is coupled to a dynamic thermal bath. The site energies $\epsilon_m(t)$ continuously and stochastically fluctuate as the surrounding molecules vibrate and move. Because the Hamiltonian $H_S(t)$ behaves as a complex, fluctuating variable (often extracted from classical molecular dynamics simulations), the TDSE cannot be solved analytically. To calculate the exciton propagation over time, we must solve this equation numerically by discretizing time into small steps. This is the core purpose of the Numerical Integration of the Schrödinger Equation (NISE) methodology [8].

2.2 NISE

To numerically solve the Time Dependent Schrödinger Equation, NISE adjusts the Hamiltonian to include a fluctuating noise term to approximate the thermal effects of the bath. In this mixed quantum-classical approximation, the system is treated explicitly via the TDSE, while the bath is approximated as a classical stochastic field. The altered Hamiltonian becomes $H = H_S + H_B + H_{SB}$ Where the H_S is equal to Equation (2.2) and the $H_B + H_{SB}$ is the term that includes the effects of the bath and the effects of the coupling between the bath and the system. NISE solves $H_B + H_{SB}$ with a classical approximation. We use the Ornstein-Uhlenbeck (OU) process driven by Brownian motion[8] as this allows us to compare with exact models that use the OU process. This transforms the Hamiltonian to:

$$H(t) = \begin{pmatrix} \epsilon_1 + \delta\epsilon_1(t) & J \\ J & \epsilon_2 + \delta\epsilon_2(t) \end{pmatrix} \tag{2.8}$$

In this simulation, we keep the excitonic coupling (J) constant to focus on how energy fluctuations at each site ($\delta\epsilon(t)$) affect the system. In a real environment, molecular vibrations cause both the

site energies and the coupling between molecules to change. However, these changes in coupling (off-diagonal disorder) are usually much smaller and less influential on the final spectrum than the site energy fluctuations (diagonal disorder). By holding J constant, we isolate the most important physical effects of the environment while keeping the simulation efficient and easier to interpret [18].

The Ornstein-Uhlenbeck (OU) process driven by Brownian motion is essentially random motion attached to a string. Typically, randomly generated trajectories have a tendency to "walk away" and drift from the mean, as shown in Figure 2.1.

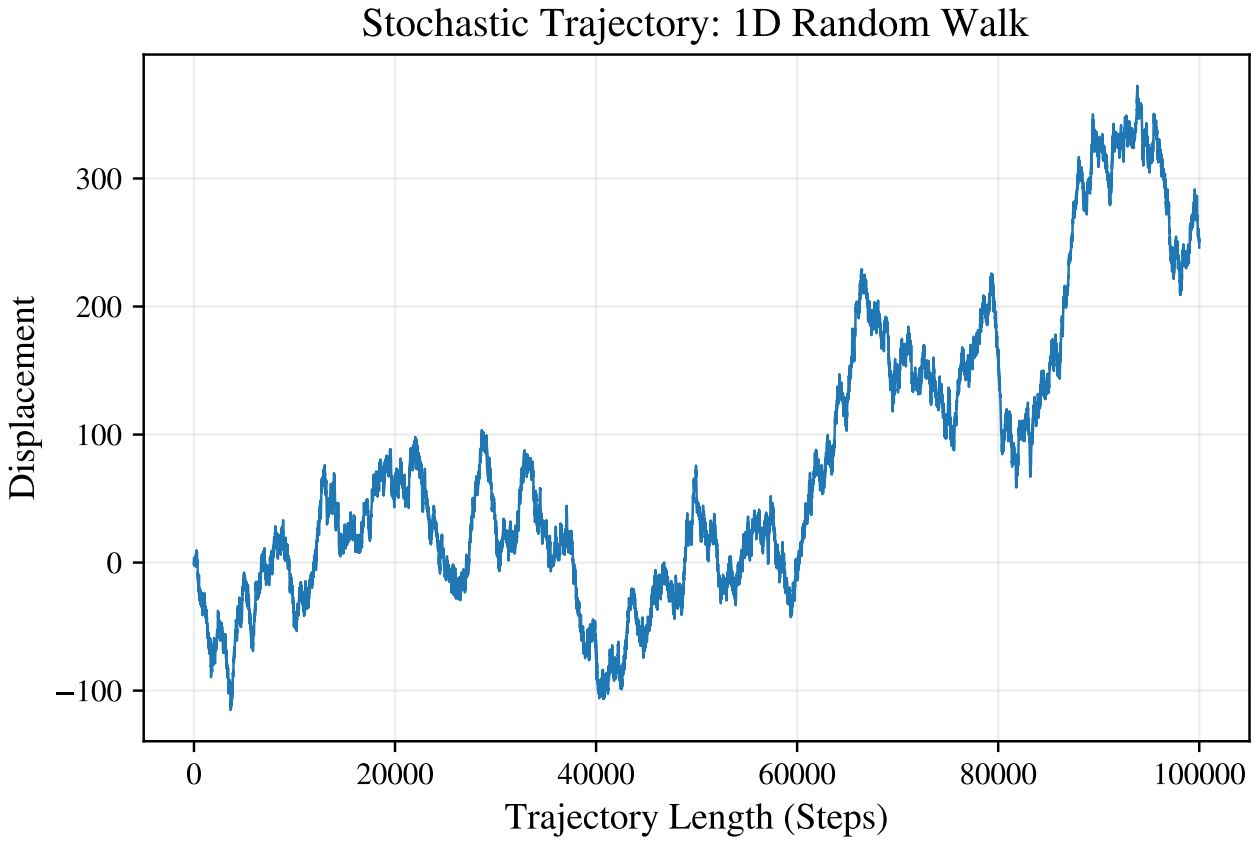

Figure 2.1. Trajectory of an unbounded Gaussian random walk ($n = 100,000$). The plot shows the cumulative sum of independent, normally distributed increments starting from an initial displacement of zero.

The OU process driven by Brownian motion adds a "mean reverting" force to the random walk. When the random walk strays too far from the mean, a restoring force pulls it back. The standard stochastic differential equation (SDE) for an OU process X_t is written as [19]:

$$dX_t = -\frac{1}{\tau}X_t dt + \sqrt{\frac{2}{\tau}}dW_t \tag{2.9}$$

Where X_t is the fluctuating variable and $-\frac{1}{\tau}X_t dt$ is the drift term. The parameter τ is the correlation time. dW_t is the purely random Brownian motion (Wiener process) driving the fluctuations. In physics, this model is also denoted as the overdamped Brownian harmonic oscillator. This describes a particle moving in a harmonic potential where the environmental damping is so strong that inertial effects are negligible. Here, the site energy is treated as being "trapped" in a local potential well, where it is constantly perturbed by thermal noise, but always pulled back toward its mean by a restorative force. Because of the restoring force, the OU process has a constant, bounded variance. This makes it perfect for modelling thermal fluctuations of molecular site energies, which jitter around a mean value but stay confined by the local molecular environment, as shown in Figure 2.2.

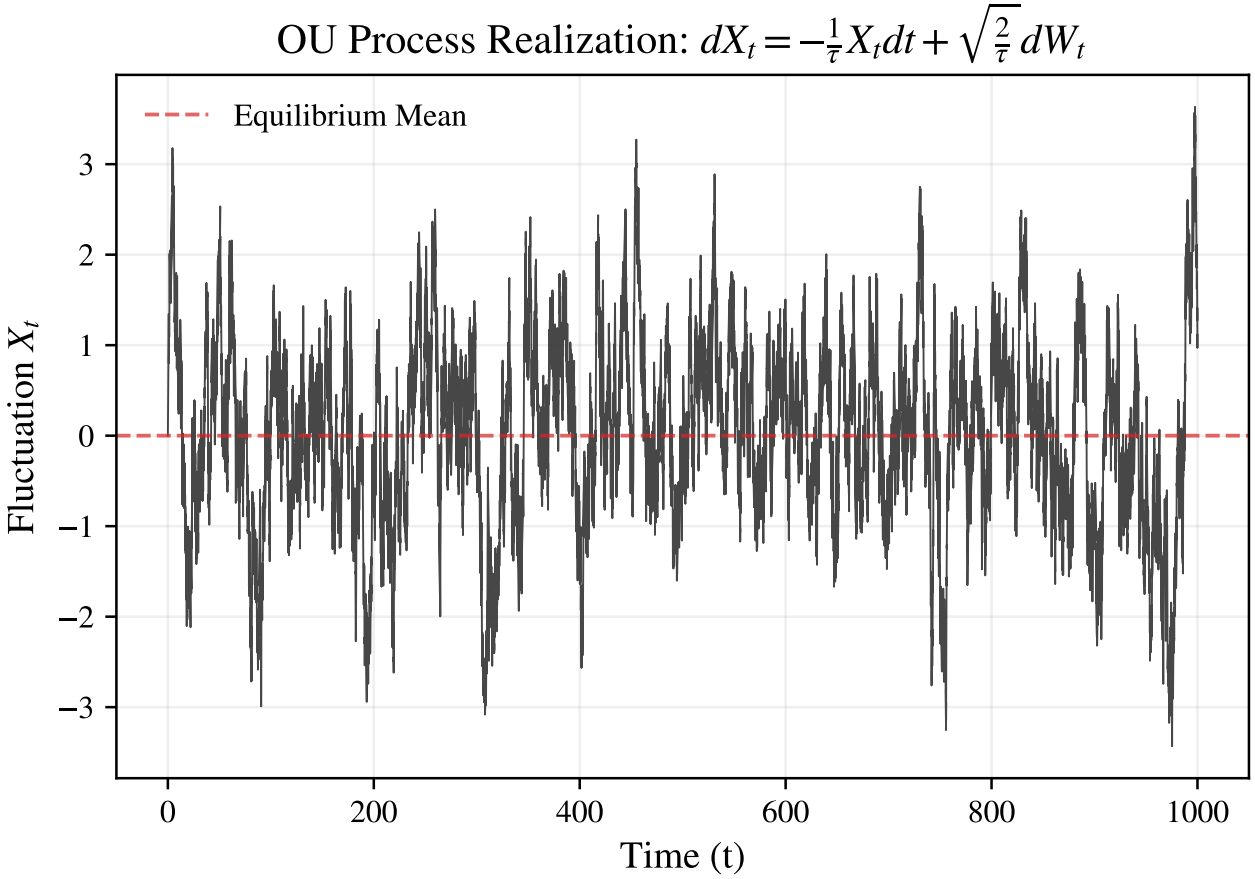

Figure 2.2. realisation of an Ornstein-Uhlenbeck process via Euler-Maruyama discretization ($\tau = 5.0, dt = 0.01$). The trajectory demonstrates mean-reverting stochastic behaviour governed by $dX_t = -\frac{1}{\tau}X_t dt + \sqrt{\frac{2}{\tau}}dW_t$. The restorative drift $-\frac{1}{\tau}X_t$ pulls the system toward $\mu = 0$, while the diffusion term maintains a stationary variance of $\sigma^2 = 1$, representing a balance between dissipative decay and stochastic fluctuations.

With the Hamiltonian discretised, we can now approximate how the quantum state evolves under the fluctuating potential. Because the continuous environmental noise is divided into discrete time intervals of length Δt, the time-dependent Hamiltonian, $H(t)$, can be effectively approximated

as being piecewise constant during each of these short intervals. Under this approximation, the Time-Dependent Schrödinger Equation can be solved step-by-step rather than analytically all at once.

For a single time interval from t to $t+\Delta t$, the evolution of the quantum system is governed by a short-time unitary evolution operator. Because the Hamiltonian is held constant over this brief window, this operator is calculated by taking the matrix exponential of the instantaneous Hamiltonian:

$$U(t+\Delta t, t) = \exp\left(-\frac{i}{\hbar}H(t)\Delta t\right) \tag{2.10}$$

By applying this localised operator to the wavefunction at time t, we obtain the state of the system at the next time step, $|\psi(t+\Delta t)\rangle = U(t, t+\Delta t)|\psi(t)\rangle$. To find the state of the exciton at any arbitrary time t_{end} , these short-time propagation steps are chained together. The total time-evolution operator becomes a time-ordered product of the individual step operators, continuously updating the wavefunction as the site energies randomly shift.

Propagating a single wavefunction only tells us what one exciton is doing under a randomly fluctuating environment. In a real molecular aggregate, we measure the macroscopic observables that represent a statistical average over a vast ensemble of molecules, each experiencing its own local bath fluctuations [20]. To mirror this, NISE uses the approach of generating many independent stochastic realisations. Calculating the squared probability amplitudes across all trajectories yields the ensemble-averaged quantum dynamics. This resulting ensemble average gives us the exciton site population data across the sites as the Hamiltonian is propagated in time [8]. The population of the site is effectively calculated as Equation 2.11.

$$P_i(t) = \frac{1}{M}\sum_{k=1}^{M} \langle\psi_k(t)|\,|i\rangle\langle i|\,|\psi_k(t)\rangle = \frac{1}{M}\sum_{k=1}^{M} |c_{i,k}(t)|^2 \tag{2.11}$$

A limitation of NISE is that it assumes the temperature goes to infinity. This is called the infinite temperature limit. In this framework, the bath is modeled as an independent, classical stochastic field that drives the quantum system. However, the bath itself experiences no "back action". In a fully quantum mechanical reality, an exciton transitioning from a high-energy localised state to a lower-energy state must spontaneously emit a quantum of energy (a phonon) into the bath to conserve energy. Conversely, an upward transition requires the system to absorb thermal energy from the bath. Because the classical noise in standard NISE possesses infinite heat capacity and its trajectory is pre-determined (or generated independently of the system's quantum state), it drives upward and downward transitions with exactly equal probability. Without the mechanism of spontaneous emission, the system cannot reach a thermal Boltzmann distribution [21]

$$P_i = \frac{1}{Z} \exp\left(-\frac{E_i}{k_B T}\right) \tag{2.12}$$

where k_B is the Boltzmann constant, T is the temperature, and Z is the partition function defined as the sum over all possible states j :

$$Z = \sum_j \exp\left(-\frac{E_j}{k_B T}\right) \tag{2.13}$$

Instead, over long propagation times, standard NISE inevitably forces the exciton population to distribute uniformly across all sites, completely disregarding their energetic differences. As such, a dimer (2-site system) will always have its site population for either site to converge to 50% regardless of the temperature of the system. The NISE propagation doesn't take temperature in as a parameter. This lack of thermodynamic relaxation is directly responsible for the inability of standard NISE to reproduce the Stokes shift.

2.3 Stokes Shift

As established, the lack of thermodynamic relaxation in standard NISE directly impedes its ability to model the Stokes shift. Phenomenologically, the Stokes shift is the energy difference between the absorption and emission spectra of a system, mathematically approximated as:

$$\Delta E_{\text{Stokes}} \approx 2\lambda \tag{2.14}$$

where λ is the reorganisation energy of the environment. The disparity between the absorption and emission energies is fundamentally governed by the Franck-Condon Principle. This principle dictates that because nuclear masses are significantly larger than those of electrons, an electronic transition occurs so rapidly that the nuclear coordinates (Q) remain essentially stationary during the event. On a potential energy diagram, this results in a "vertical" transition from the ground state to the excited state, as shown in Figure 2.3.

Immediately following this vertical excitation, the system finds itself in a non-equilibrium configuration on the excited state potential energy surface. The nuclei then undergo a process of vibrational relaxation, moving toward a new equilibrium position that minimises the energy of the excited state. This relaxation dissipates this energy, called the reorganisation energy (λ)-into the environment. Subsequent emission then occurs vertically from this new relaxed minimum back to the ground state[22], as shown in Figure 2.3.

It helps to look at the Perrin-Jablonski diagram and the resulting spectra to understand how this vibrational relaxation yields the red-shifted emission.

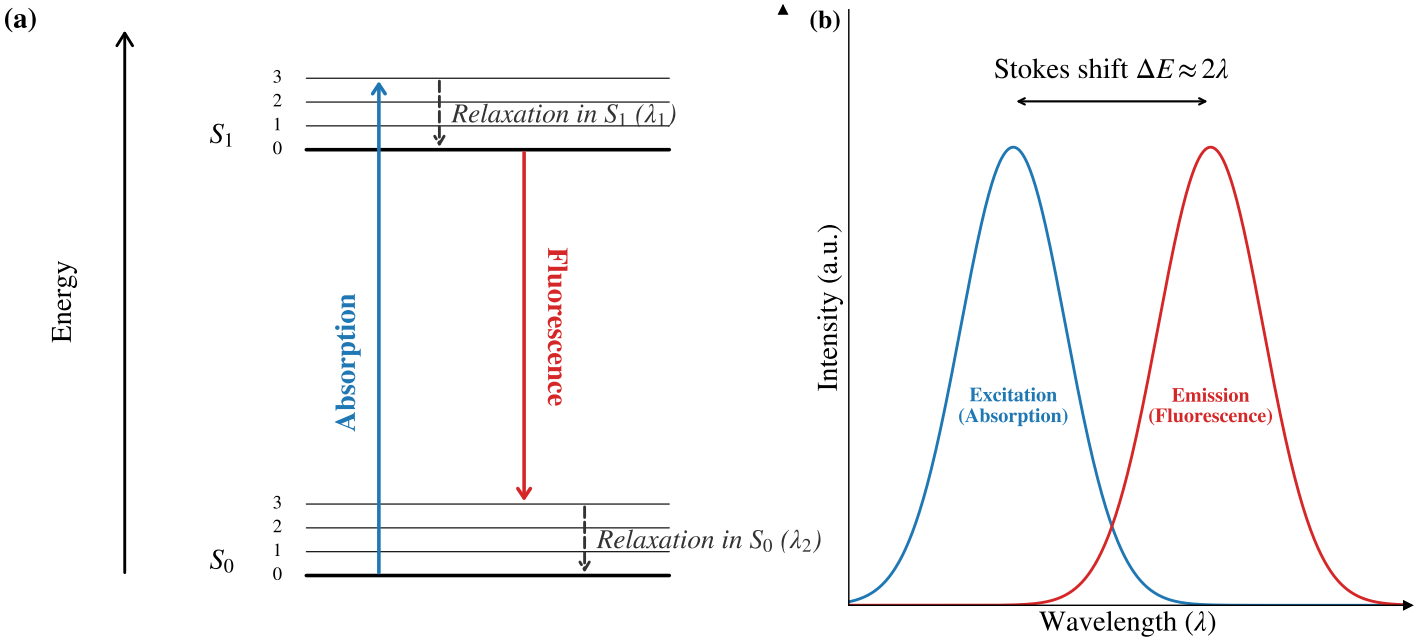

Figure 2.3. Combined Perrin-Jablonski diagram (a) and corresponding normalised excitation and emission spectra (b) illustrating the photophysical origin of the Stokes shift. In (a), absorption of a photon (blue arrow) promotes the molecule to an upper vibrational level of the first excited state (S_1). The total energy loss results from two vibrational relaxation steps (dashed grey arrows): first within the S_1 manifold (λ_1), and subsequently within the S_0 manifold (λ_2) following radiative decay (red arrow). Shown in (b), this cumulative energy loss results in an emission spectrum (red) that is red-shifted toward longer wavelengths relative to the excitation spectrum (blue), with a total Stokes shift $\Delta E \approx 2\lambda$. [23]

In the quantum mechanical picture of the Franck-Condon principle, the difference in mass between the nuclei and the electrons allows the wavefunction of the molecular state to be separated into an electronic component and a nuclear (vibrational) component (Born-Oppenheimer Approximation), as shown in Figure 2.4.

The transition probability between each vibrational level of the S_0 and S_1 states is proportional to the overlap integral between their vibrational wavefunctions. Combining these descriptions, a hybrid classical-quantum picture of the concept can be drawn (see Figure 2.4). The y-axis represents the energy of the molecule, and the x-axis represents the separation between the nuclei. The black quadratic curves illustrate the potential energy wells of the nuclear separation for the S_0 and S_1 states under the simplifying assumption that the molecule is a diatomic harmonic oscillator. The red curves depict the wavefunctions of the quantised vibrational levels (v). The bottom of each well defines the equilibrium separation distance between the nuclei (the bond length), which is characteristically longer in S_1 than S_0 due to its higher energy configuration.

At room temperature, the $v = 0$ level of the S_O state is the most populated due to the Boltzmann distribution, meaning the majority of absorption (indicated by the blue arrow) occurs from this level. This arrow is completely vertical because the nuclei are much heavier than electrons and effectively do not move during the ultrafast timescale of the electronic transition. The strongest absorption transition corresponds to the S_1 vibrational level that shares the greatest wavefunction overlap (the highest wavefunction amplitude at the location of the blue arrow) with the $v = 0$ level of S_0. Because of the larger nuclear separation in the S_1 state, this optimal overlap is never the $v = 0$ level, but rather a higher vibrational level such as $v = 2$, as shown in Figure 2.4. This specific transition defines the wavelength of the absorption spectrum's maximum.

Following absorption, this excess vibrational energy is rapidly lost, and fluorescence will almost exclusively occur from the $v = 0$ level of the S_1 state. Analogous to the absorption case, the strongest fluorescence transition (indicated by the green arrow) will be to the vibrational level of S_0 that has the greatest wavefunction overlap with the $v = 0$ level of S_1 . This downward transition defines the maximum of the fluorescence spectrum. Consequently, the fluorescence maximum will always occur at a longer wavelength than the absorption maximum, which is the fundamental origin of the Stokes shift[24].

To quantify this energy loss and the resulting shift in the spectra, we must transition from a conceptual model to a formal mathematical treatment of the site Hamiltonian.

Using the Born-Oppenheimer approximation, we can look at the site Hamiltonian coupled to a collective intramolecular nuclear coordinate, Q (which can represent bond stretches, ring deformations, or solvent reorganisation). In the classic harmonic approximation, the ground state nuclear potential is modeled as a simple harmonic oscillator:

$$V_{\text{nuc}}(Q) = \frac{1}{2} M\omega^2 Q^2 \tag{2.15}$$

where M is the effective mass of that coordinate and ω is the frequency of the mode.

Upon electronic excitation to state $|n\rangle$, a linear coupling emerges between the electronic state and the nuclear displacement. This alters the site Hamiltonian to include an interaction term:

$$H_{\text{site}}(Q) = \epsilon_n |n\rangle\langle n| + \left(\frac{P^2}{2M} + \frac{1}{2} M\omega^2 Q^2 \right) + |n\rangle\langle n| \kappa Q \tag{2.16}$$

Here, κ is the linear coupling constant between the electronic excitation and the nuclear displacement. The presence of the exciton effectively adds a linear restoring force to the nuclear coordinates, defined as $F_{\text{el}} = -\frac{\partial}{\partial Q}(\kappa Q) = -\kappa$. Consequently, the total potential energy surface for the excited state becomes

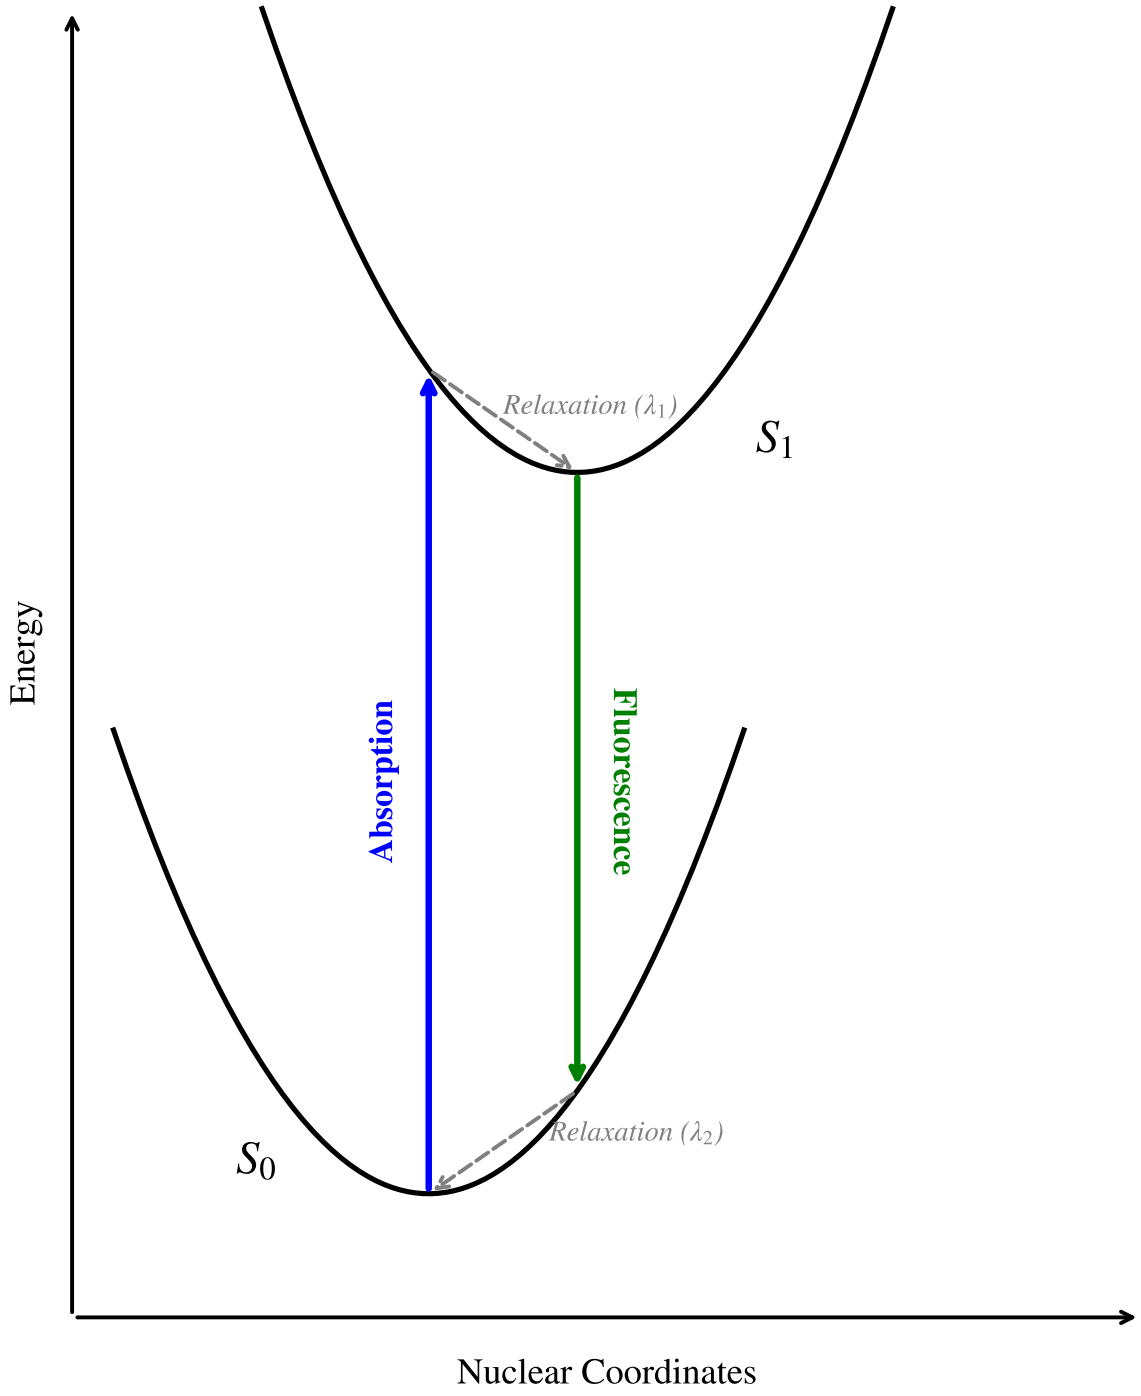

Figure 2.4. Energy coordinate diagram illustrating the Franck-Condon Principle and the physical origin of the Stokes shift using harmonic potential wells for the ground state (S_0) and excited state (S_1). Vertical absorption (blue arrow) from the S_0 equilibrium is followed by vibrational relaxation (λ_1) to the S_1 minimum. Subsequent fluorescence (green arrow) and a final relaxation step in the ground state (λ_2) complete the cycle. The non-radiative energy losses λ_1 and λ_2 result in the characteristic red-shift of emission relative to absorption. [24]

"tilted" relative to the ground state:

$$V_{\text{exc}}(Q) = \frac{1}{2}M\omega^2 Q^2 + \kappa Q + \epsilon_n \tag{2.17}$$

To find the equilibrium position of the nuclei in the excited state, we take the derivative of the excited-state potential $V_{\text{exc}}(Q)$ with respect to the coordinate Q and set it to zero:

$$\frac{dV_{\text{exc}}}{dQ} = M\omega^2 Q + \kappa = 0 \tag{2.18}$$

Solving for Q, we find the new equilibrium coordinate Q^* :

$$Q^* = -\frac{\kappa}{M\omega^2} \tag{2.19}$$

This Q^* represents the physical displacement of the nuclei (or the "bath") in response to the electronic excitation. The reorganisation energy λ is defined as the potential energy difference between the excited state at the ground-state geometry ($Q = 0$) and the excited state at its own relaxed geometry ($Q = Q^*$).

First, we evaluate the potential at the ground-state equilibrium:

$$V_{\text{exc}}(0) = \frac{1}{2}M\omega^2(0)^2 + \kappa(0) = 0 \tag{2.20}$$

Next, we evaluate the potential at the relaxed equilibrium Q^* :

$$V_{\text{exc}}(Q^*) = \frac{1}{2}M\omega^2\left(-\frac{\kappa}{M\omega^2}\right)^2 + \kappa\left(-\frac{\kappa}{M\omega^2}\right) \tag{2.21}$$

Expanding the terms:

$$V_{\text{exc}}(Q^*) = \frac{1}{2}M\omega^2\left(\frac{\kappa^2}{M^2\omega^4}\right) - \frac{\kappa^2}{M\omega^2} \tag{2.22}$$

Simplifying the first term:

$$V_{\text{exc}}(Q^*) = \frac{\kappa^2}{2M\omega^2} - \frac{\kappa^2}{M\omega^2} \tag{2.23}$$

Combining these yields the energy stabilization of the excited state:

$$V(Q^*) = -\frac{\kappa^2}{2M\omega^2} = -\lambda \tag{2.24}$$

Thus, we arrive at the explicit definition of the reorganisation energy:

$$\lambda = \frac{\kappa^2}{2M\omega^2} \tag{2.25}$$

This is the half Stokes shift using the Born-Oppenheimer Potential. The issue when propagating in NISE is that we don't have direct access to these microscopic parameters and instead we use phenomenological parameters (σ, τ, T). To bridge the gap between the microscopic reorganisation energy λ and the macroscopic variance of site energy fluctuations σ^2, we may use the equipartition theorem. This theorem connects the potential energy of a classical degree of freedom (the nuclear coordinate Q) to the thermal energy of the environment $k_B T$. In the NISE framework, we assume the site energy fluctuations $\delta\epsilon(t)$ are proportional to the displacement of the nuclear coordinate Q from its equilibrium. From the Hamiltonian in the previous section, the coupling term is $|n\rangle\langle n|\kappa Q$. This implies that the fluctuation in energy is:

$$\delta\epsilon = \kappa Q \tag{2.26}$$

The variance of the site energy, σ^2, is defined as the ensemble average of the squared fluctuations (assuming the mean $\langle Q\rangle = 0$ in the ground state):

$$\sigma^2 = \langle(\delta\epsilon)^2\rangle = \langle(\kappa Q)^2\rangle = \kappa^2\langle Q^2\rangle \tag{2.27}$$

The nuclei are modeled as harmonic oscillators. According to the equipartition theorem, in the classical limit (high temperature $k_B T \gg \hbar\omega$), every quadratic term in the Hamiltonian contributes $\frac{1}{2}k_B T$ to the average energy.

For the potential energy of the nuclear coordinate $V_{\text{nuc}}(Q) = \frac{1}{2}M\omega^2 Q^2$, we have:

$$\langle V_{\text{nuc}}(Q)\rangle = \frac{1}{2}M\omega^2\langle Q^2\rangle = \frac{1}{2}k_B T \tag{2.28}$$

From the equipartition equality above, we can isolate the term $\langle Q^2\rangle$:

$$\frac{1}{2}M\omega^2\left\langle Q^2\right\rangle = \frac{1}{2}k_B T \tag{2.29}$$

$$\langle Q^2\rangle = \frac{k_B T}{M\omega^2} \tag{2.30}$$

Now, substitute this expression for $\langle Q^2\rangle$ back into our equation for the energy variance σ^2 from

equation 2.27:

$$\sigma^2 = \kappa^2 \left(\frac{k_B T}{M\omega^2} \right) \tag{2.31}$$

Rearranging the terms slightly:

$$\sigma^2 = \left(\frac{\kappa^2}{M\omega^2} \right) k_B T \tag{2.32}$$

Recall the definition of reorganisation energy derived in the previous section:

$$\lambda = \frac{\kappa^2}{2M\omega^2} \quad \Longrightarrow \quad \frac{\kappa^2}{M\omega^2} = 2\lambda \tag{2.33}$$

Now, substitute 2λ into the expression for σ^2 :

$$\sigma^2 = (2\lambda) k_B T \tag{2.34}$$

By simply dividing both sides by $2k_B T$, we reach the final relationship used for the reorganisation energy in NISE simulations[7]:

$$\lambda = \frac{\sigma^2}{2k_B T} \tag{2.35}$$

The derivation above shows the relationship between the Born-Oppenheimer; however, a biological protein like LH2 consists of many continuous vibrational sites. To justify applying this relationship to a thermal bath, we can turn to the statistical mechanics framework formulated by Mukamel [7].

In the theory of nonlinear optical spectroscopy, the influence of the thermal bath is captured by the energy gap correlation function, $C(t)$. This two-time correlation function effectively tracks the "memory" of the environment by measuring how site energy fluctuations evolve and decay over time:

$$C(t) = \langle \epsilon(t)\epsilon(0)\rho_g \rangle \tag{2.36}$$

where $\epsilon(t)$ represents the fluctuation of the electronic energy gap at time t, and the brackets $\left\langle \ldots \rho_g \right\rangle$ denote an ensemble average taken over the ground state density matrix, ρ_g.

This complex function can be separated into a real, symmetric part $C'(t)$ representing the stochastic thermal fluctuations (noise), and an imaginary, anti-symmetric part $C''(t)$ representing the dissipation (the system's "back-action" onto the bath):

$$C(t) = C'(t) + iC''(t) \tag{2.37}$$

We define the spectral density $\tilde{C}(\omega)$. The macroscopic variance of the energy gap, σ^2 is entirely

determined by the fluctuation component at $t = 0$:

$$\sigma^2 = C'(0) = \frac{1}{\pi}\int_0^\infty d\omega \tilde{C}'(\omega) \tag{2.38}$$

The reorganisation energy λ is defined by the imaginary spectral component:

$$\lambda = \frac{1}{\pi\hbar}\int_0^\infty d\omega \frac{\tilde{C}''(\omega)}{\omega} \tag{2.39}$$

In a system that simulates the full quantum mechanical effects, the fluctuations and dissipation are not independent. They are linked by the fluctuation-dissipation theorem [25] (FDT), which states:

$$\tilde{C}'(\omega) = \coth\left(\frac{\hbar\omega}{2k_BT}\right)\tilde{C}''(\omega) \tag{2.40}$$

To link the quantum effects of the full system to the classical system used in NISE, we apply the high-temperature limit ($k_BT \gg \hbar\omega$). In this limit, the argument of the hyperbolic cotangent is small, allowing for the approximation $\coth(x) \approx 1/x$. The FDT simplifies to:

$$\tilde{C}'(\omega) \approx \left(\frac{2k_BT}{\hbar\omega}\right)\tilde{C}''(\omega) \tag{2.41}$$

Substituting this classical approximation back into the integral definition for the variance yields:

$$\sigma^2 \approx \frac{1}{\pi}\int_0^\infty d\omega \left(\frac{2k_BT}{\hbar\omega}\right)\tilde{C}''(\omega) = 2k_BT\underbrace{\left(\frac{1}{\pi\hbar}\int_0^\infty d\omega\frac{\tilde{C}''(\omega)}{\omega}\right)}_{\lambda} \tag{2.42}$$

Recognizing that the term in the parentheses is exactly the definition of the reorganisation energy λ, we arrive at:

$$\sigma^2 = 2\lambda k_BT \tag{2.43}$$

Rearranging the equation to solve for λ, we get:

$$\lambda = \frac{\sigma^2}{2k_BT} \tag{2.44}$$

2.4 HOPS

As established in the previous sections, standard NISE relies on a classical approximation of the thermal bath. While it successfully captures the fluctuating site energies that drive decoherence, it assumes the system operates in the high-temperature limit ($k_B T \gg \hbar\omega$). Because this classical bath lacks a quantum mechanical back-action mechanism, the system cannot undergo spontaneous emission to shed excess energy. This prevents the system from relaxing to a true thermal Boltzmann distribution, fundamentally restricting its ability to capture the Stokes shift.

To overcome this limitation, the environment must be treated as a fully quantum mechanical system. This is achieved using the Hierarchy of Pure States (HOPS) formalism. HOPS is a formally exact, wave-function-based method designed to model exciton transfer dynamics in open quantum systems [4].

In the HOPS framework, the exact quantum correlation function of the bath is mapped onto a complex-valued stochastic noise process. When a quantum system is coupled to this non-Markovian (memory-retaining) heat bath, the time evolution of a single stochastic realisation of the wavefunction, $|\psi(t, z^*)\rangle$, is governed by the Non-Markovian Quantum State Diffusion (NMQSD) equation [4]:

$$\frac{\partial}{\partial t}|\psi(t,z^*)\rangle = \left[-iH_S + \sum_n L_n z_n^*(t) - \sum_n L_n^\dagger \int_0^t ds\, \alpha_n(t-s)\frac{\delta}{\delta z_n^*(s)}\right]|\psi(t,z^*)\rangle \tag{2.45}$$

Here, L_n is the system-bath coupling operator for site n, $z_n^*(t)$ is the complex-valued stochastic noise derived from the environment, and $\alpha_n(t-s)$ is the exact quantum bath correlation function. Because the quantum bath retains memory of its past interactions, the equation requires an integral over the system's entire history. Inside this integral is a functional derivative with respect to the noise, $\frac{\delta}{\delta z_n^*(s)}$. While this term is physically necessary to account for the memory and back-action of the bath, it makes the equation analytically unsolvable and entirely too complex to simulate directly.

To make the problem computationally solvable, HOPS introduces a clever mathematical trick: it absorbs the problematic functional derivative by expanding the system into a hierarchy of auxiliary wavefunctions.

First, we define a tuple of indices $\mathbf{k} = (k_1, k_2, \ldots, k_N)$, where each integer k_n tracks the "depth" of the memory for a specific bath mode, and e_n represents the unit vector for that mode. A hierarchy of pure states is mathematically generated by defining a derivative operator $D_t^{(\mathbf{k})}$:

$$D_t^{(k_1 e_1 + \cdots + k_N e_N)} \equiv \prod_{n=1}^{N}\left(\int_0^t ds,\, \alpha_n(t-s)\frac{\delta}{\delta z_{n,s}^*}\right)^{k_n} \tag{2.46}$$

The actual physical state of the system acts as the root state of this hierarchy, denoted as $\left|\psi^{(\mathbf{0})}(t)\right\rangle$. By applying the functional derivative operator to this root state, we generate the auxiliary states representing deeper levels of memory [26]:

$$|\psi^{(k)}(t)\rangle = D_t^{(k)}|\psi^{(\mathbf{0})}(t)\rangle \tag{2.47}$$

Instead of calculating the massive integral from the NMQSD equation directly, the bath correlation function is approximated as a sum of exponential decays. By taking the time derivative of these auxiliary states, the complex integral is completely eliminated. This effectively "unrolls" the memory backward in time into a set of coupled ordinary differential equations (ODEs), leading to the Linear HOPS equation:

$$\begin{aligned}\frac{\partial}{\partial t}|\psi^{(k)}(t)\rangle = &\left(-iH_S - \sum_n k_n w_n + \sum_n L_n z_n^*(t)\right)|\psi^{(k)}(t)\rangle \\ &+ \sum_n k_n g_n L_n^\dagger |\psi^{(k-e_n)}(t)\rangle - \sum_n L_n |\psi^{(k+e_n)}(t)\rangle\end{aligned} \tag{2.48}$$

In this framework, each auxiliary state accounts for a portion of the system-bath entanglement. A given state at level k only communicates with the states one step higher ($\mathbf{k} + \mathbf{e}_n$) and one step lower ($\mathbf{k} - \mathbf{e}_n$) in the hierarchy. This nearest-neighbour coupling allows the system's quantum memory to cascade down through the hierarchy without requiring the computer to continually look back in time.

While the linear HOPS equation is mathematically exact, it suffers from a fatal computational flaw when tracking individual trajectories. Because the time evolution operator is non-Hermitian, individual wavefunctions are not naturally normalised. During numerical integration, the amplitude of most trajectories decays toward zero, while a few explode toward infinity, resulting in severe numerical instability.

To resolve this during simulation, a continuous normalisation shift is applied, transforming the system into the Non-linear HOPS equation:

$$\frac{\partial}{\partial t}|\psi^{(k)}(t)\rangle = (\mathcal{L}_{\text{linear}} - \Gamma_t)\,|\psi^{(k)}(t)\rangle \tag{2.49}$$

Here, $\mathcal{L}_{\text{linear}}$ represents the right-hand side of the linear HOPS equation, and Γ_t is a real-valued scalar that continuously normalises the root state ($\left(\left\langle \psi_t^{(\mathbf{0})} \mid \psi_t^{(\mathbf{0})} \right\rangle = 1\right)$) during the step-by-step time propagation. In the numerical code explicitly utilised for our simulations, this shift is calculated at each time step as [27]:

$$\begin{aligned}\Gamma_t = &\sum_n \mathrm{Re}\left[\langle L_n\rangle_t \left(z^*_{n,t} + \xi_{n,t}\right)\right] - \sum_n \left(\frac{g_n}{w_n}\right) \mathrm{Re}\left[\left(\psi_t^{(\mathbf{0})}\right)^\dagger L_n^\dagger \psi_t^{(\mathbf{e}_n)}\right] \\ &+ \sum_n \left(\frac{g_n}{w_n}\right) \mathrm{Re}\left[\left\langle L_n^\dagger \right\rangle_t \left(\psi_t^{(\mathbf{0})}\right)^\dagger \psi_t^{(\mathbf{e}_n)}\right]\end{aligned} \tag{2.50}$$

This equation structurally links the normalization back to the physics of the environment. It establishes a feedback loop between the system's expected coupling $\langle L_n\rangle_t$, the thermal noise $\xi_{n,t}$, and the first tier of auxiliary states, scaled by the bath correlation parameters g_n (the coupling parameter) and w_n (the angular frequency parameter). By continuously subtracting Γ_t, any unphysical growth or decay in the wavefunction is perfectly counterbalanced without altering the underlying physical dynamics.

Because an infinite hierarchy cannot be simulated on a computer, the auxiliary states are truncated at a maximum depth (K). The coupled ODEs are advanced through time using standard numerical solvers, such as the fourth-order Runge-Kutta method. Because a single propagation yields only one possible quantum trajectory, the simulation must be repeated across thousands of independent stochastic noise realisations. The true physical reduced density matrix, $\rho(t)$, is finally recovered by taking the ensemble average over all normalised trajectories:

$$\rho(t) = \mathbb{E}\left[|\psi(t;z)\rangle\langle\psi(t;z)|\right] \tag{2.51}$$

By utilizing formally exact quantum noise and capturing the environment's back-action via this explicit hierarchy of states, HOPS successfully bridges the gap left by strictly classical models. It properly accounts for energy dissipation into the bath, allowing the system to naturally relax to the correct thermal equilibrium and model the Stokes shift. Consequently, HOPS serves as a formally exact benchmark to validate whether the implementation of the Stokes shift within the NISE framework is physically accurate.

Since HOPS is propagated via a wavefunction, this means we can track the effective energy of the system as it propagates and track the contribution of the Stokes shift (the imaginary part) compared to the bath thermal noise (the real part). This gives us another benchmark against NISE to see if our adjustments to the NISE model by implementing the Stokes shift is making physical sense. For

instance if we take a look again at the normalised non linear HOPS equation:

$$\begin{aligned}\hbar\partial_t \left|\tilde{\psi}_t^{(k)}\right\rangle = &\left(-i\tilde{H}_{\mathrm{S}} - \vec{k}\cdot\vec{\gamma} - \Gamma_t\right)\left|\tilde{\psi}_t^{(k)}\right\rangle + \sum_n \tilde{L}_n\left(z_{n,t}^* + \xi_{n,t}\right)\left|\tilde{\psi}_t^{(k)}\right\rangle \\ &+ \sum_n^N \sum_{j_n}^{N_{\max}} k_{n,j_n}\gamma_{n,j_n}\tilde{L}_n\left|\tilde{\psi}_t^{\left(k-e_{n,j_n}\right)}\right\rangle \\ &- \sum_n^N \sum_{j_n}^{N_{\max}} \left(\frac{g_{n,j_n}}{\gamma_{n,j_n}}\right)\left(\tilde{L}_n^\dagger - \left\langle\tilde{L}_n^\dagger\right\rangle_t\right)\left|\tilde{\psi}_t^{\left(k+e_{n,j_n}\right)}\right\rangle,\end{aligned} \tag{2.52}$$

The drift term in the noise term is given by:

$$\xi_{t,n} = \frac{1}{h}\int_0^t ds\alpha_n^*(t-s)\left\langle\tilde{L}_n^\dagger\right\rangle_s, \tag{2.53}$$

With $\alpha_n(t) = (2\lambda/\beta - i\lambda\gamma)e^{-\gamma t/\hbar}$. This allows us to track the imaginary and real parts of the HOPS propagation. To explicitly track these distinct physical processes during propagation, we evaluate the timedependent, complex-valued effective site energy, $E_i(t)$, for each site i. This effective energy combines the total system Hamiltonian with the total complex bath influence (the sum of the stochastic noise and the memory drift) to give:

$$E_i(t) = H_{ii} + \sum_n \left(z_{n,t}^* + \xi_{n,t}\right)(L_n)_{ii} \tag{2.54}$$

By decomposing this effective energy into its constituent parts, we can physically and mathematically separate the environmental effects. The real component tracks the stochastic noise (pure dephasing), which represents the structural fluctuations of the energy landscape driven largely by the real, temperature-dependent part of the correlation function $(2\lambda/\beta)$:

$$\mathrm{Re}\left[E_i(t)\right] = H_{ii} + \sum_n \left(\mathrm{Re}\left[z_{n,t}^*\right] + \mathrm{Re}\left[\xi_{n,t}\right]\right)(L_n)_{ii} \tag{2.55}$$

Conversely, the imaginary component tracks the dissipative relaxation into the environment. It acts mathematically as an imaginary energetic term that directly drives the exponential decay of the state vector's probability amplitude. Governed by the imaginary part of the correlation function $(\lambda\gamma)$, the relaxation is calculated as:

$$\mathrm{Im}\left[E_i(t)\right] = \sum_n \left(\mathrm{Im}\left[z_{n,t}^*\right] + \mathrm{Im}\left[\xi_{n,t}\right]\right)(L_n)_{ii} \tag{2.56}$$

Through these separated equations, the HOPS framework allows us to independently quantify

the real-valued classical-like thermal fluctuations of the site energies versus the imaginaryvalued quantum dissipation of the system's population at any given time t.

All HOPS benchmarks in this work were conducted using a modified version of the C++ program developed by Wobben [27]. By employing this optimised implementation of the HOPS framework as a computational baseline, we generated formally exact quantum dynamics to serve as a rigorous reference. These results allowed us to validate the physical accuracy of our modified NISE model and its phenomenological treatment of the Stokes shift.

2.5 HEOM

The Hierarchical Equations of Motion (HEOM) formalism, developed by Tanimura, provides a numerically exact, non-perturbative method to describe the reduced dynamics of an open quantum system [5]. Unlike the stochastic nature of NISE or HOPS, HEOM directly propagates the reduced density matrix $\rho(t)$ by accounting for the influence of a bosonic bath through a set of coupled differential equations.

The framework begins with the standard system-bath Hamiltonian, where the environmental influence is captured by the bath correlation function $C(t)$. In the HEOM approach, this function is decomposed into a sum of exponential terms:

$$C(t) = \sum_{k=0}^{K} c_k e^{-\nu_k |t|} \tag{2.57}$$

where ν_k represents the relaxation rates and Matsubara frequencies that account for quantum thermal fluctuations [5].

To handle the non-Markovian memory of the bath, HEOM introduces a hierarchy of auxiliary density operators (ADOs), $\hat{\rho}_{\mathbf{n}}(t)$. The physical density matrix $\hat{\rho}_{\mathbf{0}}(t)$ (where $\mathbf{n} = 0$) is coupled to higher-tier ADOs that represent deeper levels of system-bath entanglement. The time evolution of the hierarchy is governed by:

$$\frac{\partial}{\partial t}\hat{\rho}_{\mathbf{n}}(t) = -\left(i\mathcal{L}_S + n\gamma + \sum_k n_k \nu_k + \hat{\Xi} \right) \hat{\rho}_{\mathbf{n}}(t) + \Phi\hat{\rho}_{\mathbf{n}+1}(t) + \Psi\hat{\rho}_{\mathbf{n}-1}(t) \tag{2.58}$$

where $\mathcal{L}_S$ is the system Liouvillian and Φ, Ψ are superoperators describing the transitions between hierarchy tiers and $\hat{\Xi}$ is equal to [5]:

$$\hat{\Xi} = -\left(\sum_{k=K+1}^{\infty} \frac{c_k}{\nu_k} \right) \frac{\hat{V}^{\times}\hat{V}^{\times}}{h^2} \tag{2.59}$$

For the molecular aggregates studied in this thesis the number of required ADOs scales factorially with the system size and hierarchy depth. To generate the exact benchmark data presented in this thesis, GPU-HEOM was used through NanoHub [28]. However, HOPS and HEOM should theoretically produce the same result. This is highlighted in Figure 2.5 where both models were simulated using the same parameters for a simplified LH2 as a dimer. The parameters for the LH2 dimer model are taken from Jansen [8].

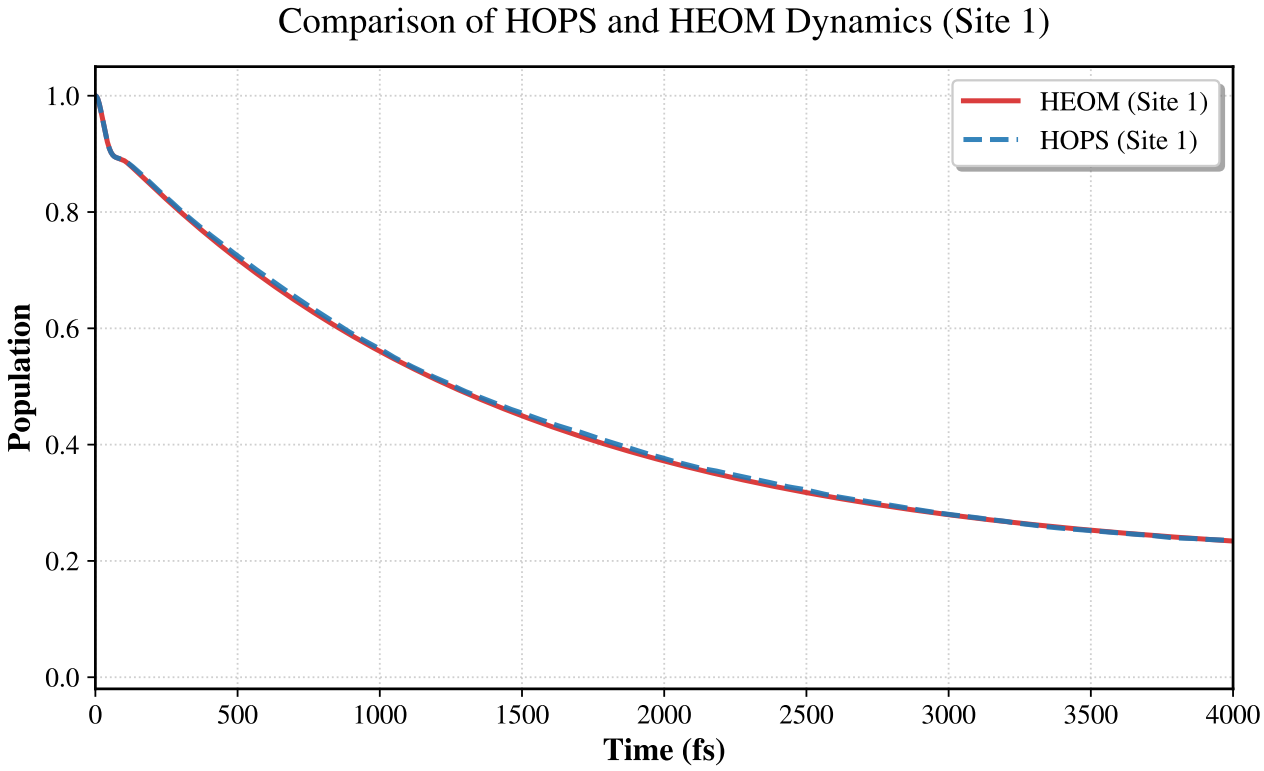

Figure 2.5. Comparison of Site 1 population dynamics obtained from the stochastic HOPS (blue solid line) and deterministic HEOM (red dashed line) methods. Simulations were performed using identical system parameters ($\Delta E = 300\,\mathrm{cm}^{-1}$, $\tau = 100\,\mathrm{fs}$, $\sigma_i = 100\,\mathrm{cm}^{-1}$, $J = 47\,\mathrm{cm}^{-1}$) for a dimer model.

One advantage with HOPS is that the parameters σ and γ can be defined independently for each site. This allows us to have a look at the dynamics of the site aggregate changes when each site has a different σ value. Furthermore, HOPS is a wavefunction based method as NISE, making the comparison easier.

Chapter 3

Methods

3.1 NISE

A NISE program was created in Python as a base to compare to other NISE models [29]. From this Python base, the Linear Reorganisation Shift and memory kernel were added. The parameters that can be adjusted in the Python NISE program are the total sites N, the total max length of the trajectory, timestep dt, number of realisations (independent stochastic trajectories) M, the coupling J cm^{-1}, the energy difference between the two sites ΔE, the σ, which represents the standard deviation of the site energy fluctuations, the Ornstein-Uhlenbeck (OU) correlation time (τ), which determines the "memory" of the environment by controlling how quickly the stochastic energy fluctuations on each site decay and lose their previous state.

In this specific implementation, the Ornstein-Uhlenbeck process is realised as a discretetime Brownian motion that incorporates a restorative memory term to simulate a bath. The code executes this by calculating a dimensionless trajectory where the state at t is a weighted sum of the previous state and a stochastic "kick," controlled by the decay factor $\alpha = e^{-\Delta t/\tau}$ and the fluctuation scaling $\beta = \sqrt{1-\alpha^2}$. This ensures that the process remains stationary and retains the desired correlation time τ across the simulation. The resulting dimensionless value $x_i(t)$ is then mapped to physical energy units and added to the site energy ϵ_i to define the fluctuating diagonal of the Hamiltonian [30]:

$$\begin{aligned} x_i(t) &= \alpha \cdot x_i(t-\Delta t) + \beta \cdot N(0,1) \\ E_i(t) &= \epsilon_i + \sigma_i \cdot x_i(t) \end{aligned} \tag{3.1}$$

3.2 Linear Reorganisation Shift

In NISE, the environment is treated as an external force that drives fluctuations in the system without experiencing any back-action from the quantum state. In this "infinite temperature" limit, the diagonal elements of the Hamiltonian are simply:

$$H_{nn}(t) = \epsilon_n + \delta E_n(t) \tag{3.2}$$

Because the bath is independent of the exciton's position, the system cannot spontaneously dissipate energy into the environment. To incorporate the Stokes shift, the model must account for the vibrational relaxation of the nuclei in response to excitation. Physically, as an exciton occupies a site, the surrounding lattice or solvent reorganises to minimise the local free energy, lowering the energy of that specific state by the reorganisation energy λ.

To simulate this effect, a mean-field correction is applied where the site energy is dynamically coupled to its own transient quantum population, $\rho_{nn}(t) = |a_n(t)|^2$. This results in the following transformation of the diagonal Hamiltonian elements[31]:

$$\epsilon_n \rightarrow \epsilon_n - \lambda \, |a_n(t)|^2 \tag{3.3}$$

The resulting instantaneous Hamiltonian becomes non-linear, as the potential landscape now depends on the probability amplitude of the wavefunction:

$$H_{nn}(t) = \epsilon_n + \delta E_n(t) - \lambda \, |a_n(t)|^2 \tag{3.4}$$

By subtracting the term $\lambda \, |a_n(t)|^2$, the exciton effectively "digs" its own potential well. As the probability of finding the exciton on site n increases, that site's energy is lowered, providing a physical mechanism for energy relaxation.

To help visualise how this relaxation unfolds over the computational trajectory, it is important to clearly define the starting conditions of the model. As illustrated in Figure 3.1, the simulation is initialised precisely at $t = 0$ immediately following vertical excitation. At this exact starting point, the exciton is already localised on the S_1 excited state, but the environment has not yet begun to reorganise; therefore, the downward shift of the site energy has just begun.

To ensure this dissipative term is consistent with the stochastic fluctuations $\delta E_n(t)$, the reorganisation energy λ is defined via the classical fluctuation-dissipation theorem:

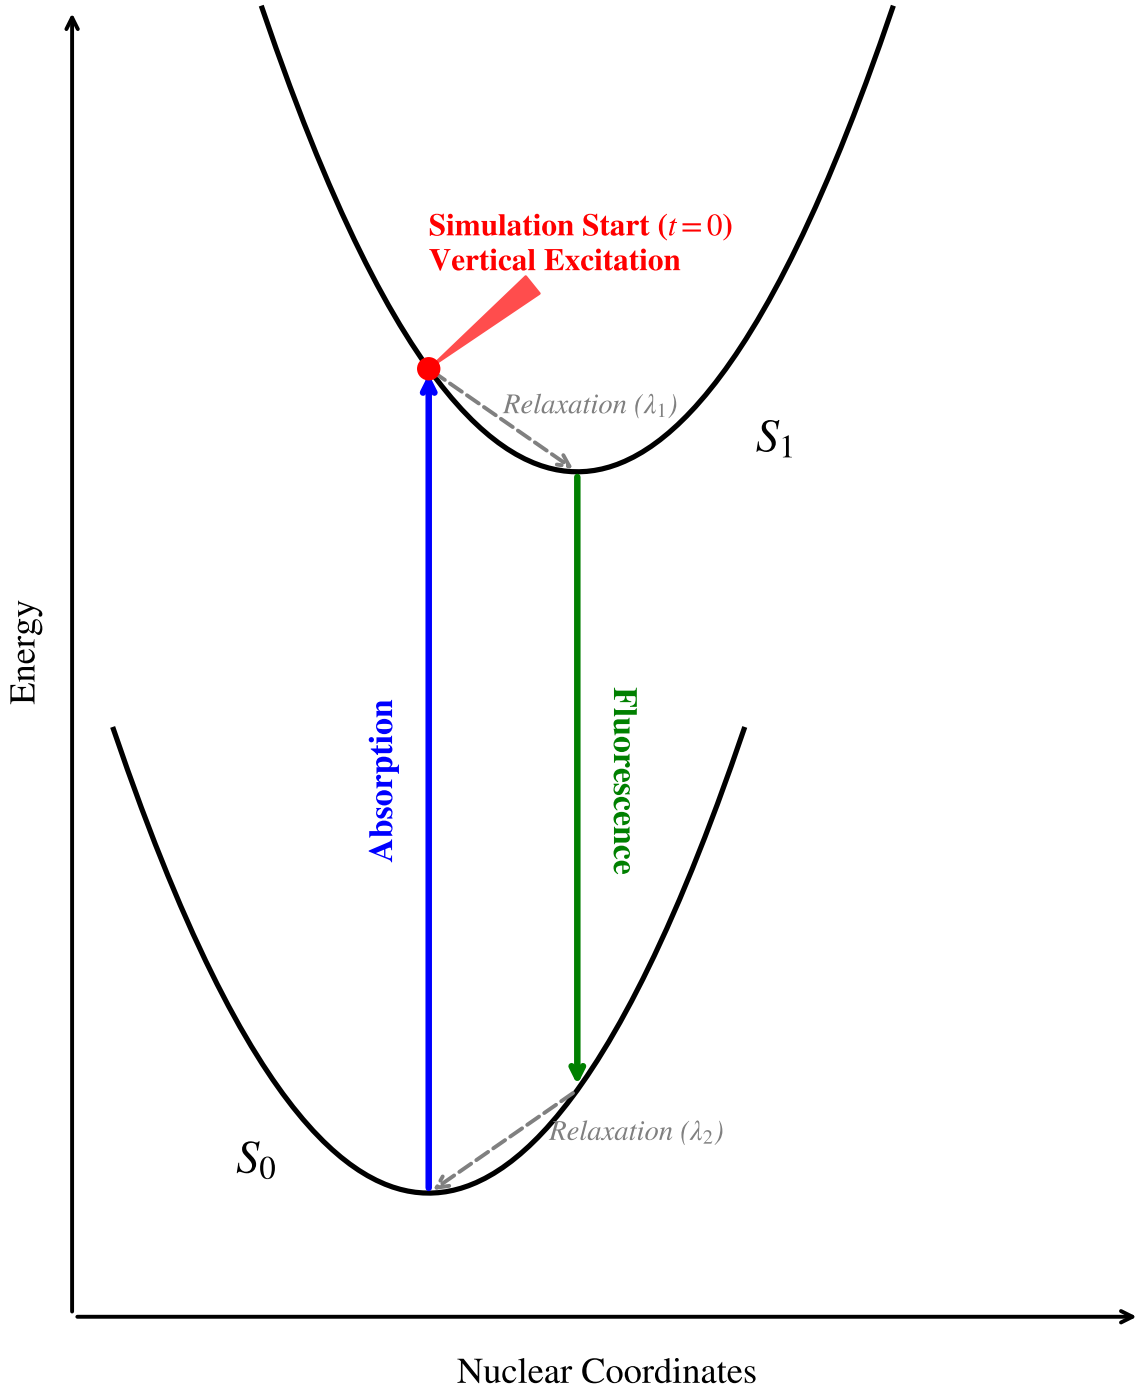

Figure 3.1. A modified diagram of Figure 2.4 highlighting the initialization point of the computational simulations. The trajectories begin exactly at $t = 0$ (red marker), representing the state where the exciton has been vertically excited to the S_1 energy surface, but prior to any physical relaxation or energetic Stokes shift.

$$\lambda = \frac{\sigma^2}{2k_B T} \tag{3.5}$$

Unlike the base NISE simulation, temperature (T) is introduced to calculate the reorganisation energy via $\lambda = \sigma^2/(2k_BT)$. This value remains constant during propagation, justified by the Franck-Condon approximation: electronic transitions occur so rapidly that nuclear coordinates remain stationary during the event. This implies the potential surface curvature and coupling strength are constant material properties. Consequently, the site energy is stabilised through the transformation $\epsilon_n \rightarrow \epsilon_n - \lambda |a_n(t)|^2$, allowing the exciton to dynamically lower the energy of its occupied state without altering the underlying environmental stiffness.

3.3 Memory Kernel Shift

While the Linear Reorganisation Shift introduces a mechanism for energy dissipation, it fundamentally assumes that the environment responds *instantaneously* to the presence of the exciton ($\epsilon_n \rightarrow \epsilon_n - \lambda |a_n(t)|^2$). Physically, this implies that the solvent molecules and nuclear coordinates possess zero inertia. In reality, the thermal bath takes time to physically reorganise around the new charge distribution, governed by the bath correlation time, τ.

To explicitly derive the time-dependent response of the bath, we must look to the exact semiclassical lineshape function, $g(t)$. Formulated by Mukamel for a system coupled to a harmonic bath[7], this function dictates the exact time-evolution of the optical energy gap in the high-temperature limit ($\beta\hbar\omega \ll 1$) [7]:

$$g(t) = \frac{2\lambda k_{\mathrm{B}} T}{\hbar} \int_0^t d\tau_2 \int_0^{\tau_2} d\tau_1 M(\tau_1) - i\lambda \int_0^t d\tau [1 - M(\tau)] \tag{3.6}$$

Here, $M(t)$ represents the memory function of the bath. For an overdamped Brownian oscillator (a Debye bath), this memory function decays exponentially over the bath correlation time τ:

$$M(t) = e^{-t/\tau} \tag{3.7}$$

The lineshape function $g(t)$ mathematically embodies the fluctuation-dissipation theorem by splitting the environmental influence into two distinct parts:

1. **The Real Part ($\propto \lambda k_BT$):** This governs the random, symmetric energetic line broadening (fluctuations). NISE models this using the real-valued Ornstein-Uhlenbeck noise, $\delta E_n(t)$.
2. **The Imaginary Part ($\propto \lambda$):** This governs the phase shift that causes the transition energy to relax (dissipation). NISE ignores this term.

If an exciton suddenly occupies a site from the ground state at $t = 0$ (a static step-function in

population), the exact energetic stabilization induced by the bath is:

$$\Delta E_{\text{step}}(t) = \frac{d}{dt}\left(-\lambda \int_0^t d\tau [1 - M(\tau)]\right) \tag{3.8}$$

By applying the fundamental theorem of calculus, the derivative cancels the integral:

$$\Delta E_{\text{step}}(t) = -\lambda[1 - M(t)] \tag{3.9}$$

Substituting the exponential memory function $M(t) = e^{-t/\tau}$ yields the step-response of the Stokes shift:

$$\Delta E_{\text{step}}(t) = -\lambda\left(1 - e^{-t/\tau}\right) \tag{3.10}$$

Physically, this is saying that upon sudden excitation, the site energy does not drop by λ instantly. Instead, it smoothly relaxes by λ over the timescale τ as the surrounding solvent progressively reorganises.

In an aggregate, the exciton moves dynamically between sites, meaning the population $P_n(t)$ is never static. Because vibrational relaxation takes time, the surrounding environment cannot respond to these shifts instantaneously. Instead, the bath acts as a linear responder that "remembers" the history of the exciton's presence. To find the current site energy, we must calculate the cumulative, delayed response of the environment to the entire history of the site's population through a convolution integral.

The impulse response function, $h(t)$, is the time derivative of the step-response:

$$h(t) = \frac{d}{dt}\left[\Delta E_{\text{step}}(t)\right] = \frac{d}{dt}\left[-\lambda\left(1 - e^{-t/\tau}\right)\right] \tag{3.11}$$

$$h(t) = -\frac{\lambda}{\tau} e^{-t/\tau} \tag{3.12}$$

Therefore, at any arbitrary time t, the dynamic Stokes shift applied to site n is the convolution of this impulse response with the continuous population history $P_n(s)$:

$$\Delta\epsilon_n(t) = \int_0^t h(t-s) P_n(s) ds \tag{3.13}$$

Substituting the impulse response yields the memory kernel:

$$\Delta\epsilon_n(t) = -\int_0^t \frac{\lambda}{\tau} e^{-(t-s)/\tau} |a_n(s)|^2 ds \tag{3.14}$$

This integral calculates the exact instantaneous reorganisation penalty based on how heavily the site has been populated in the past, exponentially weighted by how much the bath remembers that past.

Integrating this continuous memory kernel into the NISE framework transforms the system into a non-Markovian classical approximation. The modified diagonal elements of the system Hamiltonian become:

$$H_{nn}(t) = \epsilon_n + \delta E_n(t) - \int_0^t \frac{\lambda}{\tau} e^{-(t-s)/\tau} |a_n(s)|^2 ds \tag{3.15}$$

By replacing the instantaneous mean-field shift with this historical convolution, the simulation explicitly reunites the real and imaginary components of Mukamel's lineshape function within a semiclassical framework[7]. The stochastic noise supplies the temperature-dependent fluctuations (satisfying the real part), while the memory kernel provides the time-dependent dissipation (satisfying the imaginary part). Together, they satisfy the fluctuation-dissipation theorem, mathematically guaranteeing that the Hamiltonian enforces detailed balance and allowing the excitonic system to attempt to relax toward a finite-temperature Boltzmann distribution.

Chapter 4

Results

The accuracy of the NISE simulation does depend on the number of realisations calculated, as shown in Figure 4.1. All simulations in the results section utilized $M = 200,000$ realisations for both the base NISE model and NISE with the Linear Reorganisation Shift, and $M = 50,000$ realisations for the Memory Kernel Shift and HOPS benchmarks. The high number of realisations were needed as the difference between the models could be fine and the accuracy of the realisations were needed in order to determine the results. NISE and the NISE with Linear Reorganisation Shift were able to run at higher realisations due to their code being more optimised than the memory kernel method.

4.1 Average Site Energy

In this results section, we will only be testing on dimers. A dimer is two monomers linked together that we treat as a 2-site system ($N = 2$). This decision reduces the computational cost that comes with testing for large N site systems. In theory these methods can be scaled to any amount of sites, but the purpose of this paper is to test the implementation of the Stokes shift. Testing on dimers allows us to more easily see the results, and analyze the data.

When running a NISE simulation, you can collect the average site energy for each realisation. Collecting all the data and averaging each realisation by the total realisations, you can produce an average site energy over time, as shown in Figure 4.2.

The diagonal elements of the Hamiltonian are the site energies plus the Ornstein–Uhlenbeck process $H_{nn}(t) = \epsilon_n + \delta\epsilon_n(t)$, so there is no force on the Hamiltonian pushing the site energies over time as $\delta\epsilon_n(t)$ averages to 0 over time with enough realisations. When implementing the Stokes shift in NISE, the average site energies changes as the Linear Reorganisation Shift changes the site energies as it fluctuates over time. This is shown in Figure 4.3.

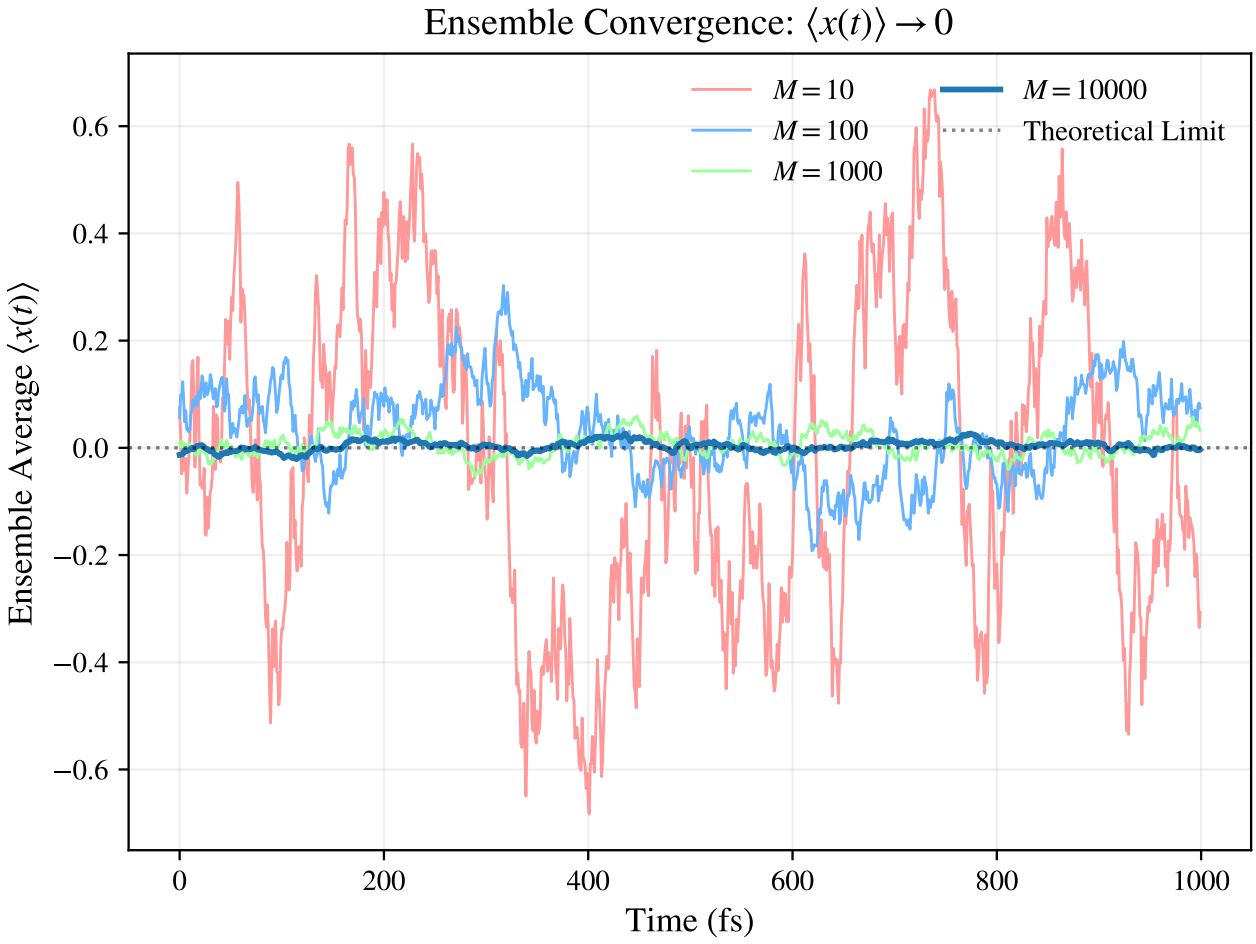

Figure 4.1. Statistical convergence of the ensemble-averaged Ornstein-Uhlenbeck (OU) process. The plot shows the average fluctuation $\langle x(t) \rangle$ for increasing numbers of realisations M, utilizing a correlation time $\tau = 50$ fs and a timestep $\Delta = 1$ fs. The larger the number of realisations M the greater the suppression of stochastic noise as the process converges toward the theoretical mean (dashed red line), a necessary condition for obtaining stable and accurate population dynamics in NISE simulations.

The simulation is initialised at $t = 0$ with the excited state localised on site 1 causing an instantaneous relaxing of energy to roughly the reorganisation energy. After that initial relaxation, both sites move towards a quasi equilibrium of approximately half the lambda reorganisation for each site. Both sites settle at half their reorganisation energy because the average population of both sites equalises at around 0.5.

We can observe an example of where site energies will cross over each other. Using the Memory Kernel Shift average energy sites, as shown in Figure 4.4. The reorganisation energy of site 2 (198.9 cm^{-1}) is greater than site 1 (100.8 cm^{-1}). Here the site energies cross over and site 2 ends up on a lower site energy than site 1 despite site 1 losing more energy initially due to the Stokes shift initial relaxation.

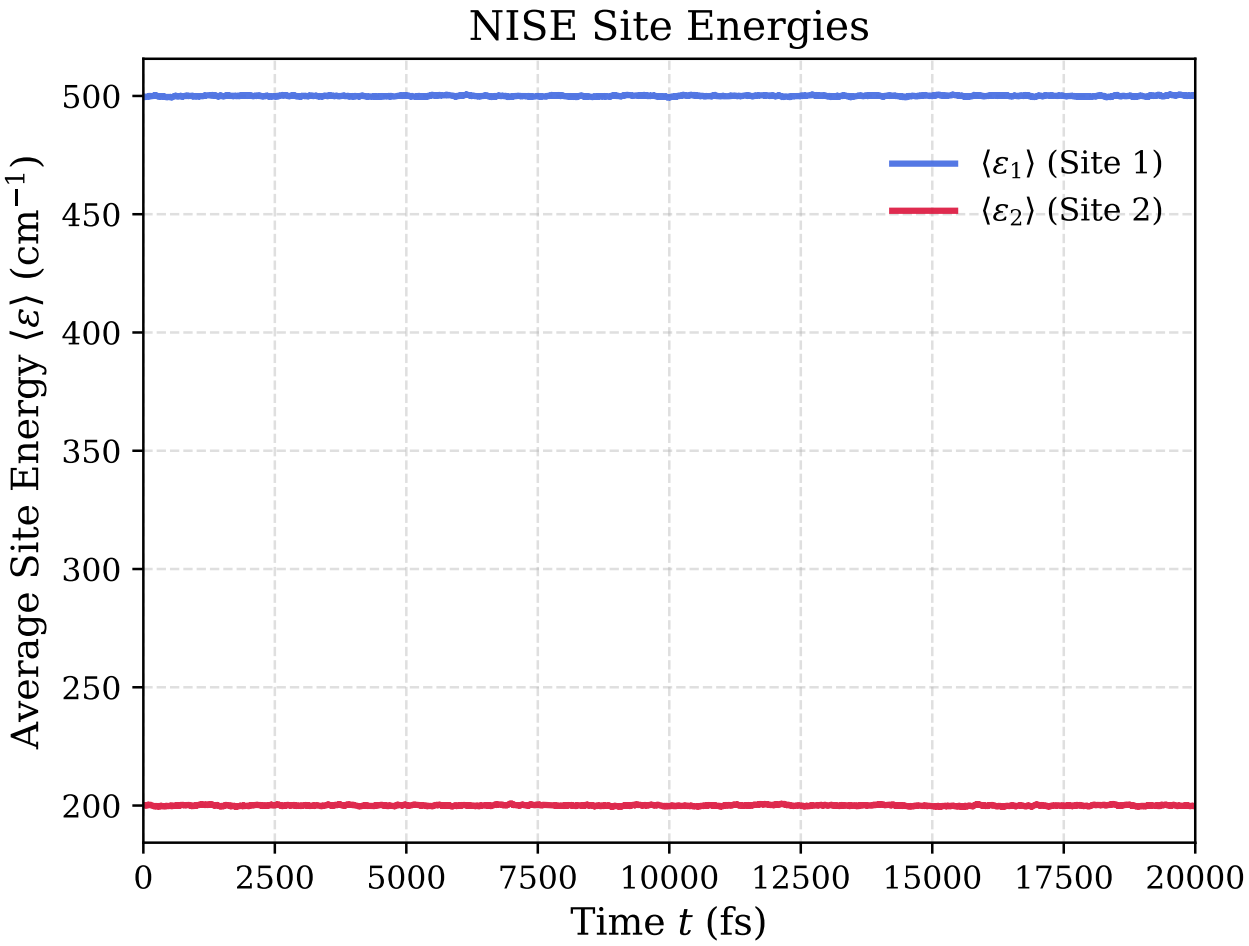

Figure 4.2. Average site energies of LH2 Dimer of a NISE simulation without the Stokes shift. $\sigma_1 = 100$ cm^{-1}, $\sigma_2 = 100$ cm^{-1}, $\tau = 100$ fs, $J = 47$ cm^{-1}, Energy Difference = $\Delta E = 300$ cm^{-1}, $T = 300$ K, $\lambda = 23.98$ cm^{-1}. Both site energies fluctuates around the initial site energy.

Although HOPS is also propagated via a wavefunction, comparing data for average site energies is not as straightforward as it is in simple NISE. In HOPS there is a real and imaginary element of the propagation. The real part handles the thermal noise fluctuations of the system while the imaginary part handles the physical relaxation of the system. To show the full effect of both of these effects it is easier to graph them separately as it is easier to see how each part impacts the site energies.

When comparing the effective energies of the HOPS and the NISE (with Stokes shift) methods there is large agreement in both models, as shown in Figure 4.6. The memory Kernel shows slightly better agreement than the linear shift in the over the longer term. All 3 models decrease the energy of site 1 by roughly the reorganisation energy (215.8cm^{-1}), which is proportional to the Stokes shift. Site 1 then stabilises higher than both NISE models with the memory kernel stabilizing higher than the linear shift. For Site 2 they stabilise in the opposite order with Linear stabilising the highest and HOPS the lowest. The explanation for why the three models diverge is due to the difference in site

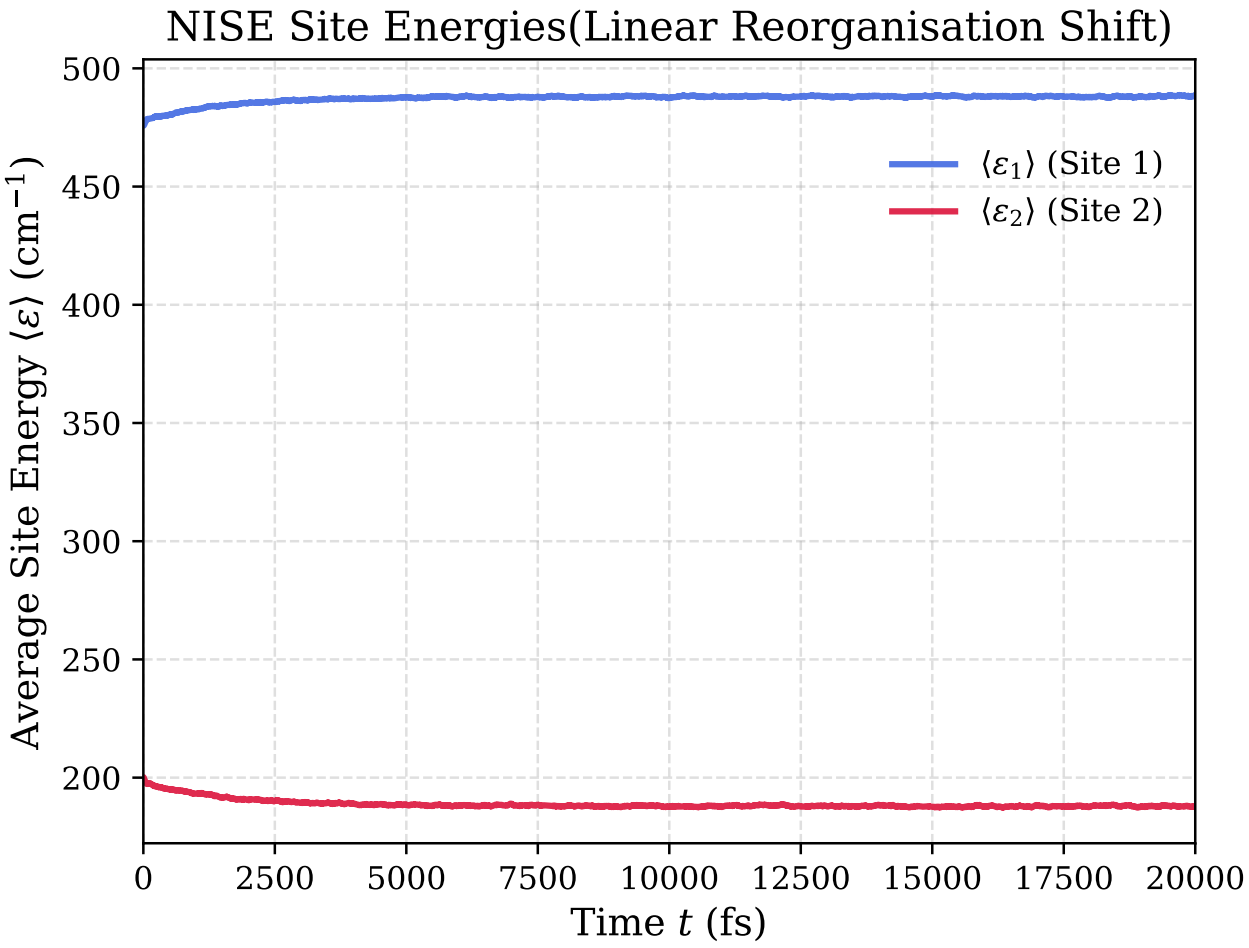

Figure 4.3. Average site energies of LH2 Dimer of a NISE simulation with the Stokes shift. $\sigma_1 = 100$ cm^{-1}, $\sigma_2 = 100$ cm^{-1}, $\tau = 100$ fs, $J = 47$ cm^{-1}, Energy Difference = $\Delta E = 300$ cm^{-1}, $T = 300$ K, $\lambda = 23.98$ cm^{-1}. Over time Site 1 shifts down by approximately 1 λ reorganisation shift (23.9 cm^{-1}) before stabilizing at around half the reorganisation shift. Site 2 evolves slowly to half the reorganisation shift below the initial starting site energy.

energy population.

At the start, between 0-200fs there is a difference between the memory kernel/HOPS and the Linear reorganisation as the Linear reorganisation starts from directly the vertical transition(500 cm^{-1} − 215.8 cm^{-1} = 284.2 cm^{-1}) and decay smoothly as it assumes the trajectory is a perfect exponential curve $E(t) = E_{eq} + \lambda e^{-t/\tau}$ so at $t = 0$ it is mathematically locked to the starting value. While HOPS and memory kernel simulates the bath at $t = 0$ they initialise the bath equilibrium with the ground state. When you excite from $t = 0$, the bath noise is sampled from a random distribution. While the average of these samples should technically match the Vertical Energy, finite sampling (running 5,000 trajectories instead of infinity) means the average at $t = 0$ will have a small statistical error (fluctuation). Additionally, memory kernels capture "Inertial" effects (the bath takes a femtosecond to start moving), whereas Linear model assumes instantaneous velocity. Zooming into the data plotted

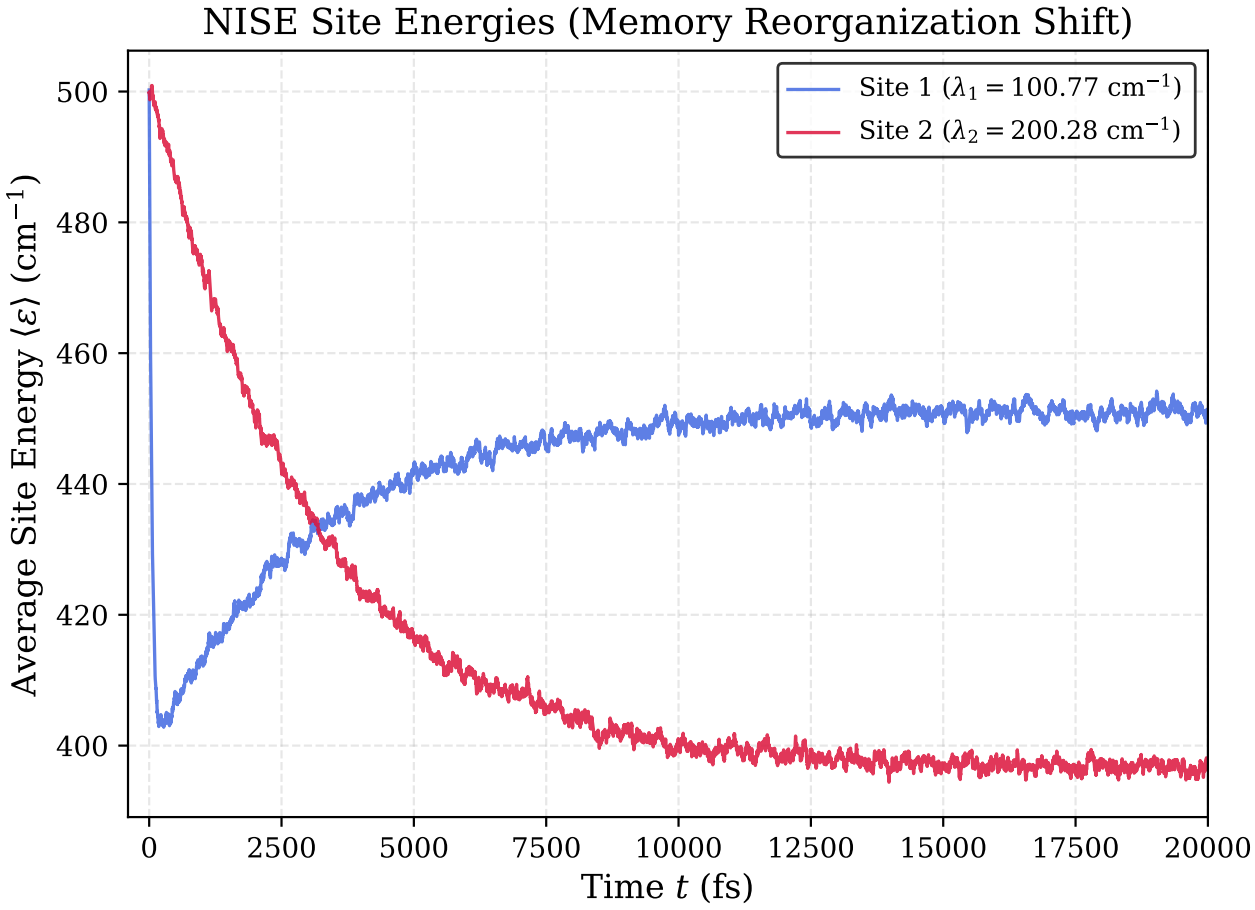

Figure 4.4. Average energy of the site in cm^{-1} over time (fs).T = 300K, τ = 50 fs with λ_1 = 100.8 cm^{-1} for Site 1 and λ_2 = 200.28 cm^{-1} for site 2. and an energy difference of 0cm^{-1}. Site 1 drops roughly the Stokes shift reorganisation energy both sites stabilise at half their respective reorganisation energy.

highlights this, as shown in Figure 4.7.

4.2 Exciton Population Dynamics

Before comparing population transfer rates, we must address an important effect introduced by the Stokes shift: it breaks the symmetry of the system. Standard NISE models the bath using classical, symmetric noise without energy loss (dissipation). Because of this symmetry, the system doesn't care whether the exciton starts on the higher-energy site or the lower-energy site. The resulting population dynamics are essentially identical and will always converge to an equal split ($P = 0.5$). However, adding the Linear Reorganisation Shift or the memory kernel introduces directional energy relaxation. This means the starting energy setup (ΔE) suddenly dictates how the system behaves. If the exciton starts on the higher-energy site, the Stokes shift lowers that site's energy over time. This closes the energy gap between the two sites, moving them closer to resonance and

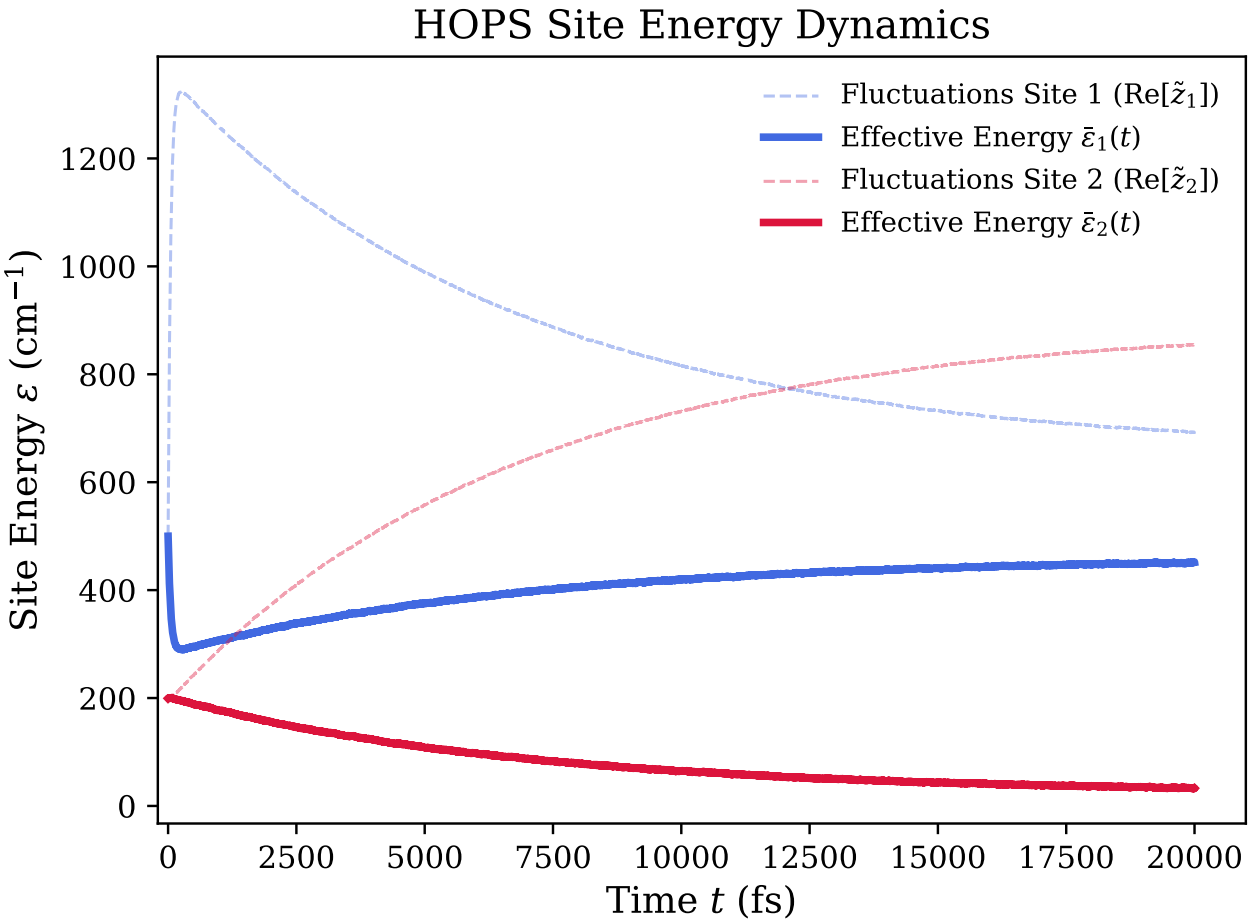

Figure 4.5. Showing the imaginary part of the effect on the Hamiltonian with the Effective Energy $\tilde{\varepsilon}_i(t)$, which is calculated as $\bar{\epsilon}_i(t) = H_{ii} - \mathrm{Im}\left[\tilde{z}_i(t)\right]$ and the real part "Stochastic Fluctuations (Re $[\tilde{z}_i]$)" is the thermal noise from the environment. $T = 300\,\mathrm{K}$, $\sigma_{ii} = 300\,\mathrm{cm}^{-1}$, $\tau = 50$ fs , $J = 10\ \mathrm{cm}^{-1}$,$\lambda = 215.82\ \mathrm{cm}^{-1}$

making it easier for the exciton to transfer. Conversely, if the exciton starts on the lower-energy site, the Stokes shift drives its energy even lower. This widens the energy gap and effectively traps the exciton, suppressing population transfer.

Therefore, to get physically accurate results and properly compare them with HOPS and HEOM benchmarks, the initial setup is critical. You must ensure that the energy offset ($+\Delta E$) is assigned to the correct site in the static Hamiltonian, so that it perfectly matches the physical conditions of the initial excitation. See Appendix A.2 for a more detailed reasoning and corresponding figures.

Comparing the population transfer between HOPS, NISE and the two new methods shows that the implemented Stokes shift in the NISE programs does produce an improvement to the population transfer predictions. When using the same parameters as used in Figure 4.7 to show the exciton population transfers we can see in Figure 4.8 that the linear shift and Memory Kernel Shifts do improve the trajectory. At the start, both NISE models with the Stokes shift improves the accuracy

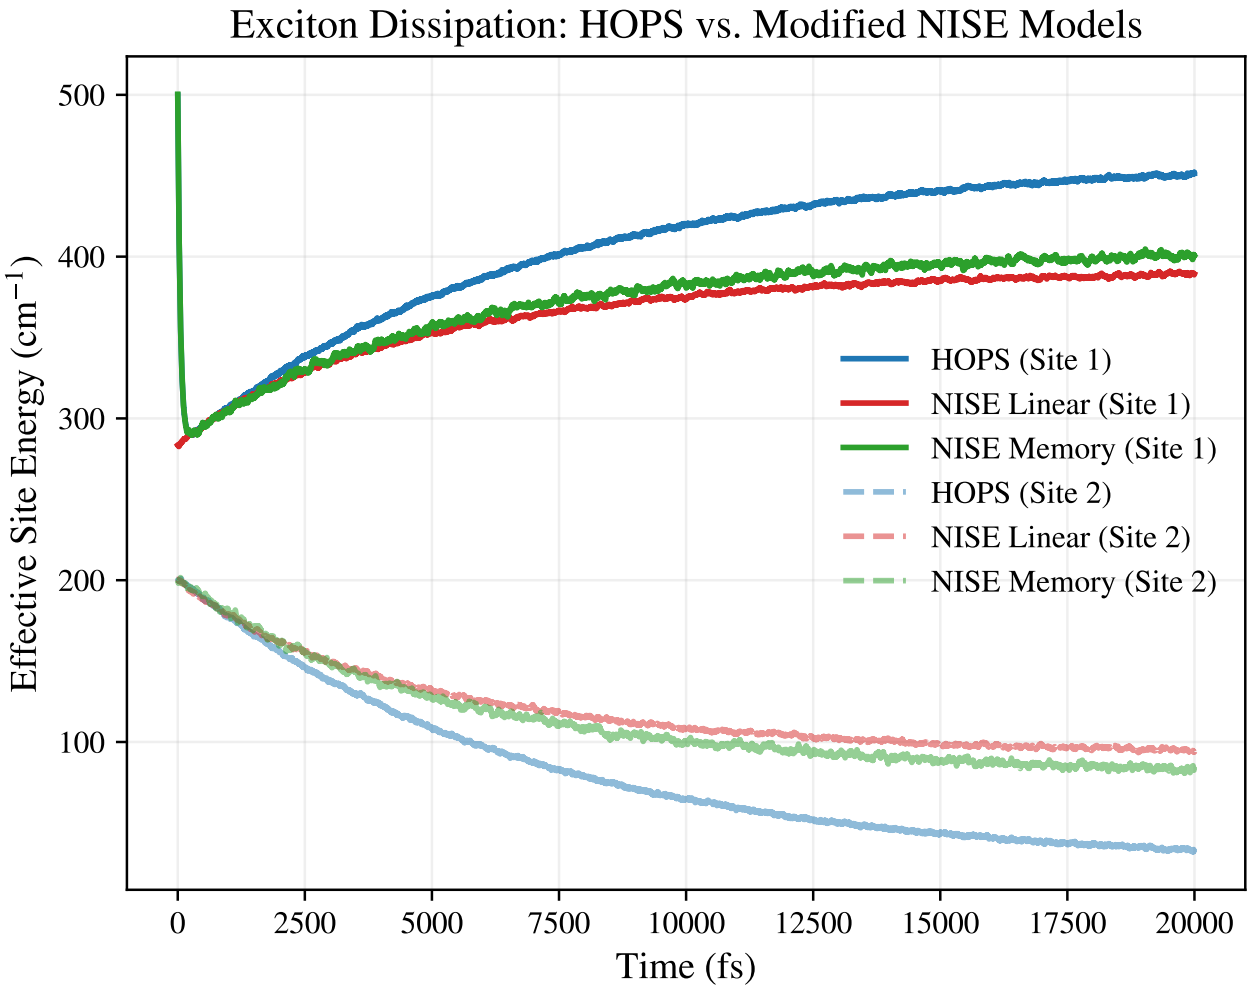

Figure 4.6. Comparing the effective energies for site 1 and site 2 of HOPS vs the NISE Linear Reorganisation Shift vs the Memory Kernel Shift. In the period of 0fs - 200fs the memory kernel and HOPS show excellent agreement, but the linear shift starts at a different position. Between 200fs-1500fs all 3 methods shown excellent agreement between the different methods. After that point they start to diverge.$\tau = 50$ fs, $\Delta E = 300$ cm^{-1}, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 300$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$ K, $\lambda = 215.82$ cm^{-1}. The starting energy for site 1 is 500cm^{-1} and site 2 is 200 cm^{-1}

and follow the HOPS trajectory, but the modified NISE models start to diverge from HOPS around 500fs. From that point onwards after 1500fs the two NISE models split with the Memory Kernel Shift showing better agreement with HOPS over the Linear reorganisation method and the unmodified NISE method. This can be shown in greater detail in Figure 4.9. The Linear Reorganisation Shift will start to merge with the NISE approximation as its rate of population transfer slows down. This is in contrast with the memory kernel that doesn't converge to the same population as NISE and continues below 0.5.

The justification for NISE going to 0.5 (along with the Linear Reorganisation Shift) is that NISE uses

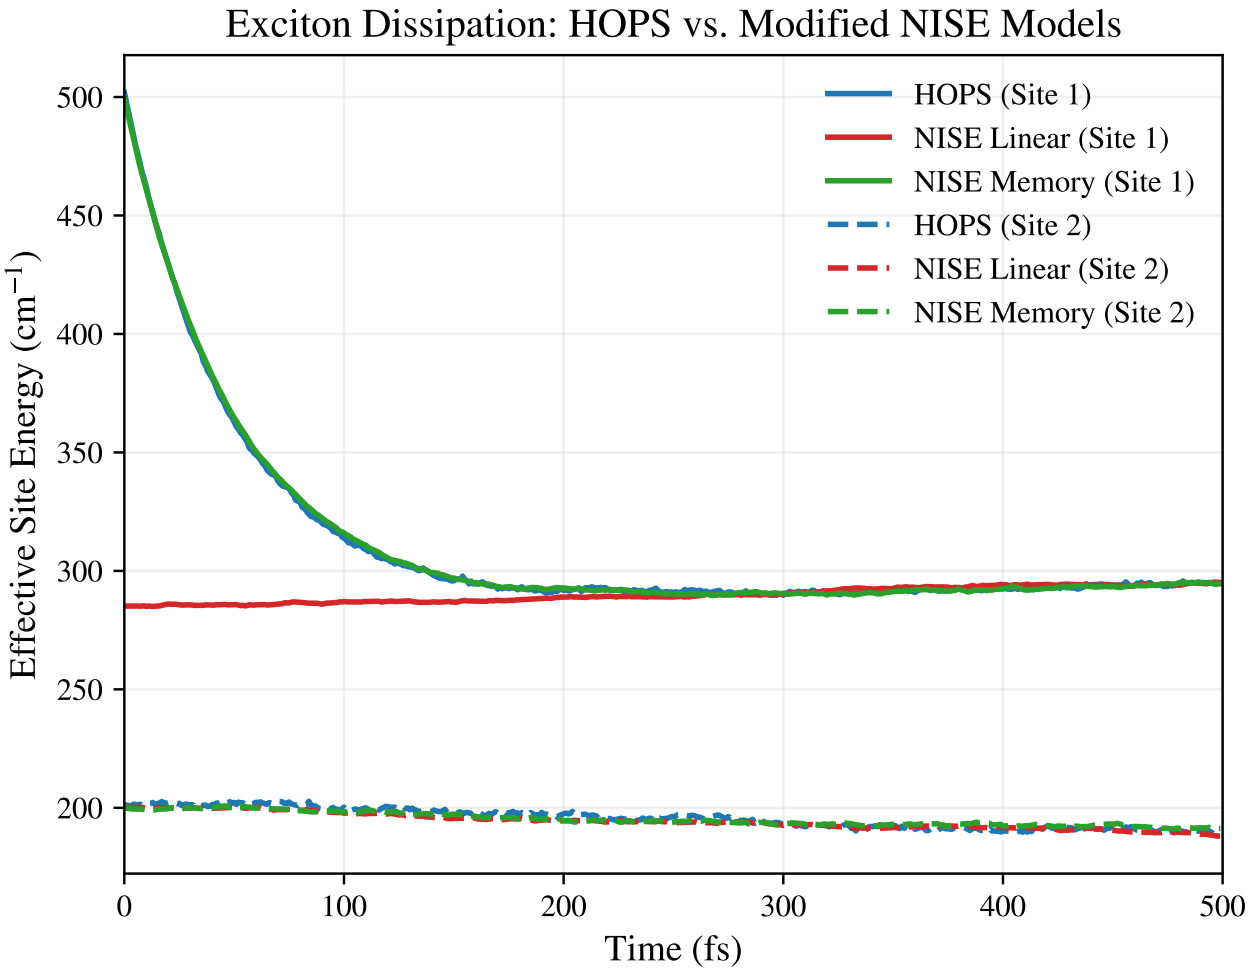

Figure 4.7. Comparing the effective energies for site 1 of HOPS vs the NISE Linear Reorganisation Shift vs the Memory Kernel Shift. The plot is zoomed in to show the first 500 fs to highlight the difference between the HOPS/memory kernel and the Linear Reorganisation Shift. $\tau = 50$ fs, $\Delta E = 300\ \text{cm}^{-1}$, $\sigma_1 = 300\ \text{cm}^{-1}$, $\sigma_2 = 300\ \text{cm}^{-1}$, $J = 10\ \text{cm}^{-1}$, $T = 300$ K, $\lambda = 215.82\ \text{cm}^{-1}$. The starting energy for site 1 is $500\ \text{cm}^{-1}$ and site 2 is $200\ \text{cm}^{-1}$

classical stochastic noise to represent the bath environment. This is symmetric and it determines the rate of transitions between the energy levels, so the rate that excitons jump down in energy ($E_1 \to E_2$) is the same rate that they jump up in energy ($E_2 \to E_1$). This means that the probability ratio:

$$\frac{P_1}{P_2} = \frac{\Gamma_{2\to1}}{\Gamma_{1\to2}} = 1 \tag{4.1}$$

is equal to 1. So $P_1 = P_2 = 0.5$, which is why the NISE model converges and finished to 0.5. The Linear Reorganisation Shift doesn't add in dissipation so it doesn't lose energy to the environment. The linear shift site energy drops according to a fixed schedule regardless of whether the exciton is actually on that site or not. The bath does not react to the exciton's presence as there is no back action.

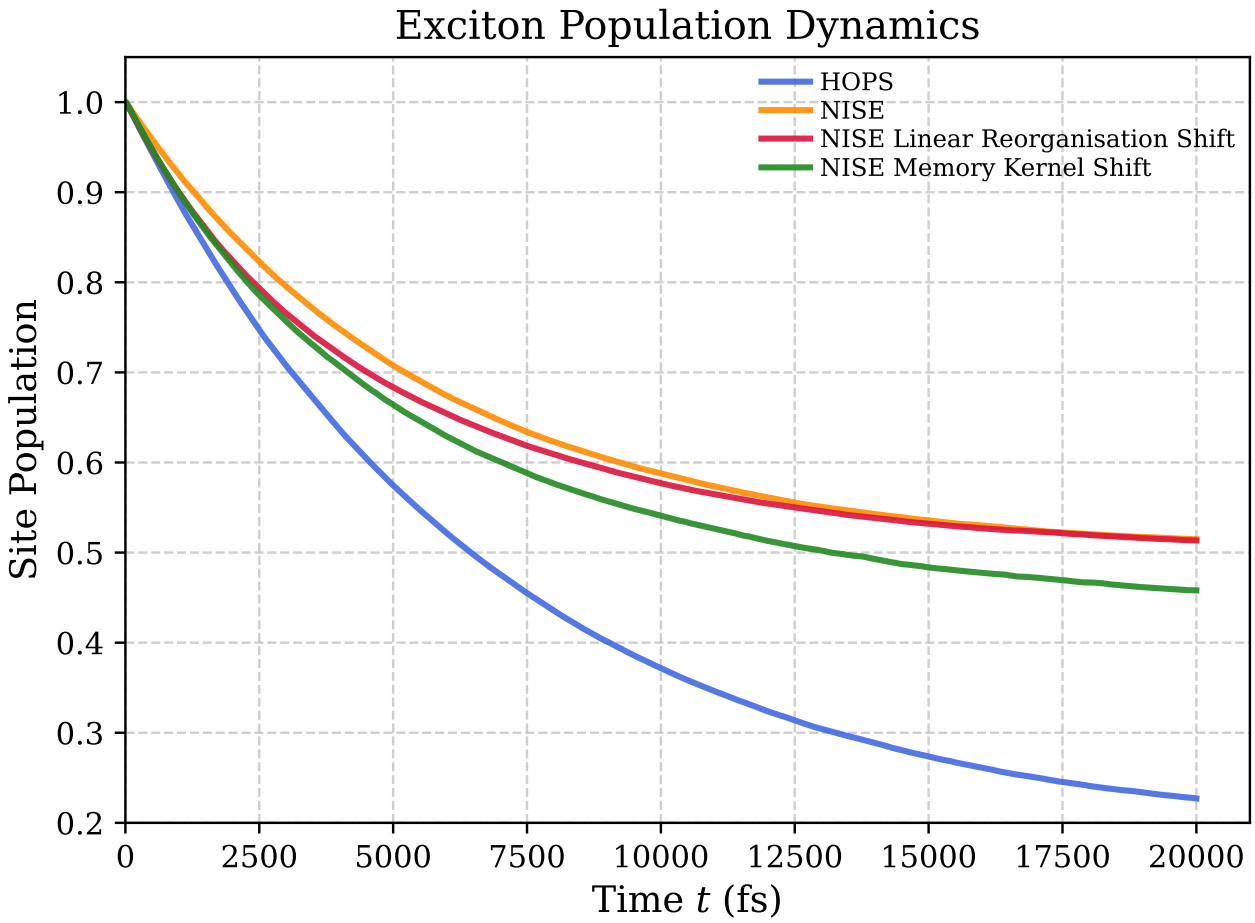

Figure 4.8. Comparing the population of the highest-energy chromophore (site basis) following initial excitation of site 1 for HOPS (blue), NISE with Linear Reorganisation Shift (red), NISE with Memory Kernel Shift (green), NISE with no added Stokes shift (orange). $\tau = 50$ fs, $\Delta E = 300$ cm^{-1}, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 300$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$ K, $\lambda = 215.82$ cm^{-1}.

This changes when looking at the memory kernel model as that introduces a dissipative force. This non linear feedback breaks the previous symmetry of NISE. If the exciton is on Site 1: P_1 is large, so the integral for $m_1(t)$ grows. This lowers the energy of Site 1 ($H_{11} \downarrow$). If the exciton moves to Site 2: P_2 becomes large, so $m_2(t)$ grows, lowering the energy of Site 2 ($H_{22} \downarrow$). This feedback adds the frictional force so instead of the site energy being a fixed value, it becomes a "well" that follows the particle. When the exciton moves to the lower energy site, in order for the exciton to jump back to site 1, it must overcome the ΔE plus the depth of the memory well it dug for itself. This makes the upward transition more difficult than the downwards transition so it allows the population to settle below 0.5. Conversely HOPS is driven towards the numerically exact solution that is given by the Boltzmann Distribution:

$$P_1 = \frac{e^{-\Delta E/k_B T}}{1 + e^{-\Delta E/k_B T}} \approx 0.19 \tag{4.2}$$

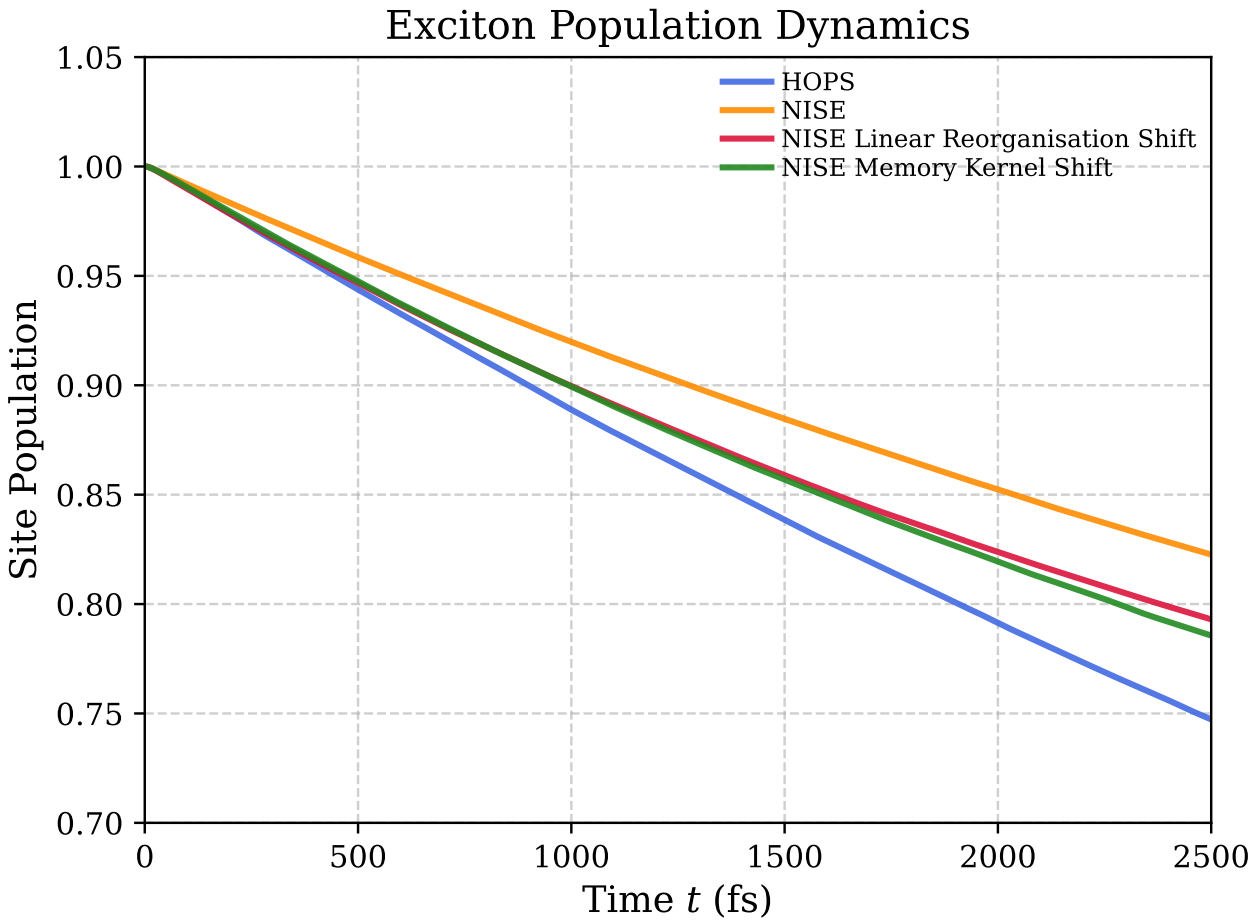

Figure 4.9. Detailed view of the first 2500 fs of population transfer from site 1. Comparisons are shown between HOPS (blue) and NISE variants utilizing linear reorganisation (red), Memory Kernel Shifts (green), and no shift (orange). Parameters: $\tau = 50$ fs, $\Delta E = 300$ cm^{-1}, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 300$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$ K, $\lambda = 215.82$ cm^{-1}.

which for these set of parameters is ≈ 0.19.

Which approximation is better is dependent on the parameters that are set. When keeping the same parameters, but changing the $\sigma_2 = 100$cm^{-1} the Linear Reorganisation Shift shows better correlation with HOPS than the memory during 0-12.5 ps, which is a contrast to when $\sigma_2 = 300$cm^{-1}, as shown in Figure 4.10. The improvement is shown in the the average energy level of the sites, but also the exciton site population. However, at around 12.5 ps, there is a crossover between the memory kernel and the linear shift. From that point onwards the memory kernel is closer to HOPS as is trending below 0.5 while NISE and NISE with linear shift is trending towards 0.5. This suggests that in the time interval of 0-12500 fs, the dynamics are primarily driven by the macroscopic reorganisation of the bath, which is more effectively captured by the smooth, pre-defined mean-field decay of the linear shift model. During this interval, the linear model's forced exponential relaxation closely mimics the initial Stokes shift seen in HOPS without the numerical 'lag' or stochastic noise inherent

in a population-dependent integral.

However, beyond this crossover point, the linear model's lack of a back-action mechanism causes it to stall at the infinite-temperature limit (P=0.5). The memory kernel's superiority in the long-time limit ($t > 12.5$ps) demonstrates that non-linear feedback is physically necessary to satisfy the fluctuation-dissipation theorem [25] and allow the system to dissipate energy into the lower-energy site, ultimately trending toward the formally exact Boltzmann distribution provided by the HOPS benchmark.

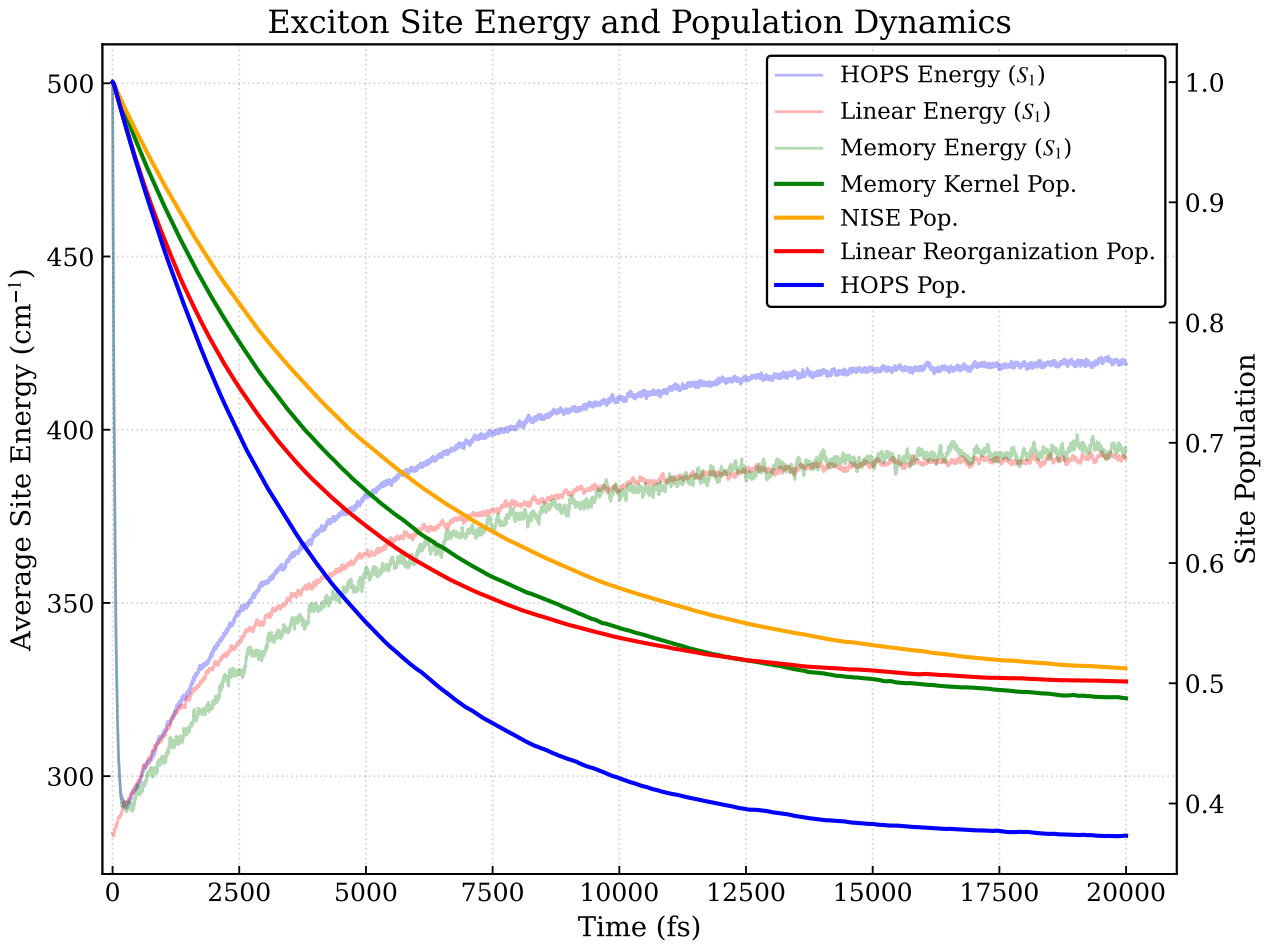

Figure 4.10. Comparing the population of the highest-energy chromophore (site basis) following initial excitation of site 1 for HOPS (blue), NISE with Linear Reorganisation Shift (red), NISE with Memory Kernel Shift (green), NISE with no added Stokes shift (orange). The effective energies for site 1 for each of these models are shown with matching colour and increased opacity. $\tau = 50$ fs, $\Delta E = 300$ cm^{-1}, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 100$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$K, $\lambda_1 = 215.82$ cm^{-1}, $\lambda_2 = 23.98$ cm^{-1}.

At very high temperatures (10^6 K) the agreement between NISE and HOPS should increase due to

agreement of the high temperature limit. Mathematically this is explained by Boltzmann factor:

$$P_1 = \frac{e^{-\beta E_1}}{e^{-\beta E_1} + e^{-\beta E_2}} = \frac{e^{-\Delta E/k_B T}}{1 + e^{-\Delta E/k_B T}} \tag{4.3}$$

As the temperature goes to infinity the Boltzmann factor reduces to e^0, which makes the equation:

$$\lim_{T\to\infty} P_1 = \frac{e^0}{1 + e^0} = \frac{1}{2} = 0.5. \tag{4.4}$$

We can also do a sanity check to make sure our methods are reproducing sane results as at this temperature the reorganisation energy is equal to $\lambda = \frac{\sigma^2}{2k_B T} \approx 0.000216\ \text{cm}^{-1}$, which should mean the Stokes shift is non existent and won't diverge compared to the NISE (which is temperature independent). We can show that this does hold true as in Figure 4.11 where the reorganisation shift is gone and both Stokes shift models replicate the original NISE model perfectly as the temperature is $T = 3 \times 10^6$ K.

Note that the HOPS program doesn't handle high temperatures well (anything above 1000 K caused issues) and can't be used while HEOM calculator via Nanohub has a max temperature of $10,000$ K. So comparison was only done against NISE (without the Stokes shift).

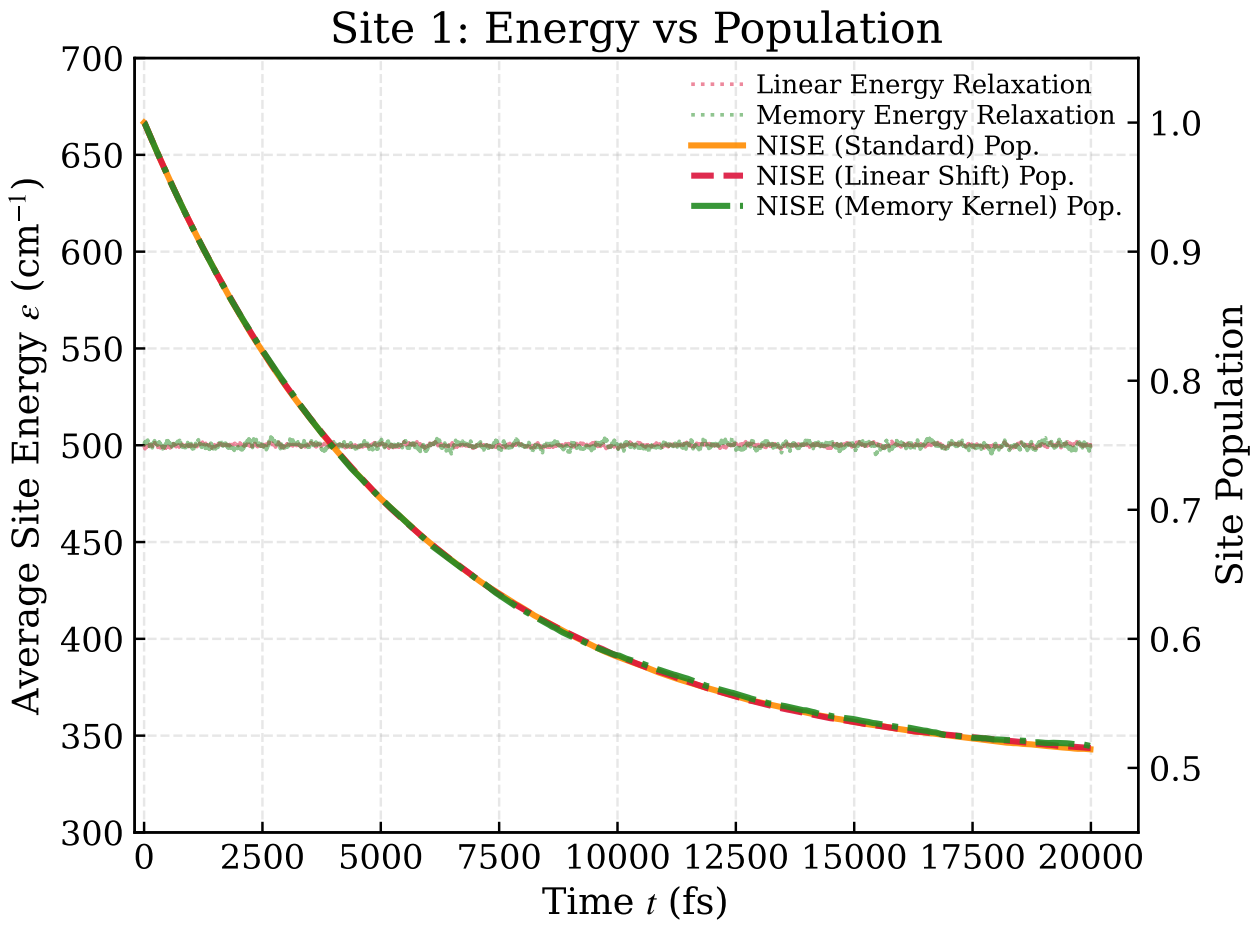

Figure 4.11. Comparing the population of the highest-energy chromophore (site basis) following initial excitation of site 1 for HOPS (blue), NISE with Linear Reorganisation Shift (red), NISE with Memory Kernel Shift (green), NISE with no added Stokes shift (orange). $\tau = 50$ fs, $\Delta E = 300$ cm^{-1}, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 300$ cm^{-1}, $J = 10$ cm^{-1}, $T = 3 \times 10^6$ K, $\lambda = 0.021582$cm^{-1}.

Chapter 5

Discussion

From the results, we can see that the implementation can lead to an improvement of the NISE simulation; however, the increase is dependent on the parameters chosen. If the parameters have a high J coupling or low σ parameters, then the added effect of the Stokes shift will be minimal. For instance, the LH2 parameters shown in Figure 4.3 shows that there is not much of a Stokes shift effect due to the large energy difference (ΔE) and the high coupling $J = 47\text{cm}^{-1}$. However, in certain scenarios where the coupling is low, and the reorganisation energy will be large, then the Stokes shift can be worth implementing. Due to the added effect of the bath memory, the memory kernel has shown to give better results due to implementing the "back action" feedback loop of the system to the bath. This allows the memory kernel to track the effective energy better than the Linear Reorganisation Shift. Occasionally in some results, the Linear reorganisation can show improved results in the early to medium transition, but at room temperature the memory kernel will have more accurate results over time.

The primary reason for attempting to improve NISE, when you have more numerically accurate models such as HEOM or HOPS, is for the scalability at large N sites. In a recent paper, a structure of 2639 molecules was simulated [32] . With this number of molecules, both NISE implementations do not significantly increase the overhead cost of the simulation.

While the propagation of the wavefunction is a global matrix operation, the feedback mechanisms are strictly local. In the Linear Shift model, the energy drop is just a simple addition to each site that follows a set schedule. In the memory kernel version, while the shift is dynamic, it is calculated site-by-site. The "memory" for Site A doesn't care what is happening at Site B; it only looks at its own history.

Because of this, adding more sites only adds a tiny bit of simple arithmetic (like addition or multiplication) for each new molecule. This is called linear scaling $\mathcal{O}(N)$. In contrast, the most expensive part

of any simulation is moving the particle the wavefunction propagation, which is much heavier. Since the computer is already "visiting" each site to update the noise and move the particle, calculating the extra back-action shift happens at the same time and is effectively "free."

Testing with a substantial amount of sites to test the computational overhead could have been done, however, the system would have to be equally optimised for all 3 methods to truly compare the computational speeds as the total system increases in sites N. While a direct comparison of computational speeds across a large number of sites (N) would have yielded valuable insights, it would not have been necessary as we can approximate the computational cost. Our primary objective is to evaluate how the inclusion of the Stokes shift modifies the NISE Hamiltonian; for this purpose, an analytical approximation of the computational cost is sufficient.

The successful implementation of the Stokes shift, particularly through the Memory Kernel Shift, expands the potential applications of the NISE framework well beyond simple exciton transport. Because this modified model can now simulate directional energy relaxation and the stabilization of lower-energy states via environmental reorganisation, it is highly suitable for modelling charge transfer dynamics. In both natural reaction centeres [12] and artificial organic photovoltaics [33], capturing the energetic "trapping" of separated charges is critical to predicting overall efficiency. Historically, accurately modelling these dissipative CT states in large molecular blends required computationally heavy methods like HOPS or HEOM to capture the necessary system-bath interactions. However, by preserving the linear $\mathcal{O}(N^2)$[9] scalability of the base NISE approach, the memory kernel method developed here provides a computationally feasible pathway to simulate these large-scale charge separation networks and the vital transition from light harvesting to stable chemical energy storage.

Chapter 6

Conclusion

In conclusion, this thesis establishes a middle ground between the numerically exact but computationally intensive HOPS and HEOM methodologies and the efficient, yet thermodynamically limited, NISE framework. By implementing the Linear Reorganisation Shift or the Memory Kernel Shift, we have demonstrated that it is possible to re-introduce the environmental "back-action" necessary to recover the Stokes shift within a mixed quantum-classical approximation. While the Linear Shift provides a simplified correction that improves short-time dynamics, it ultimately fails to deviate from the infinite-temperature limit, eventually converging to a uniform population distribution. In contrast, the Memory Kernel Shift, by accounting for the temporal response of system to the bath via a convolution integral of the population history, allows the system to satisfy the fluctuation-dissipation theorem and relax closer toward a proper Boltzmann distribution.

Crucially, these modifications preserve the $\mathcal{O}(N^2)$ scalability of the NISE method[9], ensuring that simulations of vast photosynthetic networks, such as the LH2 complex, remain feasible without the combinatorial explosion of auxiliary wavefunctions inherent in hierarchy-based methods. These results suggest that the inclusion of non-linear feedback mechanisms is a beneficial step toward accurately modelling the "downhill" energy flow in biological aggregates while maintaining the computational efficiency required for large-scale artificial light-harvesting technologies.

A direction for future research lies in the potential to combine the Memory Kernel Shift with existing frameworks for enforcing detailed balance in NISE, such as the Thermalised NISE (TNISE) methodology developed in related studies [34][8]. While the memory kernel effectively captures the dynamic Stokes shift and the time-dependent relaxation of the potential energy landscape, TNISE focuses on ensuring that long-term population transfer consistently adheres to the Boltzmann distribution. Integrating these two approaches would create a more complete semiclassical model that accounts for both the physical 'back-action' of the system onto the bath and the global requirement

for thermodynamic equilibrium. Such a hybrid method would further enhance the accuracy of NISE, bringing it even closer to the predictive power of numerically exact methods like HOPS and HEOM while maintaining its superior computational scalability for large-scale biological systems.

Appendix A

Appendix

A.1 Computational Scaling of the HOPS Method

In this section, we look at how the computational cost of the HOPS method grows as we increase the size of the environment. Since biological systems like LH2 involve many independent environmental sites, we need to know if the method remains practical as the number of sites N increases.

In the HOPS framework, the state of the system is described by a hierarchy of "auxiliary wavefunctions," denoted as $|\psi^{(\mathbf{k})}\rangle$. Each wavefunction is identified by a label $\mathbf{k} = (k_1, k_2, \dots, k_N)$, where each integer k_m represents the excitation level of the m-th environmental mode. To make the calculation possible on a computer, we have to cut this hierarchy off at a maximum depth, which we call $\mathcal{K}$. This truncation means we only keep wavefunctions where the total number of excitations across all modes is less than or equal to $\mathcal{K}$:

$$\sum_{m=1}^{N} k_m \leq \mathcal{K} \tag{A.1}$$

To find out how many wavefunctions (N_{wf}) we actually need to track, we have to count all the possible ways to arrange these excitations. This is a classic "stars and bars" counting problem. It is equivalent to figuring out how many ways we can put $\mathcal{K}$ excitations into $N + 1$ possible "bins" (where the extra bin represents any unused excitations up to the limit $\mathcal{K}$). The total number of wavefunctions is given by the binomial coefficient[26][27]:

$$N_{\text{wf}} = \binom{N + \mathcal{K}}{N} = \frac{(N + \mathcal{K})!}{N!\mathcal{K}!} \tag{A.2}$$

To see how this grows as the number of modes N gets larger, we can expand the factorials for a fixed

hierarchy depth $\mathcal{K}$:

$$N_{\text{wf}} = \frac{(N+\mathcal{K})(N+\mathcal{K}-1)\cdots(N+1)}{\mathcal{K}!} \tag{A.3}$$

If the number of modes is much larger than the depth we choose for the hierarchy ($N \gg \mathcal{K}$), then the N terms in the numerator become the most important part of the expression. Since there are $\mathcal{K}$ such terms multiplying together, the number of wavefunctions grows roughly as N to the power of $\mathcal{K}$:

$$N_{\text{wf}} \approx \mathcal{O}(N^{\mathcal{K}}) \tag{A.4}$$

Because the computer has to calculate a new time-derivative for every single one of these wavefunctions at every time step, the total calculation time is directly proportional to N_{wf}. Therefore, the computational effort required for HOPS scales polynomially with the number of modes:

$$\text{Total Cost} \propto N^{\mathcal{K}} \tag{A.5}$$

This polynomial scaling is why HOPS is considered a powerful method—it avoids the "exponential explosion" seen in some other quantum methods—but it also shows that as we want more accuracy (increasing $\mathcal{K}$), the simulation very quickly becomes much more expensive to run[5].

A.2 Reversing Energy

When the energy difference between the sites changes from $\Delta = -300$ cm^{-1} to $\Delta = +300$ cm^{-1}, the dynamics of the NISE model with the added Stokes shift change significantly when compared to NISE without the Stokes shift. Standard NISE assumes the high temperature limit ($k_B T \gg \hbar\omega$), so the effective Stokes shift is negligible. This causes a symmetry in the system, which forces the site populations to remain equal as the system propagates, as long as the J coupling and $|\Delta E|$ are equal.

In effect, this causes some unphysical realities, such as the site populations produced by the Hamiltonians in Equation A.6 being identical to each other at every timestep:

$$H \approx \begin{pmatrix} 500 & 10 \\ 10 & 200 \end{pmatrix} \approx \begin{pmatrix} 200 & 10 \\ 10 & 500 \end{pmatrix} \tag{A.6}$$

Introducing the Stokes shift, introduces directional flow into the system. Consider the example where the initial energy of site 1 (which is excited at $t = 0$) is lower than that of site 2. As can be shown in Figure A.1, standard HOPS and the Linear Reorganisation Shift start with site 1 at 500 cm^{-1} and site

2 at 200 cm^{-1}, whereas the linear reorganisation model with reversed energy starts at 200 cm^{-1} for site 1 and 500 cm^{-1} for site 2.

In this reversed scenario, an exciton jump to site 2 becomes much less likely because the Stokes shift causes site 1 to relax further down into its energy well. This deepens the energetic trap and dramatically slows down the exciton transfer rate between the sites. This effect can be seen in Figure A.2, where the site population of the Linear Reorganisation Shift propagates at the slowest rate.

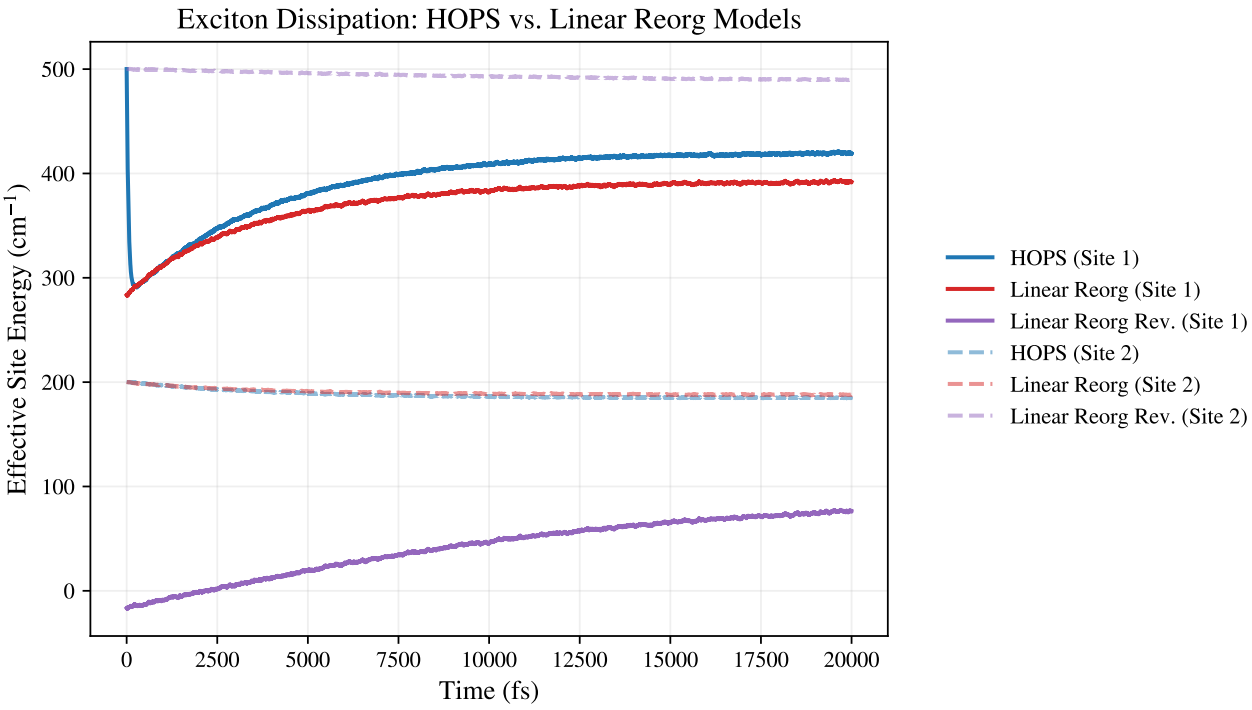

Figure A.1. Comparing the average exciton effective energy of HOPS (blue), NISE with Linear Reorganisation Shift (red), NISE with Linear Reorganisation Shift with the energy reversed (purple). The parameters are $\tau = 50$ fs, $\Delta E = 300$ cm^{-1} with the reversed energies having their site energy switched, $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 100$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$ K, $\lambda_1 = 215.82$ cm^{-1}, $\lambda_2 = 23.98$ cm^{-1}.

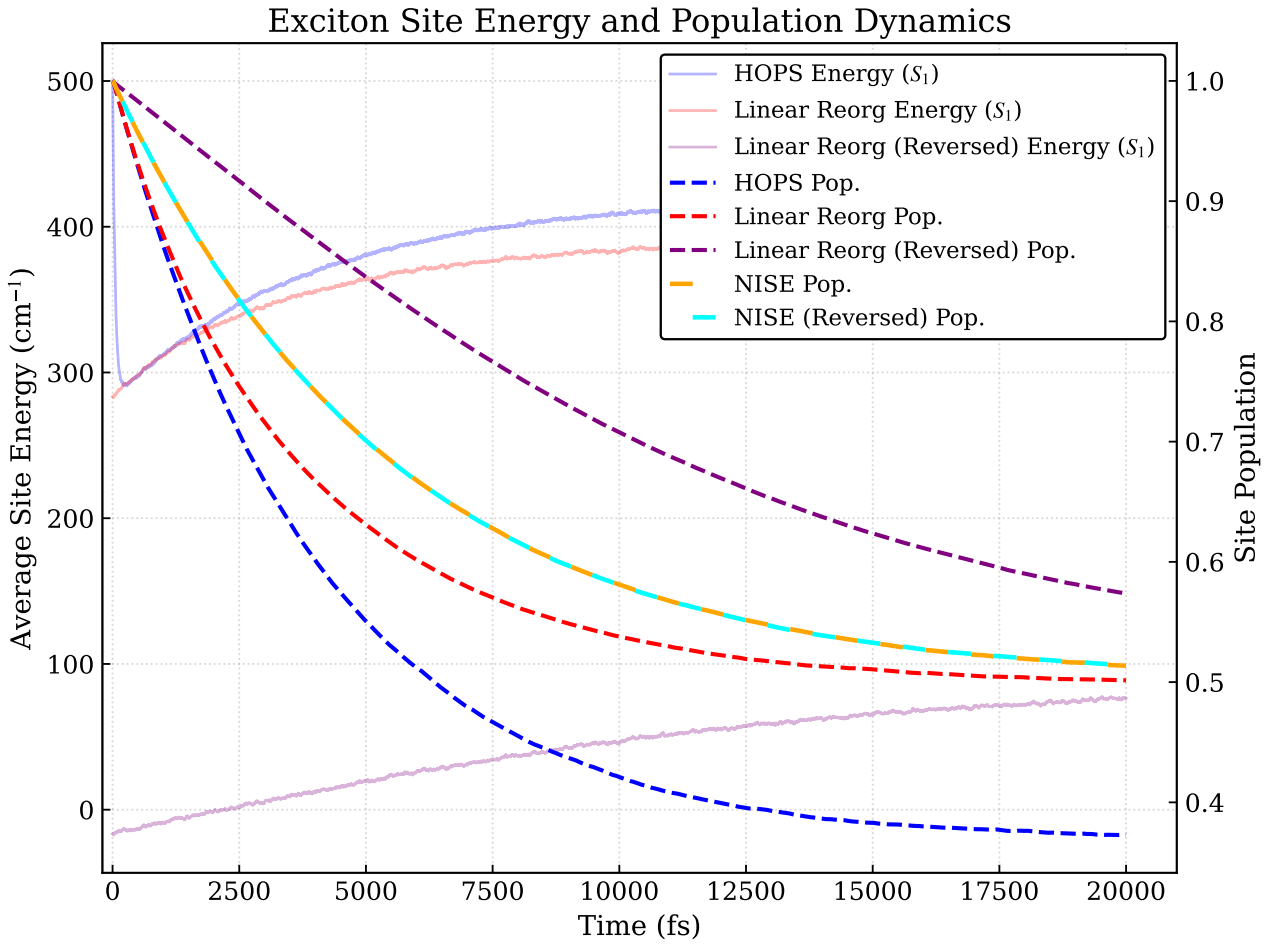

Figure A.2. Comparison of population dynamics and effective site energies. Shown are the populations of the highest-energy chromophore (site basis) following initial excitation of site 1 for HOPS (blue), NISE with Linear Reorganisation Shift (red), NISE with Linear Reorganisation Shift with reversed energies (purple), standard NISE with no Stokes shift (orange), and standard NISE with reversed energies (cyan). The corresponding effective site energies are shown with matching colours and lower opacity. Parameters: $\tau = 50$ fs, $\Delta E = 300$ cm^{-1} (for reversed cases, site 1 = 200 cm^{-1} and site 2 = 500 cm^{-1}), $\sigma_1 = 300$ cm^{-1}, $\sigma_2 = 100$ cm^{-1}, $J = 10$ cm^{-1}, $T = 300$ K, $\lambda_1 = 215.82$ cm^{-1}, $\lambda_2 = 23.98$ cm^{-1}.

Bibliography

[1] G. S. Schlau-Cohen, "Principles of light harvesting from single photosynthetic complexes," *Interface Focus*, vol. 5, no. 3, p. 20140088, 2015.

[2] C. P. van der Vegte *et al.*, "Atomistic modeling of two-dimensional electronic spectra and excited-state dynamics for a light harvesting 2 complex," *The Journal of Physical Chemistry B*, vol. 119, no. 4, pp. 1302–1313, 2015.

[3] F. Wu, D. Finkelstein-Shapiro, *et al.*, "Optical cavity-mediated exciton dynamics in photosynthetic light harvesting 2 complexes," *Nature Communications*, vol. 13, no. 1, p. 6864, 2022.

[4] R. Hartmann and W. T. Strunz, "Exact open quantum system dynamics using the hierarchy of pure states (hops)," *Journal of Chemical Theory and Computation*, vol. 13, no. 12, pp. 5834–5845, 2017.

[5] Y. Tanimura, "Numerically exact approach to open quantum dynamics: The hierarchical equations of motion (heom)," *The Journal of Chemical Physics*, vol. 153, no. 2, p. 020901, 2020.

[6] L. A. Baker and S. Habershon, "Photosynthetic pigment-protein complexes as highly connected networks: implications for robust energy transport," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 473, no. 2201, p. 20170112, 2017.

[7] S. Mukamel, *Principles of Nonlinear Optical Spectroscopy*. Oxford Series in Optical and Imaging Sciences, New York: Oxford University Press, 1995.

[8] T. L. C. Jansen, "Simple quantum dynamics with thermalization," *The Journal of Physical Chemistry A*, vol. 122, no. 1, pp. 172–183, 2018.

[9] C. Liang and T. L. C. Jansen, "An efficient n^3 scaling propagation scheme for simulating two-dimensional infrared and visible spectra," *Journal of Chemical Theory and Computation*, vol. 8, no. 5, pp. 1706–1713, 2012.

[10] Y. M. Holtkamp, E. Godinez-Ramirez, and U. Kleinekathöfer, "Spectral densities, structured noise and ensemble averaging within open quantum dynamics," *The Journal of Chemical Physics*, vol. 161, no. 13, 2024.

[11] Y. Holtkamp *et al.*, "Machine-learned correction to ensemble-averaged wave packet dynamics," *The Journal of Chemical Physics*, vol. 159, no. 9, p. 094107, 2023.

[12] F. Fassioli, R. Dinshaw, *et al.*, "Photosynthetic light harvesting: excitons and coherence," *Journal of The Royal Society Interface*, vol. 11, no. 92, p. 20130901, 2014.

[13] R. S. Knox, "Introduction to Exciton Physics," in *Collective Excitations in Solids* (B. Di Bartolo, ed.), pp. 183–245, Boston, MA: Springer US, 1983.

[14] S. C. J. Meskers, "The exciton model for molecular materials: Past, present and future?," *ChemPhysChem*, vol. 24, no. 24, p. e202300666, 2023.

[15] S. Datta and X. Marie, "Excitons and excitonic materials," *MRS Bulletin*, vol. 49, no. 9, pp. 852–861, 2024.

[16] L. Shi and A. P. Willard, "Modeling the effects of molecular disorder on the properties of frenkel excitons in organic molecular semiconductors," *The Journal of Chemical Physics*, vol. 149, no. 9, p. 094110, 2018.

[17] T. Iitaka, "Solving the time-dependent schrodinger equation numerically," *Physical Review E*, vol. 49, no. 5, pp. 4684–4690, 1994.

[18] A. S. Coolidge, H. M. James, and R. D. Present, "A study of the franck–condon principle," *The Journal of Chemical Physics*, vol. 4, no. 3, pp. 193–211, 1936.

[19] G. E. Uhlenbeck and L. S. Ornstein, "On the theory of the brownian motion," *Physical Review*, vol. 36, no. 5, pp. 823–841, 1930.

[20] V. M. Kenkre and P. Reineker, *Exciton Dynamics in Molecular Crystals and Aggregates*, vol. 94 of *Springer Tracts in Modern Physics*. Berlin: Springer-Verlag, 1982.

[21] R. W. Hakala, "Simple justification of the form of boltzmann's distribution law," *Journal of Chemical Education*, vol. 44, no. 11, p. 657, 1967.

[22] J. R. Lakowicz, *Principles of Fluorescence Spectroscopy*. New York, NY: Springer Science & Business Media, 3rd ed., 2006.

[23] A. Skallberg, *Photoemission and Characterization of Neutrophils and Nanoparticles: Energy Mapping and Elemental Composition with sub mu m Resolution*. Doctoral thesis, Linköping University, Linköping, Sweden, 2020.

[24] L. O. Kostjukova, S. V. Leontieva, and V. V. Kostjukov, "The vibronic absorption spectra and electronic states of acridine orange in aqueous solution," *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, vol. 249, p. 119302, 2021.

[25] R. Kubo, "The fluctuation-dissipation theorem," *Reports on Progress in Physics*, vol. 29, no. 1, pp. 255–327, 1966.

[26] D. Suess, W. T. Strunz, and A. Eisfeld, "Hierarchical equations for open system dynamics in fermionic and bosonic environments," *Journal of Statistical Physics*, vol. 159, no. 6, pp. 1408–1423, 2015.

[27] H. D. Wobben, "Modelling exciton transfer dynamics in open quantum systems using a hierarchy of pure states," master's thesis, University of Groningen, 2024.

[28] C. Kreisbeck and T. Kramer, "Exciton dynamics lab for light-harvesting complexes (gpu–heom)," 2013.

[29] T. L. C. Jansen, "Nise 2017: Numerical integration of the Schrodinger equation for coherent multidimensional spectra." `https://github.com/lacourjansenlab/NISE_2017`, 2017.

[30] A. Bondarenko, *Modeling of excitonic properties in tubular molecular aggregates*. PhD thesis, University of Groningen, 2019.

[31] D. S. Hoffmann, P. M. Dohmen, M. Sokolov, U. Kleinekathofer, and M. Elstner, "Exciton transfer simulations in a light-harvesting 2 complex reveal the transient delocalization mechanism," *The Journal of Physical Chemistry B*, vol. 129, no. 13, pp. 3345–3365, 2025.

[32] V. Eric, X. Li, *et al.*, "Manifestation of hydrogen bonding and exciton delocalization on the absorption and two-dimensional electronic spectra of chlorosomes," *The Journal of Physical Chemistry B*, vol. 127, no. 5, pp. 1097–1109, 2023.

[33] K. Khoussa *et al.*, "Artificial intelligence in organic photovoltaics: Predicting power conversion efficiency from the molecular chemical structure of (donor/acceptor) pairs," *Engineering Reports*, vol. 7, no. 8, p. e70334, 2025.

[34] P. Van Vliet, "Detailed balance in the numerical integration of the schrodinger equation," master's thesis, University of Groningen, 2021.