Modern simulation methods for vibrational and electronic spectroscopy

Annina Zoé Lieberherr

Lincoln College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2025

Modern simulation methods for vibrational and electronic spectroscopy

Annina Z. Lieberherr

Lincoln College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosopy

Trinity 2025

Spectroscopy offers a convenient way to collect information about molecular systems in a non-invasive way. Vibrational and electronic spectra may both be simulated by calculating response functions and Fourier transforming them to obtain spectra.

Vibrational spectra are often calculated using path integral methods, which can approximately include nuclear quantum effects in the linear response function at a linearly scaling computational cost. This thesis introduces a generalisation of f-QCMD, a promising recently developed method, to the condensed phase, enabling the calculation of the vibrational spectrum of liquid water and ice. The thesis then uses f-QCMD to analyse polariton spectra, which occur when a molecular system is placed in a microcavity. The vibrational spectrum gives a uniquely uncontroversial (compared to other properties like reaction rates) insight into cavity effects. In fact, a simple harmonic oscillator model which takes only the cavity-free spectrum and the geometry of the cavity as input can quantitatively predict cavity molecular dynamics results, implying that the cavity simply acts as an optical filter on the linear vibrational spectrum.

The thesis then moves on to the simulation of two-dimensional electronic spectra. It introduces the mean field method 'equatorial Ehrenfest' that qualitatively captures the two-dimensional electronic spectra of Frenkel exciton models, which emulate light harvesting complexes. This approach is inspired by the previously proposed polar Ehrenfest approach, which is 32 times more expensive and less accurate than equatorial Ehrenfest. Substituting Ehrenfest with a more accurate dynamics method only moderately improves the quantitative agreement with exact results, at increased computational cost.

Acknowledgements

The first person I want to thank is my supervisor David Manolopoulos. I joined Oxford when the department was just coming out of Covid. While it took some time for the social life within the department to be revived, there was never a dull day with David in the office. He has guided and supported me in my research and beyond, and I leave Oxford a stronger researcher thanks to him.

My gratitude also goes out to past and present members of the Manolopoulos group. Johan in particular has been a deep source of knowledge as well as a great friend. Lorenz has livened up the office in the short time he has been here, and has made a busy final year a pleasure. I have also enjoyed seeing numerous Part II students grow throughout their year in the group. Thank you to Seth, Chloe, Dasha, Filippo, Cole, Tom, Marcus, and Yangzhou, for their curiosity and their company.

I also want to thank my parents, who have supported me through my decisions that have brought me to this point. I don't take the freedom I have enjoyed for granted, and I would not be in the same place without their love and support. Equally, I thank my brother Moritz for his catch up calls on random week nights, and for sharing his creative spirit and his projects with me. And finally, thank you to my grandparents in Wil and my grandmother in St. Gallen.

I have enjoyed my time in Oxford tremendously, thanks to the company of a number of remarkable individuals. To my housemates (in various combinations) Lorenz, Håvard, Jeremy, Estelle and Johannes, thank you for many exciting or relaxing evenings at home, at house parties and at balls. Many joyful evenings, and lunches, were also spent with people from the Swiss society (special thanks to Viki, Amanda, Ludwig, Flurin and Lucien), with

the Sustainable Finance society (special thanks to Dennis, Marcel, Justin and Katie), with the PTCL people (special thanks to Hugh, the Tew group, particularly Daniel, Andrew, Alex and Nick, and to Ana and Zuzanna). Oxford also gave me the opportunity to test out new sports disciplines, like boxing (for a short time) with OUABC, rowing with LCBC, and Handball.

Even though it has been some time since we have lived in the same time, I want to thank people from back home: My oldest friends Lexi, Nivetha and Larissa, who have been in my life since school times; My lockdown family, Annatina, Tim, Tudor, Mischa, Dino and Marco; My uni friends Marcel, Johannes, Markus, Janik, Jakob, Kathrin, and Max.

This list would not be complete, of course, without thanking Hans, who has been at my side since (literally) my first day in Oxford after emerging from the mandatory post-travel isolation and who has supported me in more ways than I can name.

A DPhil takes so much beyond research, and I have all the people listed here (and more) to thank for making life a lot more enjoyable in and around Oxford. Much love to all of you.

Published results

This thesis contains results from published articles, listed here.

Chapter 2 is a rewritten and extended version of the article

> J. E. Lawrence, A. Z. Lieberherr, T. Fletcher, D. E. Manolopoulos. "Fast quasi-centroid molecular dynamics for water and ice," *The Journal of Physical Chemistry B*, **127** (2023)

Chapter 3 is a rewritten and extended version of the article

> A. Z. Lieberherr, S. T. E. Furniss, J. E. Lawrence, D. E. Manolopoulos."Vibrational strong coupling in liquid water from cavity molecular dynamics," *The Journal of Chemical Physics*, **158** (2023)

An article based on Chapter 4 is in preparation.

Contents

1 Introduction **1**

2 Vibrational spectroscopy **11**

2.1 Introduction 11

2.2 Path integral methods for spectroscopy 13

2.2.1 Ring polymer molecular dynamics 18

2.2.2 Thermostatted RPMD 18

2.2.3 Centroid molecular dynamics 19

2.2.4 Quasi-centroid molecular dynamics 24

2.3 f-QCMD for condensed phase systems 28

2.4 Results for water and ice 30

2.5 Summary 38

3 Vibrational strong coupling **41**

3.1 Introduction 41

3.2 Cavity molecular dynamics 44

3.3 Results 46

3.3.1 Nuclear quantum effects 46

3.3.2 More realistic simulations 48

3.3.3 A simple harmonic oscillator model 51

3.3.4 Applications to experiments 53

3.3.5 Equivalence to the transfer matrix method 58

3.4 Summary 60

4 Electronic spectroscopy **63**

4.1 Introduction . . . 63

4.2 Theory . . . 65

4.2.1 Linear spectroscopy . . . 67

4.2.2 Two-dimensional spectroscopy . . . 68

4.3 Frenkel exciton models . . . 74

4.4 Existing simulation methods . . . 77

4.5 An improved pure state Ehrenfest approach . . . 84

4.6 Results . . . 93

4.6.1 Biexciton model . . . 93

4.6.2 Fenna-Matthews-Olson complex . . . 97

4.7 Improved population dynamics . . . 100

4.7.1 Spin mapping . . . 102

4.7.2 Results . . . 104

4.8 Linear spectra . . . 107

4.9 Summary . . . 109

5 Conclusion **113**

Bibliography **129**

Appendix **131**

A Ring polymer normal modes . . . 131

B Note on spin mapping initial conditions . . . 132

Chapter 1

Introduction

Spectroscopy is the most widespread method used to probe and understand molecular systems. As early as at the beginning of the nineteenth century, diffraction gratings were used to analyse the light spectrum emitted by the sun, leading to the characterisation of the Fraunhofer lines.[1,2] Simultaneously, chemical elements started to be characterised by the light emitted upon heating the element in a flame.[3] It would take some more decades until these two approaches were connected to determine the chemical composition of the sun,[4] but it showed the power of spectroscopy early on. While spectroscopic methods were at first limited to the visible region of light, new discoveries were quickly made that opened the door to new frequency regions, such as the discovery of X-rays and microwave frequencies.[5,6] An enormous number of experimental techniques have since been added to the toolbox, and spectroscopy has found applications in a variety of fields such as synthetic chemistry,[7] medicine,[8,9] geology,[10] and exoplanet research.[11] Together with the breadth of experimental methods, a range of simulation methods has been developed for their prediction.

This thesis will focus on recent developments in the simulation of vibrational and electronic spectroscopy. Linear vibrational spectroscopy, or infrared (IR) spectroscopy, probes the vibrational frequencies of a system in its electronic ground state. The vibrational spectrum is often calculated by finding the harmonic normal mode frequencies through diagonalising the mass-weighted Hessian at the ground state geometry. This gives a stick spectrum, whereas an experimental spectrum will have vibrational peaks with non-zero widths due to homo-

geneous and inhomogeneous broadening. Homogeneous broadening arises due to the finite radiative lifetime of physical vibrations and produces line widths with a Lorentz profile where the width is inversely proportional to the vibrational life time.[12] Inhomogeneous broadening, on the other hand, occurs in collections of oscillators due to slight changes in each oscillator's environment. Due to environment effects, the vibrational frequency of each oscillator may be shifted slightly. Averaged over all oscillators, inhomogeneous broadening gives rise to a Gaussian line profile.[12,13] Inhomogeneous broadening will have a more significant impact in the condensed phase than the gas phase. At low temperatures and in a crystalline solids such as ice, all oscillators will have a similar environment, in which case the inhomogeneous broadening decreases. If both radiative processes and environmental effects occur for a given vibration, then its line shape will be the convolution of a Lorentzian profile (from the radiative processes) and a Gaussian profile (from environment effects), which combine to a Voigt profile.[12,13]

A stick spectrum obtained from the diagonalisation of the mass-weighted Hessian may be broadened artificially to resemble an experimental spectrum more closely. Even then, it is not sufficient to include only the normal mode frequencies at the ground state geometry. At any temperature there will be an ensemble of configurations of the system contributing to the vibrational spectrum. The configurations that contribute may be found according to the thermal equilibrium distribution of the system. It is unlikely that an artificial broadening procedure will reproduce the correct line widths created by the thermal distribution. Another problem of the stick spectrum approach based on the mass-weighted Hessian at the potential minimum is that any anharmonic effects in the vibrational frequencies are neglected, since only harmonic normal modes are included. (Although anharmonic effects can in principle be included through vibrational perturbation theory.[14])

Alternatively, the vibrational spectrum can be related to the Fourier transform of the dipole moment autocorrelation function.[15] The autocorrelation function's initial conditions are by definition averaged over the thermal ensemble and thus inherently capture a number of characteristics of a spectrum that are not attainable by diagonalising the Hessian at the equilibrium geometry. To calculate the dipole autocorrelation function, one would usu-

ally run a thermostatted molecular dynamics (MD) trajectory at a constant temperature to sample initial conditions, followed by MD trajectories at a constant energy, along which the autocorrelation is evaluated. When the particle number, volume and temperature are kept fixed such as for the sampling of initial conditions, we call a trajectory an NVT trajectory. A trajectory along which the energy is kept constant (together with the particle number and volume) is called an NVE trajectory.

While the autocorrelation approach will recover appropriate line widths, a classical MD simulation will still not capture the anharmonicities in the molecule's vibrations. This can be understood as follows. A classical trajectory of a particle can be thought of as a ball rolling around on a potential energy surface. It may, in particular, reach the very bottom of the well. A quantum particle would behave very differently. It cannot reach the bottom of potential energy wells, but can at most reach a zero-point energy that lies above the actual potential minimum. The thermal equilibrium distribution of a quantum particle will therefore look different from a classical particle. If the potential energy surface has typical anharmonic contributions, then the spread of the quantum particle around the potential energy minimum leads to a softening of the average potential it experiences, and to a reduced vibrational frequency compared to the classical case.

Zero-point energy (and tunnelling) are examples of nuclear quantum effects and they are neglected within classical dynamics. As a consequence, the peaks in the classical vibrational spectrum are often blue shifted with respect to the true quantum spectrum. On the other hand, it quickly becomes prohibitively expensive to run a full quantum simulation for even modest system sizes.

Fortunately, relating the vibrational spectrum to the dipole autocorrelation function allows us to employ path integral methods to calculate the autocorrelation function, which are able to include some nuclear quantum effects at a classically-scaling computational cost. Path integrals were first developed by Feynman.[16] Parrinello and Rahman later proposed to use a discretised path integral to calculate static observables such as the diamagnetic susceptibility and radial distribution functions.[17] Their path integral molecular dynamics (PIMD) method introduced the concept of the ring polymer as a representation of the discretised path in

imaginary time, and showed that classical ring polymer trajectories could exactly reproduce the static properties of the systems they considered. Note that observables obtained with PIMD, and with the other path integral methods described below, are constrained by the potential energy surface. Any approximations made to the potential energy surface will influence the results for the observables, but PIMD can be considered exact for a given potential energy surface (in the absence of identical particle exchange effects).

It is also possible to develop path integral *approximations* to the dipole autocorrelation function. The centroid molecular dynamics (CMD)[18] and ring polymer molecular dynamics (RPMD)[19] methods were long the first two methods used by the community. In RPMD, observables are evaluated by averaging them over the ring polymer beads while in CMD they are evaluated on the ring polymer centroid, which is the average Cartesian configuration of all beads. Both methods' abilities to approximate correlation functions have been thoroughly investigated, with particular interest in the short-time limit. In a general potential, RPMD has a leading error of $\mathcal{O}(t^7)$ for the position autocorrelation function[20,21], $\mathcal{O}(t^5)$ for the velocity autocorrelation function[20,21] (a simple way to see this is by noting that the velocity autocorrelation function is the second derivative of the position autocorrelation function), and $\mathcal{O}(t^3)$ for a general autocorrelation function of a non-linear local (position-dependent) operator. CMD provides lower-order approximations. Its leading error is $\mathcal{O}(t^6)$ for the position autocorrelation function, $\mathcal{O}(t^4)$ for the velocity autocorrelation function and $\mathcal{O}(1)$ for a general, nonlinear position-dependent operator. The $\mathcal{O}(1)$ error stands out, because it implies that CMD does not reproduce the correct zero-time limit of the autocorrelation function of an operator that depends non-linearly on the position (which is simply the thermal average of the square of the operator). Nevertheless, both RPMD and CMD have been shown to be suitable for the simulation of various dynamics properties, including chemical reaction rates,[22,23] diffusion coefficients[24,25] and thermal conductivity.[26] However, they both show artefacts when applied to the dipole autocorrelation function and the vibrational spectrum.[27–29]

Vibrational peaks in spectra calculated with CMD can show spurious red shifts at low temperatures in systems that contain both bending and stretching modes, which has become known as the curvature problem.[30,31] RPMD peaks remain centred around the correct fre-

quency, but can appear split into unphysical peaks due to resonances between the vibrations of the system and the artificially introduced harmonic modes of the ring polymer.[28,32] Although Rossi *et al.* showed that it is possible to remove the unphysical peaks in the RPMD spectra, with an approach called thermostatted RPMD (TRPMD),[20] the TRPMD peaks appear broadened. With all three of the established methods plagued by their own artefacts, the field was at an impasse.

Recently, quasi-centroid molecular dynamics (QCMD) entered the stage, developed by Trenins *et al.*[33] QCMD is closely related to CMD, but uses curvilinear coordinates instead of Cartesian ones which are better suited to the interplay of bending and stretching modes that CMD usually struggles with. QCMD has so far been tested on a number of gas-phase molecules,[33,34] as well as some condensed phase systems.[33] One drawback, however, was its significant computational cost, which motivated Fletcher *et al.* to develop fast QCMD (f-QCMD).[35] This new method approximates full QCMD, but it has been shown to almost perfectly reproduce QCMD results for gas-phase molecules at a reduced computational cost.[35,36]

In Chapter 2, we provide a generalisation of f-QCMD to the condensed phase. After a review of the RPMD, CMD, TRPMD, QCMD and f-QCMD approaches, we describe our f-QCMD implementation for liquid water and ice, and present its results.

Having established f-QCMD as a method with which we can study liquid water, we next focused on applying it to water in an optical microcavity. A cavity is made up of two (often gold) mirrors, such that the light trapped between them forms standing waves. In particular, the cavities are usually on the order of a few μm long. In this case, the (fundamental or overtone) frequencies of the cavity can approach and couple to physical vibrational frequencies. Vibrational strong coupling in a vacuum cavity has drawn extraordinary attention in recent years, primarily owing to claims about tuneable chemical reactions. Examples range from enhanced reactivity[37,38] and changing regioselectivity[39] to controlling the optical and electronic response of semiconductors.[40–42]

There are some effects that are well understood. For example, the coupling of the cavity to a molecular vibrational mode gives rise to two new polariton states, which are below and above the original vibrational frequency. The difference in energy between the two is the

Rabi splitting, which is predicted to scale as $\sqrt{N}$, where N is the number of molecules in the cavity. Note that there is a physical limit on the Rabi splitting. At some point the cavity is completely filled with molecules, and the Rabi splitting cannot be increased further. Also, the coupling strength between a single molecule and the cavity scales as $1/\sqrt{N}$, meaning that the cavity is only weakly coupled to any single molecule.[43] Finally it is straightforward to show that within the dipole approximation, the cavity does not affect the static properties of the system within it in the classical limit.[43]

However, there are other effects that are not so well understood. For example, there is currently no unified model for a physical process that can account for the effect of vibrational strong coupling on chemical reactions, particularly in vacuum (where the cavity is not driven by a laser).[44] Further, some experimental groups have tried and failed to reproduce cavity experiments.[45,46]

Here, we focus on a non-controversial application: the vibrational spectrum of liquid water inside a cavity. Li *et al.* showed that it is possible to calculate the cavity vibrational spectrum of liquid water with classical MD, and developed a one-dimensional harmonic oscillator model which could qualitatively (but not quantitatively) predict their results.[43,47] To determine whether nuclear quantum effects had a significant impact on the polariton spectra (and therefore may be relevant in other cavity effects), they subsequently calculated the TRPMD spectra of liquid water in a cavity together with some static equilibrium properties.[48] Just like in the classical case, they showed that static properties were the same inside or outside the cavity, and that the cavity thus had no impact whatsoever. In the vibrational spectrum, however, they noticed red shifts and broadening in the polariton peaks compared to the classical cavity spectrum. However, since they used TRPMD for their simulations, it was difficult for them to distinguish between genuine nuclear quantum effects and TRPMD artefacts in the (vibrational) spectrum.

Equipped with the tools necessary to simulate liquid water outside a cavity with f-QCMD, it was only natural to study cavity spectra with our f-QCMD implementation next. Chapter 3 describes how we introduce the cavity to our f-QCMD formalism from Chapter 2, and presents some cavity f-QCMD results. We will also show that an appropriate simple harmonic

oscillator (SHO) model can be developed that quantitatively reproduces the results of cavity MD simulations using just the cavity-free spectrum and the geometry of the cavity as input, and discuss the implications of this result.

The second half of this thesis is focused on two-dimensional electronic spectroscopy.[49,50] This is a type of spectroscopy in which three light pulses interact with a sample, giving rise to two-dimensional spectra at varying delay times. Two-dimensional electronic spectroscopy has been used extensively to study light harvesting complexes like the Fenna-Matthews-Olson complex[51,52] and light harvesting complex II,[53–55] which occur in bioluminescent bacteria. These complexes manage to harvest sunlight into chemical energy through a process where the light induces an electronic excitation, which is then funnelled from photoreceptors through a number of chromophores to the baseplate, where it is turned into chemical energy.[53] The two-dimensional spectra can be used to characterise the path by which the energy travels through the complex.[52]

Many aspects of this process remain elusive. For example, the bacteria are able to achieve very high efficiency rates for harvesting energy from light,[56,57] and it is not clear why. In early experiments, oscillations in the peak intensities were observed, which led to the hypothesis that electronic coherence could be the key to harvesting sun light efficiently.[51,58–60] However, later work revealed that these oscillations were in fact due to vibronic coherences, and that electronic coherences decay too quickly due to inhomogeneous dephasing to be involved in the light harvesting process.[61,62]

The role of nuclear quantum effects has also been debated, with some research claiming that they accelerate the transport of the exciton to the baseplate.[63–66] More recent work, however, suggests that they instead slow down the energy transport.[67] Additional insight into the dynamics of light harvesting complexes is thus still needed.

One way to study the two-dimensional spectra of light-harvesting complexes is with the hierarchical equations of motion (HEOM).[68] In simulations, the light-harvesting complexes are usually represented by Frenkel exciton models, in which the electronic excited states of the chromophores of the complex are linearly coupled to baths of harmonic oscillators. HEOM can treat these Frenkel exciton models exactly, which is one of the reasons they are frequently

used as benchmark models. These models may not be entirely appropriate to describe light harvesting complexes, however. For example, anharmonicities in the excited states or excited states that are rotated with respect to the ground state (by Duschinsky rotations) cannot be described within the Frenkel exciton framework. The computational cost of HEOM would increase in such systems, so alternative methods are needed.

It is possible to develop semiclassical approximations to the exact dynamics. For example, van der Vegte *et al.* were the first to develop a mean-field Ehrenfest approach for the simulation of two-dimensional electronic spectra.[69] Then, Atsango *et al.* developed a second, also Ehrenfest-based, method based on decomposing a coherence into four pure states.[70] Both tested their approach on Frenkel exciton models because exact benchmarks are available from HEOM. Van der Vegte *et al.* only tested their approach on one set of model parameters, which belong to the 'slow bath' regime in the language of Frenkel exciton models (i.e. the bath relaxes on a long time scale).[69] In this parameter regime, their method agreed quite closely with HEOM, but it was not clear how it would fare with different parameters. Atsango *et al.* tested their pure state approach on both slow- and fast-bath Frenkel exciton models.[70] It became clear that while their method did quite well in the slow-bath regime (similar to van der Vegte *et al.*'s observations[69]), it broke down in the fast-bath regime.[70] It also incurs a significantly increased computational effort compared to the mean classical path approach. The pure state decomposition is invoked three times (for each light-matter interaction). After the third light-matter interaction, up to $4^3 = 64$ pure state trajectories thus need to be run in order to generate a single sample of a nonlinear response function.

We will show in Chapter 4 that a small modification to the pure state Ehrenfest method gives rise to a new, equatorial Ehrenfest, method which does well in both the slow- and fast-bath parameter regimes. It is also about 32 times less expensive than Atsango *et al.*'s approach, because we run 32 times fewer pure state trajectories by taking advantage of the symmetry of the pure states and the nonlinear response functions.

To summarise, this thesis describes the development of two new simulation methods: f-QCMD for condensed phase vibrational spectroscopy, and equatorial Ehrenfest for two-dimensional electronic spectroscopy. It also presents the application of f-QCMD to liquid

water in a cavity, where it turns out that cavity f-QCMD simulations can be replaced by a simple (and faster) harmonic oscillator model. Chapter 5 summarises the main observations and conclusions in more detail, and provides perspectives on future work.

Chapter 2

Vibrational spectroscopy

2.1 Introduction

The impact of nuclear quantum effects on vibrational spectra is well understood. Due to the spread of the nuclear wave packet around a potential energy minimum, the nuclei feel a lower effective curvature than in a classical calculation, which decreases the vibrational frequency. In the spectrum, this causes a shift to longer wavelengths, or a 'red shift', in the peak positions. Path integral methods are an efficient approach to reproduce spectra with these red shifts at a low computational cost. However, the simulation of vibrational spectra through path integral methods has a long and chequered history.

PIMD was introduced in a seminal study by Parrinello and Rahman.[17] It is an algorithm to calculate thermal expectation values using a discretised Feynman path integral (also called the ring polymer).[16] Since its inception, PIMD has remained a powerful and popular method in the simulation of condensed phase systems because it scales polynomially with system size. It also allows for the exact calculation of static properties such as dielectric constants, magnetic susceptibilities and pair correlation functions.[17,71] However, PIMD itself cannot capture dynamical observables, such as vibrational spectra, chemical reaction rates and diffusion coefficients. Other methods were developed which extended the path integral formalism to dynamical observables, namely CMD,[18] and RPMD.[19] These methods find approximate time correlation functions, from which we can extract the desired dynamical observables during post-processing.[15] CMD and RPMD are well established for numerous applications,[22–26,72]

but they both struggle to describe vibrational spectra. In RPMD, a single vibrational peak is sometimes split into multiple smaller peaks because of spurious resonances between the physical and ring polymer vibrations.[28,30,32] It is possible to mitigate the spurious resonances by thermostatting the offending internal modes, an approach known as TRPMD,[20] but then the thermostat causes an artificial broadening of the peaks. CMD follows the dynamics of the ring polymer's average Cartesian configuration (the 'centroid'), which evolves on the effective potential averaged over the fluctuations of the ring polymer, and therefore avoids the spurious oscillations by design.[18] However, it was noticed early on that it suffers from a different problem, which has been termed the curvature problem.[30,31] If there is a curved feature in the potential energy surface (e.g. when there is a bending vibrational mode), then CMD peaks show an artificial shift to lower frequencies. The curvature problem is more severe at low temperatures, but already occurs in liquid water at temperatures as high as 300 K.[73]

As an alternative, Trenins *et al.* proposed to replace the centroid by a quasi-centroid, which is defined through curvilinear coordinates.[33] The coordinates can be tailored to curved features in the potential, thereby avoiding the curvature problem. They tested their QCMD method on gas-phase water and ammonia, liquid water and ice.[33,34] They showed that QCMD agrees with CMD at high temperatures, and that there is no sign of the curvature problem as the temperature is lowered. Furthermore, QCMD spectra agree well with exact results (where it is feasible to compute exact results). All of this is very encouraging. However, QCMD in its original formulation is typically 10-100 times as expensive as TRPMD or CMD.[34] To address this, Fletcher *et al.* subsequently developed f-QCMD to accelerate QCMD by pre-computing an approximate quasi-centroid potential of mean force.[35] f-QCMD has been applied to gas-phase water, ammonia, methane and hydrogen peroxide.[35,36] The f-QCMD results are virtually indistinguishable from the QCMD results for gas phase water obtained by Trenins *et al.*[33] and for gas-phase ammonia obtained by Haggard *et al.*[34] This validates the approximation made to the effective potential of mean force.

f-QCMD has so far only been applied to gas-phase systems. In this chapter, we describe its generalisation to liquid water and ice. Although we only study water, we expect that the same ideas can be used for other condensed phase systems.

Section 2.2 introduces the path integral formalism, and reviews previously developed path integral methods. Section 2.2.4 then introduces the theory behind QCMD followed by a description of different QCMD algorithms with their applications. Subsequently, Sections 2.3 and 2.4 explain the generalisation of f-QCMD to condensed phase systems and its application to liquid water and ice, respectively. Section 2.5 summarises the key developments of this chapter.

2.2 Path integral methods for spectroscopy

The vibrational spectrum can be written as[15]

$$I(\omega) = \frac{\beta}{6\varepsilon_0 cV} \int_{-\infty}^{\infty} e^{-i\omega t} C_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) \mathrm{d}t, \tag{2.1}$$

where $C_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \langle \dot{\boldsymbol{\mu}}(0) \cdot \dot{\boldsymbol{\mu}}(t) \rangle$ is the dipole derivative autocorrelation function, $\beta = 1/(k_\mathrm{B}T)$ is the inverse temperature, c is the velocity of light, V is the volume of the system and ε_0 is the vacuum permittivity. The angular brackets $\langle \cdot \rangle$ symbolise a thermal average. In fact, many dynamical properties can be expressed in terms of such time correlation functions,[15] including the diffusion coefficient (which is related to the velocity autocorrelation function[15]), the thermal conductivity (which is related to the energy flux autocorrelation function[74–76]) and chemical reaction rate constants (which are related to flux-side correlation functions[77–79]). As this thesis is focused on spectroscopy, it will be restricted to the discussion of dipole derivative autocorrelation functions.

The thermal average in Equation (2.1) can either be done in a classical or in a quantum manner. On the one hand, the classical correlation function is

$$C^{\mathrm{cl}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{(2\pi\hbar)^{3N} Z_{\mathrm{cl}}} \int \mathrm{d}\boldsymbol{p}_0 \int \mathrm{d}\boldsymbol{q}_0 e^{-\beta \mathcal{H}(\boldsymbol{p}_0, \boldsymbol{q}_0)} \dot{\boldsymbol{\mu}}(\boldsymbol{q}_0) \cdot \dot{\boldsymbol{\mu}}(\boldsymbol{q}_t), \tag{2.2}$$

where N is the number of particles in the system with positions and momenta $\boldsymbol{q}_i$ and $\boldsymbol{p}_i$ of the

ith particle, respectively, $Z_{\mathrm{cl}} = \frac{1}{(2\pi\hbar)^{3N}} \int \mathrm{d}\boldsymbol{p} \int \mathrm{d}\boldsymbol{q} e^{-\beta\mathcal{H}(\boldsymbol{p},\boldsymbol{q})}$ is the classical partition function,

$$\mathcal{H}(\boldsymbol{p},\boldsymbol{q}) = \sum_{i=1}^{N} \frac{|\boldsymbol{p}_i|^2}{2m_i} + V_{\mathrm{cl}}(\boldsymbol{q}) \tag{2.3}$$

is the classical Hamiltonian composed of kinetic energy and a potential energy V_{cl}, and m_i is the mass of particle i. We will assume throughout that we are dealing with a three-dimensional system, such that $\boldsymbol{p}_i$ and $\boldsymbol{q}_i$ hold x, y and z components.

The variables $\boldsymbol{p}_0$, $\boldsymbol{q}_0$ in Equation (2.2) are initial coordinates, which are propagated under Hamiltonian equations of motion,

$$\dot{\boldsymbol{q}}_i = \frac{\partial\mathcal{H}}{\partial\boldsymbol{p}_i} = \frac{\boldsymbol{p}_i}{m_i}, \tag{2.4a}$$

$$\dot{\boldsymbol{p}}_i = -\frac{\partial\mathcal{H}}{\partial\boldsymbol{q}_i} = -\nabla_{\boldsymbol{q}_i} V_{\mathrm{cl}}(\boldsymbol{q}). \tag{2.4b}$$

The integration can for example be done numerically using the Velocity-Verlet algorithm.[80] $\boldsymbol{q}_t$ and $\boldsymbol{p}_t$ are the coordinates after integration up to time t. On the other hand, the (Kubo-transformed) quantum correlation function is

$$\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{\beta Z_{\mathrm{qu}}} \int_0^\beta \mathrm{d}\lambda \mathrm{Tr}[e^{-\beta\hat{\mathcal{H}}} \hat{\dot{\boldsymbol{\mu}}}(0) \cdot \hat{\dot{\boldsymbol{\mu}}}(t + i\lambda)], \tag{2.5}$$

where $Z_{\mathrm{qu}} = \mathrm{Tr}[e^{-\beta\hat{\mathcal{H}}}]$ is the quantum partition function and $\hat{\mathcal{H}}$ is the Hamiltonian operator. For any operator $\hat{B}$, $\hat{B}(t) = e^{i\hat{\mathcal{H}}t/\hbar}\hat{B}e^{-i\hat{\mathcal{H}}t/\hbar}$ is the Heisenberg time evolved operator. (Note that $\hbar \neq 1$ here.) Readers may be more familiar with the following definition of the quantum correlation function,

$$C^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{Z_{\mathrm{qu}}} \mathrm{Tr}[e^{-\beta\hat{\mathcal{H}}} \hat{\dot{\boldsymbol{\mu}}}(0) \cdot \hat{\dot{\boldsymbol{\mu}}}(t)]. \tag{2.6}$$

The two quantum correlation functions are equivalent, since their Fourier transforms are related via

$$\mathcal{F}[C^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}](\omega) = \frac{\beta\hbar\omega}{1 - e^{-\beta\hbar\omega}} \mathcal{F}[\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}](\omega), \tag{2.7}$$

which can be seen by expanding the trace in the energy eigenbasis.[19] Here we will focus on the Kubo-transformed quantum correlation function in Equation (2.5) instead, because it

shares symmetry properties with the classical correlation function. It is thus more natural to compare the classical correlation function to the Kubo-transformed one. In particular they are both real,

$$\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t)^* \tag{2.8}$$

they satisfy detailed balance,

$$\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(-t), \tag{2.9}$$

and time-reversal,

$$\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(-t)^*. \tag{2.10}$$

Equations (2.8)-(2.10) also hold for $C^{\mathrm{cl}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t)$. Together, they imply that $C^{\mathrm{cl}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t)$ and $\tilde{C}^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t)$ are both real and even functions of t. The same cannot be said for $C^{\mathrm{qu}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t)$ in Equation (2.6), which only satisfies the time reversal property.

We are interested in simulating condensed phase systems with a large number of atoms. A classical calculation scales polynomially in N, but a quantum calculation scales exponentially. It therefore tends to be much harder to calculate quantum correlation functions than classical ones. But classical correlation functions lack key quantum effects, some of which are well known to have a non-negligible impact on the vibrational spectrum. Here we focus on nuclear quantum effects in particular. As discussed in Section 2.1, the spread of the nuclear wave packet around potential energy minima causes red shifts in the vibrational peaks.

Path integral methods lie somewhere between quantum and classical calculations: They include some nuclear quantum effects while avoiding full quantum simulations, and can therefore be easily applied to the condensed phase systems considered here. To do so, they map the quantum Kubo-transformed correlation function onto the classical correlation function of the ring polymer.[18,19] The ring polymer consists of P copies of the classical system, called beads, which are coupled together by harmonic springs (see Figure 2.1 for a schematic of the ring polymer and how it captures nuclear quantum effects). Its (classical) Hamiltonian is

$$\mathcal{H}_P(\boldsymbol{P},\boldsymbol{Q}) = \sum_{j=1}^{P}\left[\sum_{i=1}^{N}\left(\frac{|\boldsymbol{P}_i^{(j)}|^2}{2m_i} + \frac{1}{2}m_i\omega_P^2|\boldsymbol{Q}_i^{(j)} - \boldsymbol{Q}_i^{(j+1)}|^2\right) + V_{\mathrm{cl}}(\boldsymbol{Q}^{(j)})\right], \tag{2.11}$$

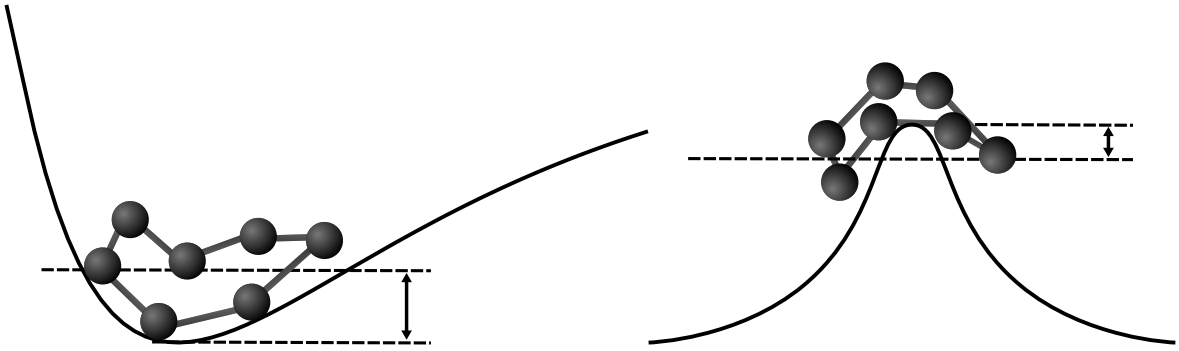

(a) (b)

Figure 2.1: Schematics of how the ring polymer captures (a) zero point energy and (b) tunnelling nuclear quantum effects. The long dashed lines represent the average energy of the beads. The average energy lies above (below) the potential energy minimum (maximum) in (a) (resp. (b)). The difference between the classical minimum (maximum) and the quantum minimum (maximum) is the nuclear quantum effect, represented by a double-sided arrow. Note that the beads are drawn above the potential energy surface to represent the ring polymer more clearly.

where $\boldsymbol{P}_i^{(j)}, \boldsymbol{Q}_i^{(j)}$ are the position and momentum coordinates of the ith atom in the jth bead, and $\boldsymbol{Q}^{(j)}$ contains all coordinates of bead j. The ring polymer is periodic (such that $\boldsymbol{Q}^{(P+1)} = \boldsymbol{Q}^{(1)}$), and the harmonic springs have a frequency $\omega_P = 1/(\beta_P \hbar)$, where $\beta_P = \frac{\beta}{P}$.

At very high temperatures, the ring polymer springs are extremely stiff and the ring polymer contracts to a single bead. This is a different way of saying that quantum mechanics coincides with classical mechanics at high temperatures. Note that, since the ring polymer comprises P copies of the classical system, it is P times more expensive to simulate than a classical system – but nevertheless a path integral simulation is usually much cheaper than a quantum simulation.

While the ultimate goal is to calculate time correlation functions, it is helpful to first consider what happens to thermal averages in the path integral formalism. Parrinello and Rahman were the first to propose this in an approach that is now known as PIMD.[17] Quantum thermal expectation values of the operator $\hat{A}$,

$$\langle \hat{A} \rangle = \frac{1}{Z_{\mathrm{qu}}} \mathrm{Tr}[e^{-\beta \hat{\mathcal{H}}} \hat{A}], \tag{2.12}$$

can be obtained within the PIMD formalism as

$$\langle \hat{A} \rangle = \lim_{P \to \infty} \langle A \rangle_{\mathrm{NVT},P}, \tag{2.13}$$

where

$$\langle A \rangle_{\mathrm{NVT},P} = \frac{1}{(2\pi\hbar)^{3NP} Z_P} \int \mathrm{d}\boldsymbol{P} \int \mathrm{d}\boldsymbol{Q} e^{-\beta_P \mathcal{H}_P(\boldsymbol{P},\boldsymbol{Q})} \overline{A}(\boldsymbol{Q}) \tag{2.14}$$

is a classical thermal average in the extended ring polymer phase space,

$$Z_P = \frac{1}{(2\pi\hbar)^{3NP}} \int \mathrm{d}\boldsymbol{P} \int \mathrm{d}\boldsymbol{Q} e^{-\beta_P \mathcal{H}_P(\boldsymbol{P},\boldsymbol{Q})} \tag{2.15}$$

is the partition function,

$$\overline{A}(\boldsymbol{Q}) = \frac{1}{P} \sum_{j=1}^{P} A(\boldsymbol{Q}^{(j)}), \tag{2.16}$$

and $A(\boldsymbol{Q}^{(j)})$ is the classical observable on bead j. By leveraging the ring polymer, we can thus obtain results at a quantum accuracy, but at a classical cost. Further, since PIMD is able to exactly capture quantum thermal averages (for systems of distinguishable particles), the formulation of quantum expectation values in terms of ring polymer classical averages is called the classical isomorphism.[17]

We would like to find a relation such as Equation (2.13) for the classical and Kubo-transformed correlation function, but unfortunately this is not possible. While it is fine to use classical methods to sample the thermal equilibrium (as in Equation (2.14)), the classical dynamics of the ring polymer cannot exactly capture the system's evolution from time 0 to time t. In particular, classical dynamics does not capture quantum coherence. We can thus at most find approximations to the Kubo-transformed quantum correlation function by classical correlation functions in the ring polymer phase space. For the systems we are interested in, such coherence effects are expected to wash out, and we therefore do not expect to face major inaccuracies in our results on their account. Various algorithms have been developed to approximate quantum correlation functions. We will next discuss a few of these methods.

2.2.1 Ring polymer molecular dynamics

The RPMD dipole derivative autocorrelation functions is[19]

$$C^{\mathrm{RPMD}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{(2\pi\hbar)^{3NP}Z_P}\int \mathrm{d}\boldsymbol{P}_0 \int \mathrm{d}\boldsymbol{Q}_0 e^{-\beta_P \mathcal{H}_P(\boldsymbol{P}_0,\boldsymbol{Q}_0)}\overline{\dot{\boldsymbol{\mu}}}(\boldsymbol{Q}_0)\cdot\overline{\dot{\boldsymbol{\mu}}}(\boldsymbol{Q}_t), \tag{2.17}$$

where $\boldsymbol{P}_0, \boldsymbol{Q}_0$ are initial conditions of the ring polymer, $\boldsymbol{P}_t, \boldsymbol{Q}_t$ are the ring polymer coordinates after propagation of the initial conditions up to time t, and

$$\overline{\dot{\boldsymbol{\mu}}}(\boldsymbol{Q}) = \frac{1}{P}\sum_{j=1}^{P}\dot{\boldsymbol{\mu}}(\boldsymbol{Q}^{(j)}). \tag{2.18}$$

RPMD works very well for a wide range of applications, including calculating chemical reaction rate constants[22] and diffusion coefficients,[19] but it is less suitable for vibrational spectra. This is because the ring polymer has internal vibrational modes which arise from the harmonic springs in the ring polymer Hamiltonian. These can be found by diagonalising the Hamiltonian in the absence of an external molecular potential, which gives

$$\mathcal{H}_{\mathrm{free},P}(\tilde{\boldsymbol{P}},\tilde{\boldsymbol{Q}}) = \sum_{k=0}^{P-1}\sum_{i=1}^{N}\left[\frac{|\tilde{\boldsymbol{P}}_i^{(k)}|^2}{2m_i} + \frac{1}{2}m_i\omega_k^2|\tilde{\boldsymbol{Q}}_i^{(k)}|^2\right], \tag{2.19}$$

with $\omega_k = 2\omega_P \sin\frac{\pi k}{P}$ and $\tilde{\boldsymbol{P}}$, $\tilde{\boldsymbol{Q}}$ are the normal mode momenta and coordinates of the free ring polymer. The transformation matrix to the normal modes is given in Appendix A. These normal modes are not physical vibrations, but rather arise from the mathematical construct of the ring polymer. However, they can become resonant with the physical modes of the system, which gives rise to spurious peaks in the vibrational spectrum.[30,32]

2.2.2 Thermostatted RPMD

TRPMD was developed to remove the spurious resonances of RPMD.[20] A Langevin thermostat is attached to the internal ring polymer modes, such that the equation of motion

becomes[20]

$$\dot{\boldsymbol{P}} = -\nabla_{\boldsymbol{Q}} V_{\mathrm{cl}}(\boldsymbol{Q}) - \Gamma \boldsymbol{P} + \sqrt{\frac{2m\Gamma}{\beta_P}}\boldsymbol{\xi}(t), \tag{2.20}$$

where Γ is a real, symmetric, positive semi-definite friction matrix and $\boldsymbol{\xi}(t)$ is a vector with independent, identically distributed entries that are sampled from a normal distribution. (Note that Equation (2.20) assumes that all masses are equal to m, but TRPMD can easily be generalised to different masses.) The thermostat does not act on the centroid by design, because we require that

$$\Gamma \boldsymbol{e} = \boldsymbol{e}^T \Gamma = 0, \tag{2.21}$$

where $\boldsymbol{e}$ is the vector with all entries equal to 1. $\boldsymbol{e}$ is also an eigenvector of $\Gamma^{\frac{1}{2}}$, and Equation (2.20) thus reduces to the thermostat-free version for the centroid. The centroid only indirectly feels the thermostat through its coupling to the internal modes via the potential V_{cl}. See Reference [20] for a detailed discussion of the friction matrix entries, which are chosen so as to optimally dampen the ring polymer normal modes. Rossi *et al.* tested TRPMD on liquid water and showed that the thermostat washes out the RPMD resonances.[20] However, they also noted that the thermostat can artificially broaden peaks.[20] While the TRPMD peak positions are accurate, and there is no sign of spurious resonances, TRPMD line widths generally cannot be trusted.[20]

2.2.3 Centroid molecular dynamics

The ring polymer centroid is the average configuration of the ring polymer beads,

$$\overline{\boldsymbol{Q}} = \frac{1}{P}\sum_{j=1}^{P} \boldsymbol{Q}^{(j)}. \tag{2.22}$$

The CMD dipole derivative autocorrelation function is

$$C^{\mathrm{CMD}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{(2\pi\hbar)^{3N} Z_{\mathrm{c}}} \int \mathrm{d}\overline{\boldsymbol{P}} \int \mathrm{d}\overline{\boldsymbol{Q}} e^{-\beta \mathcal{H}_{\mathrm{c}}(\overline{\boldsymbol{P}},\overline{\boldsymbol{Q}})} \dot{\boldsymbol{\mu}}(\overline{\boldsymbol{Q}}_0) \cdot \dot{\boldsymbol{\mu}}(\overline{\boldsymbol{Q}}_t), \tag{2.23}$$

with the centroid Hamiltonian

$$\mathcal{H}_{\mathrm{c}}(\overline{\boldsymbol{P}},\overline{\boldsymbol{Q}}) = \sum_{i=1}^{N} \frac{|\overline{\boldsymbol{P}}_i|^2}{2m_i} + V_{\mathrm{c}}(\overline{\boldsymbol{Q}}), \tag{2.24}$$

the centroid partition function $Z_{\mathrm{c}} = \frac{1}{(2\pi\hbar)^{3N}} \int \mathrm{d}\overline{\boldsymbol{P}} \int \mathrm{d}\overline{\boldsymbol{Q}} e^{-\beta\mathcal{H}_{\mathrm{c}}(\overline{\boldsymbol{P}},\overline{\boldsymbol{Q}})}$, and the centroid potential of mean force $V_c(\overline{\boldsymbol{Q}})$. Equation (2.23) is a classical correlation function, but it is fundamentally different from the classical correlation function in Equation (2.2) because the centroid evolves on the potential of mean force V_c. One way to define this potential of mean force is implicitly, by requiring that the resulting force on the centroid,

$$F(\overline{\boldsymbol{Q}}) = -\frac{\partial V_{\mathrm{c}}}{\partial \overline{\boldsymbol{Q}}} \tag{2.25}$$

is the same as the constrained average over the ring polymer,

$$F(\overline{\boldsymbol{Q}}) = -\frac{1}{\mathcal{N}} \int \mathrm{d}\boldsymbol{Q}' e^{-\beta W(\boldsymbol{Q}')} \delta(\overline{\boldsymbol{Q}}' - \overline{\boldsymbol{Q}}) \left(\frac{\partial W(\boldsymbol{Q}')}{\partial \overline{\boldsymbol{Q}}} \right) = -\left\langle \frac{\partial V_P}{\partial \overline{\boldsymbol{Q}}} \right\rangle_{\overline{\boldsymbol{Q}}}, \tag{2.26}$$

where

$$W(\boldsymbol{Q}) = V_P(\boldsymbol{Q}) + S_P(\boldsymbol{Q}) = \frac{1}{P} \sum_{j=1}^{P} V_{\mathrm{cl}}(\boldsymbol{Q}^{(j)}) + \frac{1}{2P} \sum_{j=1}^{P} \sum_{i=1}^{N} m_i \omega_P^2 |\boldsymbol{Q}_i^{(j)} - \boldsymbol{Q}_i^{(j+1)}|^2, \tag{2.27}$$

$$\mathcal{N} = \int \mathrm{d}\boldsymbol{Q}' e^{-\beta W(\boldsymbol{Q}')} \delta(\overline{\boldsymbol{Q}}' - \overline{\boldsymbol{Q}}), \tag{2.28}$$

and $\overline{\boldsymbol{Q}}'$ is the centroid of the integration variables $\boldsymbol{Q}'$. The centroid is also proportional to the $k = 0$ internal normal mode of the ring polymer (see Appendix A). The constrained average $\langle\cdot\rangle_{\overline{\boldsymbol{Q}}}$ can thus also be understood as an average over the $P-1$ remaining normal modes with $k > 0$. In particular, the spring potential S_P is independent of the centroid (which can be seen from Equation (2.19) and noting that $\omega_0 = 0$), which validates the final equality in Equation (2.26).

CMD algorithms can be separated into two groups based on whether they employ the force in Equation (2.25) (force-matching algorithms[81]) or Equation (2.26) (adiabatic CMD[82,83]). In force-matching, a data set tabulating centroid configurations against the mean field force

is compiled by running constrained PIMD simulations for each centroid configuration to calculate $\left\langle \frac{\partial V_P}{\partial \overline{\boldsymbol{Q}}} \right\rangle_{\overline{\boldsymbol{Q}}}$. The data set is then used to fit an effective potential V_c. Usually, it is assumed to be of the specific form

$$V_{\mathrm{c}}(\overline{\boldsymbol{Q}}) = V_{\mathrm{cl}}(\overline{\boldsymbol{Q}}) + \Delta V_{\mathrm{c}}(\overline{\boldsymbol{Q}}), \tag{2.29}$$

where V_{cl} is the underlying classical potential energy surface. Force-matching algorithms then only need to find the (hopefully small) correction ΔV_{c}. In the early implementations, this was done by least squares fitting on a radial grid,[81] but these days machine learning algorithms for potential energy surfaces can be repurposed to fit the correction.[84,85]

Adiabatic CMD computes the centroid forces on the fly by adiabatically separating the normal modes with $k > 0$ from the centroid normal mode. To achieve this, the masses associated with the non-centroid modes are scaled,

$$m_{i,k} = \begin{cases} m_i & k = 0 \\ \gamma^2 m_i & k > 0 \end{cases}, \tag{2.30}$$

where $\gamma > 0$ is a small number. The frequencies are shifted to new frequencies, $\omega'_{i,k} = \omega_k \sqrt{\frac{m_i}{m_{i,k}}} = \omega_k/\gamma$. As γ decreases, the frequencies increase for $k \neq 0$ and the internal modes are shifted away from the centroid mode and move on a much faster time scale.[86] The centroid now evolves on the mean field potential averaged over the internal modes.[82] Often, we choose to shift all internal modes to the same frequency Ω, which lies above the highest frequency of interest in our system,[32,83] with k-dependent scaling factors

$$\gamma_k = \frac{\omega_k}{\Omega}. \tag{2.31}$$

Further, a thermostat is attached to the internal modes to account for the canonical ensemble average of the force in Equation (2.26). In order to numerically integrate an oscillation with frequency Ω, the time step of the integration δt needs to be much smaller than the period of

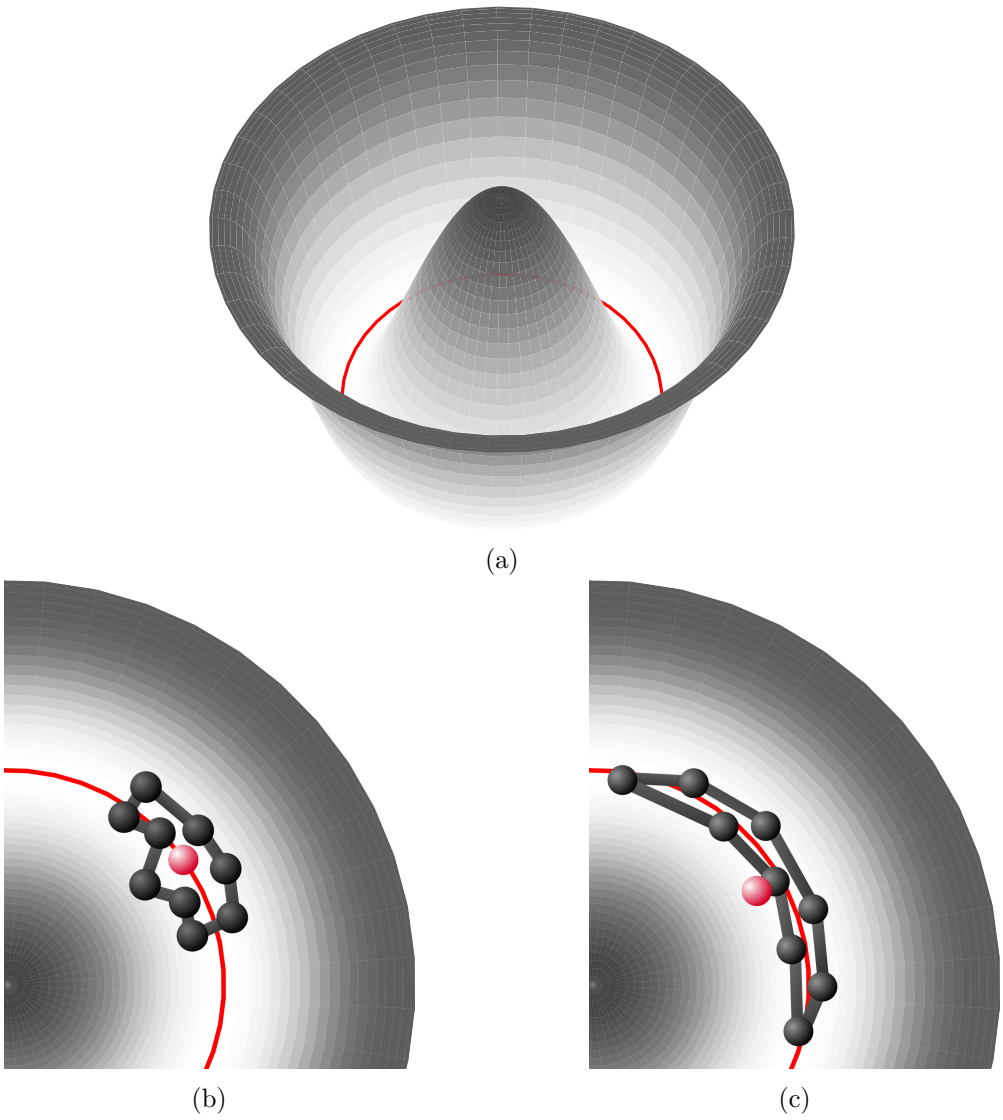

(a)

(b) (c)

Figure 2.2: An explanation of the origin of the curvature problem. (a) A potential with a minimum (highlighted in red) that is curved around the origin. (b) Example ring polymer configuration and its centroid at a high temperature. (c) As (b), but at a lower temperature. Note that this figure serves purely as a qualitative explanation of the curvature problem.

the oscillation,

$$\delta t \ll \frac{2\pi}{\Omega}. \tag{2.32}$$

We therefore need to choose Ω large enough to separate the internal modes from the system frequencies we are interested in, but not too high in order to avoid computationally expensive simulations.

Regardless of the implementation, CMD vibrational spectra are not disturbed by spurious resonances such as the ones observed in RPMD. However, they suffer from the so-called curvature problem in systems with radial and angular (bending or librational) normal modes.[30,31,73] Figure 2.2 shows a schematic for such a potential energy surface and examples of ring polymer and centroid configurations at high and low temperatures. At high temperatures, the ring polymer beads are spread around the potential minimum. As a result, so is the centroid, and it is a good representation of the ring polymer (Figure 2.2b). The situation looks very different at low temperatures (Figure 2.2c). Here, the ring polymer beads tend to lie close to the potential minimum, but the centroid lies closer to the origin at a higher potential energy. This is the origin of the curvature problem: The centroid can visit unphysical (high energy) configurations, such as the one shown in Figure 2.2c, while none of the beads are at unphysical configurations. References [33] and [34] discuss in detail how the ring polymer beads spread out along the angular coordinates and give rise to an artificial instanton. As a result of the ring polymer spreading out, the centroid potential of mean force softens at the centroid configuration. This softening becomes increasingly pronounced at low temperatures and leads to a spurious red shift in the CMD spectrum. The red shift is more pronounced at low temperatures, but it is already present in liquid water at 300 K and thus poses a serious issue.[30,31,73]

For a while, there was therefore no artefact-free path integral method for the calculation of vibrational spectra. Recently, QCMD was proposed as a novel alternative. For a number of problems that the path integral methods discussed above have struggled with, QCMD has proven very well suited. We will next explain the theory behind QCMD, and some approximations that have been introduced to find more efficient implementations.

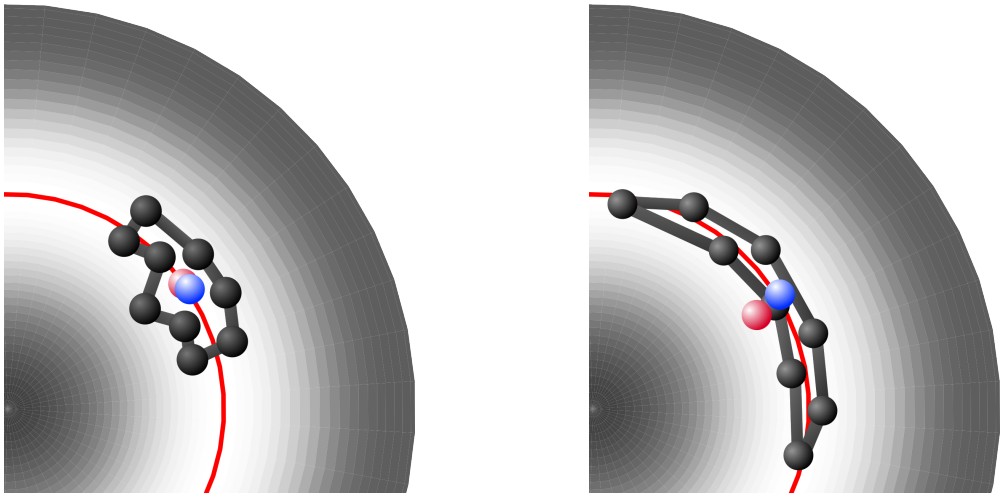

(a) (b)

Figure 2.3: Centroid (red) and quasi-centroid (blue) configurations at (a) high and (b) low temperatures for the same potential as in Figure 2.2.

2.2.4 Quasi-centroid molecular dynamics

QCMD was proposed as an alternative to CMD to address the curvature problem. The quasi-centroid is defined in terms of curvilinear coordinates, which may be tailored to any curved potential features. For the two-dimensional potential shown in Figure 2.2, it is natural to define the curvilinear coordinates in terms of polar coordinates. Figure 2.3 shows the quasi-centroid and centroid configurations for ring polymer configurations at high and low temperatures. Clearly, the quasi-centroid is a better representation of the ring polymer configuration. In the low-temperature case, where the centroid may visit unphysical configurations, the quasi-centroid remains at a configuration that is a good representation of the beads. From this observation, it can be expected that QCMD avoids the CMD curvature problem.

The curvilinear coordinates are tailored to the system in question. We therefore cannot define them in general, beyond saying that they usually contain bond lengths, bond angles and dihedral angles. Let us thus write $\boldsymbol{\xi}$ for the set of curvilinear quasi-centroid coordinates. Then, let $\tilde{\boldsymbol{Q}}(\boldsymbol{\xi})$ be the Cartesian quasi-centroid configuration which corresponds to the curvilinear

degrees of freedom $\boldsymbol{\xi}$. In QCMD, correlation functions are calculated via

$$C^{\mathrm{QCMD}}_{\dot{\boldsymbol{\mu}}\dot{\boldsymbol{\mu}}}(t) = \frac{1}{(2\pi\hbar)^{3N} Z_{\mathrm{qc}}} \int \mathrm{d}\tilde{\boldsymbol{P}} \int \mathrm{d}\tilde{\boldsymbol{Q}} e^{-\beta \mathcal{H}_{\mathrm{qc}}(\tilde{\boldsymbol{P}},\tilde{\boldsymbol{Q}})} \dot{\boldsymbol{\mu}}(\tilde{\boldsymbol{Q}}_0) \cdot \dot{\boldsymbol{\mu}}(\tilde{\boldsymbol{Q}}_t), \tag{2.33}$$

with

$$\mathcal{H}_{\mathrm{qc}}(\tilde{\boldsymbol{P}},\tilde{\boldsymbol{Q}}) = \sum_{i=1}^{N} \frac{|\tilde{\boldsymbol{P}}_i|^2}{2m_i} + V_{\mathrm{qc}}(\tilde{\boldsymbol{Q}}) \tag{2.34}$$

and the quasi-centroid partition function $Z_{\mathrm{qc}} = \frac{1}{(2\pi\hbar)^{3N}} \int \mathrm{d}\tilde{\boldsymbol{P}} \int \mathrm{d}\tilde{\boldsymbol{Q}} e^{-\beta \mathcal{H}_{\mathrm{qc}}(\tilde{\boldsymbol{P}},\tilde{\boldsymbol{Q}})}$. In analogy to CMD, we write the force on the quasi-centroid as

$$F(\tilde{\boldsymbol{Q}}) \simeq -\left\langle \frac{\partial V_P}{\partial \tilde{\boldsymbol{Q}}} \right\rangle_{\boldsymbol{\xi}}, \tag{2.35}$$

where V_P is defined through Equation (2.27), the constrained average is

$$\langle \cdot \rangle_{\boldsymbol{\xi}} = \frac{1}{\mathcal{N}} \int \mathrm{d}\boldsymbol{Q}' e^{-\beta W(\boldsymbol{Q}')} \, (\cdot) \, \delta(\boldsymbol{\xi}' - \boldsymbol{\xi}), \tag{2.36}$$

$\boldsymbol{\xi}'$ are the quasi-centroid coordinates averaged over the ring polymer beads $\boldsymbol{Q}'$, W is defined in Equation (2.27), and

$$\mathcal{N} = \int \mathrm{d}\boldsymbol{Q}' e^{-\beta W(\boldsymbol{Q}')} \delta(\boldsymbol{\xi}' - \boldsymbol{\xi}). \tag{2.37}$$

Equation (2.35) is an approximation, although it looks almost the same as its CMD equivalent Equation (2.26) (which is exact). This is because the spring potential may not be independent of the quasi-centroid.

To assess the accuracy of this approximation and QCMD in general, Trenins *et al.* developed an adiabatic QCMD algorithm and applied it to gas-phase and liquid water, and hexagonal ice.[33,87] In a subsequent study, they also applied it to gas-phase ammonia.[34] Adiabatic QCMD is similar to adiabatic CMD, in the sense that the quasi-centroid is adiabatically separated from the internal modes and propagated on the mean field averaged potential. The QCMD spectra are consistent with exact quantum results (where it was possible to do a full quantum simulation) or embedding methods such as the local monomer approach.[73,88,89] This remains true in cases where the CMD curvature problem is severe, with the exception of the

bending vibration of gas-phase ammonia along the umbrella inversion mode. The reason for this is discussed in detail by Haggard *et al.*[34] They conclude that QCMD is not expected to do well for vibrations that are not well approximated classically to begin with, or that are highly anharmonic. Nevertheless, it is promising in many systems. Further, adiabatic QCMD relies on the existence of a function $\tilde{\boldsymbol{Q}}(\boldsymbol{\xi})$, which assumes that we can always find a Cartesian representation for a given set of quasi-centroid coordinates. However, it is possible for this problem to be overdetermined, in which case a solution $\tilde{\boldsymbol{Q}}$ may not exist for all $\boldsymbol{\xi}$. It is not yet clear how to run adiabatic QCMD in this case.

Another, more pressing issue with adiabatic QCMD is that the adiabatic separation parameter γ has to be very small to ensure complete adiabatic separation of the quasi-centroid from the internal modes. This means that adiabatic QCMD is 10-100 times more expensive than RPMD, TRPMD or (partially adiabatic) CMD.[34]

A fast implementation

Fletcher *et al.* developed f-QCMD as an efficient alternative to adiabatic QCMD.[35] They precompute the quasi-centroid potential of mean force V_{qc}, such that only a single classical MD simulation is needed to compute the correlation function in Equation (2.33). We will describe the algorithm based on its application to water, and touch on extensions to other molecules below. In gas-phase water, there are three curvilinear quasi-centroid coordinates: The O-H bond lengths r_1 and r_2, and the H-O-H bond angle, θ_{12}. For a ring polymer configuration $\boldsymbol{Q}$, the quasi-centroid values of these are defined through

$$\overline{r}_i(\boldsymbol{Q}) = \frac{1}{P}\sum_{j=1}^{P} r_i(\boldsymbol{Q}^{(j)}), \tag{2.38a}$$

$$\overline{\theta}_{12}(\boldsymbol{Q}) = \frac{1}{P}\sum_{j=1}^{P} \theta_{12}(\boldsymbol{Q}^{(j)}). \tag{2.38b}$$

We can also define their distribution functions,

$$\rho_r^{\mathrm{exact}}(r_i) = \frac{1}{4\pi r_i^2}\left\langle \delta\left(r_i - \overline{r}_i(\boldsymbol{Q})\right)\right\rangle_{\mathrm{NVT},P}, \tag{2.39a}$$

$$\rho_\theta^{\text{exact}}(\theta_{12}) = \frac{1}{\sin\theta_{12}} \left\langle \delta\left(\theta_{12} - \overline{\theta}_{12}(\boldsymbol{Q})\right)\right\rangle_{\text{NVT},P}, \tag{2.39b}$$

which are easy to calculate in a PIMD simulation. To find V_{qc}, we first write it as

$$V_{\text{qc}}(\tilde{\boldsymbol{Q}}) = V_{\text{cl}}(\tilde{\boldsymbol{Q}}) + \Delta V_{\text{qc}}(\tilde{\boldsymbol{Q}}), \tag{2.40}$$

where V_{cl} is the classical potential energy surface (for example the Partridge-Schwenke potential energy surface for gas-phase water[90]). Note the similarity to Equation (2.29) in the force matching approach to CMD. The key assumption in f-QCMD is that ΔV_{qc} is separable in the curvilinear coordinates,

$$\Delta V_{\text{qc}}(\tilde{\boldsymbol{Q}}) \approx \Delta V_r(r_1(\tilde{\boldsymbol{Q}})) + \Delta V_r(r_2(\tilde{\boldsymbol{Q}})) + \Delta V_\theta(\theta_{12}(\tilde{\boldsymbol{Q}})), \tag{2.41}$$

To find the quasi-centroid corrections, we use the distribution functions of the quasi-centroid coordinates (Equation (2.39)). These distribution functions will be equivalent to the classical averages under V_{qc},

$$\rho_r(r_i) = \frac{1}{4\pi r_i^2}\langle\delta(r_i - \overline{r}_i(\boldsymbol{Q}))\rangle_{\text{NVT},1} \tag{2.42a}$$

$$\rho_\theta(\theta_{12}) = \frac{1}{\sin\theta_{12}}\langle\delta(\theta_{12} - \overline{\theta}_{12}(\boldsymbol{Q}))\rangle_{\text{NVT},1}, \tag{2.42b}$$

provided we have found the correct V_{qc}. The iterative Boltzmann inversion (IBI) finds the one-dimensional corrections ΔV_r and ΔV_θ as follows.[91] First, we run a PIMD simulation to obtain the distribution functions ρ_r^{exact} and $\rho_\theta^{\text{exact}}$ according to Equation (2.39). This only needs to be done once. We then initialise the corrections, $\Delta V_r^{(0)} = 0$ and $\Delta V_\theta^{(0)} = 0$, such that the initial guess for the quasi-centroid effective potential is just the classical potential itself. For $n \geq 0$, the update to the corrections is

$$\Delta V_r^{(n+1)}(r_i) = \Delta V_r^{(n)}(r_i) - \frac{1}{\beta}\log\frac{\rho_r^{\text{exact}}(r_i)}{\rho_r^{(n)}(r_i)}, \tag{2.43a}$$

$$\Delta V_\theta^{(n+1)}(\theta_{12}) = \Delta V_\theta^{(n)}(\theta_{12}) - \frac{1}{\beta}\log\frac{\rho_\theta^{\text{exact}}(\theta_{12})}{\rho_\theta^{(n)}(\theta_{12})}, \tag{2.43b}$$

until convergence, where $\rho_r^{(n)}$ and $\rho_\theta^{(n)}$ are the distribution functions resulting from the corrected potential in the nth iteration, $V_{\mathrm{qc}}^{(n)} = V_{\mathrm{cl}} + \Delta V_{\mathrm{qc}}^{(n)}$.

The extension of f-QCMD beyond water molecules is completely straightforward, as long as a suitable set of curvilinear coordinates can be found. Fletcher *et al.* applied f-QCMD to gas-phase water, ammonia and methane.[35] In these calculations, the curvilinear coordinates comprise the bond lengths between the central atom (O, N or C) and the hydrogens, as well as the angles between the hydrogens. Fletcher then also extended the formalism to include dihedral angles, which allowed the application to hydrogen peroxide.[36] For gas-phase water and ammonia, the f-QCMD spectra agree to graphical accuracy with Trenins *et al.*'s[33] full QCMD spectra. This implies that the approximation in Equation (2.41) is justified, and that f-QCMD is a good approximation for QCMD in the gas phase. But it had not been tested in the condensed phase.

2.3 f-QCMD for condensed phase systems

To generalise f-QCMD to the condensed phase, we assume that the quasi-centroid potential of mean force can be separated into an intra- and intermolecular part, such that Equation (2.40) is extended to

$$V_{\mathrm{qc}}(\boldsymbol{r}) \approx V_{\mathrm{cl}}(\boldsymbol{r}) + \Delta V_{\mathrm{intra}}(\boldsymbol{r}) + \Delta V_{\mathrm{inter}}(\boldsymbol{r}). \tag{2.44}$$

We will describe the algorithm for liquid water and ice, but it is straightforward to extend to other materials. The intramolecular correction $\Delta V_{\mathrm{intra}}$ contains contributions from the intramolecular O-H distances and H-O-H bond angles, allowing us to adopt the procedure described for gas-phase water in Section 2.2.4. We thus only need to specify how to find $\Delta V_{\mathrm{inter}}$. To this end, we focus on the pairwise O-O, O-H and H-H distances, with classical distribution functions

$$g_{\mathrm{OO}}(r) = \frac{1}{4\pi r^2 N\rho} \left\langle \sum_{j>i}^{N} \delta(r - |\boldsymbol{r}_{\mathrm{O},i} - \boldsymbol{r}_{\mathrm{O},j}|) \right\rangle_{\mathrm{NVT},1}, \tag{2.45a}$$

$$g_{\mathrm{OH}}(r) = \frac{1}{4\pi r^2 N\rho}\left\langle \sum_{j>i}^{N} \delta(r - |\boldsymbol{r}_{\mathrm{O},i} - \boldsymbol{r}_{\mathrm{H},j}|)\right\rangle_{\mathrm{NVT},1}, \tag{2.45b}$$

$$g_{\mathrm{HH}}(r) = \frac{1}{4\pi r^2 N\rho}\left\langle \sum_{j>i}^{N} \delta(r - |\boldsymbol{r}_{\mathrm{H},i} - \boldsymbol{r}_{\mathrm{H},j}|)\right\rangle_{\mathrm{NVT},1}, \tag{2.45c}$$

where $\rho = N/V$ is the bulk number density of water molecules, $\boldsymbol{r}_{\mathrm{O},i}$ and $\boldsymbol{r}_{\mathrm{H},i}$ are the coordinates of the oxygen and hydrogen atoms in the ith water molecule, respectively. Equation (2.45) are the classical intermolecular distribution functions, but we can also obtain the exact quantum distribution functions needed for the IBI iterations by using PIMD estimators for the pairwise distances. These are the most expensive part of the f-QCMD algorithm. The distribution functions are usually calculated as histograms, and many samples are needed to obtain smooth distributions which give suitable potential corrections. We can, however, take advantage of recently developed low variance force estimators to obtain the distribution functions with smaller statistical errors.[92,93]

Having specified the quasi-centroid coordinates, we are ready to describe the quasi-centroid potential of mean force in more detail. The intramolecular correction is

$$\Delta V_{\mathrm{intra}}(\boldsymbol{r}) = \sum_{i=1}^{N} \left[\Delta V_r(r_{\mathrm{OH}_1,i}) + \Delta V_r(r_{\mathrm{OH}_2,i}) + \Delta V_\theta(\theta_i)\right], \tag{2.46}$$

where $r_{\mathrm{OH}_1,i}$, $r_{\mathrm{OH}_2,i}$ and θ_i are the intramolecular quasi-centroid coordinates of the ith water molecule, and the intermolecular correction is

$$\Delta V_{\mathrm{inter}}(\boldsymbol{r}) = \sum_{j>i}^{N} \sum_{\mathrm{X}\in i} \sum_{\mathrm{Y}\in j} \Delta V_{\mathrm{XY}}(|\boldsymbol{r}_{\mathrm{X},i} - \boldsymbol{r}_{\mathrm{Y},j}|). \tag{2.47}$$

Here, X and Y loop over all possible O-O, O-H and H-H pairs in the molecules i, j. Both $\Delta V_{\mathrm{intra}}$ and $\Delta V_{\mathrm{inter}}$ can be found via the IBI.[91,94] For $\Delta V_{\mathrm{inter}}$, we set

$$\Delta V_{\mathrm{XY}}^{(n+1)}(r) = \Delta V_{\mathrm{XY}}^{(n)}(r) - \frac{1}{\beta}\log\frac{g_{\mathrm{XY}}^{\mathrm{exact}}(r)}{g_{\mathrm{XY}}^{(n)}(r)} \tag{2.48}$$

in the nth iteration (starting from $\Delta V_{\mathrm{XY}}^{(0)} = 0$). $g_{\mathrm{XY}}^{(n)}(r)$ is the pairwise distribution function obtained from a classical evolution with the current correction $\Delta V_{\mathrm{qc}}^{(n)}$. The intramolecular

corrections $\Delta V_{\text{intra}}^{(n)}$ are found with the update rules in Equations (2.43).

There are heuristic arguments that the IBI always converges,[91] and progress has been made towards proving so.[95] However, it has been previously observed[94] that close to convergence, the IBI iterations start to oscillate, and can even become unstable. We observed the same behaviour in our iterations. To avoid convergence issues, the IBI update is sometimes scaled by a prefactor outside the logarithm in Equation (2.43) and (2.48).[96] However, we found that a much better approach is to regularise the IBI update rule as

$$\Delta V_{\alpha}^{(n+1)}(r) = \Delta V_{\alpha}^{(n)}(r) - \frac{1}{\beta} \log \frac{g_{\alpha}^{\text{exact}}(r) + \varepsilon g_{\alpha}^{\infty}}{g_{\alpha}^{(n)}(r) + \varepsilon g_{\alpha}^{\infty}}, \tag{2.49}$$

where α refers to any of the inter- or intramolecular quasi-centroid degrees of freedom, ε is a small number, and $g_{\alpha}^{\infty} = \max_r \{g_{\alpha}^{\text{exact}}(r), g_{\alpha}^{(n)}(r)\}$ ensures that ε remains dimensionless. The regularisation of the IBI is shown in Figure 2.4 for the O-H pairwise distance distributions. With increasing ε, the magnitude of the correction is damped, which should mitigate the oscillating behaviour close to convergence. At the same time, any discontinuities in ΔV_{OH} that arise from statistical noise in the low sampling region of the O-H distribution functions are smoothed out. This allows us to differentiate the potential corrections to evaluate the nuclear force. The regularisation in Equation (2.49) is by no means a new idea, having already been suggested in the paper introducing IBI by Soper,[91] but it has not been widely used.

2.4 Results for water and ice

For liquid water and ice, we choose the q-TIP4P/F potential as the underlying classical potential V_{cl}.[29] The intermolecular part of the q-TIP4P/F potential,

$$V_{\text{inter}}(\boldsymbol{r}) = \sum_{j>i}^{N} \left[4\varepsilon \left(\left(\frac{\sigma}{|\boldsymbol{r}_{\text{O},i} - \boldsymbol{r}_{\text{O},j}|} \right)^{12} - \left(\frac{\sigma}{|\boldsymbol{r}_{\text{O},i} - \boldsymbol{r}_{\text{O},j}|} \right)^{6} \right) + \sum_{m \in i} \sum_{n \in j} \frac{q_m q_n}{r_{mn}} \right], \tag{2.50}$$

is composed of Lennard-Jones interactions between the oxygen atoms and electrostatic interactions of auxiliary partial charges placed on the water molecules. The Lennard-Jones parameters ε and σ set the energy and length scales of the intermolecular interactions, re-

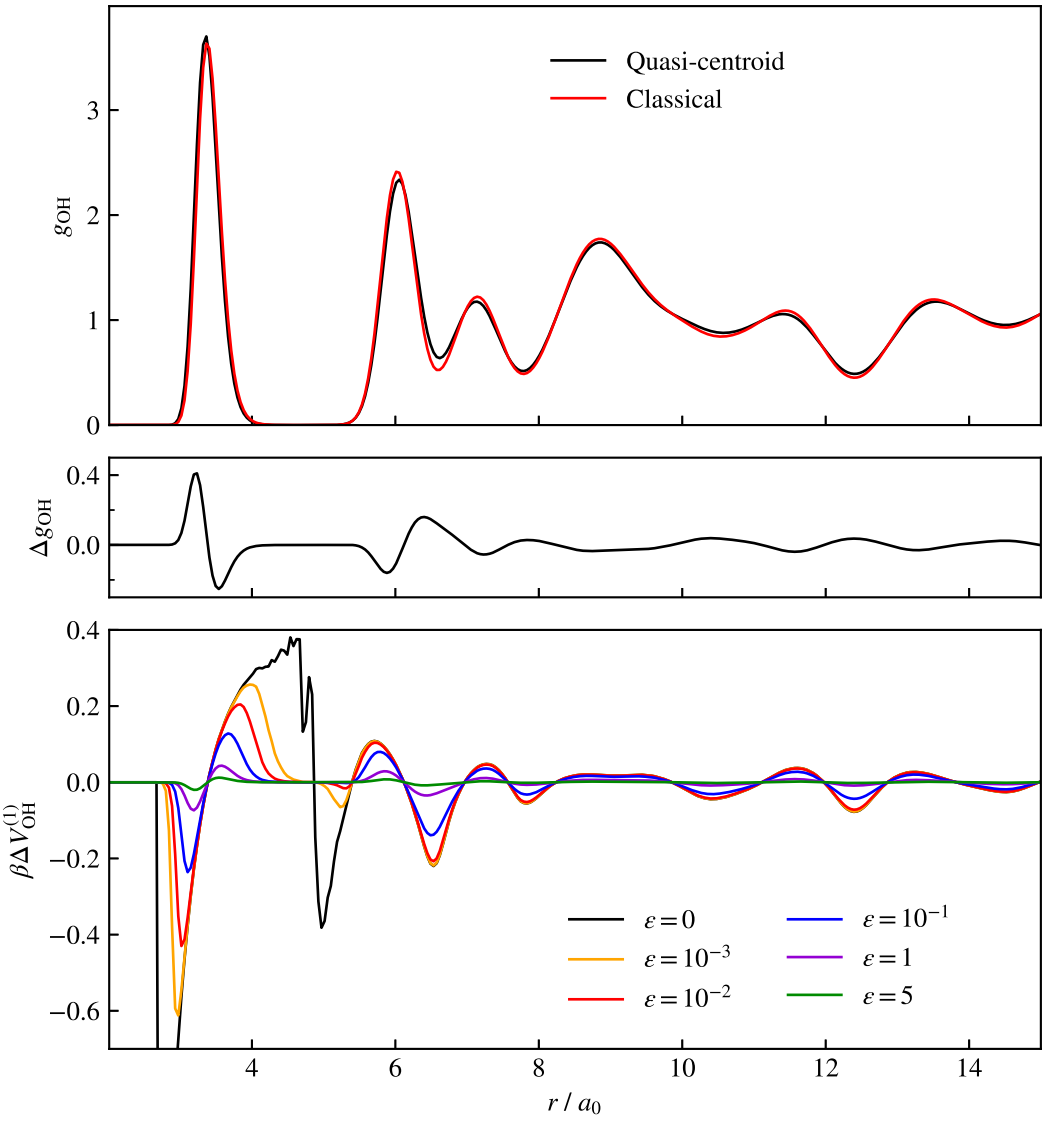

Figure 2.4: Regularisation procedure for the first iteration of the IBI update to the intermolecular correction ΔV_{OH}. Starting from the classical and the exact quasi-centroid distributions (top panel), the correction to the potential is shown in the bottom panel at increasing regularisation strengths ε. The middle panel shows the difference between the two distribution functions. a_0 is the Bohr radius.

spectively. The indices n and m enumerate the partial charges q_m, q_n that are placed on the ith and jth water molecules, respectively, separated by the pairwise distance r_{mn}. To find the intermolecular quasi-centroid potential of mean force, we use the regularised IBI procedure as outlined in Section 2.3.

For the intramolecular part we employ an additional trick. The q-TIP4P/F intramolecular potential for a given molecule includes a fourth order polynomial in the O-H bond length,

$$V_r(r) = k_2(r - r_{\mathrm{eq}})^2 + k_3(r - r_{\mathrm{eq}})^3 + k_4(r - r_{\mathrm{eq}})^4, \tag{2.51}$$

and a quadratic polynomial in the H-O-H bond angle,

$$V_\theta(\theta) = k_\theta(\theta - \theta_{\mathrm{eq}})^2, \tag{2.52}$$

with the equilibrium O-H bond length r_{eq}, the equilibrium H-O-H angle θ_{eq}, and constants k_2, k_3, k_4, and k_θ. The full intramolecular potential is

$$V_{\mathrm{intra}}(\boldsymbol{r}) = \sum_{i=1}^{N} \left[V_r(r_{\mathrm{OH}_1,i}) + V_r(r_{\mathrm{OH}_2,i}) + V_\theta(\theta_i) \right]. \tag{2.53}$$

Although we could use the regularised IBI routine as described in Section 2.3, we can simplify our approach further by assuming that the intramolecular quasi-centroid potential of mean force has the same form as Equations (2.51)-(2.53). After each IBI update we then simply refit the polynomial coefficients to $V_r + \Delta V_r^{(n)}$ and $V_\theta + \Delta V_\theta^{(n)}$, where $\Delta V_r^{(n)}$ and $\Delta V_\theta^{(n)}$ are obtained through Equation (2.43). The fitting ensures that the resulting quasi-centroid potential is smooth, so we get reasonable MD forces throughout the simulation.

Our simulation contains 216 water molecules in a cubic box with side length 35.24 a_0 for water, and 96 water molecules in an orthorhombic box with side lengths 25.62, 29.58 and 27.89 a_0 for ice, where a_0 is the Bohr radius. We used $P = 32$ beads at 300 K and $P = 64$ beads at 150 K for the PIMD simulation. During the IBI, we used $\varepsilon = 1$ and $\varepsilon = 5$ at 300 and 150 K, respectively. For both temperatures, we iterated the regularised IBI 30 times, after which the distribution functions obtained with the effective potential were in

perfect agreement with the exact PIMD distribution functions. The converged quasi-centroid distribution functions for ice are shown in Figure 2.5 to underline this point.

This is a good time to discuss the uniqueness of the solution we find with our IBI iterations. The IBI is guaranteed to find the correct potential of mean force if the target potential itself is a pairwise potential.[97] However, if three and higher body-order interactions are present, then uniqueness is not guaranteed. For example, Evans has argued that in this case, higher order distribution functions would be necessary to unambiguously determine the effective potential.[98] It is unclear what the shape of the correction ΔV_{qc} we are targeting is, but it is easily conceivable that it contains terms beyond pairwise ones. Indeed, we observe small differences between the final effective potentials found by different IBI runs with different regularisation parameters ε. However, these differences did not have an observable effect on the spectrum.

Figure 2.6 compares the f-QCMD spectrum of liquid water with the classical spectrum, the adiabatic QCMD spectrum, and the spectrum calculated with a recently developed path integral coarse grained method (Te PIGS).[84] In Te PIGS, the CMD potential of mean force V_{c} is precomputed at a higher temperature, and then used at all lower temperatures.[84] (For example, in liquid water, the potential is precomputed at 600 K.[84]) The adiabatic QCMD results were calculated with an improved torque estimator[87] and provided to us by George Trenins.[33] The Te PIGS results were provided by Venkat Kapil.[84]

We focus on the fundamental region of the spectrum, which is composed of a libration band at around 600 cm^{-1}, a bending band at around 1600 cm^{-1}, and a stretch band at around 3500 cm^{-1}. At 300 K, our f-QCMD results are nearly on top of the adiabatic QCMD results, except for a small red shift in the stretch band and small blue shifts in the libration and bend bands. All shifts are on the order of tens of cm^{-1}. For liquid water, our separable assumption in Equations (2.44), (2.46) and (2.47) is thus adequate, and the nuclear quantum effects are captured for the most part by pairwise interactions.

The extent to which nuclear quantum effects are relevant can be determined by looking at the red shifts in the f-QCMD spectrum with respect to the classical spectrum. These are most pronounced for the high frequency stretch band. The red shifts decrease in the bending

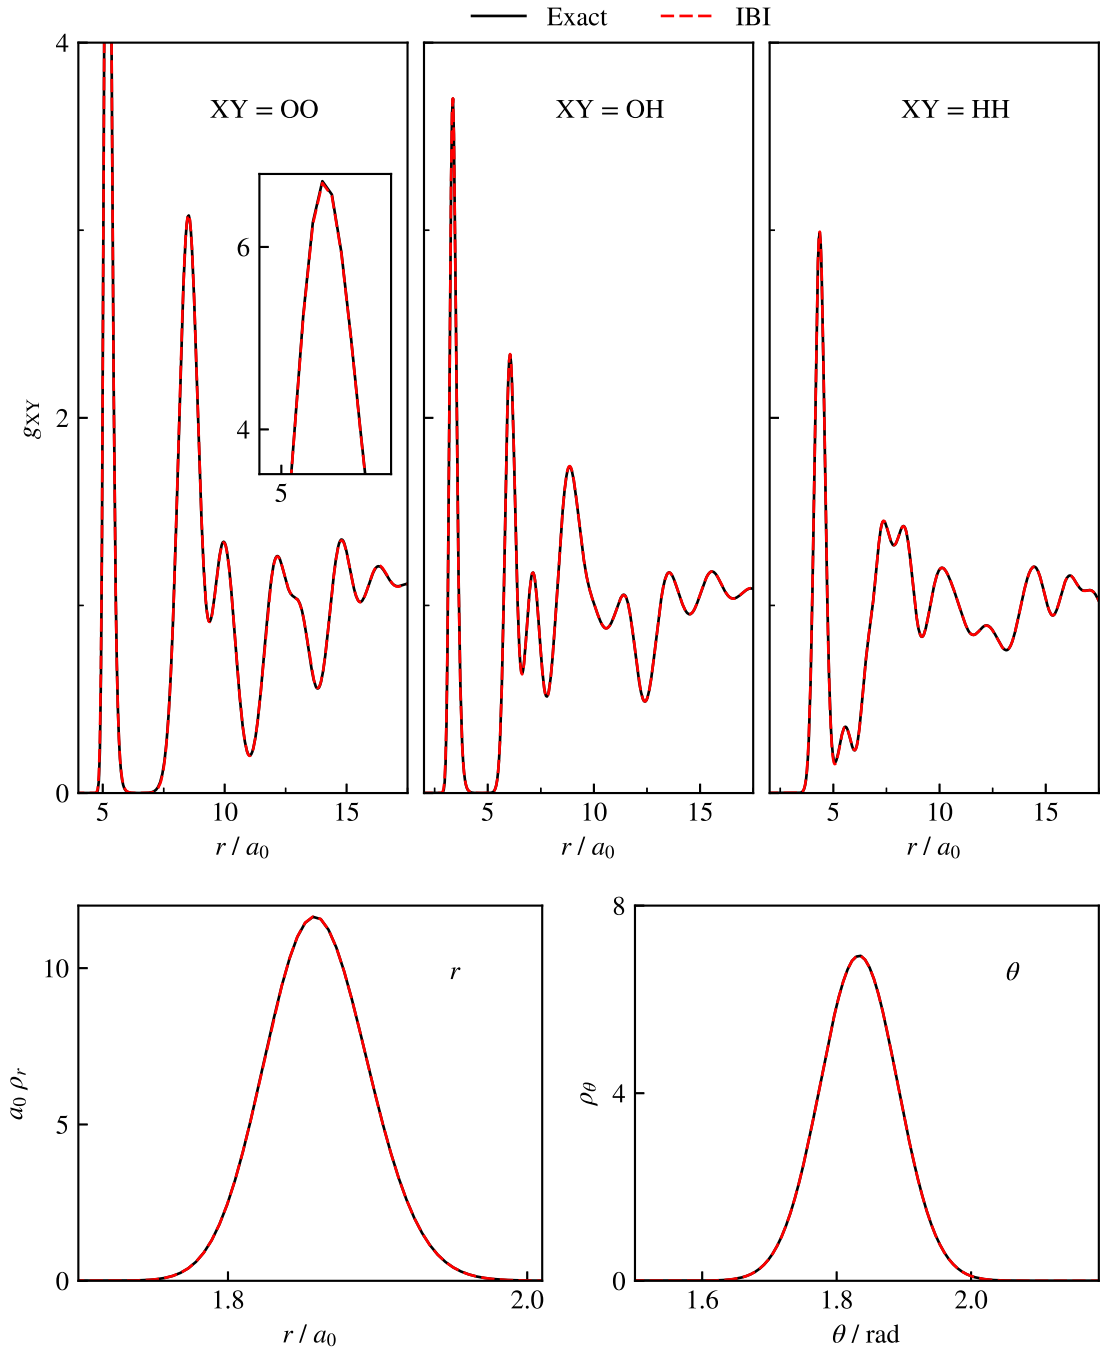

Figure 2.5: Comparison of exact quasi-centroid inter- and intramolecular distribution functions in ice to the distribution functions obtained from a classical MD simulation on the quasi-centroid potential of mean force found with IBI.

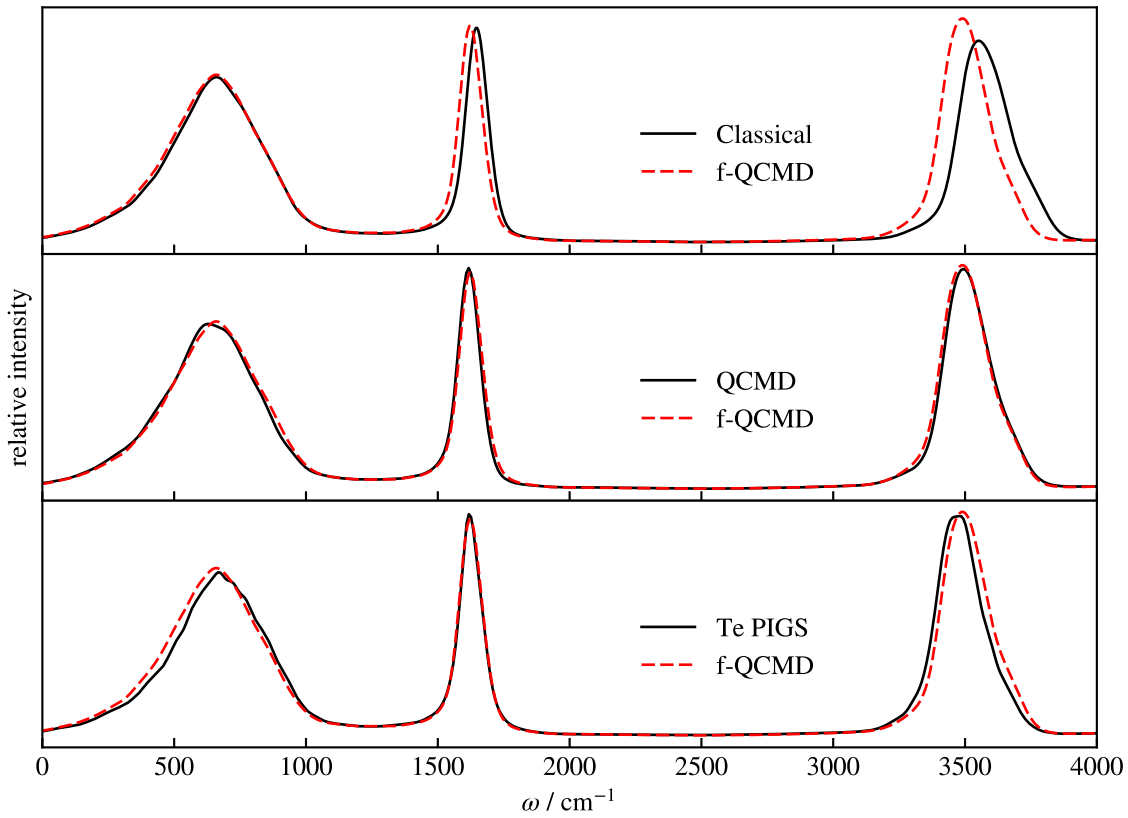

Figure 2.6: Comparison of classical MD (top), adiabatic QCMD (middle) and Te PIGS (bottom) with f-QCMD for the vibrational spectrum of liquid water at 300 K.

band and almost completely disappear in the libration band, where the thermal energy $k_\mathrm{B}T$ (ca. 200 cm^{-1} at 300 K) is comparable to the vibrational frequency, and the vibration is therefore effectively fully classical. The Te PIGS spectra are again similar to the f-QCMD spectra, except for a red shift in the stretch band and a very slight blue shift in the libration band. Because of the uneven shape of the libration band, we suspect that its slight blue shift is due to insufficient convergence of the Te PIGS correlation function, which will mostly affect the low frequency regime of the spectrum. The red shift in the stretch peak, however, could be a sign of the curvature problem, which can manifest already at 600 K.[33]

Figure 2.7 compares the vibrational spectra for hexagonal ice at 150 K. The main difference to the water spectra is that the peaks are now significantly narrower and have more structure, in particular the O-H stretch peak. We can also see a tiny peak at around 250 cm^{-1}, which corresponds to an intermolecular O-O stretch. The difference between f-QCMD and adiabatic QCMD is now more noticeable than at 300 K. There is a small red shift in the

frequency of the stretch band, the intensity of the bend band is noticeably lower in f-QCMD, and both the bend and libration bands are slightly blue shifted compared to QCMD. The most likely explanation is that the separable approximation to the quasi-centroid potential of mean force breaks down at lower temperatures. It is nevertheless worth briefly outlining other potential reasons for the disagreement. It could, for example, be that the adiabatic QCMD results carry some artefacts from their implementation. A comparison between the exact quasi-centroid distribution functions and the distribution functions obtained with adiabatic QCMD in Figure 2.8 shows some subtle differences. It is possible that the adiabatic QCMD results are either not fully converged with respect to the adiabatic separation, or that the torque estimator used to evaluate the forces in adiabatic QCMD is not entirely accurate. However, the QCMD peak positions have previously been tested for convergence with respect to the adiabatic separation[99] Also, QCMD and CMD agree in the low frequency region,[87] which suggests that the torque estimator is fine. The adiabatic QCMD peaks are therefore probably reliable, and the error is more likely to lie in the separable approximation we have used in f-QCMD.

In contrast, the Te PIGS results are remarkably close to our f-QCMD results. This is slightly surprising, since the two methods obtain their potentials of mean force in completely different ways. At first glance, there is no reason they should produce almost identical spectra. A possible explanation is the following. At 300 K, the adiabatic QCMD and f-QCMD results agree well (Figure 2.6). This indicates that the difference between the effective potential and the classical potential is dominated by pairwise interactions. The Te PIGS effective potential is computed at an even higher temperature, and we can imagine that it will therefore also be dominated by pairwise interactions. As we move towards lower temperatures, the pairwise separable approximation becomes worse. However, Te PIGS reuses the elevated temperature potential of mean force for all temperatures, so it continues to employ a pairwise potential. F-QCMD recomputes the effective potential for each new temperature, but we have *a priori* restricted it to have a separable form. By this argument, f-QCMD and Te PIGS end up with similar approximations to the effective potential, and they could thus result in similar spectra. Finally, we observe again shifts in the f-QCMD peaks compared to the classical spectrum

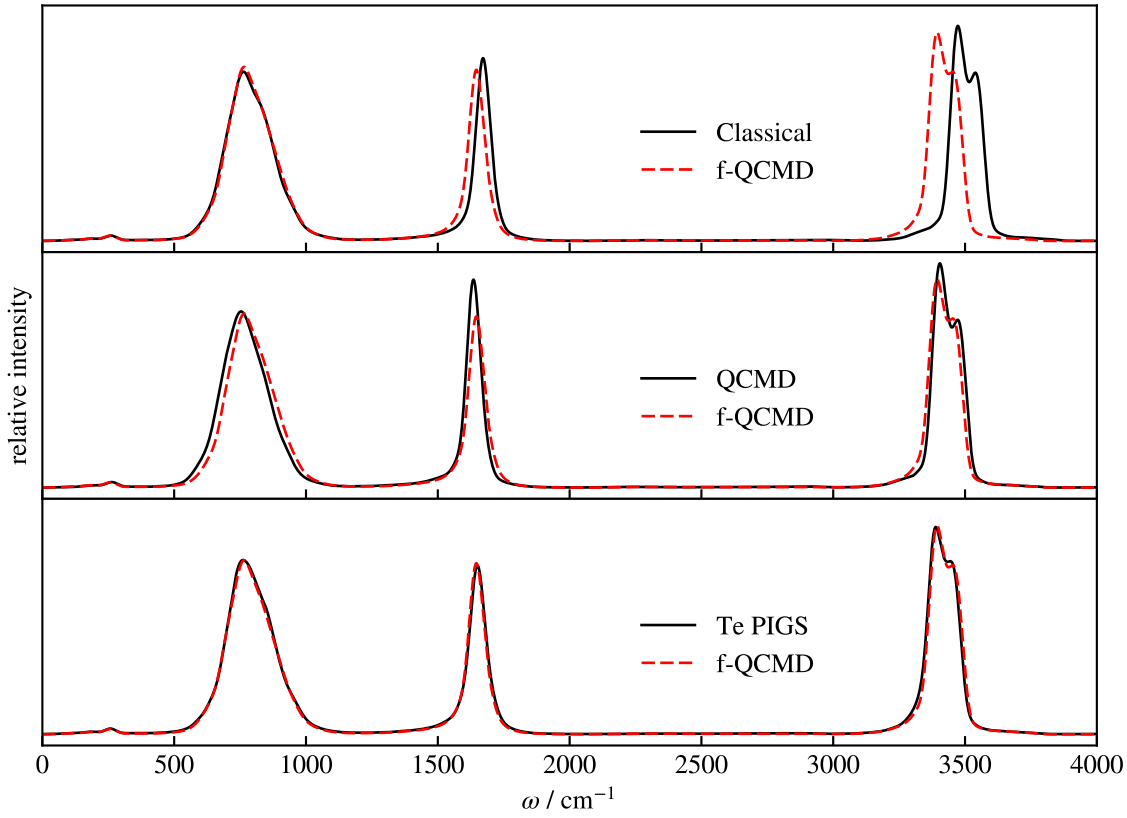

Figure 2.7: Like Figure 2.6, but for hexagonal ice at 150 K.

due to nuclear quantum effects. They cause a red shift of about 100 cm^{-1} in the stretch peak and about 50 cm^{-1} in the bend peak. There is no red shift in the libration peak (although adiabatic QCMD does show a slight red shift with respect to the classical spectrum). The shape of all peaks, however, is the same in classical and f-QCMD spectra. In fact, the peak shapes are very similar in all of the methods discussed here. This is a drastically different picture than we would have seen about a decade ago. The state-of-the-art methods at the time like (T)RPMD and CMD all have their own issues when it comes to vibrational spectra, which we outlined in Section 2.2. We are now reaching a consensus on what the 'correct' spectra of water and ice are in the fundamental region below 4,000 cm^{-1}, which was not possible previously.

2.5 Summary

This chapter has outlined the extension of f-QCMD to the condensed phase. There are three stages to an f-QCMD calculation. First, we define a set of intra- and intermolecular quasi-centroid coordinates and run a PIMD simulation to obtain their distribution functions. The coordinates are tailored to the system of interest and are thus defined on a case-by-case basis. Second, we find an approximate quasi-centroid potential of mean force through the IBI by assuming that any corrections on top of the underlying classical potential are separable in the quasi-centroid coordinates. Third, we run a classical MD simulation on the quasi-centroid potential of mean force to calculate the dipole autocorrelation function according to Equation (2.33), with which we obtain the vibrational spectrum upon Fourier transformation.

We have applied this procedure to liquid water and ice. For the first stage, we defined the intramolecular O-H bond lengths and H-O-H bond angles and the intermolecular pair-wise O-O, O-H and H-H distances as quasi-centroid coordinates. We then showed that the convergence of the IBI iterations needed in the second stage can be made more stable by regularising the IBI update, Equation (2.49). Our regularisation has the additional benefit of producing smooth potentials of mean force which are amenable to MD simulations. While such a regularisation was proposed as early as in the initial publication on IBI,[91] it did not

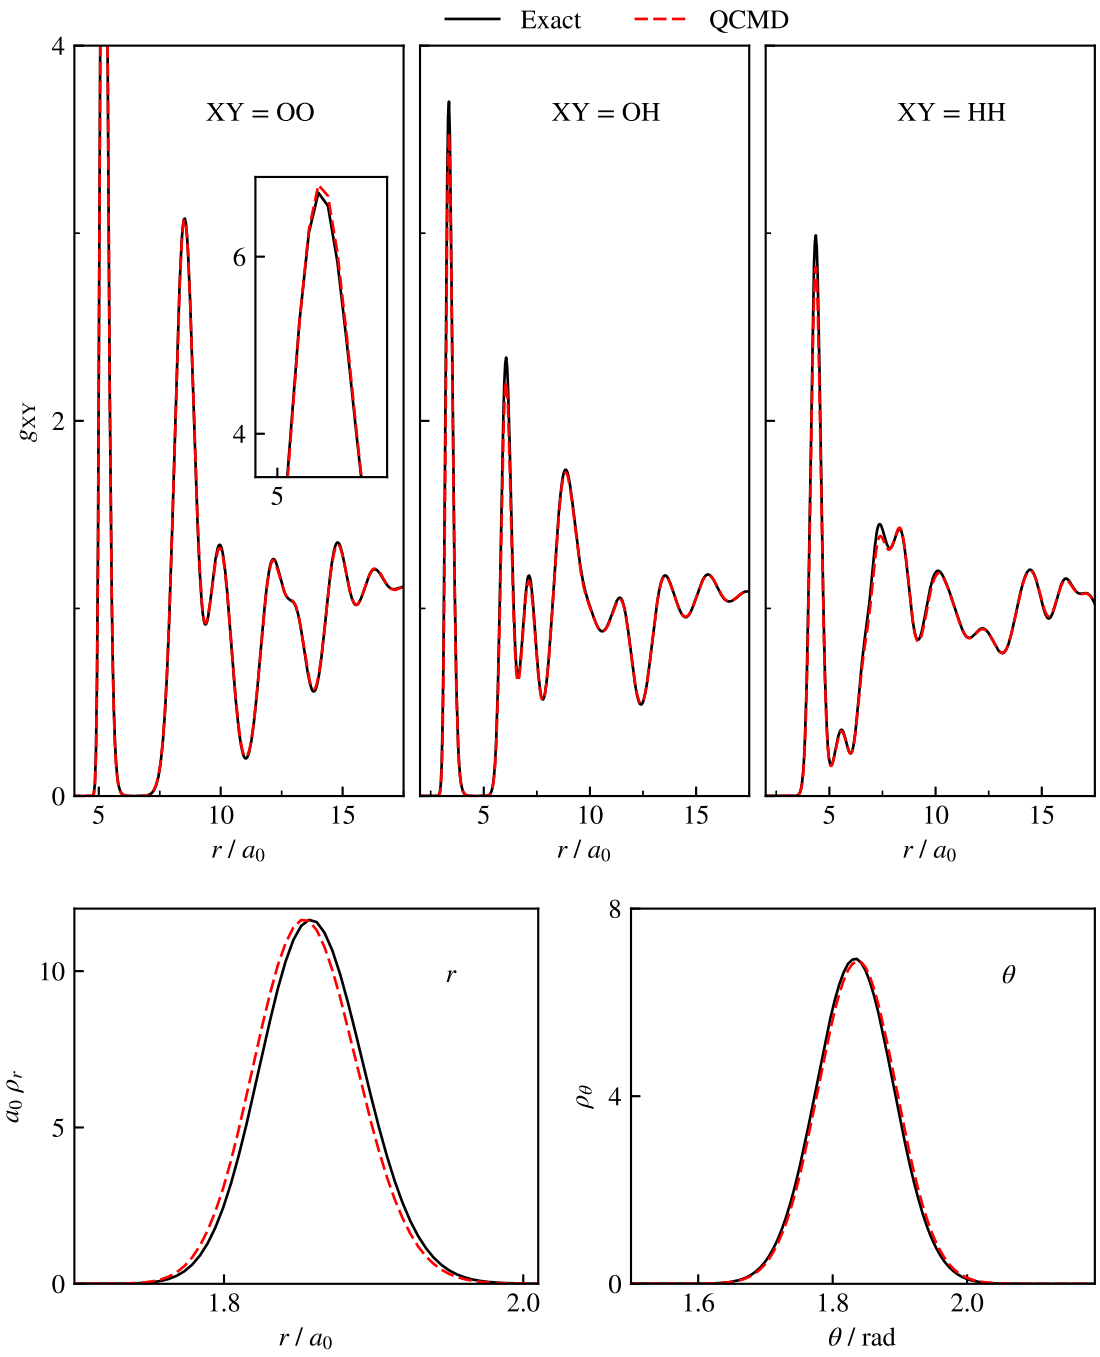

Figure 2.8: Exact quasi-centroid and QCMD inter- and intramolecular distribution functions for ice at 150 K. The exact results were obtained by us from a PIMD simulation, and the QCMD results were provided by George Trenins.

seem to gain traction. Instead, the iterations were controlled by an overall scaling factor in front of the correction term,[96] which does not remove discontinuities in the potential of mean force which arise due to statistical noise in the distribution functions.

We continued by comparing the f-QCMD spectrum with classical, adiabatic QCMD and Te PIGS results for water at 300 K and ice at 150 K. F-QCMD agrees better with adiabatic QCMD in water than in ice, which suggests that the separable assumption on the quasi-centroid potential of mean force (Equations (2.44)-(2.47)) is better at higher temperatures. Te PIGS also agrees well with f-QCMD, except for minor shifts in the liquid water spectrum, which we suspect are either due to insufficient convergence or an early manifestation of the curvature problem. In general, although there are small shifts between the quantum methods, they agree quite well in the line shapes. We are thus now reaching a consensus between path integral methods. This is in stark contrast to the situation a decade ago, when the state-of-the-art methods TRPMD, RPMD and CMD could produce wildly different spectra. This allows us to finally quantify the impact of nuclear quantum effects by comparing f-QCMD to the classical spectrum. For both water and ice, the classical and quantum line shapes are similar, and the only nuclear quantum effect is an anharmonic red shift in the frequency of the O-H stretching band.

As a caveat, we should note that although we are now confident that recently developed methods like QCMD variants and Te PIGS work well for the fundamental region of the spectrum, they are less well suited for the overtone region where they severely underestimate intensities. This has been extensively discussed elsewhere,[100,101] where perturbative correction factors have been derived to rescale the overtone bands. However, since this is not straightforward to do if bands overlap, the reliable treatment of overtones remains an open problem.

Chapter 3

Vibrational strong coupling in liquid water

3.1 Introduction

There has been considerable recent interest in understanding molecular processes inside a microcavity (Figure 3.1). Experiments suggest that the presence of the microcavity can significantly influence observables like (linear and nonlinear) infrared spectra,[102] chemical rates[37,39,103], and conductivity and photoconductivity of semiconductors.[40,41,104] The underlying mechanism of many of these effects is not well understood, however, and the observed effects remain hotly debated.[42,45,46] Changes to chemical reaction rates in particular have stirred up a frenzy because of their potential applications. They are also surprising, since the coupling between the light and any single molecule in the cavity is weak,[105] and since any collective excitation of chemical reactions to their transition state would have a prohibitively high activation energy.[44] In contrast, we understand the effect of the cavity on the linear infrared spectrum much better: When the cavity is in resonance with one of the vibrational modes of the system, two new light-matter states form, which we call the polariton states. The energy difference between the polariton states is called the Rabi splitting. The polariton states appear in the infrared spectrum as two narrow peaks, as shown in Figure 3.2. Since the cavity interacts with the collective dipole moment of the system within it, it is not surprising

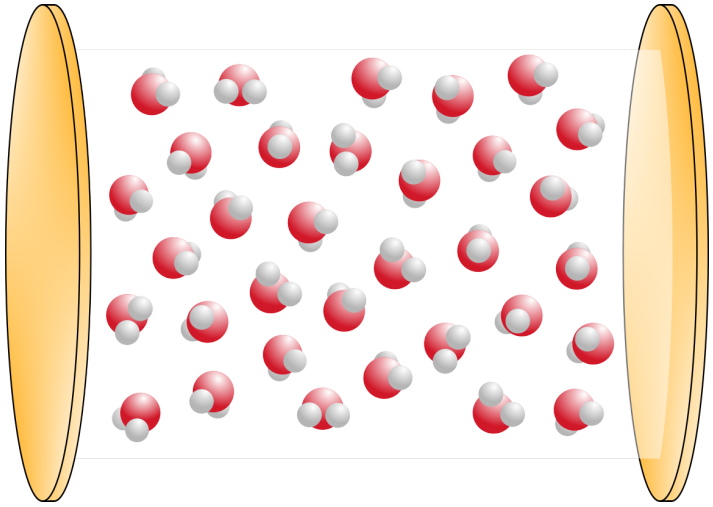

Figure 3.1: An example of a cavity made from gold mirrors and filled with water. In the applications considered in this section, the spacing between the gold mirrors is usually on the order of a few micrometers.

to see such a pronounced effect in this observable.

Li *et al.* developed a classical MD method to describe the effect of the cavity,[47] and applied it to the vibrational spectrum of liquid water.[43] In the same study, they calculated static equilibrium properties such as the dielectric constant of liquid water and the pair distribution functions we encountered earlier (Equation (2.45)). They finished with a discussion of the cavity's effect on dynamical properties such as the velocity and orientation autocorrelation functions. For the vibrational spectrum, the O-H stretching peak splits into two narrow polariton peaks as expected from the discussion above. Li *et al.* also developed a one-dimensional harmonic oscillator model which qualitatively (but not quantitatively) predicts the polariton peak positions and their intensities with varying coupling strengths.[43] For the static equilibrium properties, they showed analytically that within a classical treatment, no property of the system can be modified through the presence of the cavity, which they then confirmed through simulations. They did not observe any significant modifications to the velocity autocorrelation function and only a tiny modification to the orientational autocorrelation function when the cavity was introduced. For the latter, they found that small peaks appear in its Fourier transform in the same places as the polariton peaks in the IR spectrum. However, the orientational autocorrelation function is a single-molecule property. We expect any effect of the cavity to wash out in a simulation that takes into

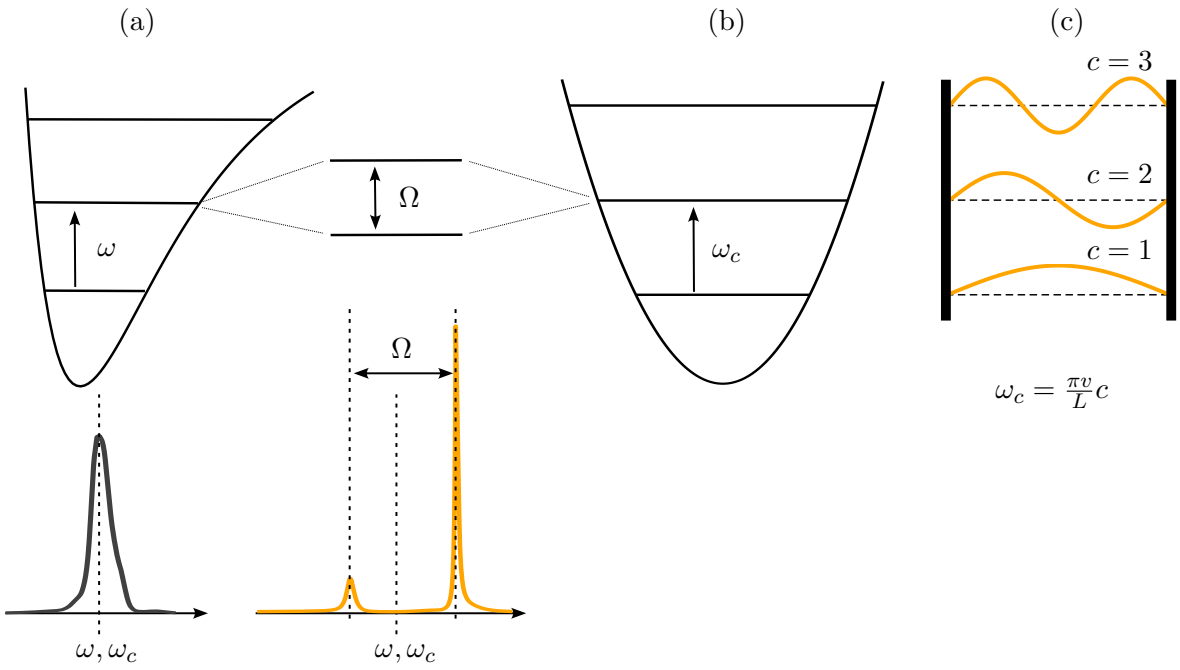

Figure 3.2: Coupling between (a) the system vibrational modes and (b, c) the cavity photons. The vacuum cavity harmonic oscillator potential is shown in panel (b) and the standing waves of the lowest three cavity modes are shown in panel (c). The polariton states are drawn between the system and cavity panels (a) and (b). The cavity frequencies are $\omega_c = \frac{\pi v}{L} c$, where $v = c_0/n_\infty$ is the speed of light in the cavity, c_0 is the speed of light in vacuum and n_∞ is the refractive index of liquid water. The integer index c enumerates the cavity modes.
The system's infrared spectra inside and outside the cavity are drawn in orange and grey, respectively, below the relevant panels. The two polariton states arising from the interaction between the system and the cavity cause the vibrational peak to split into two narrow polariton peaks, which are separated by the Rabi splitting Ω. The cavity frequency ω_c and the system frequency ω are chosen to be equal here, and are represented by a dotted line in the spectra.

account more than the 216 water molecules they considered.

Li *et al.* concluded that, while the cavity has a significant effect on the IR spectrum, it seems to be completely irrelevant for equilibrium properties and single-molecule dynamical properties in the limit of classical nuclear motion.[43] To test whether these conclusions hold up under the inclusion of nuclear quantum effects, they followed up with a study of liquid water in a cavity using TRPMD.[48] Compared to the classical case, they did not notice significant changes to the dielectric constant or the Rabi splitting upon the introduction of nuclear quantum effects. However, the TRPMD polariton peaks associated with the water O-H stretch were found to experience a red shift and almost double in width compared to the

classical ones. The red shift is unsurprising. It is a known consequence of nuclear quantum effects, as we have discussed extensively in Chapter 2. What they could not unequivocally determine, however, is whether the broadening in the polariton peaks is due to the usual broadening introduced by TRPMD (discussed in Section 2.2.2), or whether it is a genuine effect brought on by the cavity. Having just developed an f-QCMD algorithm for liquid water, we are in the perfect position to shine light on this question.

3.2 Cavity molecular dynamics

The Hamiltonian for an MD simulation of a system in an optical cavity is[47]

$$\mathcal{H} = \mathcal{H}_{\text{sys}} + \mathcal{H}_{\text{cav}}, \tag{3.1}$$

with the system Hamiltonian

$$\mathcal{H}_{\text{sys}} = \sum_{i=1}^{N} \frac{|\boldsymbol{p}_i|^2}{2m_i} + V_{\text{cl}}(\boldsymbol{q}), \tag{3.2}$$

and the cavity Hamiltonian

$$\mathcal{H}_{\text{cav}} = \sum_{c,d} \left[\frac{p_{c,d}^2}{2} + \frac{1}{2}\omega_c^2 \left(q_{c,d} + \frac{\mu_d(\{\boldsymbol{q}_i\})}{\omega_c \sqrt{\varepsilon_0 V_{\text{cav}}}} \right)^2 \right]. \tag{3.3}$$

$\boldsymbol{q}_i$ and $\boldsymbol{p}_i$ are the position and momentum of the ith of the total N molecules in the cavity, and μ_d is the d component of the molecular dipole moment, where $d \in \{x, y\}$ and we assume that the cavity mirrors are placed along the z direction. The double sum in Equation (3.3) goes over the cavity modes c with frequencies ω_c, and over the polarisation directions d. All that remains is to specify the coupling $\frac{1}{\sqrt{\varepsilon_0 V_{\text{cav}}}}$, where V_{cav} is the volume of the cavity. To this end, we note first that the dipole moment μ_d ought to be normalised by $\frac{1}{\sqrt{V_{\text{sys}}}}$, where V_{sys} is the system volume. One way to see this is to note that the dielectric constant, which is an intensive material property, is related to the dipole variance divided by V_{sys}.[106–108] In general, $V_{\text{sys}} \leq V_{\text{cav}}$, for example if the mirrors are coated with silica so as to avoid interactions with the substrate. With all this in mind, we can define the ratio between the system and cavity

volume, $R = V_{\text{sys}}/V_{\text{cav}}$, and the scaled, system size-independent dipole moment

$$\nu_d(\{\boldsymbol{q}_i\}) = \frac{\mu_d(\{\boldsymbol{q}_i\})}{\sqrt{\varepsilon V_{\text{sys}}}}, \tag{3.4}$$

with which the cavity Hamiltonian can be written as

$$\mathcal{H}_{\text{cav}} = \sum_{c,d} \left[\frac{p_{c,d}^2}{2} + \frac{1}{2}\omega_c^2 \left(q_{c,d} + \sqrt{R}\frac{\nu_d(\{\boldsymbol{q}_i\})}{\omega_c} \right)^2 \right]. \tag{3.5}$$

This has two advantages. Firstly, using the scaled dipole moment $\boldsymbol{\nu}(\{\boldsymbol{q}_i\})$ is convenient for our simulations, which use periodic boundary conditions. We fix R, and increase the simulation box size until our results are converged. Scaling the dipole moment in Equation (3.4) makes sure that the box size does not impact our simulations. Secondly, it makes obvious that there is a physical limit to the coupling strength, since it constrains $0 < R \leq 1$. As we will show below, $R = 1$ provides an upper limit to the Rabi splitting we can physically expect. To obtain the cavity spectrum, we compute the Fourier transform of the dipole autocorrelation function just as in Chapter 2,

$$I(\omega) = \frac{\beta}{4\varepsilon_0 c V_{\text{sys}}} \int_{-\infty}^{\infty} e^{-i\omega t} \sum_{d=x,y} \langle \dot{\mu}_d(0)\dot{\mu}_d(t)\rangle \mathrm{d}t. \tag{3.6}$$

In contrast to Equation (2.1), we do not include the z direction along which the cavity mirrors are placed. Since the cavity only interacts with the x and y components of the dipole moment, the spectrum along the z direction would simply recover the cavity-free spectrum.

From here on, a classical cavity MD simulation proceeds just like any classical MD simulation, with additional degrees of freedom for the cavity photons and with equations of motion derived from the total Hamiltonian (Equation (3.1)). A cavity TRPMD simulation is also straightforward. We simply add cavity degrees of freedom to the ring polymer.[47,48] For cavity f-QCMD simulations, we will reuse the cavity-free quasi-centroid potentials of mean force for liquid water obtained in Chapter 2. In doing this, we implicitly define the cavity quasi-centroid coordinates to coincide with the Cartesian centroids, and note that the cavity does not impact the radial and angular distribution functions of the water molecules.[43] For

simplicity, we will label the f-QCMD results as 'QCMD' from here on.

3.3 Results

3.3.1 Nuclear quantum effects

We started by computing the classical cavity spectrum for liquid water, and comparing it to Li *et al.*'s results.[43] Both used the q-TIP4P/F force field for the molecular potential in Equation (3.2). We increased R until our Rabi splitting (720 cm^{-1}) was close to their reported value (715 cm^{-1}), which we achieved for $R = 2.25$. Like any $R > 1$, this is clearly unphysical, since it corresponds to the case where the volume of water inside the cavity exceeds the cavity volume. We will consider more realistic values in the coming sections, but for the sake of comparing like with like we will use $R = 2.25$ in this section. We then computed TRPMD, RPMD and QCMD spectra for this parameter. The results are shown in Figure 3.3. To start, we observe that the classical line widths are also similar to Li *et al.*'s: 18 cm^{-1} and 59 cm^{-1} (us) compared to 24 cm^{-1} and 60 cm^{-1} (them) for the lower and upper polariton peaks, respectively

Let us now turn to the TRPMD results, Firstly, both inside and outside the cavity, the spectra are red shifted by around 100 cm^{-1} with respect to the classical spectrum due to the usual anharmonic red shift. We measured a TRPMD Rabi splitting of 727 cm^{-1}, with upper and lower polariton line widths of 78 cm^{-1} and 134 cm^{-1}, respectively. This is close to the values reported by Li *et al.*, which are 720 cm^{-1} for the Rabi splitting, 65 cm^{-1} for the upper and 129 cm^{-1} for the lower polariton line width.[48] In TRPMD, the polariton bands are thus about twice as broad as in classical MD. Using the fact that the broadening is less severe outside the cavity (it comes out to around a 30% broadening), Li *et al.* argued that part of the broadening may be due to nuclear quantum effects in the cavity.[48]

The QCMD results shed light on this suggestion. While the peak positions are the same as in TRPMD, the line widths are noticeably smaller. They are in fact close to the classical values, with an upper polariton line width of 18 cm^{-1} and a lower polariton line width of 44 cm^{-1}. We have also computed the RPMD spectra inside and outside the cavity. The cavity-free spectrum is known to show signs of a spurious resonance between the O-H stretch

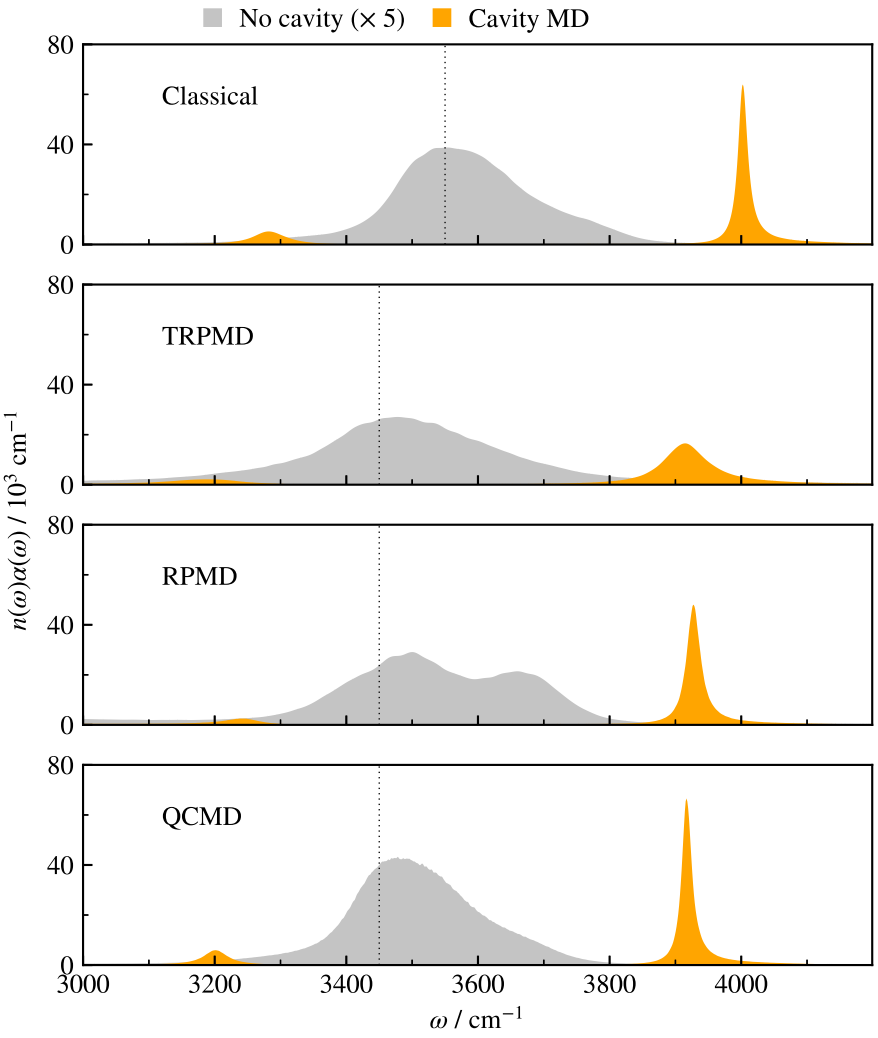

Figure 3.3: Comparison of the infrared spectrum of liquid water outside (grey) and inside (orange) the cavity, as obtained from classical MD and three different path integral methods. A single cavity mode is coupled to the O-H stretch at a frequency that is indicated by the dotted line (3,550 cm^{-1} for the classical simulation, 3,450 cm^{-1} for the rest). The intensities of the cavity-free spectra are scaled by a factor 5 for an easier comparison.

and the ring polymer internal modes,[20] which gives rise to a double hump. However, it seems that at least the upper polariton fortuitously is not perturbed by resonances (the lower polariton is smeared out, making it barely visible). By themselves, RPMD results should be taken with a grain of salt, but here we see the same red shift as TRPMD and QCMD. Together with our conclusions from Chapter 2, we are confident that this is the correct peak position at our level of theory. The RPMD line width of the upper polariton is 25 cm^{-1}. This is much closer to the classical and QCMD line widths, and serves as another piece of evidence that the observed broadening in TRPMD is purely an artefact. It also justifies our choice of the QCMD quasi-centroid potentials of mean force. From here on, we will only present QCMD spectra, and consider more realistic R values.

3.3.2 More realistic simulations

Figure 3.4 shows cavity QCMD spectra for a number of different cavity setups at the highest physical coupling strength obtained with $R = 1$. Figure 3.5 shows the same but for $R = 0.5$. When only a single cavity mode is resonant with a molecular vibration (panels (a)-(c)), the cavity spectrum is relatively straightforward to interpret. The peak that is resonant with the cavity splits into an upper and a lower polariton. There is very little impact on the remaining peaks, except for slight shifts away from the polariton peaks (visible in panels (b) and (c) of Figure 3.4), which can be thought of as repulsive interactions with the polariton states. These interactions are hardly noticeable at lower coupling strengths (Figure 3.5), though.

The Rabi splitting is known to scale as $\sqrt{N}$, where N is the number of molecules in the cavity. Since R is proportional to N, the Rabi splitting also scales as $\sqrt{R}$. We can verify this by comparing Figure 3.4 to Figure 3.5, where the Rabi splitting is indeed about a factor $\sqrt{2}$ smaller.

The multimode spectra in panels (d) of Figures 3.4 and 3.5 show more complex features, and it is difficult to assign peaks to the physical vibrations. The most noticeable difference between the multimode spectra at different coupling strengths in Figures 3.4 and 3.5 is that the highest frequency peak is significantly narrower and more blue shifted for $R = 1$ than for $R = 0.5$. However, this could be related to the fact that we do not include cavity frequencies

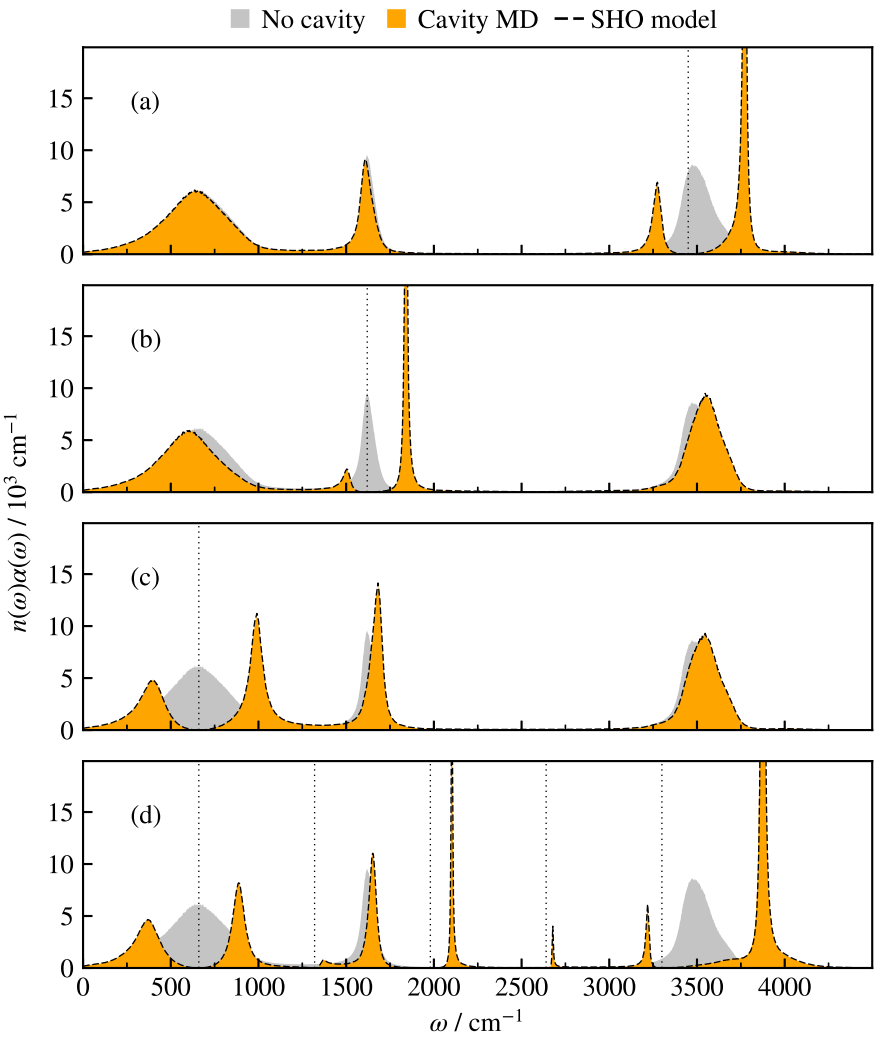

Figure 3.4: Cavity QCMD spectra of liquid water with the cavity tuned to different fundamental vibrational frequencies at the maximal coupling strength $R = 1$. Panels (a)-(c) assume there is just a single cavity mode. Panel (d) shows a multi-mode spectrum, which includes four overtones on top of the fundamental frequency, which is tuned to the libration band. The dotted lines indicate the cavity frequencies at (a) 3,450 cm^{-1}, (b) 1,620 cm^{-1}, (c) 660 cm^{-1} and (d) $c \cdot 660$ cm^{-1}, where $c = 1, \ldots, 5$.

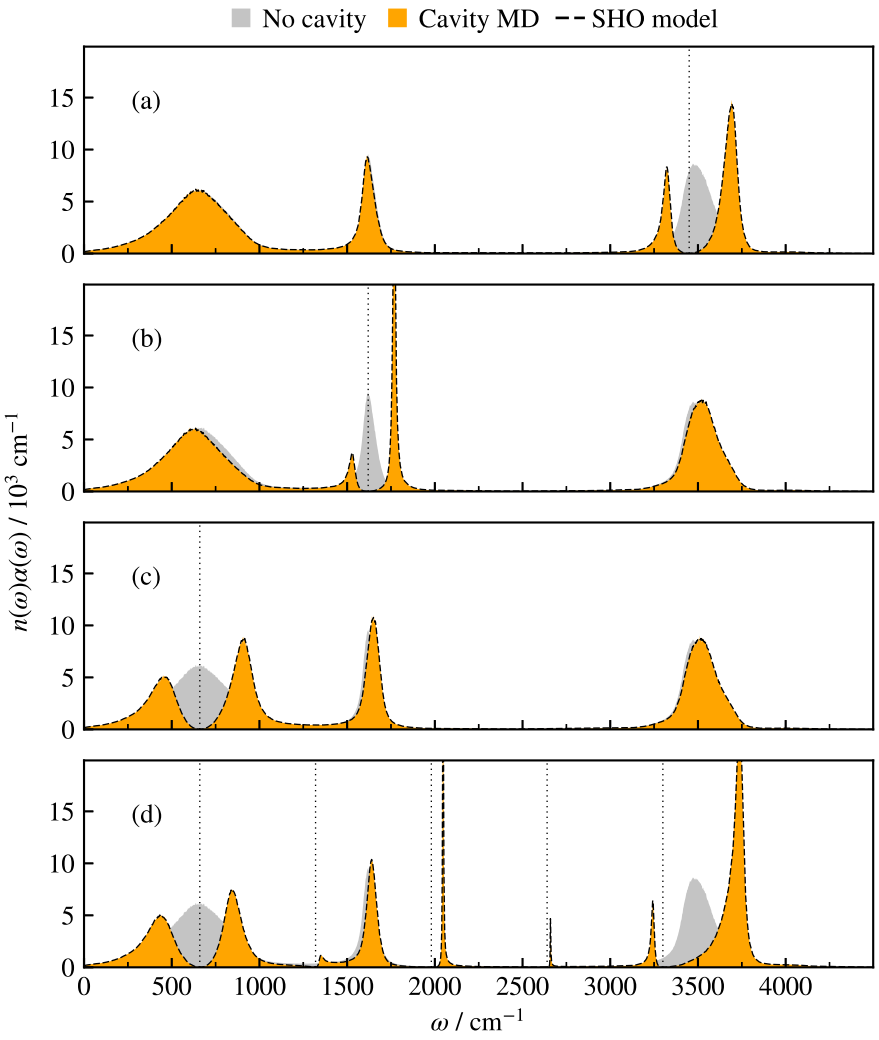

Figure 3.5: As Figure 3.4, but for $R = 0.5$.

above 4,000 cm^{-1}. All of the effects we observe in the single or multi-mode cavity QCMD spectra are consistent with theoretical expectations. However, we will next discuss how we can avoid doing any MD simulation, and instead predict the cavity spectra to graphical accuracy using a simple one-dimensional harmonic oscillator (SHO) model.

3.3.3 A simple harmonic oscillator model

To build our model, we first consider the cavity-free spectrum, which is obtained via

$$I(\omega) = \frac{\beta}{2c} \int_{-\infty}^{\infty} e^{-i\omega t} \langle \dot{\nu}_x(\{\boldsymbol{q}_i(0)\}) \dot{\nu}_x(\{\boldsymbol{q}_i(t)\}) \rangle \mathrm{d}t, \tag{3.7}$$

where we have substituted the dipole moment in Equation (3.6) with the scaled dipole moment from Equation (3.4), and we consider only the x component of the dipole moment in order to simplify the discussion. Consider next a grid of frequencies ω_i, $i = 1, \ldots, N$. We assume that the cavity-free spectrum is available on such a grid, and define the pre-limit delta functions,

$$\delta_i(\omega) = \begin{cases} \frac{1}{\Delta\omega_i} & \omega \in (\frac{\omega_i + \omega_{i-1}}{2}, \frac{\omega_{i+1} + \omega_i}{2}) \\ 0 & \text{else} \end{cases}, \tag{3.8}$$

where $\Delta\omega_i = (\omega_{i+1} - \omega_{i-1})/2$ and with $\omega_0 = \max(0, 2\omega_1 - \omega_2)$, $\omega_{N+1} = 2\omega_N - \omega_{N-1}$. Using these, we discretise the spectrum as

$$I(\omega) = \sum_{i=1}^{N} \Delta\omega_i I(\omega_i) [\delta_i(\omega) + \delta_i(-\omega)]. \tag{3.9}$$

Next, consider a system of N one-dimensional harmonic oscillators with the same frequencies ω_i and with the Hamiltonian

$$\mathcal{H}_{\mathrm{sys}}(\{p_i, q_i\}) = \sum_{i=1}^{N} \left(\frac{p_i^2}{2} + \frac{1}{2}\omega_i^2 q_i^2 \right) \tag{3.10}$$

and with a linear dipole moment

$$\nu = \sum_{i=1}^{N} \nu_i q_i. \tag{3.11}$$

It is straightforward to derive an analytic expression for the dipole derivative autocorrelation function for this model system,

$$\langle \dot{\nu}(0)\dot{\nu}(t)\rangle = \frac{1}{\beta}\sum_{i=1}^{N}\nu_i^2\cos(\omega_i t), \tag{3.12}$$

which also allows us to find for the vibrational spectrum,

$$I(\omega) = \frac{\pi}{2c}\sum_{i=1}^{N}\nu_i^2[\delta(\omega-\omega_i)+\delta(\omega+\omega_i)] \approx \frac{\pi}{2c}\sum_{i=1}^{N}\nu_i^2[\delta_i(\omega)+\delta_i(-\omega)]. \tag{3.13}$$

Equations (3.13) and (3.9) are equal if we set

$$\nu_i = \sqrt{\frac{2c\Delta\omega_i I(\omega_i)}{\pi}}, \tag{3.14}$$

which determines the dipole moments of the harmonic oscillators. Finally, we couple the harmonic oscillators to the cavity with these dipole moments. The cavity Hamiltonian is then

$$\mathcal{H}_{\text{cav}}(\{q_i\},q_c,p_c) = \sum_{c=1}^{M}\left[\frac{p_c^2}{2}+\frac{1}{2}\omega_c^2\left(q_c+\sqrt{R}\sum_{i=1}^{N}\frac{\nu_i q_i}{\omega_c}\right)^2\right], \tag{3.15}$$

and we can write the full Hamiltonian of this SHO model as

$$\mathcal{H}(\mathbf{p},\mathbf{q}) = \frac{1}{2}\mathbf{p}^T\mathbf{p}+\frac{1}{2}\mathbf{q}^T\mathbf{H}\mathbf{q}, \tag{3.16}$$

with $\mathbf{p} = (\{p_i\},\{p_c\})$, $\mathbf{q} = (\{q_i\},\{q_c\})$ and the Hessian matrix elements $H_{ij} = \omega_i^2\delta_{ij} + MR\nu_i\nu_j$, $H_{ic} = H_{ci} = \sqrt{R}\omega_c\nu_i$ and $H_{cc'} = \omega_c^2\delta_{cc'}$. Diagonalising the Hessian matrix with an orthogonal matrix $\mathbf{C}$

$$\mathbf{C}^T\mathbf{H}\mathbf{C} = \tilde{\boldsymbol{\omega}}, \tag{3.17}$$

and transforming the dipole moment into the eigenbasis,

$$\tilde{\boldsymbol{\nu}} = \mathbf{C}^T\boldsymbol{\nu}, \tag{3.18}$$

gives the predicted cavity spectrum of the SHO model as

$$I(\omega) = \frac{\pi}{2c} \sum_{i=1}^{N} \tilde{\nu}_i [\tilde{\delta}_i(\omega) + \tilde{\delta}_i(-\omega)], \tag{3.19}$$

with the pre-limit delta functions for the new frequencies defined in analogy to Equation (3.8),

$$\tilde{\delta}_i(\omega) = \begin{cases} \frac{1}{\Delta\tilde{\omega}_i} & \omega \in \left[\frac{\tilde{\omega}_i + \tilde{\omega}_{i-1}}{2}, \frac{\tilde{\omega}_{i+1} + \tilde{\omega}_i}{2}\right] \\ 0 & \text{else} \end{cases}, \tag{3.20}$$

$\Delta\tilde{\omega}_i = (\omega_{i+1} - \omega_{i-1})/2$, $\tilde{\omega}_0 = \max(0, 2\tilde{\omega}_1 - \tilde{\omega}_2)$, and $\tilde{\omega}_{N+1} = 2\tilde{\omega}_N - \tilde{\omega}_{N-1}$. In Equation (3.18), $\boldsymbol{\nu} = (\{\nu_i\}, \{0\})$ is the dipole vector with entries from Equation (3.14) that has been padded with M zeros to be compatible with the matrix $\mathbf{C}$. We already included the SHO model's predictions in Figures 3.4 and 3.5 as black dashed lines, using the cavity-free QCMD spectrum as input to the discretisation in Equation (3.9). The model agrees to graphical accuracy with the cavity QCMD simulations in all panels and at both interaction strengths. It thus inherits its accuracy directly from the input spectrum: We input the result of a cavity-free QCMD simulation using the q-TIP4P/F potential, so the SHO model returns the analogous cavity spectrum. This means that there is no need for cavity MD simulations. With the model described here we can perfectly predict cavity MD spectra, at a fraction of the cost of a cavity MD simulation.

3.3.4 Applications to experiments

Another advantage of the model described in Section 3.3.3 is that we have complete freedom in choosing the input spectrum. A particularly attractive option is the experimental spectrum. It is the exact (quantum mechanical) cavity-free spectrum and avoids all of the assumptions inherent to MD, like the approximation of the true potential energy surface, finite size effects, and approximate treatment of the nuclear quantum effects. The experimental spectrum of liquid water is available on a grid of frequencies.[109] We can then follow the procedure outlined in Section 3.3.3 to discretise the spectrum and to obtain a prediction for cavity spectra. The only thing left to do is to specify where to place the cavity modes. From an experimental

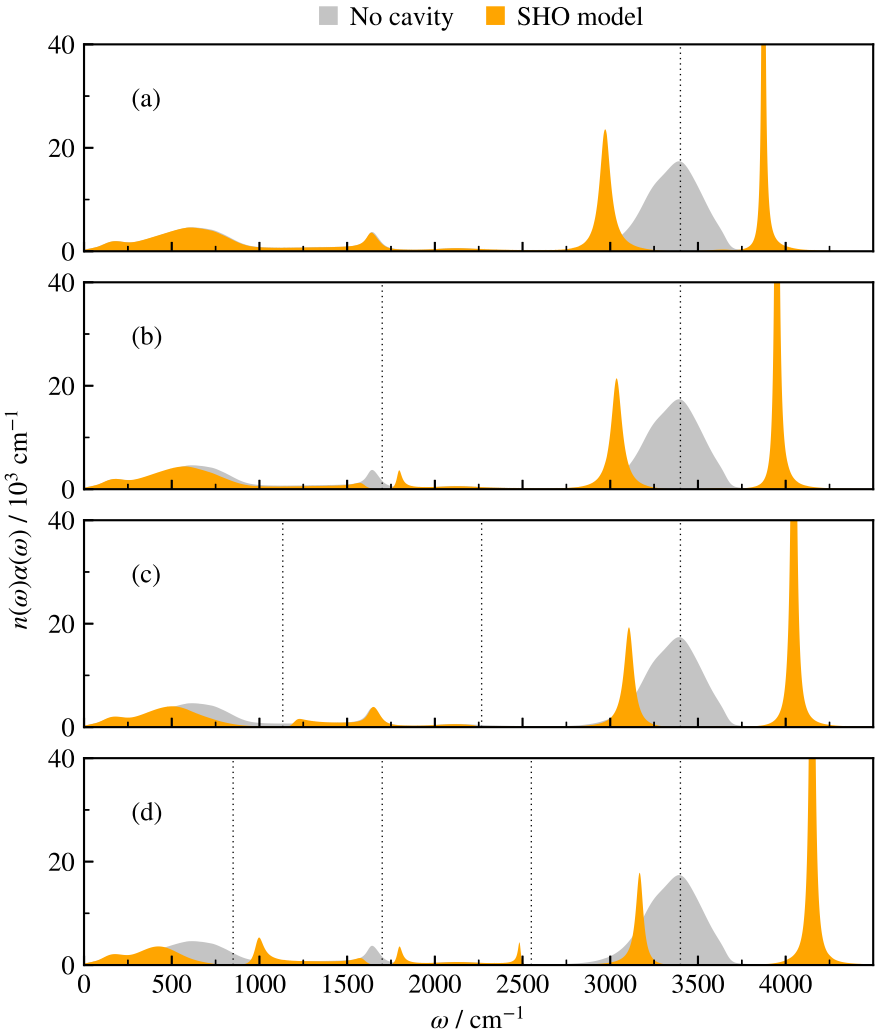

Figure 3.6: Impact of the cavity on the experimental liquid water spectrum as predicted from the SHO model using $R = 1$. The cavity modes with (a) $c = 1$, (b) $c = 2$, (c) $c = 3$ and (d) $c = 4$ are resonant with the OH stretch, as described in the text. The cavity frequencies are indicated with dotted vertical lines, and we only include modes with frequencies up to $3,500$ cm^{-1}.

point of view, the O-H stretch is the most interesting degree of freedom. It has a much higher intensity than the bend or libration peaks in the experimental spectrum, and is therefore expected to result in a larger Rabi splitting. Therefore, we choose either the fundamental cavity mode or one of the overtones to be resonant with the O-H stretch. This consideration was less relevant when we were studying water with the q-TIP4P/F model: The intensity of the O-H stretch peak is much smaller than in the experimental spectrum, and the cavity has a comparably large effect on all fundamental modes. Given that we are now in a position to make realistic predictions about the cavity spectrum, we view the simulations presented in Section 3.3.2 mostly as forming our initial understanding of single- and multi-mode cavity spectra. Figure 3.6 shows the SHO model's predictions for a coupling strength of $R = 1$ when the experimental cavity-free spectrum is used as input. The Rabi splittings are similar in all four panels, which is consistent with observations in experiment.[45]

There have been numerous experiments on liquid water in a cavity.[38,45,110] For example, Fukushima *et al.* reported upper and lower polariton peak positions as a function of the (fundamental) cavity frequency ω_c.[110] This is precisely the information we need (together with the geometry of their cavity) to assess our predictions that use the cavity-free water spectrum[109] as input. We use $R = 0.96$ throughout, which corresponds to a cavity spacing of 1.0 μm and accounts for the 0.02 μm thick silica coatings of the mirrors. Note that, in principle, the ratio R changes slightly when the mirrors are moved to implement different cavity frequencies, but we have not taken this into account in our simulations. We have also not taken into account any changes to the spectra due to the presence of the silica layers. Still, the agreement of our simulations with experiment shown in Figure 3.7 is remarkable. Our predictions are practically on top of the experimental lower polariton peak positions. The agreement in the upper polariton is slightly worse, where our predictions start off close to experiments at lower cavity wavenumber ω_c, but drift away from the experimental data when ω_c increases. On the one hand, this could be because we only include a single cavity mode in our predictions, which is a crude approximation. Including overtones would shift the upper polariton down, and it is possible that this would improve agreement. On the other hand, it is possible that the reported experimental values suffer from some artefacts (see,

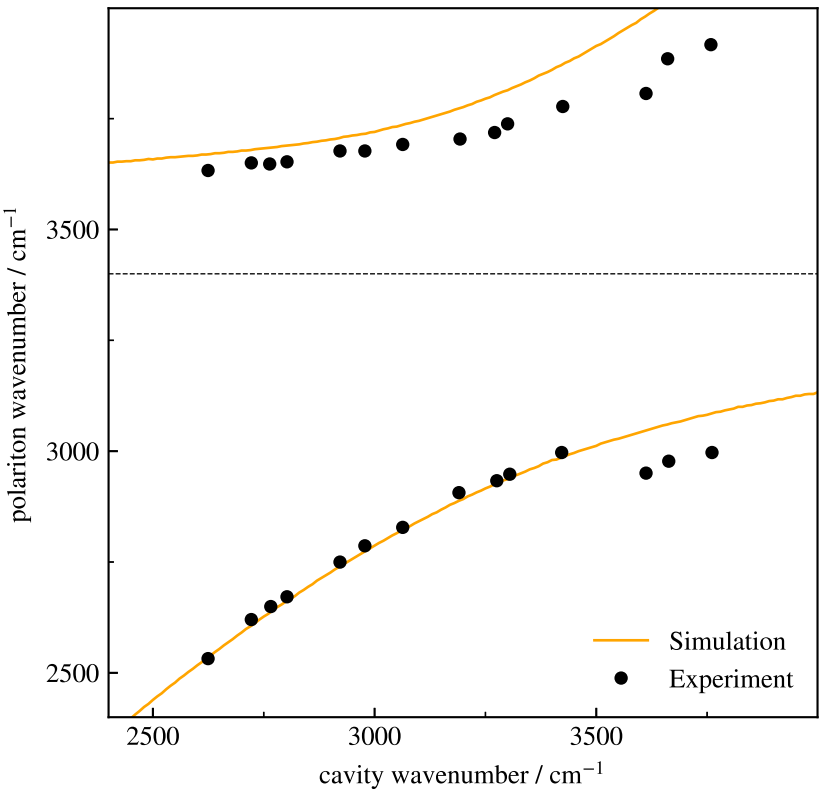

Figure 3.7: Comparison of experimental and simulated polariton peak positions at different cavity frequencies for a cavity filled with liquid water. The cavity-free O-H stretching frequency (3,400 cm^{-1}) is indicated by a dotted line.

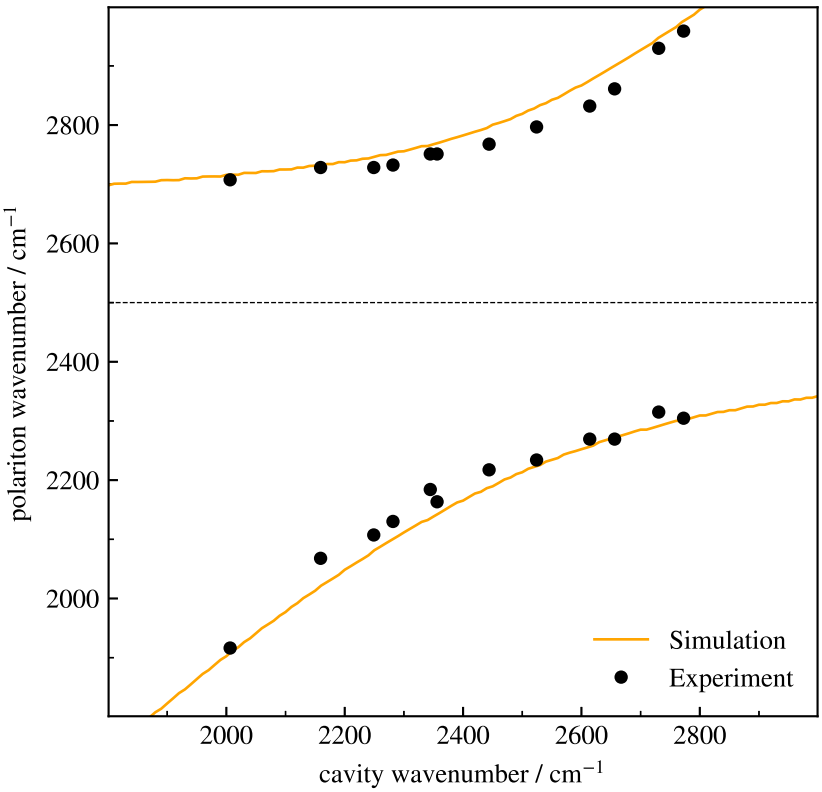

Figure 3.8: As Figure 3.7 but for deuterated water. The dashed line indicates the cavity-free O-D stretch frequency (2,500 cm^{-1}).

e.g., the data points for $\omega_c > 3,500\ \mathrm{cm}^{-1}$).

Luckily, Fukushima *et al.* repeated their experiments with deuterated water, so we can test the SHO model on a second system.[110] The cavity-free deuterated water experimental spectrum is also available on a frequency grid.[111] Equipped with this, it is easy to compute SHO results and compare them to Fukushima *et al.*'s measurements, which we present in Figure 3.8. The agreement between our predictions and experiments is significantly improved compared to Figure 3.7. They are now near perfect. While we cannot exclude artefacts due to the number of cavity modes included in our simulation, our results show that the SHO model provides a convenient and cheap way to predict cavity spectral measurements.

3.3.5 Equivalence to the transfer matrix method

The SHO model uses the geometry of the cavity (to determine R), and the cavity-free infrared spectrum as input. It turns out that this is equivalent to the transfer matrix method,[112,113] which is well established and has frequently been used by experimentalists to verify their cavity measurements.[37,39,42,45,46] The transfer matrix method uses as input the complex refractive index $N(\omega) = n(\omega) + i\kappa(\omega)$. $n(\omega)$ is the frequency dependent refractive index and $\kappa(\omega)$ is the extinction coefficient. The cavity-free spectrum can be written as $I(\omega) = n(\omega)\alpha(\omega)$, where $\alpha(\omega) = \frac{2\omega\kappa(\omega)}{c}$ is the absorption coefficient. We can relate $N(\omega)$ and $I(\omega)$ by noting that the refractive index and the extinction coefficient are related by Kramers-Kronig relations.[114] This allows us to obtain both $n(\omega)$ and $\kappa(\omega)$ from the spectrum $I(\omega)$, for example via the following iterative procedure. Starting from $n^{(0)}(\omega) = n_\infty$, where $n_\infty = \lim_{\omega\to\infty} n(\omega)$ is the material's refractive index, we update the extinction coefficient and refractive index in the jth iteration as

$$\kappa^{(j+1)}(\omega) = \frac{cI(\omega)}{2\omega n^{(j)}(\omega)}, \tag{3.21a}$$

$$n^{(j+1)}(\omega) = \mathrm{Re}\left[n_\infty + \frac{2}{\pi}\int_0^\infty \frac{\omega'\kappa^{(j+1)}(\omega')}{\omega'^2 - \omega^2 + i\varepsilon}\mathrm{d}\omega'\right], \tag{3.21b}$$

where ε is a small number. Adding $i\varepsilon$ to the denominator shifts the poles away from the real axis, so that we can easily do the integral numerically. Equation (3.21b) is one of the Kramers-Kronig relations, and Equation (3.21a) ensures that we reproduce a given vibrational

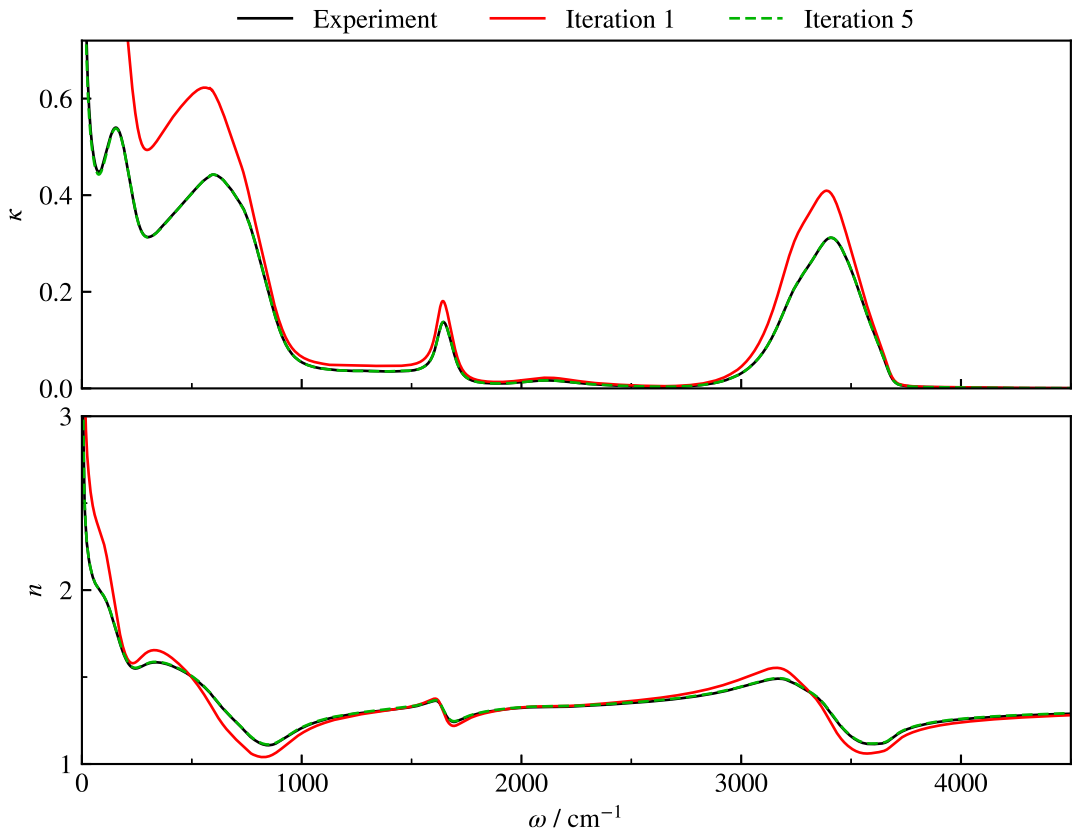

Figure 3.9: Comparison of extinction coefficient (top panel) and refractive index (bottom panel) obtained from the iterative procedure outlined in Equations (3.21) and the text to experiments.

spectrum. For liquid water, the spectrum $I(\omega)$, the refractive index $n(\omega)$ and the extinction coefficient $\kappa(\omega)$ are all available on frequency grids with a high accuracy,[109] making it the ideal test for this approach. As Figure 3.9 shows, after only five iterations the predicted $n^{(5)}(\omega)$ and $\kappa^{(5)}(\omega)$ agree up to graphical accuracy with experiments.

The SHO model and the transfer matrix method are only equivalent if we consider just a single material. A cavity can have many additional layers, like the mirrors and the silica coatings. The transfer matrix method can easily be generalised for any cavity setup, whereas this is more difficult to do with the SHO model.

3.4 Summary

In this chapter, we applied f-QCMD to liquid water inside a cavity. This system had previously been studied with classical MD and TRPMD by Li *et al.*[43,48] They observed a red shift and a broadening of the polariton peaks in the TRPMD spectrum compared to the spectrum obtained with classical MD. Using our newly developed f-QCMD method, we were able to show that the broadening is an artefact of TRPMD and not a real nuclear quantum effect. The red shift, however, persists in the f-QCMD polariton spectra, and it is a true nuclear quantum effect. We also observe anharmonic red shifts in the cavity-free spectrum, as we discussed extensively in Chapter 2. There is thus nothing particularly surprising about the cavity spectrum. We then computed cavity f-QCMD spectra for cavity setups where a single cavity mode is resonant with either the O-H stretch, the bend, or the libration bands, as well as a multi-mode cavity. In all of those, our cavity f-QCMD simulations produce results that are consistent with experimental observations.

We then presented the SHO model to predict cavity spectra. The model discretises the cavity-free spectrum into uncoupled one-dimensional harmonic oscillators, which are subsequently coupled to the cavity. The total Hamiltonian for this system can be diagonalised in a matter of minutes to predict a cavity spectrum, whereas an f-QCMD simulation can take days. Using the cavity-free f-QCMD spectrum as input, we have shown that this SHO model reproduces all of our cavity f-QCMD simulations to graphical accuracy. It is also much more flexible, since we can use any cavity-free spectrum as input – in particular, we can input the experimental spectrum. We can therefore make experimental predictions with the SHO model for a variety of cavity setups. Water in a cavity has been extensively studied,[38,45,110] and experimental data was available for us to compare with our predictions.[110] Our SHO model agrees well with cavity experiments. The agreement is better in deuterated water than liquid water, which could either be due to our neglect of higher frequency cavity modes or to artefacts in the experiments.

The striking simplicity of the SHO model together with its accurate predictions of cavity spectra casts some doubt on other polariton experiments. The effects we observe in the linear vibrational spectra are captured entirely without invoking any chemical or structural

changes to the contents of the cavity. As Schwennicke *et al.* showed, the cavity is nothing but an optical filter for the linear spectrum.[115] They could reproduce polariton spectra without the need for a cavity, but with the use of a filtered laser. However, they point out that to describe some processes like spontaneous Rahman scattering, a more sophisticated quantum electrodynamics treatment is needed.

On the one hand, it could therefore be argued that cavity effects will only become noticeable in higher-order measurements, such as two-dimensional infrared spectra.[102,116] On the other hand, there have already been reports about issues with reproducing cavity experiments.[45,46] While it may be possible to observe cavity effects in a pumped cavity, it seems unlikely to us that claims of vacuum light field effects on chemical reaction rates in liquid samples will ever find an acceptable explanation.

We closed this chapter by pointing out that the SHO model is equivalent to the transfer matrix method of applied optics,[112,113] which is much better established and more flexible. Given that cavity MD results can be predicted with the SHO model, and that the SHO model input is equivalent to that of the transfer matrix method, there is really nothing that cavity MD has to offer for equilibrium simulations of the spectrum. However, since its introduction, cavity MD has also been used to explore non-equilibrium cavity effects.[117] In this context it is more powerful than the transfer matrix method and it could potentially reveal some real effects of the cavity on the system it contains.

Chapter 4

Electronic spectroscopy

4.1 Introduction

Within the wide range of spectroscopic methods, two-dimensional spectroscopy has emerged as a particularly useful experimental tool. It offers insight into nonadiabatic dynamics, including energy and charge transfer.[50,116,118] This poses a challenge to theory, since nonadiabatic processes necessitate the simulation of both nuclear and electronic degrees of freedom simultaneously, as opposed to the previous Chapters 2 and 3, where we could largely ignore the electronic component of the problem by approximating it with the q-TIP4P/F force field for liquid water. A zoo of different propagation methods and approximations exists to describe nonadiabatic dynamics. In applications to large biomolecules such as the ones we are interested in, semiclassical approaches, where the system is separated into a classical bath and a quantum subsystem, have emerged as the most popular approaches. Often, the quantum subsystem comprises all electronic degrees of freedom, and the nuclei are represented by the classical bath. Semiclassical methods can be differentiated by how they evaluate the effect of the electronic degrees of freedom on the nuclear bath.

The Ehrenfest method, for example, employs a mean field approach and averages the nuclear force over the current electronic state.[119] As a consequence, it is not able to describe branching of the wave packet when two electronic states cross. Further, while it gives good results at short times, Ehrenfest converges to an overheated equilibrium, and is therefore generally unable to describe long-time population dynamics.[120] The most common alterna-

tive to Ehrenfest is trajectory surface hopping, where nuclei evolve on a single adiabatic state and hop stochastically when two adiabatic states approach each other.[121] Since the hops are done stochastically, they are not reversible. Furthermore, the electronic wave function and nuclear forces can become inconsistent, resulting in the violation of detailed balance.[122,123] Surface hopping also cannot appropriately describe decoherence, which occurs in wave packet branching. Various decoherence corrections have been developed to account for this, but they have to be introduced ad hoc.[124–128] As a further alternative, mapping methods map the electronic states onto classical variables, such that nuclei and electrons are treated on the same footing.[129–131] The force on the nuclei is averaged over the mapping degrees of freedom. Although they are generally better suited for population dynamics than Ehrenfest,[132] mapping approaches are still mean field and therefore cannot describe wave packet splitting. The field had thus been left in a somewhat unsatisfying place, with none of the above methods being able to both capture accurate population dynamics and describe wave packet branching. A promising way to unify the benefits of mapping and surface hopping approaches has been recently introduced with the mapping approach to surface hopping (MASH).[122,123,133] In MASH, the nuclei evolve on a single electronic surface at any one time, as in surface hopping, but hops are decided deterministically based on the electronic mapping variables, recovering detailed balance.

Although we will keep the other methods in mind, here we will first focus on methods based on Ehrenfest dynamics. This is primarily motivated by a desire to minimise computational cost: Ehrenfest is the least expensive of the above methods since there is no need to sample the electronic wave function (unlike in mapping) and there is no stochasticity to the evolution (unlike in surface hopping). We can also typically use a larger time step than in surface hopping, because there are no abrupt jumps in the force. Even though we expect the results to be affected by the fact that Ehrenfest converges to an overheated equilibrium , we anticipate that this will have a minor effect at shorter times and that Ehrenfest will therefore reproduce the major features of the spectra reliably.

Ehrenfest has previously been used to investigate two-dimensional electronic spectra by calculating the dynamics of a coherence to obtain third order response functions. Van der

Vegte *et al.* first applied Ehrenfest dynamics to the simulation of electronic spectra.[69] They used a mean classical path approximation, where the nuclei evolve under an Ehrenfest force induced by the states that appear in the coherence being propagated. Their method agrees well with exact results obtained with HEOM for systems where the classical bath relaxes slowly. We are however interested in systems with baths which may have shorter relaxation times. Atsango *et al.* followed up with a different formulation and considered a wider range of parameters.[70] They decomposed coherences into four pure states, which were then propagated according to the wave function formulation of Ehrenfest.[119] While their approach performs well in model systems with a slow relaxation time, it deteriorates severely as the relaxation time decreases.

After a review of the theory of electronic spectroscopy in Section 4.2 and Frenkel exciton models in Section 4.3, this chapter provides an in-depth analysis of Atsango *et al.*'s pure state Ehrenfest approach in Section 4.4. We reveal that their choice of 'polar' pure states give rise to unphysical contributions to the quantum subsystem's evolution. However, we also propose and analyse an alternative, 'equatorial', set of pure states in Sections 4.5-4.8, for which no such unphysical contributions occur, and for which we can incidentally also reduce the computational effort by a factor 32. Section 4.9 summarises the main results.

4.2 Theory

The Hamiltonian for a one-dimensional system irradiated by light is

$$\hat{\mathcal{H}}(t) = \hat{\mathcal{H}}_{\text{mol}} + \mathcal{H}_{\text{light}}(t) + \hat{\mathcal{H}}_{\text{int}}(t), \tag{4.1}$$

where $\hat{\mathcal{H}}_{\text{mol}}$ is the system Hamiltonian, $\mathcal{H}_{\text{light}}(t)$ is the classical Hamiltonian of the light field, and $\hat{\mathcal{H}}_{\text{int}}$ is the interaction Hamiltonian. This Hamiltonian is closely related to the cavity Hamiltonian in Chapter 3 (Equation (3.1)). One main difference is that the total Hamiltonian is now an operator since we are dealing with multiple electronic states. In Chapter 3, the cavity Hamiltonian has been projected onto the ground state.[47]

The interaction with the light's electric field E is

$$\hat{\mathcal{H}}_{\text{int}}(t) = \hat{\mu}E(r,t), \tag{4.2}$$

where we have applied the dipole approximation in assuming that the dipole moment operator $\hat{\mu}$ does not depend on the position r and already in Equation (4.1) by assuming the light to be classical. It is often useful to decompose the dipole moment operator as $\hat{\mu} = \hat{\mu}_+ + \hat{\mu}_-$, where $\hat{\mu}_+$ is an excitation operator and $\hat{\mu}_-$ is a de-excitation operator.

The dynamics induced by light-matter interactions are probed via the polarisation

$$P(t) = \text{Tr}[\hat{\mu}\hat{\rho}(t)], \tag{4.3}$$

where $\hat{\rho}(t)$ is the density matrix and satisfies the Liouville-von Neumann equation,

$$\frac{\partial}{\partial t}\hat{\rho}(t) = -i[\hat{\mathcal{H}}(t), \hat{\rho}(t)]. \tag{4.4}$$

(Note that $\hbar = 1$ here.) The polarisation is a macroscopic quantity that captures electric field fluctuations related to absorption, stimulated emission and other responses of the system to the light, as we will describe shortly. Unfortunately, Equation (4.4) is rather difficult to solve generally. One would usually try to find an eigenbasis for the Hamiltonian, within which the time evolution is trivial (this approach is also known as exact diagonalisation).[134] But the size of the eigenbasis grows exponentially with the number of particles in a system, which makes diagonalising the Hamiltonian reminiscent of searching for needles in a haystack. Even worse, the eigenbasis may change with time, since the light-matter Hamiltonian in Equation (4.1) is time-dependent, and the diagonalisation may have to be updated every few time steps. It is not hard to see that this quickly becomes unfeasible and that a compromise is needed.

Instead of attempting to find the full polarisation, we instead expand it perturbatively up to nth order in the light-matter interaction $\hat{H}_{\text{int}}$.

$$P(t) \approx P^{(1)}(t) + P^{(2)}(t) + \cdots + P^{(n)}(t), \tag{4.5}$$

where

$$P^{(k)}(t) = \int_0^\infty \mathrm{d}t_k \cdots \int_0^\infty \mathrm{d}t_1 S^{(k)}(t_k, t_{k-1}, \ldots, t_1) E(r, t - t_k) \cdots E(r, t - t_k \cdots - t_1). \quad (4.6)$$

Equation (4.6) is a result from time-dependent perturbation theory. For the sake of conciseness it is stipulated here, but it is also derived in detail in textbooks, for example in Reference [118]. The kth order response function is

$$S^{(k)}(t_k, t_{k-1}, \ldots, t_1) = i^k \Theta(t_1) \cdots \Theta(t_k) \mathrm{Tr}[\hat{\mu} \mathcal{G}(t_k)[\hat{\mu}, \mathcal{G}(t_{k-1})[\hat{\mu}, \cdots \mathcal{G}(t_1)[\hat{\mu}, \hat{\rho}_0]]]], \quad (4.7)$$

where $\mathcal{G}(t)(\cdot) = e^{-i\hat{\mathcal{H}}t}(\cdot)e^{i\hat{\mathcal{H}}t}$ is the time propagator and t_i is the time between the ith and $(i+1)$th light-matter interaction. All t_i are non-negative to respect causality. This is clear from the integration boundaries in Equation (4.6), but is additionally imposed by the Heaviside functions $\Theta(t_i)$ in Equation (4.7). Note in particular that the dipole moment operator appears $k+1$ times in the kth order response function. The first k appearances correspond to interactions between the system and the light field. The final dipole moment operator, however, arises from the trace in Equation (4.3) and hence is of different character.

4.2.1 Linear spectroscopy

The linear response function is[118]

$$S^{(1)}(t) = i\Theta(t)\mathrm{Tr}[\hat{\mu}\mathcal{G}(t)[\hat{\mu}, \hat{\rho}_0]] = -2\Theta(t)\mathrm{Im}[\hat{\mu}\mathcal{G}(t)(\hat{\mu}\hat{\rho}_0)], \quad (4.8)$$

and the linear spectrum is its Fourier transform,[118]

$$I(\omega) = \mathrm{Im} \int_0^\infty e^{-i\omega t} S^{(1)}(t) \mathrm{d}t. \quad (4.9)$$

Only the imaginary part of the Fourier transform is related to the extinction coefficient and contributes to the spectrum.[118]

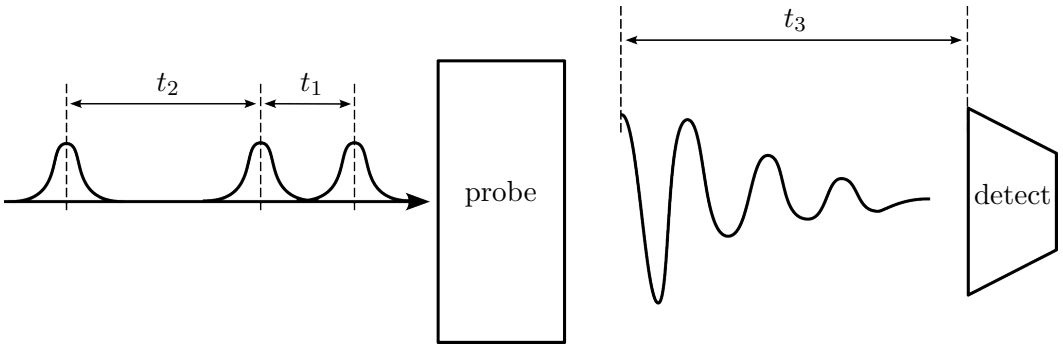

Figure 4.1: Experimental setup for time-domain two-dimensional spectroscopy. Three staggered ultrashort pulses interact with the probe, with delay times t_1 and t_2 between them. After the interaction with the light pulses, the probe emits a signal, which is detected after a third delay time t_3. We assume here that all laser pulses have the same incident angle.

4.2.2 Two-dimensional spectroscopy

In time-domain two-dimensional spectroscopy, three very short pulses interact with the system (see Figure 4.1). The corresponding third order response function is[118]

$$S^{(3)}(t_1,t_2,t_3) = -i\Theta(t_1)\Theta(t_2)\Theta(t_3)\mathrm{Tr}[\hat{\mu}\mathcal{G}(t_3)[\hat{\mu},\mathcal{G}(t_2)[\hat{\mu},\mathcal{G}(t_1)[\hat{\mu},\rho_0]]]]. \tag{4.10}$$

The three commutators expand to a total of eight terms, such that

$$S^{(3)}(t_1,t_2,t_3) = -i\sum_{j=1}^{4}[R_j(t_1,t_2,t_3) - R_j^*(t_1,t_2,t_3)], \tag{4.11}$$

where

$$R_1(t_1,t_2,t_3) = \mathrm{Tr}[\hat{\mu}\mathcal{G}(t_3)(\mathcal{G}(t_2)(\mathcal{G}(t_1)(\hat{\mu}\rho_0)\hat{\mu})\hat{\mu})], \tag{4.12a}$$

$$R_2(t_1,t_2,t_3) = \mathrm{Tr}[\hat{\mu}\mathcal{G}(t_3)(\mathcal{G}(t_2)(\hat{\mu}\mathcal{G}(t_1)(\rho_0\hat{\mu}))\hat{\mu})], \tag{4.12b}$$

$$R_3(t_1,t_2,t_3) = \mathrm{Tr}[\hat{\mu}\mathcal{G}(t_3)(\hat{\mu}\mathcal{G}(t_2)(\mathcal{G}(t_1)(\rho_0\hat{\mu})\hat{\mu}))], \tag{4.12c}$$

$$R_4(t_1,t_2,t_3) = \mathrm{Tr}[\hat{\mu}\mathcal{G}(t_3)(\hat{\mu}\mathcal{G}(t_2)(\hat{\mu}\mathcal{G}(t_1)(\hat{\mu}\rho_0)))]. \tag{4.12d}$$

The separate terms look very complicated at first sight, but we can greatly simplify them with the rotating wave approximation (RWA).

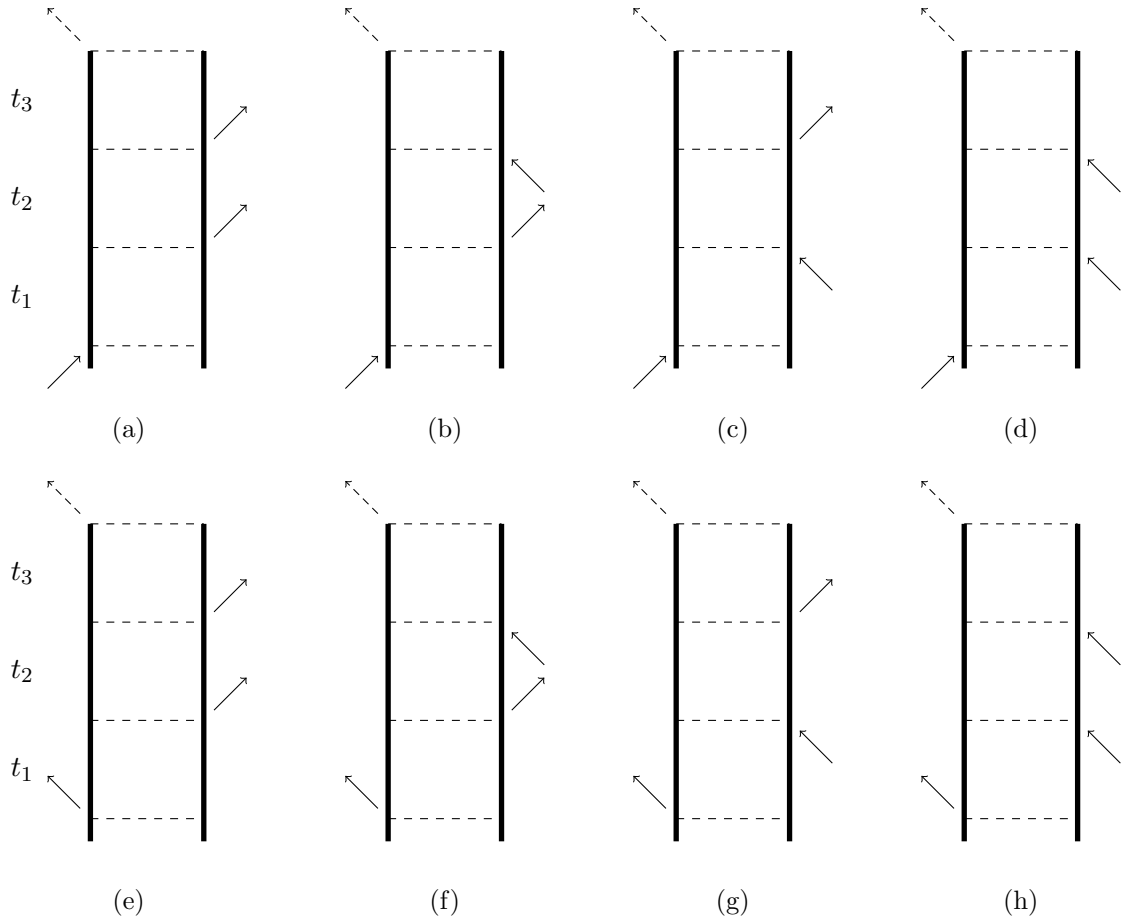

Figure 4.2: The eight possible Feynman diagrams that arise from the three electric field interactions in the response function R_1. The arrows are placed on the side which the relevant dipole moment operator acts from. They also denote the direction of the electric field propagation (see text). The final arrow is drawn with dashed lines, since it corresponds to the measurement of the polarisation (Equation (4.3)). Further, it corresponds to the emission of the spectroscopic signal, and therefore points away from the system.

The rotating wave approximation

Figure 4.2 shows the possible Feynman diagrams for the response function R_1. The arrows represent the electric fields in Equation (4.3). An arrow pointing to the right represents a right-propagating field, $E_i(r,t) = E_i(t)e^{i(k_i r - \omega_i t)}$, while an arrow pointing to the left represents a left-propagating field, $E_i(r,t) = E_i(t)e^{i(-k_i r + \omega_i t)}$. The RWA allows us to establish a direct relationship between the electric fields and electronic excitations. Let us first focus

only on the t_1 integral in Equation (4.6). The contribution of the response function R_1 is

$$\cdots \int_0^\infty \mathrm{d}t_1 R_1(t_3, t_2, t_1) E(r, t - t_3 - t_2 - t_1) \cdots . \tag{4.13}$$

Let us now consider what happens to the initial condition $\hat{\mu}\rho_0$ (see Equation (4.12a)) during the evolution through t_1. Since $\rho_0 = |g\rangle\langle g|$ is the ground state density matrix, $\hat{\mu}$ can only induce an excitation. For the simplicity of the argument, let us assume that there is only one excited state $|s\rangle$ with a fixed energy and write the energy difference between $|g\rangle$ and $|s\rangle$ as ω_{sg}. The coherence $|s\rangle\langle g|$ accumulates a phase during the t_1 evolution,

$$\mathcal{G}(t_1)(\hat{\mu}\rho_0) \propto e^{-i\omega_{sg}t_1}|s\rangle\langle g|. \tag{4.14}$$

We can write the integral over t_1 as

$$\begin{aligned}
&\int_0^t \mathcal{G}(t_1)(\hat{\mu}\rho_0) E(r, t - t_3 - t_2 - t_1)\mathrm{d}t_1 \\
&\qquad \propto \int_0^t \left(e^{i(kr - \omega(t - t_1))} + e^{-i(kr - \omega(t - t_1))} \right) e^{-i\omega_{sg}t_1}|s\rangle\langle g|\mathrm{d}t_1 \\
&\qquad = \left[e^{i(kr-\omega t)} \int_0^t e^{-i(\omega_{sg} - \omega)t_1}\mathrm{d}t_1 + e^{-i(kr-\omega t)} \int_0^t e^{-i(\omega_{sg} + \omega)t_1}\mathrm{d}t_1 \right] |s\rangle\langle g|.
\end{aligned} \tag{4.15}$$

For simplicity, the first step in Equation (4.15) assumes a very simple form of the electric field, $E(r,t) = e^{i(kr-\omega t)} + e^{-i(kr-\omega t)}$ and we have dropped the terms containing $e^{\pm i\omega(t_3+t_2)}$ for clarity of the final expression. (The dropped terms would appear as prefactors together with $e^{i(k_1 r - \omega_1 t)}$.) Equation (4.15) should be seen as a proof of concept, but a more detailed discussion of the RWA may be found in References [118] and [135].

The integrand $e^{-i(\omega_{sg}+\omega)t_1}$ oscillates very quickly compared to the integrand $e^{-i(\omega_{sg}-\omega)t_1}$. During the integration it will effectively cancel out to zero. We can therefore neglect it, such that only the first integral in Equation (4.15) survives, which arose from the right-propagating term of the electric field. This is the RWA. Since the first electric field interaction in the response function R_1 acts on the left of the system density matrix (see Figure 4.2), only the first row of Feynman diagrams survives when we invoke the RWA for the t_1 evolution. We can analyse the t_2 and t_3 evolution in the same way,[135] where we would come to the conclusion

that only the Feynman diagram in Figure 4.2c survives, and that there is only one possible path for the response function R_1 through the excitation manifolds that survives the RWA.

Consider instead the response function R_2, with the initial condition $\rho_0\hat{\mu}$. Because the dipole moment operator now acts from the right, the phase accumulated during the t_1 evolution has opposite sign, with

$$\mathcal{G}(t_1)(\rho_0\hat{\mu}) \propto e^{i\omega_{sg}t_1}|g\rangle\langle s|. \tag{4.16}$$

If we insert this into the polarisation, we make the exact opposite conclusion as we did for R_1. The sign in the exponential in Equation (4.14) would be inverted, which means that only the left-propagating term of the electric field survives within the RWA (and analogously for the t_2 and t_3 evolution).

These observations can be generalised to the following rules for Feynman diagrams: Under the RWA, arrows that point *towards* the diagrams must correspond to an excitation, while arrows pointing *away* from the diagram must correspond to a de-excitation. At the end, only six Feynman diagrams survive out of the total eight response functions in Equation (4.12). Operators that are associated with arrows on the right side of the Feynman diagram act on the left (i.e. on the bra state in the density matrix). An excitation in the bra state is thus caused by a $\hat{\mu}_-$ operator. Conversely, a $\hat{\mu}_+$ operator causes a de-excitation in the bra state. This observation can be reformulated into a second rule, which says that a left-pointing arrow is associated with the operator $\hat{\mu}_-$, while a right-pointing arrow is associated with the operator $\hat{\mu}_+$.

The correspondence between arrows in the Feynman diagrams with excitations and de-excitations of the system allows us to greatly simplify the response functions in Equation (4.12). Response functions within the RWA will be denoted with the Greek letter 'Φ', and we will adjust the numbering to agree with literature that is relevant to this work.[70,136] Using the excitation and de-excitation operators μ_+ and μ_-, these are

$$R_2 \rightsquigarrow \Phi_1 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\mathcal{G}(t_2)[\hat{\mu}_+\mathcal{G}(t_1)(\rho_0\hat{\mu}_-)]\hat{\mu}_+]], \tag{4.17a}$$

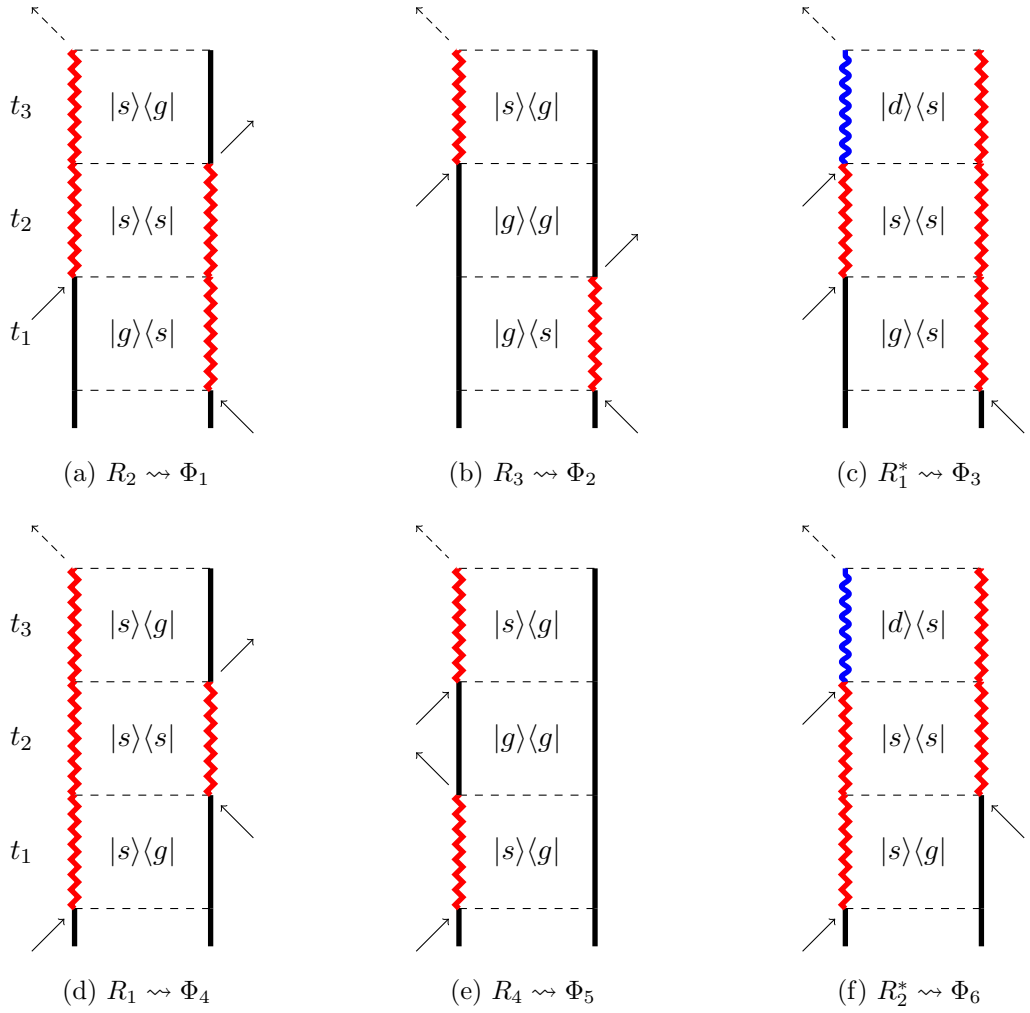

(a) $R_2 \rightsquigarrow \Phi_1$ (b) $R_3 \rightsquigarrow \Phi_2$ (c) $R_1^* \rightsquigarrow \Phi_3$

(d) $R_1 \rightsquigarrow \Phi_4$ (e) $R_4 \rightsquigarrow \Phi_5$ (f) $R_2^* \rightsquigarrow \Phi_6$

Figure 4.3: Feynman diagrams of the nonlinear response functions that survive the RWA. The caption indicates which response function from Equation (4.12) each diagram originates from, and which response function in Equation (4.17) it corresponds to. The colours indicate the active excitation manifold of the ket and bra states of the density matrix: g is the ground state (black, straight), s is the singly excited (red, zigzag) and d is the doubly excited (blue, wavy) manifold. Note that in contrast to Figure 4.2, each arrow only corresponds to either a $\hat{\mu}_+$ or $\hat{\mu}_-$ operator.

$$R_3 \rightsquigarrow \Phi_2 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\hat{\mu}_+\mathcal{G}(t_2)[\mathcal{G}(t_1)(\rho_0\hat{\mu}_-)\hat{\mu}_+]]], \tag{4.17b}$$

$$R_1^* \rightsquigarrow \Phi_3 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\hat{\mu}_+\mathcal{G}(t_2)[\hat{\mu}_+\mathcal{G}(t_1)(\rho_0\hat{\mu}_-)]]], \tag{4.17c}$$

$$R_1 \rightsquigarrow \Phi_4 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\mathcal{G}(t_2)[\mathcal{G}(t_1)(\hat{\mu}_+\rho_0)\hat{\mu}_-]\hat{\mu}_+]], \tag{4.17d}$$

$$R_4 \rightsquigarrow \Phi_5 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\hat{\mu}_+\mathcal{G}(t_2)[\hat{\mu}_-\mathcal{G}(t_1)(\hat{\mu}_+\rho_0)]]], \tag{4.17e}$$

$$R_2^* \rightsquigarrow \Phi_6 = \mathrm{Tr}[\hat{\mu}_-\mathcal{G}(t_3)[\hat{\mu}_+\mathcal{G}(t_2)[\mathcal{G}(t_1)(\hat{\mu}_+\rho_0)\hat{\mu}_-]]], \tag{4.17f}$$

all of which depend on the delay times (t_1, t_2, t_3). Figure 4.3 shows the corresponding Feynman diagrams. The path of the initial density matrix through the excitation manifolds is now well-defined. In particular, we note that during t_2, the system is within a single excitation manifold (either the ground state or the singly excited manifold) whereas during t_1 and t_3 it is in a coherence, which is especially easy to see from the Feynman diagrams in Figure 4.3. This is why t_2 is sometimes also referred to as the population time.

These response functions can be related to experimental signals. Φ_1-Φ_3 constitute the rephasing response function,

$$R_{\mathrm{rp}} = \Phi_1 + \Phi_2 - \Phi_3, \tag{4.18}$$

and Φ_4-Φ_6 the non-rephasing response function,

$$R_{\mathrm{nr}} = \Phi_4 + \Phi_5 - \Phi_6, \tag{4.19}$$

where the terms 'rephasing' and 'non-rephasing' refer to the direction in which the signal is emitted. For the rephasing terms, the wave vector resulting from the first three interactions with the electric field is $k_{\mathrm{rp}} = -k_1 + k_2 + k_3$ whereas for the non-rephasing terms it is $k_{\mathrm{nr}} = k_1 - k_2 + k_3$, which can be seen from Figure 4.3. If k_{rp} and k_{nr} point in different directions, then the signal could in principle be measured separately by placing a detector along the different directions. Here, we assume that all laser pulses are collinear (as shown in Figure 4.1), in which case rephasing and non-rephasing signals emit in the same direction. In such a setup, the signals cannot be discerned experimentally, although we can of course always separate them in a simulation if we calculate each of the six response functions in Equation (4.17)

separately. The Feynman diagrams can also be related to physical processes: Φ_1 and Φ_4 correspond to stimulated emission, Φ_2 and Φ_5 correspond to ground state bleaching, and Φ_3 and Φ_6 correspond to excited state absorption.

Obtaining two-dimensional spectra

We can obtain the spectrum from the rephasing and non-rephasing response functions as[137,138]

$$S(\omega_3, t_2, \omega_1) = \mathrm{Re} \int_0^\infty \mathrm{d}t_1 \int_0^\infty \mathrm{d}t_3 e^{i\omega_3 t_3} \left[e^{i\omega_1 t_1} R_{\mathrm{nr}}(t_1, t_2, t_3) + e^{-i\omega_1 t_1} R_{\mathrm{rp}}(t_1, t_2, t_3) \right]. \quad (4.20)$$

Pump-probe spectra (also called transient absorption spectra) are also third order measurements. They correspond to the case where $t_1 = 0$, and are easily obtained once we calculate the response functions R_{nr} and R_{rp} by integrating $S(\omega_3, t_2, \omega_1)$ over t_1.

4.3 Frenkel exciton models

To assess the accuracy of our methods, we will study Frenkel exciton model Hamiltonians,

$$\hat{\mathcal{H}}(\hat{\boldsymbol{p}}, \hat{\boldsymbol{q}}) = \begin{pmatrix} \hat{H}^g(\hat{\boldsymbol{q}}) & 0 & 0 \\ 0 & \hat{H}^s(\hat{\boldsymbol{q}}) & 0 \\ 0 & 0 & \hat{H}^d(\hat{\boldsymbol{q}}) \end{pmatrix} + \hat{H}_{\mathrm{B}}(\hat{\boldsymbol{p}}, \hat{\boldsymbol{q}}), \quad (4.21)$$

with a harmonic bath of M oscillators for each of the N electronic sites,

$$\hat{H}_{\mathrm{B}}(\hat{\boldsymbol{p}}, \hat{\boldsymbol{q}}) = \sum_{j=1}^{M} \sum_{n=1}^{N} \left(\frac{\hat{p}_{jn}^2}{2} + \frac{1}{2} \omega_j^2 \hat{q}_{jn}^2 \right). \quad (4.22)$$

$\hat{\mathcal{H}}(\hat{\boldsymbol{p}}, \hat{\boldsymbol{q}})$ is block-diagonal in the ground, singly excited and doubly excited manifolds, with respective Hamiltonians $\hat{H}^g$, $\hat{H}^s$ and $\hat{H}^d$. The ground state energy ε_0 is fixed at zero, and therefore $\hat{H}^g = \varepsilon_0 |0\rangle\langle 0| = 0$. The Hamiltonian of the singly excited manifold is

$$\hat{H}^s(\hat{\boldsymbol{q}}) = \hat{H}_{\mathrm{S}}^s + \hat{H}_{\mathrm{SB}}^s(\hat{\boldsymbol{q}}), \quad (4.23)$$

where

$$\hat{H}_{\mathrm{S}}^{s} = \sum_{n=1}^{N} \varepsilon_n |n\rangle\langle n| + \sum_{m\neq n}^{N} J_{nm} |n\rangle\langle m|, \tag{4.24}$$

describes the subsystem, and

$$\hat{H}_{\mathrm{SB}}^{s}(\hat{\boldsymbol{q}}) = \sum_{j=1}^{M}\sum_{n=1}^{N} c_j \hat{q}_{jn} |n\rangle\langle n|, \tag{4.25}$$

is the system-bath coupling term. In Equations (4.24) and (4.25), ε_n is the site energy, and J_{nm} is the coupling between sites n and m. Doubly excited states are simultaneous excitations of two different electronic sites. Since it does not matter which order they are excited in, the states $|nm\rangle$, $m > n$, can serve as a basis for $\hat{H}^d$, such that

$$\hat{H}^d(\hat{\boldsymbol{q}}) = \hat{H}_{\mathrm{S}}^{d} + \hat{H}_{\mathrm{SB}}^{d}(\hat{\boldsymbol{q}}), \tag{4.26}$$

with

$$\begin{aligned} \hat{H}_{\mathrm{S}}^{d} = &\sum_{m>n}^{N} (\varepsilon_n + \varepsilon_m) |nm\rangle\langle nm| \\ &+ \left[\sum_{\ell>m>n}^{N} \left(J_{m\ell} |nm\rangle\langle n\ell| + J_{n\ell} |nm\rangle\langle m\ell| + J_{nm} |n\ell\rangle\langle m\ell| \right) + \text{c.c.} \right], \end{aligned} \tag{4.27}$$

and

$$\hat{H}_{\mathrm{SB}}^{d}(\hat{\boldsymbol{q}}) = \sum_{m>n}^{N}\sum_{j=1}^{M} c_j (\hat{q}_{jn} + \hat{q}_{jm}) |nm\rangle\langle nm|. \tag{4.28}$$

The dipole moment raising operator couples the ground state to the singly excited manifold, and the singly excited manifold to the doubly excited manifold. The coupling between the ground state and the singly excited states is

$$\hat{\mu}_{+}^{sg} = \sum_{n=1}^{N} \mu_n |n\rangle\langle 0|, \tag{4.29}$$

and the coupling between the single and double excitation manifolds is

$$\hat{\mu}_+^{ds} = \sum_{m>n}^{N} \left(\mu_m |nm\rangle\langle n| + \mu_n |nm\rangle\langle m|\right). \tag{4.30}$$

The dipole moment raising operator is thus

$$\hat{\mu}_+ = \begin{pmatrix} 0 & 0 & 0 \\ \hat{\mu}_+^{sg} & 0 & 0 \\ 0 & \hat{\mu}_+^{ds} & 0 \end{pmatrix}, \tag{4.31}$$

and the full dipole moment operator is $\hat{\mu} = \hat{\mu}_+ + \hat{\mu}_-$, with $\hat{\mu}_- = (\hat{\mu}_+)^\dagger$.

For simplicity, we use identical baths for all sites in Equations (4.22), (4.25) and (4.28), which is why c_j does not have a site subscript. We also assume that their frequencies follow a Debye power spectrum,

$$J(\omega) = \frac{2\lambda\omega_c\omega}{\omega^2 + \omega_c^2} \tag{4.32}$$

where $\lambda = \frac{1}{\pi}\int_0^\infty \frac{J(\omega)}{\omega}\mathrm{d}\omega$ is the reorganisation energy and ω_c is the bath's characteristic frequency. Using a Debye bath allows us to obtain exact results with HEOM,[68] which is instrumental in benchmarking the methods described further below.

In order to treat the bath spectral density numerically, we discretise it into frequencies ω_j and couplings c_j,

$$J(\omega) \approx \frac{\pi}{2}\sum_{j=1}^{M} \frac{c_j^2}{\omega_j}\delta(\omega - \omega_j). \tag{4.33}$$

There are different ways of choosing the parameters ω_j, c_j, but a particularly elegant one is to choose them based on the quadrature rule for the reorganisation energy integral.[139,140] To this end, consider the reorganisation energy

$$\lambda = \frac{1}{\pi}\int_0^\infty \frac{J(\omega)}{\omega}\mathrm{d}\omega = \frac{2\lambda}{\pi}\int_0^\infty \frac{1}{1+x^2}\mathrm{d}x, \tag{4.34}$$

where $x = \frac{\omega}{\omega_c}$. The integral over x can either be solved by contour integration or by substi-

tuting $x = \tan\theta$. The substitution renders a trivial integral,

$$\int_0^\infty \frac{1}{1+x^2}\mathrm{d}x = \int_0^{\pi/2} \mathrm{d}\theta, \tag{4.35}$$

which is solved by the midpoint rule with M quadrature points, $\theta_j = \frac{(j-\frac{1}{2})\pi}{2M}$, $j = 1, \ldots, M$. The points θ_j determine the frequency grid $\omega_j = \omega_c \tan\theta_j$. For the weights c_j, we require that the quadrature must reproduce the correct reorganisation energy and that all modes contribute equally to the reorganisation energy, such that

$$\lambda = \frac{1}{2}\sum_{j=1}^{M} \frac{c_j^2}{\omega_j^2} \tag{4.36}$$

and

$$\frac{c_j^2}{\omega_j^2} = \frac{2\lambda}{M}. \tag{4.37}$$

To summarise, the discretised frequencies ω_j and c_j are

$$\omega_j = \omega_c \tan\frac{(j-\frac{1}{2})\pi}{2M}, \tag{4.38a}$$

$$c_j = \sqrt{\frac{2\lambda}{M}}\omega_j. \tag{4.38b}$$

Figure 4.4 shows examples of the quadrature points ω_j for $M = 1, 5, 10$ overlaid with the original Debye spectrum. If M is odd, the $\frac{M+1}{2}$th point is placed at the maximum of the spectral density (at $\omega = \omega_c$). The remaining quadrature points are distributed evenly between the left and right sides of the maximum, although it should be noted that the quadrature points are spaced much further apart on the right side, due to the long tail of $J(\omega)$. If M is even, there is no quadrature point at ω_c, but the quadrature points are still split evenly on either side of ω_c.

4.4 Existing simulation methods

Frenkel exciton models are a classic example of an open quantum system, where an electronic subsystem is coupled to a nuclear bath. Sometimes it is feasible to treat these systems

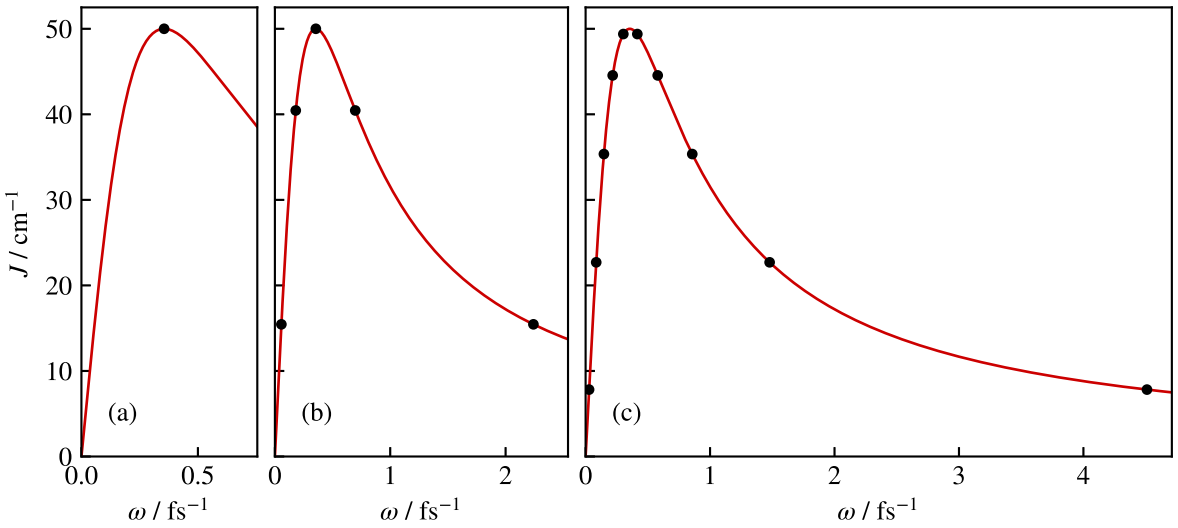

Figure 4.4: Quadrature points for the discretisation of a Debye bath with (a) $M = 1$, (b) $M = 5$ and (c) $M = 10$ bath modes. The bath has a relaxation time $\tau_c = 17.7$ fs (where $\tau_c = \frac{2\pi}{\omega_c}$) and reorganisation energy $\lambda = 50\ \mathrm{cm}^{-1}$.

fully quantum mechanically. In the case of Frenkel exciton models, choosing a Debye power spectrum (Equation (4.32)) for the nuclear bath allows us to use the HEOM method,[68] which in principle gives exact results. However, HEOM is limited in its applicability. Like any exact method its computational cost increases significantly with system size. Semiclassical methods pose a viable alternative: They treat the nuclear degrees of freedom classically, such that only a much smaller subsystem, usually containing the electronic degrees of freedom, needs a quantum treatment. To this end, the bath operators $\hat{\boldsymbol{p}}$ and $\hat{\boldsymbol{q}}$ are replaced by classical momenta and positions $\boldsymbol{p}$ and $\boldsymbol{q}$, which evolve under the equations of motion

$$\begin{aligned} \dot{\boldsymbol{p}} &= F(\boldsymbol{q}, c), \\ \dot{\boldsymbol{q}} &= \boldsymbol{p}. \end{aligned} \tag{4.39}$$

The nuclear force F may depend on the nuclear positions and the electronic wave function $|c\rangle$. The wave function $|c\rangle$ evolves under the Schrödinger equation,

$$|\dot{c}\rangle = -i\hat{\mathcal{H}}|c\rangle. \tag{4.40}$$

We further approximate the quantum trace by a mixed quantum-classical trace,

$$\mathrm{Tr}[\cdot] \approx \frac{1}{(2\pi)^{NM}} \int \mathrm{d}\boldsymbol{p} \int \mathrm{d}\boldsymbol{q}\ \mathrm{Tr}_{\mathrm{el}}[\cdot], \tag{4.41}$$

where $\mathrm{Tr}_{\mathrm{el}}$ is the trace over the electronic degrees of freedom only. The initial condition is a mixed quantum-classical density matrix,

$$\hat{\rho}_0(\boldsymbol{p}, \boldsymbol{q}) = |0\rangle \rho_{\mathrm{B}}(\boldsymbol{p}, \boldsymbol{q}) \langle 0|, \tag{4.42}$$

where

$$\rho_{\mathrm{B}}(\boldsymbol{p}, \boldsymbol{q}) = \prod_{j,n} \frac{\sigma_j}{\pi} \exp\left[-\frac{2\sigma_j}{\omega_j} \left(\frac{p_{jn}^2}{2} + \frac{\omega_j^2 q_{jn}^2}{2} \right) \right] \tag{4.43}$$

is the Wigner distribution of the bath modes, and $\sigma_j = \tanh\left(\frac{\beta\omega_j}{2}\right)$.[141] The Wigner distribution accounts for the additional spread of the nuclear wave packet compared to the classical case. The simulations from here on proceed similarly to the vibrational spectroscopy simulations described in Chapters 2 and 3, except for one key difference. Previously, all dynamics happened on the ground state potential energy surface. With the multiple electronic states present in the Frenkel exciton model, we need to specify how the electronic wave function impacts the nuclear dynamics. We can thus characterise different semiclassical methods by how they approximate the nuclear force.

In the Ehrenfest method, the nuclear force is averaged over the electronic degrees of freedom. The force is

$$F(\boldsymbol{q}) = -\langle c | \nabla_{\boldsymbol{q}} \hat{\mathcal{H}}(\boldsymbol{p}, \boldsymbol{q}) | c \rangle. \tag{4.44}$$

for a normalised wave function $|c\rangle$ and

$$F(\boldsymbol{q}) = -\mathrm{Tr}[\nabla_{\boldsymbol{q}} \hat{\mathcal{H}}(\boldsymbol{p}, \boldsymbol{q}) \hat{\rho}] \tag{4.45}$$

for a density matrix $\hat{\rho}$. The density matrix formulation of Ehrenfest can be derived as a mean field approximation to the quantum-classical Liouville equation.[142] While it is an extremely popular method, it is not without its flaws. For example, the Ehrenfest equations of motion

do not conserve the true Boltzmann distribution,

$$\rho(\boldsymbol{p}, \boldsymbol{q}, c) \propto \langle c|e^{-\beta\hat{\mathcal{H}}}|c\rangle, \tag{4.46}$$

but instead conserve,

$$\rho_{\mathrm{Ehr}}(\boldsymbol{p}, \boldsymbol{q}, c) \propto e^{-\beta\langle c|\hat{\mathcal{H}}|c\rangle}. \tag{4.47}$$

The two distributions are generally not identical, which is why Ehrenfest performs poorly for population dynamics except in the short-time limit (which is often referred to as the 'overheating problem').[120] Further, Ehrenfest cannot describe wave packet branching: Imagine two potential energy surfaces that cross at a point $q^{\ddagger}$. A wave packet approaching the crossing point from the left on the lower potential energy surface will split into two wave packets upon crossing $q^{\ddagger}$, which evolve separately on each potential energy surface. In Ehrenfest, however, the wave packet instead evolves on an average of the two potential energy surfaces.

To address the overheating problem in Ehrenfest, mapping approaches[129,131,143] and trajectory surface hopping[121] have been developed. In mapping approaches, the wave function $|c\rangle$ is mapped onto a classical phase space, such that both nuclear and electronic degrees of freedom are consistently treated classically.[129,131,143] They are generally better suited to calculating population dynamics than Ehrenfest, but struggle to describe wave packet branching just like Ehrenfest.

Conversely, in surface hopping the nuclear force is always evaluated on a single adiabatic state (the 'active' state).[121] The active state is updated stochastically throughout the simulation through 'hops' between adiabatic states. Surface hopping thus has no problem describing wave packet branching: it is built in by design. However, the active state in surface hopping and the state that actually has the highest population in the wave function can become inconsistent. As a result, surface hopping does not produce accurate population dynamics, and may give wrong equilibrium populations.

Recently, a mapping approach to surface hopping (MASH) has been introduced that promises to combine the strengths of mapping and surface hopping approaches.[136,144] In MASH, the wave function is mapped onto a classical phase space, but the force is determined

by a single adiabatic surface. Hops between adiabatic states happen deterministically based on the mapping variables. By virtue of its similarity to surface hopping, it is easy to see that it can describe wave packet branching. Further, it has been shown to work very well for describing population dynamics[136,144] and chemical rate constants.[133]

Here, we will use Ehrenfest, because it is simple to implement and we anticipate it to converge very quickly with respect to the total trajectory number. This was our primary concern when choosing an appropriate method for the present work. We want to develop a practical method which can be applied to a wide range of systems. Since the third order response functions are calculated on three-dimensional time grids, the cost for a single trajectory is already considerable, not to mention the cost of a fully converged calculation. We are aware of the Ehrenfest method's shortcomings, and will discuss the results accordingly below. However, its low computational expense compared to mapping, surface hopping and MASH tipped the scale in its favour here: there is no need to sample the electronic wave function (in contrast to mapping methods) and there is no stochasticity in its evolution (in contrast to surface hopping). The instantaneous hops between the adiabatic states also usually require a smaller time step in surface hopping methods and MASH than in Ehrenfest, which would further increase the cost of such calculations.

Although Ehrenfest is at face value very simple, its application to calculate the two-dimensional spectrum is not trivial. The initial condition in the nonlinear response functions is either $\hat{\mu}_+\rho_0$ or $\rho_0\hat{\mu}_-$ (see Equation (4.17)). Both of these objects are not proper density matrices since they have a trace of zero and cannot be propagated with the Ehrenfest equations of motion.[142] Another way to see this is by noting that the Ehrenfest force (Equation (4.45)) induced by a coherence $|a\rangle\langle b|$ is imaginary (or zero, if the Hamiltonian is real), which is clearly unphysical. (In our case, the force induced by $\hat{\mu}_+\rho_0$ or $\rho_0\hat{\mu}_-$ is always zero since the Hamiltonian is block-diagonal in the excitation manifolds.) Two workarounds have been proposed: First, van der Vegte *et al.* developed a mean classical path approximation to the evolution of a coherence.[69] They average the force between the coherence's bra and ket states, such that the force on the nuclei induced by a coherence $|a\rangle\langle b|$ is

$$F_{\mathrm{MCP}} = -\frac{1}{2}(\langle a|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|a\rangle + \langle b|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|b\rangle), \tag{4.48}$$

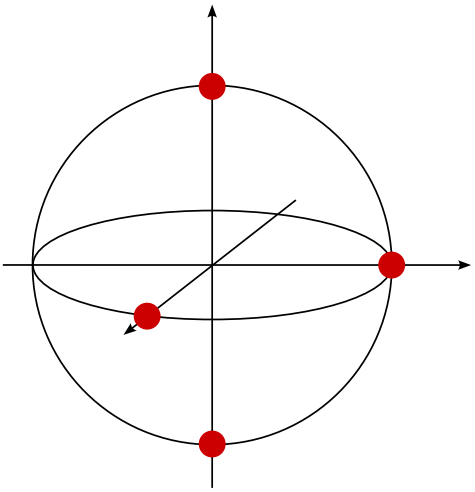

Figure 4.5: Bloch sphere representation of the polar pure states.

assuming that both $|a\rangle$ and $|b\rangle$ are normalised wave functions. They showed that their method agrees very well with exact HEOM results for a Frenkel biexciton model ($N = 2$) with relaxation time $\tau_c = 220\,\text{fs}$, but did not explore other parameter regimes. Second, Atsango *et al.* proposed to decompose a coherence $|a\rangle\langle b|$ into four pure states,[70]

$$|a\rangle\langle b| = \sum_{j=0}^{3} w_j |j\rangle\langle j|, \tag{4.49}$$

which are then propagated with the Ehrenfest force in Equation (4.44). They chose the pure states to be $|j\rangle \equiv |\psi_j\rangle$, with

$$|\psi_0\rangle = |a\rangle, \qquad w_0^{(p)} = -\frac{1+i}{2}, \tag{4.50a}$$

$$|\psi_1\rangle = |b\rangle, \qquad w_1^{(p)} = -\frac{1+i}{2}, \tag{4.50b}$$

$$|\psi_2\rangle = \frac{1}{\sqrt{2}}(|a\rangle + |b\rangle), \qquad w_2^{(p)} = 1, \tag{4.50c}$$

$$|\psi_3\rangle = \frac{1}{\sqrt{2}}(|a\rangle + i|b\rangle), \qquad w_3^{(p)} = i. \tag{4.50d}$$

Because of the placement of the states on the Bloch sphere (see Figure 4.5), we shall refer to this as the polar decomposition.

If $|a\rangle$, $|b\rangle$ are normalised then so are the pure states $|\psi_i\rangle$, and they can directly be used to evaluate the Ehrenfest force (Equation (4.44)). However, consider the initial condition of

a rephasing response function, where $|a\rangle \equiv |0\rangle$ and $|b\rangle \equiv |\mu\rangle$, and $|\mu\rangle = \mu_+|0\rangle$. While the ground state is normalised, we cannot say the same about the excited state superposition $|\mu\rangle$, which has norm $\langle\mu|\mu\rangle = \sum_n \mu_n^2$. As a remedy, Atsango *et al.* normalise the pure states *a posteriori* before evaluating the Ehrenfest force. However, there are reasons to believe that this is not the best procedure. For example, it is known that in the limit of uncoupled states, averaging the dynamics evenly between the ground and excited states reduces to the exact solution. This is also known as the Wigner averaged classical limit.[145,146] In the polar decomposition, if the norm $\langle\mu|\mu\rangle \gg 1$, the excited states dominate the Ehrenfest force and the dynamics effectively takes place entirely in the singly excited manifold (except for the dynamics of $|\psi_0\rangle$, which has no contributions from $|\mu\rangle$). As a result, polar Ehrenfest fails to reproduce the Wigner averaged classical limit.

Further, polar Ehrenfest carries a significant computational cost. The decomposition in Equation (4.50) is invoked every time a dipole operator acts on the system, which happens three times for each of the non-linear response functions in Equation (4.17). The first dipole interaction generates four pure states, which are propagated through t_1. Each of these four states is acted on by the second dipole operator, and for each of them the resulting term is decomposed into four further states. We therefore need to propagate a total of 16 states through t_2, and – by the same logic – a total of 64 states through t_3. At the very end, the response function is obtained by tracing over all pure state density matrices with the final dipole de-excitation operator and summing over the pure states with the appropriate weights. Atsango *et al.* argue that their approach is less *ad hoc* than the mean classical path approach, and that they thus accept a steep increase in the computational cost.[70] We shall reach a different conclusion below.

Atsango *et al.* applied the polar Ehrenfest method to two Frenkel biexciton models with $\tau_c = 300$ fs and $\tau_c = 17.7$ fs, respectively.[70] The method worked very well for the first model, but showed serious problems in the second model, where the bath relaxes on a fast time scale. Here, their two-dimensional spectra exhibit spurious oscillations that worsen with longer delay times t_2. They did not find an explanation for their observations.

To provide an explanation, we need to analyse the dynamics induced by the polar decom-

position. Figure 4.6 shows the evolution of the wave function during the calculation of Φ_1. The figure not only shows the exponentially increasing size of the calculation with each dipole interaction, but it also highlights a more fundamental issue with the polar decomposition. From the Feynman diagram (Figure 4.3a), we already know exactly what the wave function evolution for Φ_1 is supposed to look like: Since it starts off in the ground state, the system should be in a coherence between ground and singly excited states during t_1. During t_2, it should evolve purely in the singly excited manifold, and during t_3 it returns to a coherence between ground and singly excited states. The wave function dynamics generated by the polar decomposition paint a very different picture. In all time intervals there are wave function contributions that we do not expect based on the Feynman diagram. This is especially striking in the t_2 evolution, where there are contributions from the ground and doubly excited states. These could be the reason behind the spurious oscillations that Atsango *et al.* saw in their two-dimensional spectrum.[70] We will now demonstrate that these oscillations can be eliminated with a simple modification to the polar Ehrenfest method.

4.5 An improved pure state Ehrenfest approach

The polar decomposition is only one of many decompositions that satisfy Equation (4.49). Instead, we propose an alternative decomposition with pure states $|j\rangle \equiv |\phi_j\rangle$,

$$|\phi_j\rangle = \frac{1}{\sqrt{2}}(|\overline{a}\rangle + e^{ij\frac{\pi}{2}}|\overline{b}\rangle), \qquad w_j^{(e)} = \frac{N_{ab}}{2}e^{ij\frac{\pi}{2}}, \tag{4.51}$$

where $|\overline{a}\rangle = c_a^{-\frac{1}{2}}|a\rangle$, $|\overline{b}\rangle = c_b^{-\frac{1}{2}}|b\rangle$ and $N_{ab} = (c_a c_b)^{\frac{1}{2}}$. The normalisation coefficients are $c_a = \langle a|a\rangle$ and $c_b = \langle b|b\rangle$. We call this the equatorial decomposition, since the pure states lie on the equator of the Bloch sphere (see Figures 4.5 and 4.7 for a representation of polar and equatorial pure states on the Bloch sphere, respectively). We have normalised the states $|a\rangle$, $|b\rangle$ before building the pure states $|\phi_j\rangle$ instead of normalising the states *a posteriori* as in the polar decomposition. This way, the equatorial decomposition is more reminiscent of the Wigner averaged classical limit than the polar decomposition, and the pure states are

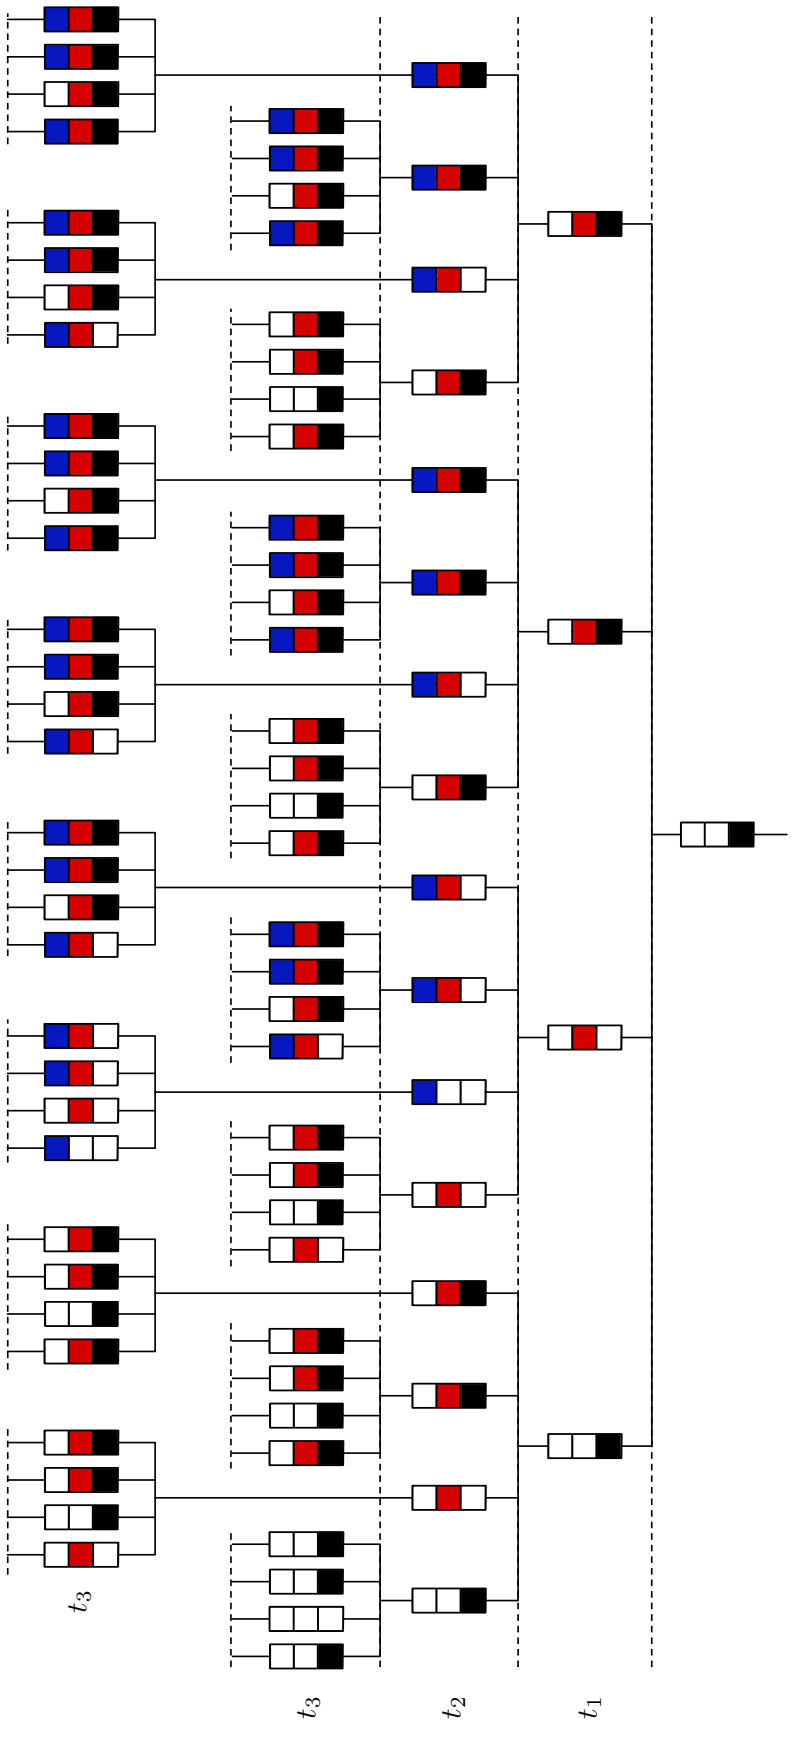

Figure 4.6: Wave function dynamics during the calculation of Φ_1 with polar Ehrenfest. There are three relevant excitation manifolds, so we represent each wave function by three boxes whose colours match the ones in the Feynman diagrams in Figure 4.3. The boxes are coloured in if there are non-zero coefficients in the corresponding manifold. The 64 states that are propagated through t_3 do not fit comfortably on a single row, and are instead presented in a staggered manner.

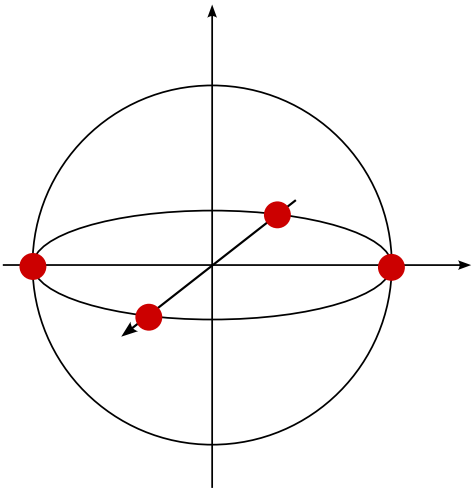

Figure 4.7: Bloch sphere representation of the equatorial pure states.

invariant to overall scaling of the dipole moments. The weights $w_j^{(e)}$ satisfy

$$\sum_{j=0}^{3} w_j^{(e)} = \sum_{j=0}^{3} w_j^{(e)} e^{ij\frac{\pi}{2}} = 0, \tag{4.52}$$

and

$$\sum_{j=0}^{3} w_j^{(e)} e^{-ij\frac{\pi}{2}} = 2N_{ab}. \tag{4.53}$$

If the Hamiltonian is block-diagonal and $|\overline{a}\rangle$ and $|\overline{b}\rangle$ belong to different manifolds, such that $\langle\overline{a}|\hat{\mathcal{H}}|\overline{b}\rangle = 0$, the pure state evolution turns out to be particularly simple. In this case, the Ehrenfest force

$$\begin{aligned} F_j(\boldsymbol{q}) &= -\langle\phi_j|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|\phi_j\rangle \\ &= -\frac{1}{2}(\langle\overline{a}|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|\overline{a}\rangle + \langle\overline{b}|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|\overline{b}\rangle) \equiv F(\boldsymbol{q}) \end{aligned} \tag{4.54}$$

is independent of j and all pure states follow identical nuclear trajectories. This force is actually identical to the force within the mean classical path approximation, since Equation (4.48) assumes the states $|a\rangle$, $|b\rangle$ to be normalised. Our method thus sits somewhere between polar Ehrenfest and the mean classical path approximation.

Note that the equatorial Ehrenfest method is not entirely equivalent to the Wigner averaged classical limit. The dipole moment raising operator $\hat{\mu}_+$ generates the state $|\mu\rangle$, which is a superposition of all excited states. It would only be equivalent to the Wigner averaged

classical limit if there were only one state in the excited manifold, but in general we propagate a superposition. Still, we expect that this will be more accurate than the polar approach.

Based on Equation (4.54), we can easily write down the time evolution of the equatorial pure states,

$$|\phi_j(t)\rangle = \frac{1}{\sqrt{2}}(|\overline{a}(t)\rangle + e^{ij\frac{\pi}{2}}|\overline{b}(t)\rangle), \tag{4.55}$$

where $|\overline{x}(t)\rangle = \mathcal{T}e^{-i\int_0^t \mathcal{H}(\boldsymbol{q}_\tau)\mathrm{d}\tau}|\overline{x}\rangle$, $\boldsymbol{q}_\tau$ evolves under the force $F(\boldsymbol{q}_\tau)$ from Equation (4.54) and $\mathcal{T}$ is the time ordering operator. This observation, combined with Equations (4.52) and (4.53), allows us to evaluate the pure state sum after a propagation explicitly,

$$\mathcal{G}(t)|a\rangle\langle b| = \sum_{j=0}^{3} w_j^{(e)}|\phi_j(t)\rangle\langle\phi_j(t)| = |a(t)\rangle\langle b(t)|. \tag{4.56}$$

In particular, we only need to propagate one of the four pure states because we can project out $|a(t)\rangle$ and $|b(t)\rangle$ using projection operators on their respective manifolds. For example,

$$|a(t)\rangle = \sqrt{2c_a}P_a|\phi_0(t)\rangle, \tag{4.57a}$$

$$|b(t)\rangle = \sqrt{2c_b}P_b|\phi_0(t)\rangle, \tag{4.57b}$$

where P_a and P_b are the projection operators on the a and b manifolds, respectively.

Let us now consider how the resummation impacts the calculation of the response functions. During the t_1 and t_3 time evolution, the density matrix is in a coherence between excitation manifolds, and we can resum the pure states after the propagation. As a consequence, we can invoke Equation (4.57), and reduce the cost of the t_1 and t_3 evolution by a factor of four each. During t_2, the dynamics occur within a single state manifold. Here, we cannot apply the resummation trick and just accept that we need to propagate four pure states. However, there is a different property that we can take advantage of: In the equatorial decomposition, the coherences $|a\rangle\langle b|$ and $|b\rangle\langle a|$ give rise to the same four pure states but with complex conjugated weights, such that

$$|b\rangle\langle a| = \sum_{j=0}^{3} w_j^{(e)*}|\phi_j\rangle\langle\phi_j|, \tag{4.58}$$

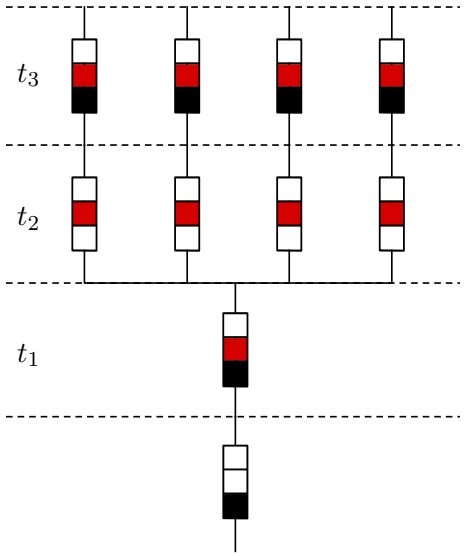

Figure 4.8: As Figure 4.6, but for the equatorial decomposition.

where $|\phi_j\rangle$ are the pure states from Equation (4.51). For each of the non-rephasing response functions, the object propagated through t_2 is exactly the adjoint of the one propagated in the rephasing analogue. This is easiest to see by comparing the column-wise pairs of Feynman diagrams in Figure 4.3. The non-rephasing response functions can therefore be obtained for free during the calculation of the rephasing ones, saving us an additional factor 2. Overall, we save a factor 32 compared to polar Ehrenfest (but our approach is still about twice as expensive as the mean classical path approach of van der Vegte *et al.*[69]). The wave function dynamics now also agree with what we would expect from the Feynman diagrams (Figure 4.8).

We will next derive an expression for the response functions Φ_3 (and thereby also for Φ_6 due to Equation (4.58)) in detail. All remaining response functions can be derived by following the same steps, so we will keep their discussion brief. We start by propagating the state $|\phi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |\overline{\mu}\rangle)$ through t_1, where $|\overline{\mu}\rangle = c_\mu^{-\frac{1}{2}}|\mu\rangle$ is a superposition of excited states, with $c_\mu = \sum_n \mu_n^2$. Because $|0\rangle$ and $|\overline{\mu}\rangle$ belong to different manifolds, we can invoke Equation (4.56), by which

$$\mathcal{G}(t_1)(\rho_0\mu_-) = |0(t_1)\rangle\langle\mu(t_1)| \tag{4.59}$$

and

$$|0(t_1)\rangle = \sqrt{2}P_g|\phi(t_1)\rangle, \tag{4.60a}$$

$$|\mu(t_1)\rangle = \sqrt{2c_\mu}P_s|\phi(t_1)\rangle, \tag{4.60b}$$

where $P_g = |0\rangle\langle 0|$ is the projection on the ground state and $P_s = \sum_{n=1}^{N} |n\rangle\langle n|$ is the projection on the singly excited manifold. The nuclear force is

$$F(\boldsymbol{q}) = -\frac{1}{2}(\langle 0|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|0\rangle + \langle\overline{\mu}|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|\overline{\mu}\rangle), \tag{4.61}$$

following Equation (4.54). Next, the dipole moment operator acts on Equation (4.59) and the result is decomposed into four new pure states,

$$\mu_+\mathcal{G}(t_1)(\rho_0\mu_-) = \mu_+|0(t_1)\rangle\langle\mu(t_1)| = \sum_{j=0}^{3} w_{s,j}^{(e)}|\phi_{s,j}^{(t_1)}\rangle\langle\phi_{s,j}^{(t_1)}|. \tag{4.62}$$

with

$$|\phi_{s,j}^{(t_1)}\rangle = \frac{1}{\sqrt{2}}\left(\overline{\mu_+|0(t_1)\rangle} + e^{ij\frac{\pi}{2}}\overline{|\mu(t_1)\rangle}\right). \tag{4.63}$$

The subscript 's' emphasises that they belong to the singly excited manifold. Here, we denote the normalisation within the pure states by overlining them, such that

$$\overline{\mu_+|0(t_1)\rangle} = c_\mu^{-\frac{1}{2}}\mu_+|0(t_1)\rangle, \tag{4.64a}$$

$$\overline{|\mu(t_1)\rangle} = c_\mu^{-\frac{1}{2}}|\mu(t_1)\rangle. \tag{4.64b}$$

The normalisation constant c_μ is the same for the bra and ket states, and in particular independent of t_1. The weights are $w_{s,j}^{(e)} = \frac{e^{ij\frac{\pi}{2}}}{2}c_\mu$. Note that the pure states $|\phi_{s,j}^{(t_1)}\rangle$ themselves are not normalised, since the states $\mu_+|0(t_1)\rangle$ and $|\mu(t_1)\rangle$ are not orthogonal, so we actually propagate the normalised state $\overline{|\phi_{s,j}^{(t_1)}\rangle}$ through t_2, and multiply the norm back in at the end of the t_2 propagation to obtain $|\phi_{s,j}^{(t_1)}(t_2)\rangle$. The third dipole moment operator then acts on each pure state separately, but for each creates only a single state to propagate through t_3, because we can project out the desired terms at the end by invoking Equation (4.56) a second

time. We thus define

$$|\phi_{\mathrm{esa},j}^{(t_2,t_1)}\rangle = \frac{1}{\sqrt{2}}\left[\overline{\mu_+|\phi_{s,j}^{(t_1)}(t_2)\rangle} + \overline{|\phi_{s,j}^{(t_1)}(t_2)\rangle}\right], \tag{4.65}$$

with the normalised states

$$\overline{\mu_+|\phi_{s,j}^{(t_1)}(t_2)\rangle} = \frac{1}{\sqrt{c_{1,j}^{(t_2,t_1)}}}\mu_+|\phi_{s,j}^{(t_1)}(t_2)\rangle, \tag{4.66a}$$

$$\overline{|\phi_{s,j}^{(t_1)}(t_2)\rangle} = \frac{1}{\sqrt{c_{2,j}^{(t_1)}}}|\phi_{s,j}^{(t_1)}(t_2)\rangle. \tag{4.66b}$$

The subscript 'esa' indicates that the target is an excited state absorption response function and the subscript 'j' indicates which of the four pure states from Equation (4.63) the state originated from. In Equation (4.66), we defined the normalisation constants

$$c_{1,j}^{(t_1,t_2)} = \langle\phi_{s,j}^{(t_1)}(t_2)|\mu_-\mu_+|\phi_{s,j}^{(t_1)}(t_2)\rangle \tag{4.67}$$

and (the t_2-independent)

$$c_{2,j}^{(t_1)} = \langle\phi_{s,j}^{(t_1)}|\phi_{s,j}^{(t_1)}\rangle, \tag{4.68}$$

because we will reuse them below. All that remains is to apply the final dipole moment operator and to trace over the electronic degrees of freedom. The final result for Φ_3 is thus

$$\Phi_3(t_3,t_2,t_1) = \int \mathrm{d}\boldsymbol{p}\int \mathrm{d}\boldsymbol{q}\rho_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q})\sum_{j=0}^{3} w_{s,j}^{(e)} n_{\mathrm{esa},j}^{(t_2,t_1)}\langle\phi_{\mathrm{esa},j}^{(t_2,t_1)}(t_3)|P_s\mu_-P_d|\phi_{\mathrm{esa},j}^{(t_2,t_1)}(t_3)\rangle, \tag{4.69}$$

where P_s and P_d are the projections on the singly and doubly excited manifold, respectively, and $n_{\mathrm{esa},j}^{(t_2,t_1)} = 2\sqrt{c_{1,j}^{(t_2,t_1)}c_{2,j}^{(t_1)}}$ collects the normalisation factors that arise from projecting the final terms out of $|\phi_{\mathrm{esa},j}^{(t_2,t_1)}(t_3)\rangle$ (compare to Equation (4.57)). Φ_6 can be obtained from the same expression as Φ_3 by swapping the weights $w_{s,j}^{(e)}$ with their complex conjugates.

The stimulated emission response function Φ_1 has the same Feynman diagram as Φ_3 up until the third dipole moment interaction, meaning its derivation is completely identical to the discussion above, up to but excluding Equation (4.65). We only need to define new states

for the t_3 propagation,

$$|\phi_{\mathrm{se},j}^{(t_2,t_1)}\rangle = \frac{1}{\sqrt{2}}\left[\,\overline{|\phi_{s,j}^{(t_1)}(t_2)\rangle} + \overline{\mu_-|\phi_{s,j}^{(t_1)}(t_2)\rangle}\,\right], \tag{4.70}$$

with the normalised states

$$\overline{\mu_-|\phi_{s,j}^{(t_1)}(t_2)\rangle} = \frac{1}{\sqrt{c_{3,j}^{(t_2,t_1)}}}\mu_-|\phi_{s,j}^{(t_1)}(t_2)\rangle \tag{4.71}$$

and $\overline{|\phi_{s,j}^{(t_1)}(t_2)\rangle}$ from Equation (4.66b). Note that we have defined an additional normalisation constant

$$c_{3,j}^{(t_2,t_1)} = \langle\phi_{s,j}^{(t_1)}(t_2)|\mu_+\mu_-|\phi_{s,j}^{(t_1)}(t_2)\rangle. \tag{4.72}$$

The final result is then

$$\Phi_1(t_3,t_2,t_1) = \int \mathrm{d}\boldsymbol{p}\int \mathrm{d}\boldsymbol{q}\rho_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q})\sum_{j=0}^{3} w_{s,j}^{(e)} n_{\mathrm{se},j}^{(t_2,t_1)}\langle\phi_{\mathrm{se},j}^{(t_2,t_1)}(t_3)|P_g\mu_-P_s|\phi_{\mathrm{se},j}^{(t_2,t_1)}(t_3)\rangle \tag{4.73}$$

where $n_{\mathrm{se},j}^{(t_2,t_1)} = 2\sqrt{c_{2,j}^{(t_1)}c_{3,j}^{(t_2,t_1)}}$. By swapping the weights with their complex conjugates, we obtain Φ_4.

For Φ_2, the analogue of Equation (4.63) is

$$|\phi_{g,j}^{(t_1)}\rangle = \frac{1}{\sqrt{2}}\left[|0(t_1)\rangle + e^{ij\frac{\pi}{2}}\overline{\mu_-|\mu(t_1)\rangle}\right], \tag{4.74}$$

with

$$\overline{\mu_-|\mu(t_1)\rangle} = \frac{1}{\sqrt{c_4^{(t_1)}}}\mu_-|\mu(t_1)\rangle, \tag{4.75}$$

$c_4^{(t_1)} = \langle\mu(t_1)|\mu_+\mu_-|\mu(t_1)\rangle$ and weights $w_{g,j}^{(e)} = \frac{e^{ij\frac{\pi}{2}}}{2}\sqrt{c_4^{(t_1)}}$. The states in Equation (4.74) are proportional to the ground state, but they are not properly normalised (since $|0(t_1)\rangle$ and $\overline{\mu_-|\mu(t_1)\rangle}$ are not orthogonal). Just as in the derivation of Φ_3, we propagate the normalised state $\overline{|\phi_{g,j}^{(t_1)}\rangle}$ through t_1, and multiply the normalisation constant back in before proceeding

with the t_3 evolution. For each j, the state for the t_3 propagation is

$$|\phi_{\mathrm{gb},j}^{(t_2,t_1)}\rangle = \frac{1}{\sqrt{2}}\left[\overline{\mu_+|\phi_{g,j}^{(t_1)}(t_2)\rangle} + \overline{|\phi_{g,j}^{(t_1)}(t_2)\rangle}\right], \tag{4.76}$$

with the normalised states

$$\overline{\mu_+|\phi_{g,j}^{(t_1)}(t_2)\rangle} = \frac{1}{\sqrt{c_{5,j}^{(t_2,t_1)}}}\mu_+|\phi_{g,j}^{(t_1)}(t_2)\rangle, \tag{4.77a}$$

$$\overline{|\phi_{g,j}^{(t_1)}(t_2)\rangle} = \frac{1}{\sqrt{c_{6,j}^{(t_1)}}}|\phi_{g,j}^{(t_1)}(t_2)\rangle. \tag{4.77b}$$

To this end we needed to define two final normalisation constants,

$$c_{5,j}^{(t_2,t_1)} = \langle\phi_{g,j}^{(t_1)}(t_2)|\mu_-\mu_+|\phi_{g,j}^{(t_1)}(t_2)\rangle \tag{4.78}$$

and

$$c_{6,j}^{(t_1)} = \langle\phi_{g,j}^{(t_1)}|\phi_{g,j}^{(t_1)}\rangle. \tag{4.79}$$

We can now write down the expressions for the ground state bleaching response functions,

$$\Phi_2(t_3,t_2,t_1) = \int \mathrm{d}\boldsymbol{p}\int \mathrm{d}\boldsymbol{q}\rho_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q})\sum_{j=0}^{3} w_{g,j}^{(e)} n_{\mathrm{gb},j}^{(t_2,t_1)}\langle\phi_{\mathrm{gb},j}^{(t_2,t_1)}(t_3)|P_g\mu_-P_s|\phi_{\mathrm{gb},j}^{(t_2,t_1)}(t_3)\rangle \tag{4.80}$$

with $n_{\mathrm{gb},j}^{(t_2,t_1)} = 2\sqrt{c_{5,j}^{(t_2,t_1)}c_{6,j}^{(t_1)}}$ and Φ_5 is obtained by replacing the weights with their complex conjugates.

It would be natural to wonder at this point whether there is another set of pure states that would have other beneficial properties. The key development of this chapter is the summation trick (Equation (4.56)), which is only possible because the pure states lie on the equator of the Bloch sphere and thus follow identical nuclear trajectories. Generally, pure states obtained from any horizontal cut through the Bloch sphere follow identical nuclear trajectories. But equal contributions between ground and excited states only occur along the equator. Without summing the pure states, we would populate the wrong excitation manifolds and produce dynamics that are inconsistent with the Feynman diagrams. While

there are infinitely many pure state decompositions, we believe that the equatorial splitting is optimal.

The final expressions for the response functions in Equations (4.69), (4.73) and (4.80) look complicated because of the various weights and normalisations to keep track of, but they are easy to calculate in a simulation. With the publication of the article based on this chapter, an implementation will be available online.[147] After calculating the response functions, we smoothen them in the t_1 and t_3 dimensions with a Hann window before Fourier transforming them to get the two-dimensional spectra. To assess the accuracy of our approach, we will apply it to two Frenkel exciton models: The fast bath biexciton model previously studied by Atsango *et al.*,[70] and a seven state model of the Fenna-Matthews-Olson light-harvesting complex.[53,138,148]

4.6 Results

4.6.1 Biexciton model

The first model is a biexciton model with two electronic states with site energies $\varepsilon_1 = 50\ \mathrm{cm}^{-1}$, $\varepsilon_2 = -50\ \mathrm{cm}^{-1}$ and site-site coupling $J = 100\ \mathrm{cm}^{-1}$. The transition dipole moments of the two sites are anti-aligned, such that $\mu_1/\mu_2 = -5$. The relaxation time is $\tau_c = 17.7$ fs and the reorganisation energy is $\lambda = 35\ \mathrm{cm}^{-1}$. This is the same system considered by Atsango *et al.*, where polar Ehrenfest showed strong spurious oscillations.[70]

Figure 4.9 compares the exact two-dimensional spectra of the biexciton model to the spectra obtained from polar and equatorial Ehrenfest methods. We are interested in relative peak intensities, and can therefore rescale the spectrum arbitrarily. Here, we choose to normalise each spectrum by its largest absolute magnitude over all delay times t_2. The HEOM results were obtained with `pyrho`,[149] using the convergence parameters $L = 15, K = 0$. It is clear from Figure 4.9 that the equatorial Ehrenfest spectra are smoother than the polar Ehrenfest ones, and are close to the exact result upon initial inspection. This supports our previous hypothesis that the poor performance of polar Ehrenfest can be traced back to the wave function evolution shown in Figure 4.6. But we cannot yet make a quantitative assessment of polar and equatorial Ehrenfest by only considering the two-dimensional spectra.

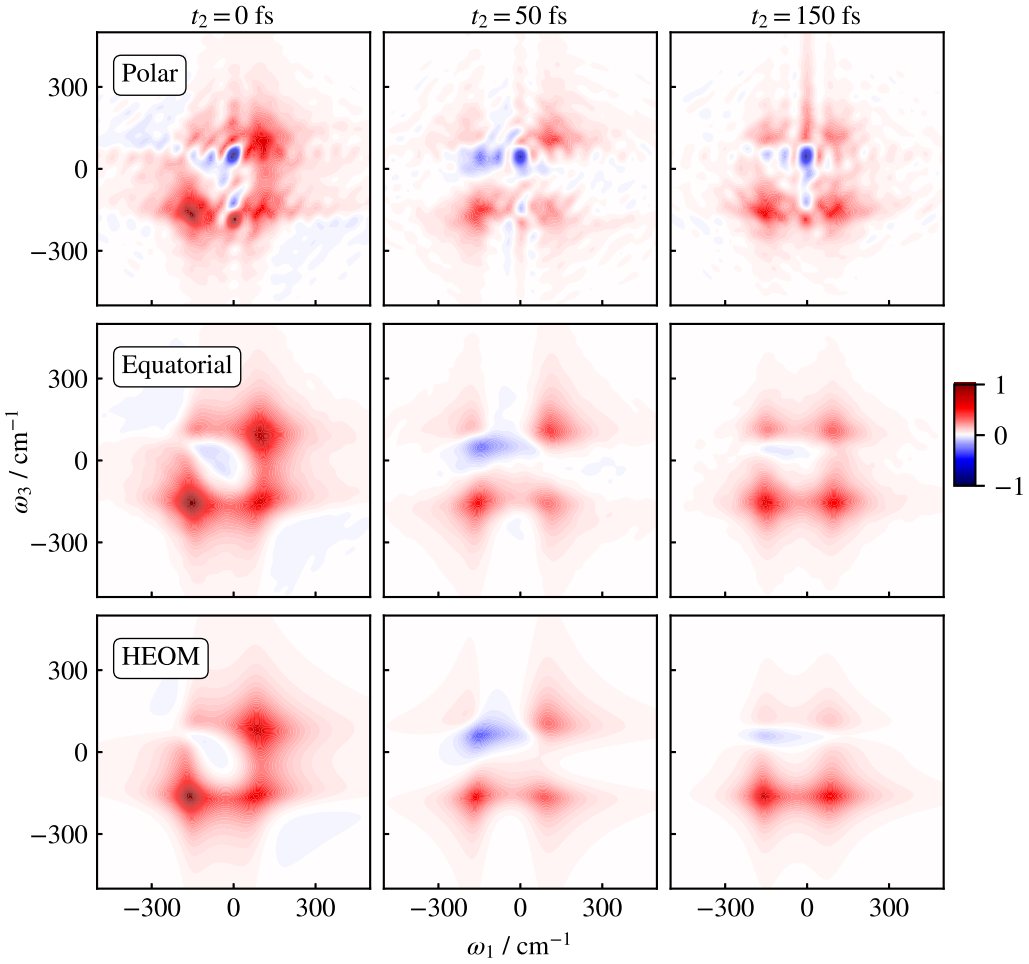

Figure 4.9: Comparison of HEOM, equatorial and polar Ehrenfest two-dimensional electronic spectra for the biexciton model at different delay times.

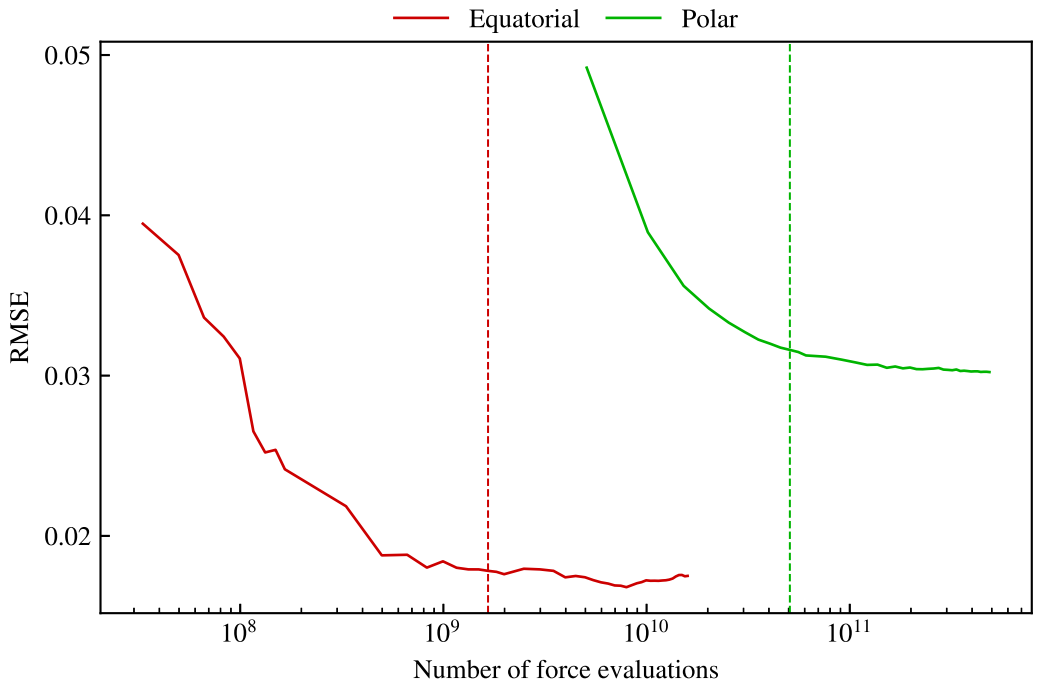

Figure 4.10: RMSE of Ehrenfest spectra with respect to exact HEOM results, shown here against the total number of force evaluations necessary for obtaining all six response functions. Note that only the three rephasing response functions need to be calculated in equatorial Ehrenfest, since the non-rephasing ones can be obtained for free as explained in the text following Equation (4.58). The dashed vertical lines correspond to running 10,000 trajectories for each point (t_1, t_2, t_3) on the three-dimensional time grid.

Instead, we calculate the root mean squared error (RMSE) in the spectrum,

$$\mathrm{RMSE} = \left(\frac{1}{N_{t_2}(\omega_+ - \omega_-)^2} \sum_{t_2=0}^{N_{t_2}} \int_{\omega_-}^{\omega_+} \mathrm{d}\omega_1 \int_{\omega_-}^{\omega_+} \mathrm{d}\omega_3 |I_{\mathrm{HEOM}}(\omega_1, t_2, \omega_3) - I_X(\omega_1, t_2, \omega_3)|^2 \right)^{1/2}. \tag{4.81}$$

N_{t_2} corresponds to the number of t_2 values at which the spectrum is calculated, ω_+ and ω_- correspond to the lowest and highest frequencies at which we obtain the spectrum, and X either refers to polar or equatorial Ehrenfest. Figure 4.10 compares the RMSEs of the Ehrenfest methods. We plot the error against the number of force evaluations because these are the limiting factor in *ab initio* calculations. The equatorial Ehrenfest RMSE is shifted left and downward compared to the polar Ehrenfest RMSE, implying that it is both cheaper and more accurate than polar Ehrenfest for this model. This is fully consistent with expectations: It is cheaper because we need fewer pure states in equatorial Ehrenfest, and we only need

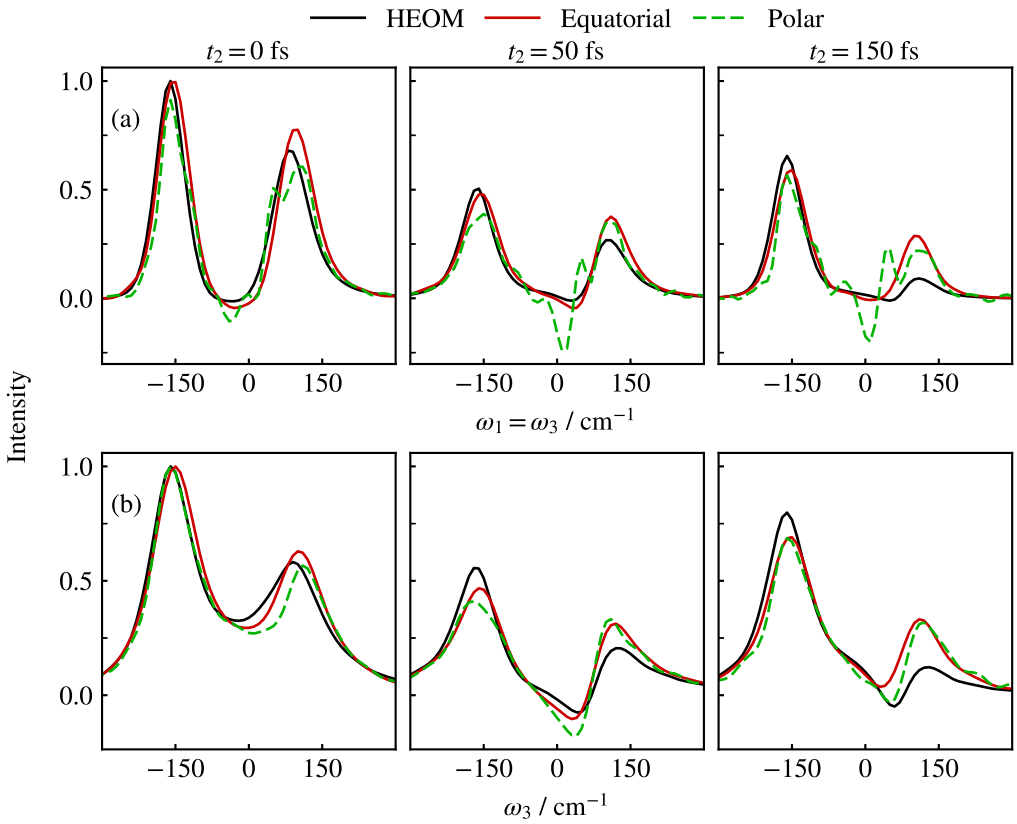

Figure 4.11: (a) Diagonal slices through two-dimensional spectra and (b) pump-probe spectra for the biexciton model.

to calculate the rephasing response functions. It is more accurate because the wave function dynamics are consistent with the Feynman diagram.

To get an even better picture of the quality of the Ehrenfest spectra, we show the spectra along the diagonal frequency slice (where $\omega_1 = \omega_3$) together with the pump-probe spectra, which were calculated by integrating the two-dimensional spectra over ω_1, in Figure 4.11. While the spurious oscillations in polar Ehrenfest are very noticeable in the diagonal slices and cause a large-amplitude negative intensity between the two main peaks, they disappear in the pump-probe spectra where the two Ehrenfest methods are in close agreement. Any remaining disagreement is likely due to the different normalisation conventions of the two methods.

Equatorial Ehrenfest gives smooth spectra throughout, making it more useful for subsequent analyses. However, both Ehrenfest methods overestimate the relative peak intensity of the higher-frequency peak around 80 cm^{-1} compared to HEOM. The agreement particularly

deteriorates with increasing population time t_2. This is not too surprising, since Ehrenfest is well-known to give wrong population dynamics[120] and the peak intensities are related to the populations of the excited states. Similar conclusions have been made before but it has been unclear how to overcome this limitation without going to significantly more sophisticated (and computationally expensive) methods. We will return to this question below. First, we will test equatorial Ehrenfest on a larger system.

4.6.2 Fenna-Matthews-Olson complex

The Fenna-Matthews-Olson (FMO) complex is a light-harvesting complex in green sulphur bacteria. It is composed of seven bacteriochlorophyll molecules which funnel energy from the base plate to the reaction centre.[53] Later, an eighth bacteriochlorophyll was discovered in the complex,[54,55] but here we will use a seven site model studied by Kramer *et al.*,[138] since it has been more widely studied and a HEOM benchmark exists for its two-dimensional spectrum.[144]

The singly excited Hamiltonian of the FMO electronic subsystem is

$$\hat{H}_{\mathrm{S}}^{s} = \begin{pmatrix} 200 & -87.7 & 5.5 & -5.9 & 6.7 & -13.7 & -9.9 \\ -87.7 & 320 & 30.8 & 8.2 & 0.7 & 11.8 & 4.3 \\ 5.5 & 30.8 & 0 & -53.5 & -2.2 & -9.6 & 6.0 \\ -5.9 & 8.2 & -53.5 & 110 & -70.7 & -17.0 & -63.3 \\ 6.7 & 0.7 & -2.2 & -70.7 & 270 & 81.1 & -1.3 \\ -13.7 & 11.8 & -9.6 & -17.0 & 81.1 & 420 & 39.7 \\ -9.9 & 4.3 & 6.0 & -63.3 & -1.3 & 39.7 & 230 \end{pmatrix} \mathrm{cm}^{-1} + 12,210\ \mathrm{cm}^{-1} \hat{I}_N, \tag{4.82}$$

where $\hat{I}_N$ is the $N \times N$ identity matrix (in this case, $N = 7$). Like in the biexciton model, the ground state energy is fixed at $\varepsilon_0 = 0$, such that $\hat{H}_{\mathrm{S}}^{g} = 0$. Each site is coupled to a harmonic bath with a Debye spectral density with reorganisation energy $\lambda = 35\ \mathrm{cm}^{-1}$ and relaxation time $\tau_c = 50$ fs.

The FMO complex is a three-dimensional system, which has two consequences: Firstly,

the dipole moment operator is now a vector $\hat{\boldsymbol{\mu}} = (\hat{\mu}_x, \hat{\mu}_y, \hat{\mu}_z)$ with three distinct Cartesian components. Each component is built from the dipole vectors of the seven bacteriochlorophyll molecules, which point in the directions[138]

$$\boldsymbol{d}_1 = \Big(-0.741, \quad -0.561, \quad -0.3696\Big), \tag{4.83a}$$

$$\boldsymbol{d}_2 = \Big(-0.857, \quad 0.504, \quad -0.107\Big), \tag{4.83b}$$

$$\boldsymbol{d}_3 = \Big(-0.197, \quad 0.957, \quad -0.211\Big), \tag{4.83c}$$

$$\boldsymbol{d}_4 = \Big(-0.799, \quad -0.534, \quad -0.277\Big), \tag{4.83d}$$

$$\boldsymbol{d}_5 = \Big(-0.737, \quad 0.656, \quad 0.164\Big), \tag{4.83e}$$

$$\boldsymbol{d}_6 = \Big(-0.135, \quad -0.879, \quad 0.457\Big), \tag{4.83f}$$

$$\boldsymbol{d}_7 = \Big(-0.495, \quad -0.708, \quad -0.503\Big). \tag{4.83g}$$

Within the notation of the Frenkel exciton models, each of the three Cartesian components is obtained through the definitions in Equation (4.29) and (4.30) by setting $\mu_{n,\alpha} = d_{n,\alpha}$ for $\alpha = x, y, z$, thereby taking the magnitude of each bacteriochlorophyll dipole, which simply affects the overall scaling of the spectra, to be 1. Here, we have added the additional subscript α to account for the multiple dimensions.

Secondly, we need to take into account different orientations of the FMO sample. To this end, we rotationally average the response function or (equivalently) average over possible light polarisations. We assume that all light pulses are collinear, which greatly simplifies the discussion. Hein *et al.* found appropriate sets of light polarisations by demanding that the average of the product of two dipole operators (in the case of linear spectra) or the product of four dipole operators (in the case of two-dimensional spectra) is reproduced correctly by averaging over the polarisation directions.[148] We use their convention here. For the linear spectrum, we can simply use the canonical basis vectors $\boldsymbol{e}_x = (1, 0, 0)$, $\boldsymbol{e}_y = (0, 1, 0)$ and $\boldsymbol{e}_z = (0, 0, 1)$. For the nonlinear spectra, we average over ten vectors, $\boldsymbol{e}_{1,2,3,4} = (\pm 1, \pm 1, 1)$, $\boldsymbol{e}_{5,6} = (0, \pm\frac{1}{\phi}, \phi)$, $\boldsymbol{e}_{7,8} = (\pm\frac{1}{\phi}, \phi, 0)$ and $\boldsymbol{e}_{9,10} = (\pm\phi, 0, \frac{1}{\phi})$. The dipole moment operator is

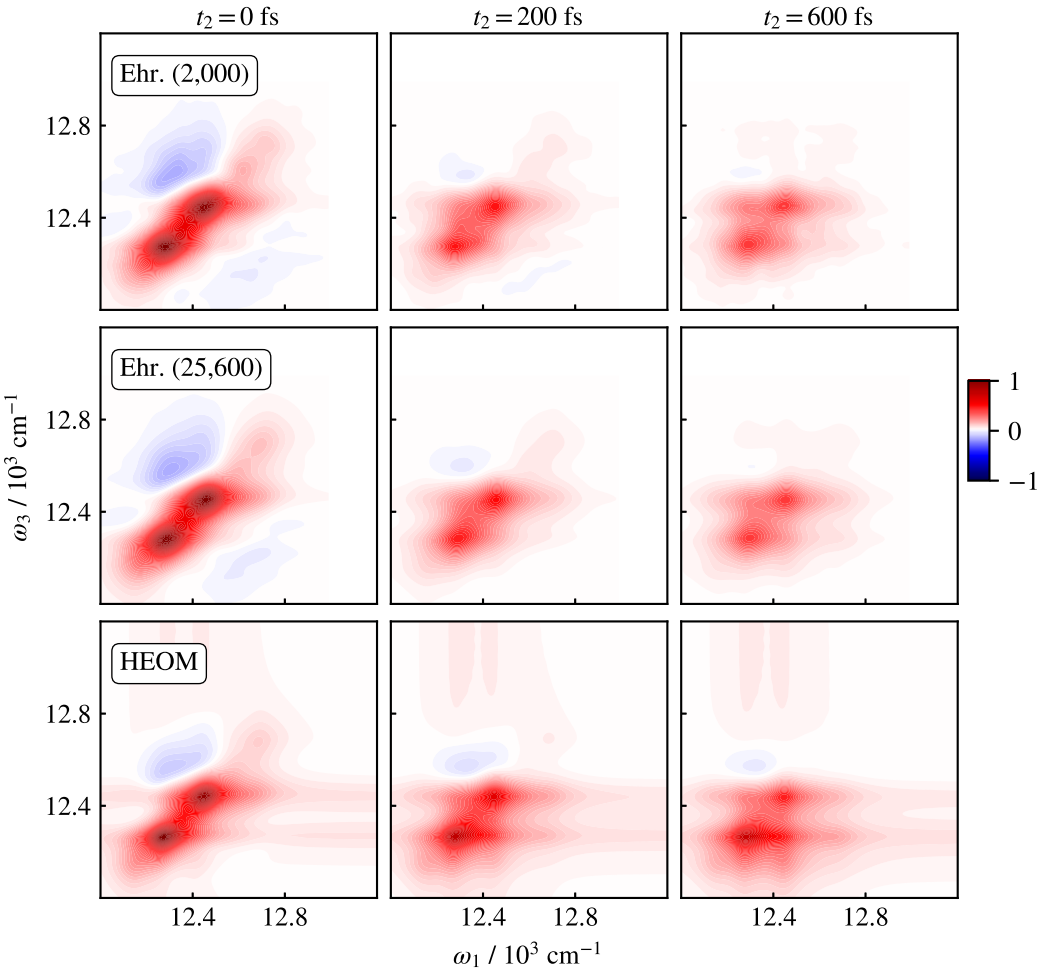

Figure 4.12: HEOM and equatorial Ehrenfest two-dimensional electronic spectra for the FMO complex. The numbers in the Ehrenfest panels correspond to the number of trajectories that were averaged over to produce the spectra.

projected onto each of the light polarisations,

$$\hat{\mu}_i = \hat{\boldsymbol{\mu}} \cdot \boldsymbol{e}_i, \tag{4.84}$$

where $i = (x, y, z)$ for linear spectra and $i = 1, \ldots, 10$ for two-dimensional spectra.

With seven electronic states, the FMO model is computationally much harder to solve for an exact method like HEOM than the biexciton model, but it is doable. The HEOM results shown here were generated by Jonathan Mannouch and have been in part presented elsewhere.[144] We compare two-dimensional HEOM and equatorial Ehrenfest spectra at different

convergence levels in Figure 4.12. Because polar Ehrenfest is significantly more expensive than equatorial Ehrenfest, and because of its other drawbacks as discussed above, we did not compute polar Ehrenfest results for the FMO complex.

Let us first focus on the two bottom rows which compare fully converged Ehrenfest spectra to HEOM results. Similar to the biexciton model, the equatorial Ehrenfest spectra agree well visually with the HEOM spectra. The two top rows compare fully converged Ehrenfest spectra averaged over 57,000 trajectories with coarse spectra averaged over 2,000 trajectories. The spectrum at 2,000 trajectories is already a good approximation of the fully converged spectrum. This is partly due to the windowing of the response functions prior to applying the two-dimensional Fourier transform, but it is nevertheless important from a practical perspective. Applications may be limited in the total trajectory number that can be run in a reasonable time. Even at low trajectory numbers, equatorial Ehrenfest is thus able to provide almost the same physical insight as a fully converged calculation.

However, already from the two-dimensional plots it is apparent that there is a spread in the HEOM spectra that is not captured by equatorial Ehrenfest. To gain better insight into our predictions, we show the diagonal slices through the FMO spectra and the pump-probe spectra in Figure 4.13. The diagonal slices agree well at zero delay time, $t_2 = 0$. At longer delay times, the Ehrenfest peaks remain in the right places but their intensities diverge from HEOM intensities. For example, in the HEOM spectra at $t_2 = 600$ fs, the lower-frequency peak around 12,300 cm^{-1} has a noticeably higher intensity at longer delay times t_2 than the high-frequency peak around 12,500 cm^{-1}. Ehrenfest completely fails to capture this. Instead, it predicts nearly equal intensities of the two peaks in the diagonal slices, and even an inverted trend in the pump-probe spectra. As in the biexciton model, we suspect that this is due to the overheated Ehrenfest population dynamics.

4.7 Improved population dynamics

Having identified the overheating in Ehrenfest as the main source of error in the two-dimensional spectra, we will now explore whether it is possible to improve on Ehrenfest with a method that is better suited for population dynamics. We focus on the t_2 evolution,

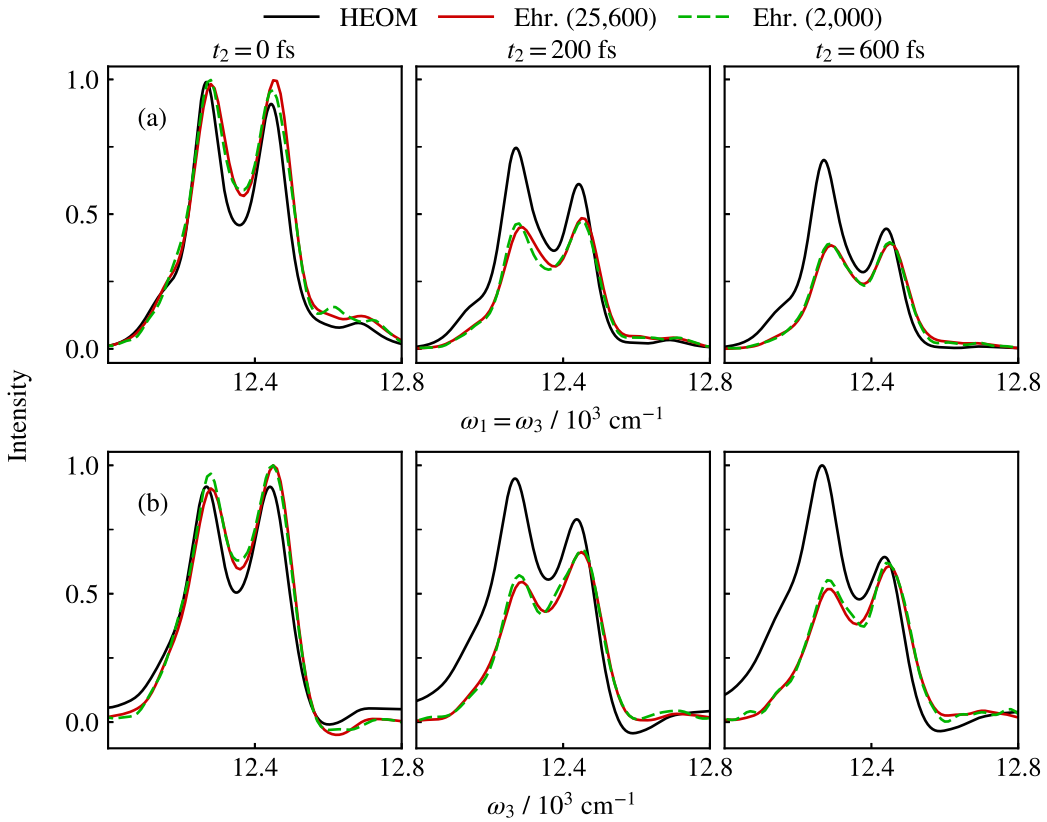

Figure 4.13: Comparing diagonal slices through the two-dimensional spectra (top) and the pump-probe spectra obtained by integrating the two-dimensional spectra over the ω_1 axis (bottom) for the coarse (green) and fully converged (red) equatorial Ehrenfest, and exact HEOM results.

since we expect it to be most affected by errors in the population dynamics. A first idea is to simply substitute the t_2 Ehrenfest dynamics with spin mapping dynamics.[143,150] We chose spin mapping because we can easily extend our existing code to run spin mapping instead of Ehrenfest, and because spin mapping is more accurate for the FMO model than other mapping approaches.[143] Just like in Ehrenfest, we cannot straightforwardly propagate a coherence in spin mapping. We will therefore use the same equatorial pure states for the t_2 propagation. This section discusses any updates that need to be made to accommodate the change in methods.

It could be argued that a method like MASH would be even better suited than spin mapping, since it is designed to reproduce the correct equilibrium populations.[122,123,151] However, the computational cost of MASH is expected to be much higher than the cost of spin mapping because it needs a smaller time step. If the spin mapping results improve on Ehrenfest, we could follow it up with a MASH modification.

4.7.1 Spin mapping

Originally, spin mapping was introduced through the Stratonovich-Weyl representation, which maps the electronic wave function onto vectors in the Bloch space, whose entries are found through the $\mathfrak{su}(N)$ generators.[143,150] For example, for $N = 2$ the Bloch space is a sphere in three dimensions. However, the Bloch space representation becomes complicated if there are more than two electronic states: In this case, the Bloch vector entries are subject to a number of additional constraints.[152,153] It is then easier to formulate spin mapping in terms of the usual electronic wave function $|c\rangle$. The resulting spin mapping density matrix is

$$\hat{\rho}_{\mathrm{sm}} = \sqrt{N+1}|c\rangle\langle c| - \frac{\gamma_{\mathrm{W}}}{2}\hat{I}_N, \tag{4.85}$$

where $\gamma_{\mathrm{W}} = \frac{2}{N}(\sqrt{N+1} - 1)$. Only the case where the wave function $|c\rangle$ is in the singly excited manifold is relevant here. This is because during the t_2 propagation, the system is either in the ground state or in the singly excited manifold. If it is in the ground state (which is the case for the ground state bleaching response functions Φ_2 and Φ_5), then $N = 1$, in which case spin mapping and Ehrenfest give the same results, and we do not need to make

any modifications. We can thus focus on the Frenkel exciton Hamiltonian within the singly excited manifold,

$$\hat{H}(\boldsymbol{p},\boldsymbol{q}) = \hat{H}^s_{\mathrm{S}} + \hat{H}^s_{\mathrm{SB}}(\boldsymbol{q}) + H_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q}), \tag{4.86}$$

which contains the bath Hamiltonian $H_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q})$, the singly excited system Hamiltonian $\hat{H}^s_{\mathrm{S}}$, and the system-bath coupling within the singly excited manifold, $\hat{H}^s_{\mathrm{SB}}(\boldsymbol{q})$ (see Equations (4.22), (4.24) and (4.25) for their definitions). Together with the density matrix in Equation (4.85), the spin mapping Hamiltonian is then

$$H(\boldsymbol{p},\boldsymbol{q},c) = \mathrm{Tr}[\hat{\rho}_{\mathrm{sm}}\hat{H}(\boldsymbol{p},\boldsymbol{q})] = \sqrt{N+1}\langle c|\hat{H}^s(\boldsymbol{q})|c\rangle - \frac{\gamma_{\mathrm{W}}}{2}\mathrm{Tr}[\hat{H}^s(\boldsymbol{q})] + H_{\mathrm{B}}(\boldsymbol{p},\boldsymbol{q}), \tag{4.87}$$

with $\hat{H}^s(\boldsymbol{q}) = \hat{H}^s_{\mathrm{S}} + \hat{H}^s_{\mathrm{SB}}(\boldsymbol{q})$ as defined in Equation (4.23). Spin mapping differs from Ehrenfest in three key respects: First, the electronic wave function coefficients are sampled from a semiclassical distribution, as opposed to being initialised in a deterministic manner. Secondly, the equations of motions are different, which is clear from comparing the density matrix in Equation (4.85) to the Ehrenfest density matrix, $\hat{\rho}_{\mathrm{Ehr}} = |c\rangle\langle c|$. And thirdly, we need to specify how the final density matrix should be used in the t_3 propagation.

Initialisation We use focused initial conditions, with which we can restrict the sampling to states which reproduce a given pure state population. Let us assume that we want to initialise a pure state $|c\rangle$, with $c_n = 1$ and $c_{m\neq n} = 0$. To this end, we sample a vector $|\tilde{c}\rangle$, where $\tilde{c}_n = \sqrt{2+\gamma_{\mathrm{W}}}\; e^{i\phi_n}$, $\tilde{c}_{m\neq n} = \sqrt{\gamma_{\mathrm{W}}}\; e^{i\phi_m}$, and all ϕs are sampled uniformly between 0 and 2π. The need for sampling the ϕs is one reason why our spin mapping modification is more expensive than Ehrenfest.

Moreover, we generalised focused sampling to a 'doubly focused sampling' scheme to accurately reproduce a given density matrix as initial condition. The scheme is described in App. B.

Propagation The only thing that changes between Ehrenfest and spin mapping in terms of the propagation is the force on the nuclei, which is

$$\begin{aligned} F(\boldsymbol{q}, c) &= -\mathrm{Tr}[\hat{\rho}_{\mathrm{sm,el}} \nabla_{\boldsymbol{q}} \hat{H}(\boldsymbol{p}, \boldsymbol{q}, c)] \\ &= -\left[\sqrt{N+1}\langle c|\nabla_{\boldsymbol{q}} \hat{H}^s(\boldsymbol{q})|c\rangle - \frac{\gamma_{\mathrm{W}}}{2}\mathrm{Tr}[\nabla_{\boldsymbol{q}} \hat{H}^s(\boldsymbol{q})] + \nabla_{\boldsymbol{q}} H_{\mathrm{B}}(\boldsymbol{p}, \boldsymbol{q})\right] \end{aligned} \tag{4.88}$$

in spin mapping. The wave function and the nuclear positions $\boldsymbol{q}$ follow the same equations of motion as before (in Equations (4.39) and (4.40)).

Measuring the density matrix and setup for t_3 evolution At the end of the t_2 propagation, we calculate the density matrix in Equation (4.85). The density matrix is then acted on by μ_+ either from the left (for Φ_3 and Φ_6) or from the right (for Φ_1 and Φ_4). For the t_3 propagation with Ehrenfest, we need to separate the result into a bra and a ket state (see. Equation (4.65) and the preceding paragraph). However, the spin mapping density matrix is not a pure state, so we diagonalise the density matrix before the action of $\hat{\mu}_+$. This creates N states in the singly excited manifold, which are then ready for the t_3 evolution. The fact that we have to propagate N pure states through t_3 for each pure state propagated through t_2 is another reason why spin mapping is more expensive than Ehrenfest.

4.7.2 Results

Figure 4.14 compares the population dynamics for the FMO model for the different methods considered here. Spin mapping improves the Ehrenfest population dynamics. This is most noticeable at long times, but there is also some improvement at shorter times. On the time scales we are interested in (t_2 up to 600 fs), we thus expect some improvement from spin mapping. Spin mapping spectra are presented in Figure 4.15 and compared to HEOM and equatorial Ehrenfest spectra. The diagonal slices and pump-probe spectra offer the insight we need, so we only present those here.

At $t_2 = 0$ the spin mapping spectra seem to accumulate additional intensity between the two dominant peaks. This may be due to a convergence issue, since Ehrenfest and spin mapping should be indistinguishable (once the sampling in spin mapping has been fully

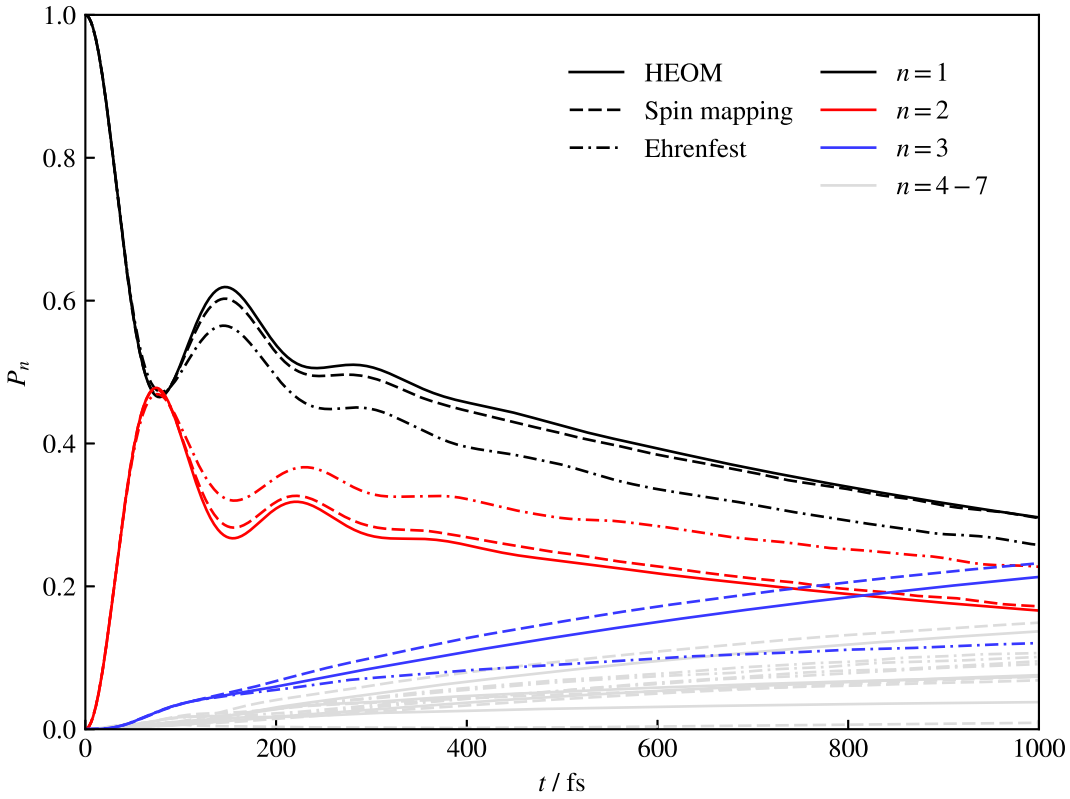

Figure 4.14: HEOM, spin mapping, and Ehrenfest population dynamics for the FMO model. On the time scale we are interested in, the first three states carry the majority of the population. In the interest of clarity, the remaining states are therefore only shown in grey. Coloured versions can be found by comparing Figure 6 of Reference [143] and Figure 4 in Reference [122].

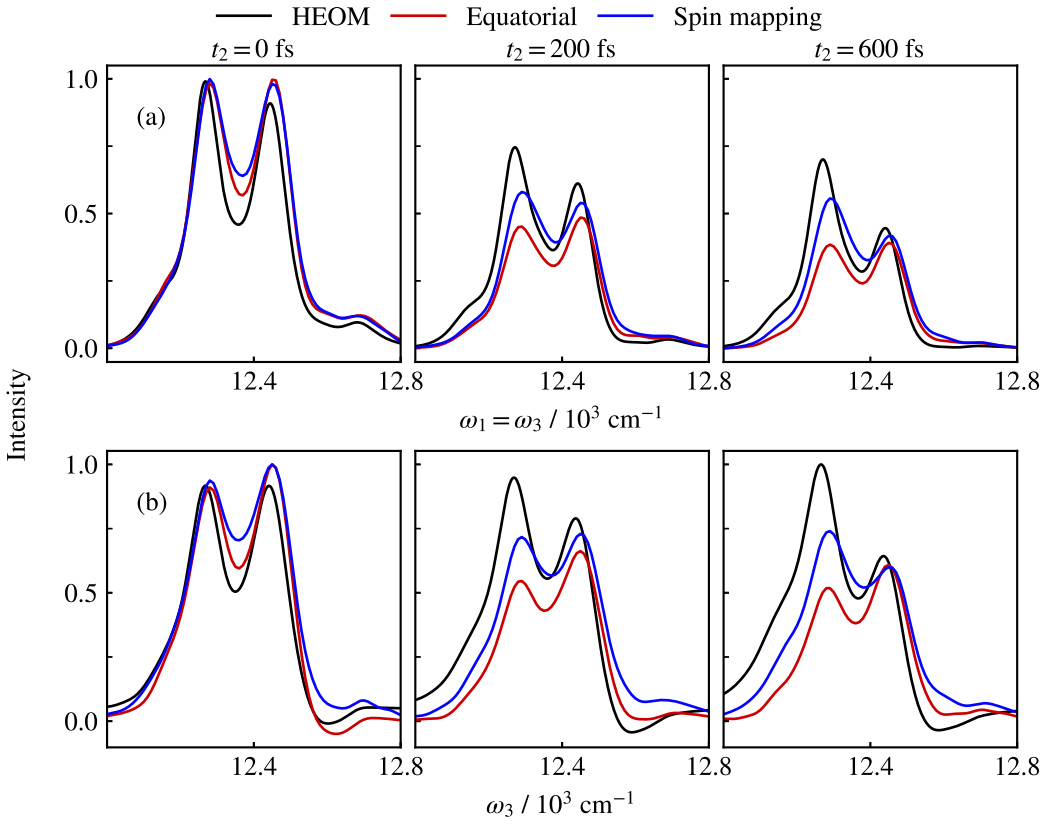

Figure 4.15: Comparison of diagonal slices through the two-dimensional spectra (top) and pump-probe spectra (bottom) obtained with spin mapping to equatorial Ehrenfest and HEOM. The (fully converged) Ehrenfest and HEOM results are the same as in Figure 4.13.

t_2 / fs	HEOM	Ehrenfest	Spin mapping
0	0.91	1.02	0.98
200	0.81	1.06	0.93
600	0.61	1.03	0.75

Table 4.1: Ratio of the intensity of the high-frequency peak to the intensity of the low-frequency peak along the diagonal of the two-dimensional FMO spectrum for different dynamics methods.

converged) for $t_2 = 0$. At the intermediate delay time ($t_2 = 200$ fs), spin mapping only offers modest (if any) improvement over Ehrenfest, but the improvement gets better at longer delay times ($t_2 = 600$ fs). This can be seen, for example, from the relative peak intensities which are also summarised in Table 4.1. Clearly, the relative peak intensities are better in spin mapping, even though it does not reproduce the HEOM spectra quantitatively. It would be interesting to see how an even better dynamics method (like MASH) would do. We expect that most of the spin mapping methodology outlined above could be generalised to MASH, the main changes beyond the change in the dynamics being the use of MASH focused initial conditions and MASH observables to measure the density matrix.

4.8 Linear spectra

While we have focused on the nonlinear spectra, we can straightforwardly apply our methodology to the linear response function to target the linear spectrum (Equations (4.8) and (4.9)). In polar Ehrenfest, four pure states are propagated up to time t in order to calculate the linear response function. However, the polar states $|\psi_0\rangle = |0\rangle$ and $|\psi_1\rangle = |\mu\rangle$ do not contribute to the final result, since they vanish in the final trace with μ_-. Only the mixed states $|\psi_2\rangle$ and $|\psi_3\rangle$ contribute, such that the final result is

$$S^{(1)}(t) = \mathrm{Im}\left(\mathrm{Tr}[\mu_-|\psi_2(t)\rangle\langle\psi_2(t)|] + i\mathrm{Tr}[\mu_-|\psi_3(t)\rangle\langle\psi_3(t)|]\right) = \mathrm{Im}\mathrm{Tr}[\mu_-|\mu(t)\rangle\langle 0(t)|], \quad (4.89)$$

where the time evolution is given by

$$|\psi_i(t)\rangle = \mathcal{T} e^{i\int_0^t \mathcal{H}(\boldsymbol{q}'_\tau)\mathrm{d}\tau}|\psi_i\rangle. \quad (4.90)$$

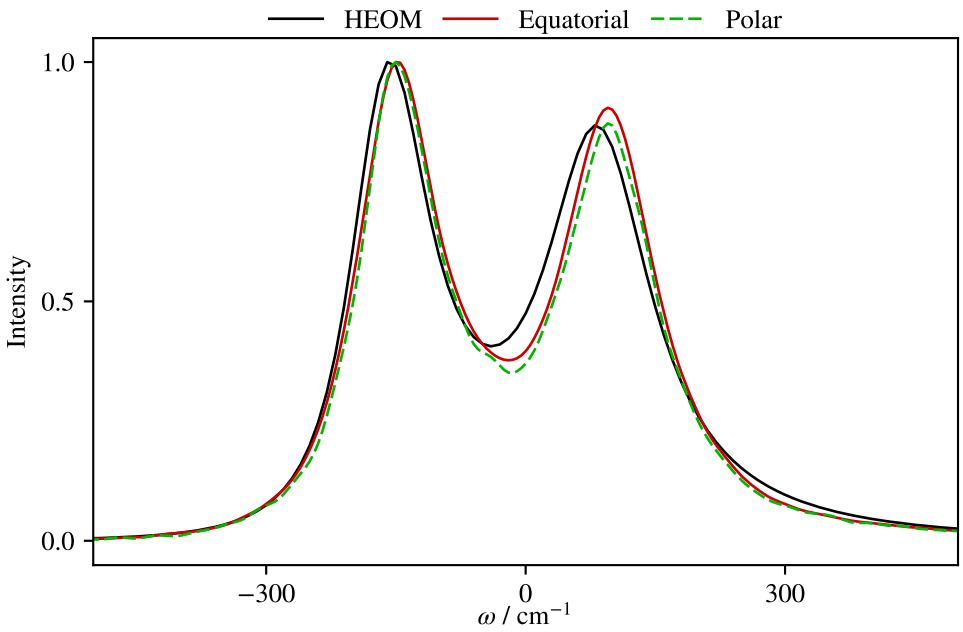

Figure 4.16: Linear spectrum of the biexciton model, comparing equatorial and polar Ehrenfest with exact HEOM results. There are subtle differences between the spectra due to the different normalisation conventions of the pure states in the two methods.

This is almost the same answer as we recover from the equatorial decomposition (see Equation (4.59)), except that the nuclear degrees of freedom $\boldsymbol{p}'_\tau, \boldsymbol{q}'_\tau$ evolve under the force

$$F_{\mathrm{pol}}(\boldsymbol{q}) = -\frac{1}{1+c_\mu}(\langle\mu|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|\mu\rangle + \langle 0|\nabla_{\boldsymbol{q}}\hat{\mathcal{H}}|0\rangle), \tag{4.91}$$

with the state $|\mu\rangle = \mu_+|0\rangle$ and coefficient $c_\mu = \sum_n \mu_n^2$ from above. The force arising from the polar pure states is thus subtly different from the one arising from the equatorial pure states (see Equation (4.61)), which manifests in small differences in the linear spectrum. We have also verified this numerically. Figure 4.16 shows that the linear spectra for the biexciton model differ slightly between the two Ehrenfest methods. However, the differences are a lot smaller than in the two-dimensional spectra. This is likely because the biexciton model's dipole moments are relatively small in magnitude, and the polar force is similar enough to the equatorial force. For larger transition dipole moments, we would expect to see a larger difference between the two methods. In contrast, equatorial pure state Ehrenfest is entirely equivalent to the mean classical path approach for linear spectra.[69] We have already shown that the force in equatorial Ehrenfest (Equation (4.54)) is identical to the mean classical path

force (Equation (4.48)) in the propagation of a coherence $|a\rangle\langle b|$. The linear spectrum involves a single coherence propagation, and so the pure state decomposition is only implicitly invoked when we use the resummation in Equation (4.59).

4.9 Summary

In this chapter, we discussed recent advances in the semiclassical simulation of two-dimensional electronic spectra. Our goal was to develop a computationally inexpensive method. We thus chose to focus on semiclassical mean-field (Ehrenfest) dynamics,[142] which tends to need fewer trajectories than other approaches but is known to converge to an overheated equilibrium.[120] Ehrenfest has been previously studied in the context of electronic spectroscopy. Van der Vegte *et al.* developed the mean classical path approach, where a coherence is propagated under the average force of its ket and bra states. They tested their method on a two-state Frenkel exciton model. In Frenkel exciton models, a quantum system is coupled to a Debye bath of classical oscillators. The relaxation time of the oscillators determines if the system can be appropriately described in a classical limit (if the relaxation time is slow) or whether coherence effects play some role (if the relaxation time is fast). The mean classical path method has been shown to be well suited to biexciton models with slow baths at short delay times, but it performs worse at longer delay times t_2. Van der Vegte *et al.* attributed this to the convergence of Ehrenfest to an overheated equilibrium.[69,120]

Later, Atsango *et al.* proposed the polar Ehrenfest approach,[70] where they decomposed a coherence into pure states before propagating them with conventional Ehrenfest dynamics. They argued that their approach is less *ad hoc* than the mean classical path method, and applied it to slow and fast bath Frenkel biexciton models. Polar Ehrenfest does well with a slow bath, but the results again decrease in quality at longer delay times, probably due to the overheated Ehrenfest dynamics. In contrast, the polar Ehrenfest spectra severely deteriorate for a fast bath biexciton model, where they show spurious oscillations already at very short delay times.

One obvious drawback of the polar Ehrenfest method is the computational cost, which

scales exponentially with the number of dipole interactions. However, this is not its most serious drawback. Through careful analysis, we showed that the wave function dynamics in polar Ehrenfest contains unphysical contributions, in particular during the population time t_2 (see Figure 4.6 and its discussion in the text). The reason for this is the choice of polar pure states, which is only one of many possible decompositions. We proposed an alternative equatorial decomposition, and started by analysing it with the insights accumulated from our analysis of polar Ehrenfest. Firstly, the choice of pure states and weights reduces the computational effort significantly – by a factor 32 – through a resummation trick (Equation (4.56)) and by taking advantage of the symmetry of the response functions. But secondly, the equatorial decomposition together with the resummation trick means that the wave function dynamics agree with physical expectations. There are no longer any unphysical contributions.

We then compared equatorial Ehrenfest to polar Ehrenfest by computing two-dimensional spectra for the fast bath biexciton model, where polar Ehrenfest fails. The equatorial Ehrenfest spectra are much better behaved, and are close to exact HEOM spectra. However, just as observed previously, our results become worse at longer delay times, where the relative diagonal peak intensities diverge from the HEOM ones. Since the diagonal peaks are related to the populations of the excited states, we conclude once again that the known issues in Ehrenfest population dynamics plague electronic spectra.

The reduced computational cost allowed us to look at systems with more than two electronic states, so we also computed equatorial Ehrenfest spectra for a seven state model of the FMO complex. Already with as few as 2,000 trajectories, the results are close to fully converged, and are qualitatively close to HEOM results. Quantitatively, they suffer from the same problems as other Ehrenfest approaches and get the population dynamics wrong at long delay times.

To overcome the errors introduced by Ehrenfest, we replaced the propagation through the delay time t_2 with spin mapping. Spin mapping observables (together with the correct sampling and propagation routines) are much better suited to population dynamics. However, we needed to add an additional step at the end of the t_2 evolution (before starting the t_3 evolution), which increases the computational effort by a factor of N, where N is the number

of singly excited states. For the FMO model, the relative peak intensities obtained with spin mapping improve on Ehrenfest dynamics. However, the absolute intensities are only modestly improved, and the spin mapping spectra are still some way from the exact results. Further studies, ideally with comparisons to experiment, are needed to determine whether the increased computational cost is worth it, or whether Ehrenfest approaches themselves already provide sufficient physical insight. To make a useful contribution to this field, such studies will need to go beyond Frenkel exciton models, for example to allow for excited states that have different fundamental harmonic frequencies than the ground state or are rotated with respect to the ground state (which are known as Duschinsky rotations), or to include anharmonic potential energy surfaces.

Finally, we explained how our approach is closely related to the mean classical path method, and even reduces exactly to it for linear spectra (whereas polar Ehrenfest is not entirely equivalent to the mean classical path approach because of a different normalisation convention for the pure states). The mean classical path approach is thus no less *ad hoc* than pure state Ehrenfest. In fact it is even computationally less expensive than equatorial Ehrenfest (by about a factor of 2), and is likely to be about as accurate as equatorial Ehrenfest for two-dimensional spectra.

Chapter 5

Conclusion

This thesis has described recent developments in the simulation of vibrational and electronic spectroscopy. The first half, comprising Chapters 2 and 3, focused on the application of path integral methods to vibrational spectroscopy.

In Chapter 2, we reviewed the chequered history of path integral methods for vibrational spectroscopy. We then focused on the recently developed QCMD method, and in particular on an approximate, faster variant called f-QCMD. In previous work, f-QCMD had been applied to gas-phase systems, but it had not been generalised to any liquid or solid phase systems. We proposed that f-QCMD can be generalised to the spectroscopy of condensed phases by separating the quasi-centroid potential of mean force into intra- and intermolecular contributions, which were added as corrections to the classical interaction potential. To find the intramolecular contributions to the potential of mean force, we followed the strategy that had previously been used for gas-phase water.[35] The intermolecular corrections were found with IBI, by fitting the radial distribution functions obtained with the potential of mean force to exact quasi-centroid radial distribution functions obtained from a short PIMD simulation. A key development in obtaining the quasi-centroid potential of mean force was the regularisation of the IBI update rule. This improves not only the convergence of the IBI, which is notorious for oscillations that over- or undershoot the true target,[94] but also gives smooth effective potentials which can be readily used to evaluate MD forces.

It is difficult to test any QCMD variant generally, since the quasi-centroid degrees of

freedom need to be freshly specified for each system. However, well converged adiabatic QCMD spectra were available for liquid water and hexagonal ice at 300 K and 150 K,[33,87] respectively, which allowed us to get a meaningful assessment of the approximations made in f-QCMD. We also included a comparison to Te PIGS, which precomputes a centroid potential of mean force at an effective (high) temperature, and assumes that the same potential may be used at lower temperatures. Comparing all three path integral methods, we noted that only minor differences occur in the peak shapes and that the peaks are shifted only very slightly from each other (on the order of tens of cm^{-1}). In spite of these differences, the three methods seem to converge more-or-less to the same answer. This is in stark contrast to the wildly different predictions of vibrational spectra with the established methods RPMD, CMD and TRPMD, and we believe that we are finally seeing the 'true' path integral spectrum for liquid water and ice. We also compared the f-QCMD spectra (which can be considered representative of all path integral spectra, as just discussed) to classical spectra to ascertain the impact of nuclear quantum effects on the spectrum. Nuclear quantum effects do not change the peak line shape, but they do cause anharmonic red shifts in the peak positions which are more pronounced in the higher frequency peaks.

Our methodology can easily be extended to other liquids or solids, provided a suitable set of intra- and intermolecular quasi-centroid coordinates can be defined. In the gas phase, f-QCMD performed just as well for molecules including ammonia, methane and hydrogen peroxide, as it did for water.[35,36] The same is likely to be true in the condensed phase, which could be tested (for example) in future calculations on liquid ammonia.

London and co-workers have recently incorporated our f-QCMD method in DL_POLY Quantum 2.1, a computational suite for the simulation of dynamical properties and vibrational spectra.[154] They also recently developed h-CMD, a hybrid method for the simulation of heterogeneous systems such as water molecules in zeolite cavities, by treating the water molecules with f-QCMD and the zeolite framework with a fast version of (Cartesian) CMD (f-CMD).[155] The implementation of f-CMD and f-QCMD in freely available software packages will certainly help the wider application of these methods, in particular because the f-CMD implementation can be combined with centroid potentials of mean force obtained from the

Te PIGS method of Kapil and co-workers.[84]

In Chapter 3, we used the f-QCMD approximation developed in Chapter 2 to investigate vibrational strong coupling (polariton spectroscopy) in liquid water. Having reviewed previous work,[43,47,48] we began by demonstrating that for a range of cavity setups f-QCMD performs just as well as it does in the cavity-free simulations of Chapter 2.

However, the most important development in this chapter was the introduction of the SHO model, which replaces a complicated molecular potential by a set of uncoupled harmonic oscillators. The SHO model predicts cavity spectra by coupling these oscillators to the cavity through their linear dipole moments, which are entirely determined by the cavity-free infrared spectrum. Because the cavity is just another harmonic oscillator, SHO model predictions for the cavity spectrum can be obtained through a simple matrix diagonalisation. The SHO model perfectly predicts the cavity f-QCMD spectra, provided we use the cavity-free f-QCMD spectrum as input. This holds for all cavity setups for which we computed spectra, meaning that the SHO model can replace cavity f-QCMD simulations altogether. This is generalisable: the SHO model makes any cavity simulation of the infrared spectrum redundant, since we can use the relevant cavity-free spectrum as input.

With this realisation in mind, we argued that the experimental spectrum is the most sensible input. It does not suffer from errors due to approximations made to the underlying potential energy surfaces, unlike any simulated input spectrum would. Since experimental data for the cavity-free spectra of water and deuterated water,[109,111] as well data on cavity experiments is available online,[110] we could immediately assess the SHO predictions on real examples. Given the simple nature of the model, it gives very good predictions. The agreement with cavity measurements is almost quantitative in deuterated water, and slightly worse in water.

The quality of the SHO model's prediction may appear unexpected at first sight. However, we then showed that the input to the SHO model is equivalent to the input to the transfer matrix method, which is already widely used by experimentalists to interpret their cavity spectra, making the SHO model's accurate predictions somewhat less surprising. We argue that our findings could hint at the fact that there is nothing fundamentally special or novel

about cavity spectra. The cavity simply provides a different way of 'looking' at the system enclosed within. Together with questions about the reproducibility of cavity experiments,[45,46] our results casts some doubt on reported vacuum cavity effects (specifically on the rates of chemical reactions in liquid samples).

Recently, Yuen-Zhou, Weichman and their co-workers have published a review[115] titled "When do molecular polaritons behave like optical filters?", which settles the debate. In many situations (such as the linear spectroscopy studied by us), the cavity simply acts as an optical filter and its effects can be explained using classical linear optics. It is only in some situations, like spontaneous Rahman scattering, that cavity effects may be non-trivial and a more sophisticated quantum electrodynamics treatment is warranted.

In Chapter 4, the focus shifted to electronic spectroscopy, and in particular two-dimensional electronic spectroscopy. We started by discussing the theoretical background and deriving expressions for the response functions relevant for two-dimensional spectra. We then zeroed in on two works using semiclassical Ehrenfest methods to compute the response functions, which we call the mean classical path[69] and polar pure state Ehrenfest[70] (or 'polar Ehrenfest') approaches, respectively. We then proposed the equatorial Ehrenfest method, which is obtained through a simple modification to the polar pure states method. Equatorial Ehrenfest both reduces the computational cost as well as improves the accuracy of polar Ehrenfest by taking advantage of the symmetries of the equatorial pure states and the nonlinear response functions. It also serves as a conceptual bridge between the mean classical path and polar Ehrenfest methods. For linear spectra, equatorial Ehrenfest is even equivalent to the mean classical path method. There is thus no argument to be made that polar Ehrenfest is more rigorous or less *ad hoc* than the mean classical path, as has been claimed elsewhere.[70] Quite the opposite is true: Polar Ehrenfest is less accurate and significantly more expensive than the mean classical path approach.

We suspect, in agreement with previous work[69,70] that the main issue in the two-dimensional spectra computed with Ehrenfest methods is the propagation through the second of the three propagation times, t_2, which is also called the population time. In a condensed phase system (or a system-bath model of such a system), Ehrenfest converges to an overheated equilibrium,

which causes significant errors in the Ehrenfest population dynamics. This is supported by the observation that the spectra become worse with increasing t_2. To overcome this, we suggested to replace the propagation through t_2 by spin mapping. After describing the necessary modifications, we showed that such a spin mapping-based method improves the relative peak intensities on the diagonals of the two-dimensional spectra, but still fails to reproduce exact results. Also, the computational cost is increased by a factor N, where N is the number of singly excited states.

Whether the additional cost is worth it will depend on the relevant application. We thus need to better understand which properties are extracted from experimental two-dimensional spectra. To do so, we will need to design models that closely emulate experiments. Throughout Chapter 4, we considered only Frenkel exciton models, for which we could obtain exact HEOM results for benchmarking. However, these models are not expected to be good representations of experiments: Realistic systems will have a more complicated excited state structure that includes effects such as Duschinsky rotations, electronic state-dependent vibrational frequencies and anharmonicities, and we do not know how the methods discussed here will fare for them. But at least they can in principle be applied to more realistic systems, which exact quantum mechanical methods such as HEOM currently cannot.

While the work described in Chapter 4 has not yet been published, a draft preprint has been shared with some members of the community. Reichman and co-workers have used our equatorial pure state decomposition to study (linear) polaron spectral functions.[156] They also show results obtained with MASH estimators.[156] The idea of using MASH estimators for the t_2 propagation in two-dimensional spectroscopy, as Reichman and co-workers have done for the linear case, would be interesting to investigate in order to see whether we can recover more of the exact spectral amplitudes than we could with the Ehrenfest and spin mapping discussed in Chapter 4. However, the MASH results in Reference [156] are not very encouraging, and an entirely different approach may be necessary. How best to calculate nonlinear spectra with a simulation method that scales well in systems of increasing complexity remains an open question.

Bibliography

(1) Wollaston, W. H. *Philosophical transactions of the Royal Society of London* **1802**, 365–380.

(2) Fraunhofer, J. *Annalen der Physik* **1817**, *56*, 264–313.

(3) Melvill, T. *Journal of the Royal Astronomical Society of Canada, Vol. 8, p. 231* **1914**, *8*, 231.

(4) Kirchhoff, G., *Untersuchungen über das Sonnenspectrum und die Spectren der chemischen Elemente*; Dümmler: 1862; Vol. 1.

(5) Röntgen, W. C., *Ueber eine neue Art von Strahlen:(Vorläufige Mittheilung).* Verlag der Stahel'schen k. Hof-u. Univers.-Buch-u. Kunsthandlung: 1895.

(6) Bryant, J. *IEEE Transactions on Microwave Theory and Techniques* **1988**, *36*, 830–858.

(7) Pretsch, E.; Bühlmann, P.; Affolter, C.; Badertscher, M., *Spektroskopische Daten zur Strukturaufklärung organischer Verbindungen*; Springer Berlin Heidelberg: 2001.

(8) Mansfield, P. Snap-Shot MRI, Nobel Prize Lecture, 2003.

(9) Lauterbur, P. C. All Science is Interdisciplinary – from Magnetic Moments to Molecules to Men, Nobel Prize Lecture, 2003.

(10) Hawthorne, F. C. In *Spectroscopic Methods in Mineralogy and Geology*; De Gruyter: 1988, pp III–III.

(11) Aronson, E.; Waldén, P. *Astronomy & Astrophysics* **2015**, *578*, A133.

(12) Krainov, V.; Smirnov, B. M., *Atomic and Molecular Radiative Processes: With Applications to Modern Spectroscopy and the Greenhouse Effect*; Springer International Publishing: 2019.

(13) Klingshirn, C. F., *Semiconductor Optics*; Springer Berlin Heidelberg: 2012.

(14) Schweer, S. M.; Nejad, A.; Suhm, M. A. *Physical Chemistry Chemical Physics* **2022**, *24*, 26449–26457.

(15) McQuarrie, D. A., *Statistical Mechanics*; Harper & Row: 1976.

(16) Feynman, R. P.; Hibbs, A. R., *Quantum Mechanics and Path Integrals*; McGraw Hill, New York: 1965.

(17) Parrinello, M.; Rahman, A. *The Journal of Chemical Physics* **1984**, *80*, 860–867.

(18) Cao, J.; Voth, G. A. *The Journal of Chemical Physics* **1994**, *100*, 5106–5117.

(19) Craig, I. R.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2004**, *121*, 3368–3373.

(20) Rossi, M.; Ceriotti, M.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2014**, *140*, 234116.

(21) Braams, B. J.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2006**, *125*, 124105.

(22) Craig, I. R.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2005**, *122*, 084106.

(23) Geva, E.; Shi, Q.; Voth, G. A. *The Journal of Chemical Physics* **2001**, *115*, 9209–9222.

(24) Pavese, M.; Voth, G. A. *Chemical Physics Letters* **1996**, *249*, 231–236.

(25) Miller, T. F.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2005**, *122*, 184503.

(26) Sutherland, B. J.; Moore, W. H. D.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2021**, *154*, 174104.

(27) Ohta, Y.; Ohta, K.; Kinugawa, K. *The Journal of Chemical Physics* **2004**, *120*, 312–320.

(28) Shiga, M.; Nakayama, A. *Chemical Physics Letters* **2008**, *451*, 175–181.

(29) Habershon, S.; Markland, T. E.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2009**, *131*, 024501.

(30) Witt, A.; Ivanov, S. D.; Shiga, M.; Forbert, H.; Marx, D. *The Journal of Chemical Physics* **2009**, *130*, 194510.

(31) Ivanov, S. D.; Witt, A.; Shiga, M.; Marx, D. *The Journal of Chemical Physics* **2010**, *132*, 031101.

(32) Habershon, S.; Fanourgakis, G. S.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2008**, *129*, 074501.

(33) Trenins, G.; Willatt, M. J.; Althorpe, S. C. *The Journal of Chemical Physics* **2019**, *151*, 054109.

(34) Haggard, C.; Sadhasivam, V. G.; Trenins, G.; Althorpe, S. C. *The Journal of Chemical Physics* **2021**, *155*, 174120.

(35) Fletcher, T.; Zhu, A.; Lawrence, J. E.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2021**, *155*, 231101.

(36) Fletcher, T. Fast quasi-centroid molecular dynamics, PhD thesis, University of Oxford, 2023.

(37) Lather, J.; Bhatt, P.; Thomas, A.; Ebbesen, T. W.; George, J. *Angewandte Chemie International Edition* **2019**, *58*, 10635–10638.

(38) Vergauwe, R. M. A.; Thomas, A.; Nagarajan, K.; Shalabney, A.; George, J.; Chervy, T.; Seidel, M.; Devaux, E.; Torbeev, V.; Ebbesen, T. W. *Angewandte Chemie International Edition* **2019**, *58*, 15324–15328.

(39) Thomas, A.; Lethuillier-Karl, L.; Nagarajan, K.; Vergauwe, R. M. A.; George, J.; Chervy, T.; Shalabney, A.; Devaux, E.; Genet, C.; Moran, J.; Ebbesen, T. W. *Science* **2019**, *363*, 615–619.

(40) Orgiu, E.; George, J.; Hutchison, J. A.; Devaux, E.; Dayen, J. F.; Doudin, B.; Stellacci, F.; Genet, C.; Schachenmayer, J.; Genes, C.; Pupillo, G.; Samorì, P.; Ebbesen, T. W. *Nature Materials* **2015**, *14*, 1123–1129.

(41) Nagarajan, K.; George, J.; Thomas, A.; Devaux, E.; Chervy, T.; Azzini, S.; Joseph, K.; Jouaiti, A.; Hosseini, M. W.; Kumar, A.; Genet, C.; Bartolo, N.; Ciuti, C.; Ebbesen, T. W. *ACS Nano* **2020**, *14*, 10219–10225.

(42) Di Virgilio, L.; Geuchies, J. J.; Kim, H.; Krewer, K.; Wang, H.; Grechko, M.; Bonn, M. *Light: Science & Applications* **2023**, *12*, 183.

(43) Li, T. E.; Subotnik, J. E.; Nitzan, A. *Proceedings of the National Academy of Sciences* **2020**, *117*, 18324–18331.

(44) Campos-Gonzalez-Angulo, J. A.; Poh, Y. R.; Du, M.; Yuen-Zhou, J. *The Journal of Chemical Physics* **2023**, *158*, 230901.

(45) Imperatore, M. V.; Asbury, J. B.; Giebink, N. C. *The Journal of Chemical Physics* **2021**, *154*, 191103.

(46) Wiesehan, G. D.; Xiong, W. *The Journal of Chemical Physics* **2021**, *155*, 241103.

(47) Li, T. E.; Nitzan, A.; Subotnik, J. E. *The Journal of Chemical Physics* **2020**, *152*, 234107.

(48) Li, T. E.; Nitzan, A.; Hammes-Schiffer, S.; Subotnik, J. E. *Journal of Physical Chemistry Letters* **2022**, *13*, 3890–3895.

(49) Brixner, T.; Stenger, J.; Vaswani, H. M.; Cho, M.; Blankenship, R. E.; Fleming, G. R. *Nature* **2005**, *434*, 625–628.

(50) Schlau-Cohen, G. S.; Calhoun, T. R.; Ginsberg, N. S.; Read, E. L.; Ballottari, M.; Bassi, R.; van Grondelle, R.; Fleming, G. R. *The Journal of Physical Chemistry B* **2009**, *113*, 15352–15363.

(51) Engel, G. S.; Calhoun, T. R.; Read, E. L.; Ahn, T.-K.; Mančal, T.; Cheng, Y.-C.; Blankenship, R. E.; Fleming, G. R. *Nature* **2007**, *446*, 782–786.

(52) Read, E. L.; Schlau-Cohen, G. S.; Engel, G. S.; Wen, J.; Blankenship, R. E.; Fleming, G. R. *Biophysical Journal* **2008**, *95*, 847–856.

(53) Hauska, G.; Schoedl, T.; Remigy, H.; Tsiotis, G. *Biochimica et Biophysica Acta (BBA) - Bioenergetics* **2001**, *1507*, 260–277.

(54) Ben-Shem, A.; Frolow, F.; Nelson, N. *FEBS Letters* **2004**, *564*, 274–280.

(55) Tronrud, D. E.; Wen, J.; Gay, L.; Blankenship, R. E. *Photosynthesis Research* **2009**, *100*, 79–87.

(56) Nagy, A.; Prokhorenko, V.; Miller, R. *Current Opinion in Structural Biology* **2006**, *16*, 654–663.

(57) Gilmore, J.; McKenzie, R. H. *The Journal of Physical Chemistry A* **2008**, *112*, 2162–2176.

(58) Pisliakov, A. V.; Mančal, T.; Fleming, G. R. *The Journal of Chemical Physics* **2006**, *124*, 234505.

(59) Cheng, Y.-C.; Fleming, G. R. *The Journal of Physical Chemistry A* **2008**, *112*, 4254–4260.

(60) Ishizaki, A.; Fleming, G. R. *Proceedings of the National Academy of Sciences* **2009**, *106*, 17255–17260.

(61) Wilkins, D. M.; Dattani, N. S. *Journal of Chemical Theory and Computation* **2015**, *11*, 3411–3419.

(62) Jonas, D. M. *Annual Review of Physical Chemistry* **2018**, *69*, 327–352.

(63) Mančal, T. *The Journal of Physical Chemistry B* **2013**, *117*, 11282–11291.

(64) Reppert, M.; Brumer, P. *The Journal of Chemical Physics* **2018**, *149*, 234102.

(65) Renger, T. *The Journal of Physical Chemistry B* **2021**, *125*, 6406–6416.

(66) Kundu, S.; Dani, R.; Makri, N. *Science Advances* **2022**, *8*, eadd0023.

(67) Runeson, J. E.; Manolopoulos, D. E. *Science Advances* **2025**, *11*, eadw4798.

(68) Tanimura, Y. *The Journal of Chemical Physics* **2020**, *153*, 020901.

(69) Van der Vegte, C. P.; Dijkstra, A. G.; Knoester, J.; Jansen, T. L. C. *The Journal of Physical Chemistry A* **2013**, *117*, 5970–5980.

(70) Atsango, A. O.; Montoya-Castillo, A.; Markland, T. E. *The Journal of Chemical Physics* **2023**, *158*, 074107.

(71) Allen, M. P.; Tildesley, D. J., *Computer Simulation of Liquids*; Oxford University Press: 1987.

(72) Pérez de Tudela, R.; Aoiz, F. J.; Suleimanov, Y. V.; Manolopoulos, D. E. *The Journal of Physical Chemistry Letters* **2012**, *3*, 493–497.

(73) Rossi, M.; Liu, H.; Paesani, F.; Bowman, J.; Ceriotti, M. *The Journal of Chemical Physics* **2014**, *141*, 181101.

(74) Green, M. S. *The Journal of Chemical Physics* **1954**, *22*, 398–413.

(75) Kubo, R. *Journal of the Physical Society of Japan* **1957**, *12*, 570–586.

(76) Kubo, R.; Yokota, M.; Nakajima, S. *Journal of the Physical Society of Japan* **1957**, *12*, 1203–1211.

(77) Yamamoto, T. *The Journal of Chemical Physics* **1960**, *33*, 281–289.

(78) Miller, W. H. *The Journal of Chemical Physics* **1974**, *61*, 1823–1834.

(79) Miller, W. H.; Schwartz, S. D.; Tromp, J. W. *The Journal of Chemical Physics* **1983**, *79*, 4889–4898.

(80) Verlet, L. *Physical Review* **1967**, *159*, 98–103.

(81) Hone, T. D.; Izvekov, S.; Voth, G. A. *The Journal of Chemical Physics* **2005**, *122*, 054105.

(82) Cao, J.; Voth, G. A. *The Journal of Chemical Physics* **1994**, *101*, 6168–6183.

(83) Hone, T. D.; Rossky, P. J.; Voth, G. A. *The Journal of Chemical Physics* **2006**, *124*, 154103.

(84) Musil, F.; Zaporozhets, I.; Noé, F.; Clementi, C.; Kapil, V. *The Journal of Chemical Physics* **2022**, *157*, 181102.

(85) Loose, T. D.; Sahrmann, P. G.; Voth, G. A. *Journal of Chemical Theory and Computation* **2022**, *18*, 5856–5863.

(86) Tuckerman, M. E.; Berne, B. J.; Rossi, A. *The Journal of Chemical Physics* **1991**, *94*, 1465–1469.

(87) Trenins, G.; Haggard, C.; Althorpe, S. C. *The Journal of Chemical Physics* **2022**, *157*, 174108.

(88) Wang, Y.; Bowman, J. M. *The Journal of Chemical Physics* **2011**, *134*, 154510.

(89) Wang, Y.; Bowman, J. M. *The Journal of Chemical Physics* **2012**, *136*, 144113.

(90) Partridge, H.; Schwenke, D. W. *The Journal of Chemical Physics* **1997**, *106*, 4618–4639.

(91) Soper, A. *Chemical Physics* **1996**, *202*, 295–306.

(92) Rotenberg, B. *The Journal of Chemical Physics* **2020**, *153*, 150902.

(93) Coles, S. W.; Mangaud, E.; Frenkel, D.; Rotenberg, B. *The Journal of Chemical Physics* **2021**, *154*, 191101.

(94) Reith, D.; Pütz, M.; Müller-Plathe, F. *Journal of Computational Chemistry* **2003**, *24*, 1624–1636.

(95) Hanke, M. *Journal of Statistical Physics* **2017**, *170*, 536–553.

(96) Hadley, K. R.; McCabe, C. *The Journal of Chemical Physics* **2010**, *132*, 134505.

(97) Henderson, R. *Physics Letters A* **1974**, *49*, 197–198.

(98) Evans, R. *Molecular Simulation* **1990**, *4*, 409–411.

(99) Trenins, G. Private Communication, 2023.

(100) Plé, T.; Huppert, S.; Finocchi, F.; Depondt, P.; Bonella, S. *The Journal of Chemical Physics* **2021**, *155*, 104108.

(101) Benson, R. L.; Althorpe, S. C. *The Journal of Chemical Physics* **2021**, *155*, 104107.

(102) Xiang, B.; Ribeiro, R. F.; Dunkelberger, A. D.; Wang, J.; Li, Y.; Simpkins, B. S.; Owrutsky, J. C.; Yuen-Zhou, J.; Xiong, W. *Proceedings of the National Academy of Sciences* **2018**, *115*, 4845–4850.

(103) Thomas, A.; George, J.; Shalabney, A.; Dryzhakov, M.; Varma, S. J.; Moran, J.; Chervy, T.; Zhong, X.; Devaux, E.; Genet, C.; Hutchison, J. A.; Ebbesen, T. W. *Angewandte Chemie* **2016**, *128*, 11634–11638.

(104) Krainova, N.; Grede, A. J.; Tsokkou, D.; Banerji, N.; Giebink, N. C. *Physical Review Letters* **2020**, *124*, 177401.

(105) Yang, P.-Y.; Cao, J. *The Journal of Physical Chemistry Letters* **2021**, *12*, 9531–9538.

(106) Onsager, L. *Journal of the American Chemical Society* **1936**, *58*, 1486–1493.

(107) Kirkwood, J. G. *The Journal of Chemical Physics* **1939**, *7*, 911–919.

(108) Adams, D. J. *Nature* **1981**, *293*, 447–449.

(109) Bertie, J. E.; Lan, Z. *Applied Spectroscopy* **1996**, *50*, 1047–1057.

(110) Fukushima, T.; Yoshimitsu, S.; Murakoshi, K. *Journal of the American Chemical Society* **2022**, *144*, 12177–12183.

(111) Bertie, J. E.; Ahmed, M. K.; Eysel, H. H. *The Journal of Physical Chemistry* **1989**, *93*, 2210–2218.

(112) Centurioni, E. *Applied Optics* **2005**, *44*, 7532.

(113) Német, N.; White, D.; Kato, S.; Parkins, S.; Aoki, T. *Physical Review Applied* **2020**, *13*, 064010.

(114) Fox, M., *Optical Properties of Solids*, 2nd ed.; Oxford University Press: 2010.

(115) Schwennicke, K.; Koner, A.; Pérez-Sánchez, J. B.; Xiong, W.; Giebink, N. C.; Weichman, M. L.; Yuen-Zhou, J. *Chemical Society Reviews* **2025**, *54*, 6482–6504.

(116) Hamm, P.; Zanni, M., *Concepts and Methods of 2D Infrared Spectroscopy*; Cambridge University Press: 2011.

(117) Li, T. E.; Nitzan, A.; Subotnik, J. E. *Nature Communications* **2022**, *13*, 4203.

(118) Mukamel, S., *Principles of Nonlinear Optical Spectroscopy*; Oxford University Press: 1999.

(119) Sawada, S.-I.; Nitzan, A.; Metiu, H. *Physical Review B* **1985**, *32*, 851–867.

(120) Parandekar, P. V.; Tully, J. C. *Journal of Chemical Theory and Computation* **2006**, *2*, 229–235.

(121) Tully, J. C. *The Journal of Chemical Physics* **1990**, *93*, 1061–1071.

(122) Runeson, J. E.; Manolopoulos, D. E. *The Journal of Chemical Physics* **2023**, *159*, 094115.

(123) Mannouch, J. R.; Richardson, J. O. *The Journal of Chemical Physics* **2023**, *158*, 104111.

(124) Schwartz, B. J.; Bittner, E. R.; Prezhdo, O. V.; Rossky, P. J. *The Journal of Chemical Physics* **1996**, *104*, 5942–5955.

(125) Fang, J.-Y.; Hammes-Schiffer, S. *The Journal of Physical Chemistry A* **1999**, *103*, 9399–9407.

(126) Wong, K. F.; Rossky, P. J. *The Journal of Chemical Physics* **2002**, *116*, 8418–8428.

(127) Jasper, A. W.; Truhlar, D. G. *The Journal of Chemical Physics* **2005**, *123*, 064103.

(128) Jain, A.; Alguire, E.; Subotnik, J. E. *Journal of Chemical Theory and Computation* **2016**, *12*, 5256–5268.

(129) Meyera), H.-D.; Miller, W. H. *The Journal of Chemical Physics* **1979**, *70*, 3214–3223.

(130) Meyer, H.-D.; Miller, W. H. *The Journal of Chemical Physics* **1979**, *71*, 2156–2169.

(131) Stock, G.; Thoss, M. *Physical Review Letters* **1997**, *78*, 578–581.

(132) Miller, W. H.; Cotton, S. J. *Faraday Discussions* **2016**, *195*, 9–30.

(133) Lawrence, J. E.; Mannouch, J. R.; Richardson, J. O. *The Journal of Physical Chemistry Letters* **2024**, *15*, 707–716.

(134) Weiße, A.; Fehske, H. In *Computational Many-Particle Physics*, Fehske, H., Schneider, R., Weiße, A., Eds.; Springer Berlin Heidelberg: 2008, pp 529–544.

(135) Hamm, P. Principles of Nonlinear Optical Spectroscopy: A Practical Approach, 2005.

(136) Mannouch, J. R.; Richardson, J. O. *The Journal of Chemical Physics* **2022**, *156*, 024108.

(137) Schlau-Cohen, G.; Ishizaki, A.; Fleming, G. R. *Chemical Physics* **2011**.

(138) Kramer, T.; Noack, M.; Reinefeld, A.; Rodríguez, M.; Zelinskyy, Y. *Journal of Computational Chemistry* **2018**, *39*, 1779–1794.

(139) Wang, H.; Thoss, M. *The Journal of Physical Chemistry A* **2003**, *107*, 2126–2136.

(140) Berkelbach, T. C.; Reichman, D. R.; Markland, T. E. *The Journal of Chemical Physics* **2012**, *136*, 034113.

(141) Wang, H.; Sun, X.; Miller, W. H. *The Journal of Chemical Physics* **1998**, *108*, 9726–9736.

(142) Grunwald, R.; Kelly, A.; Kapral, R. In *Energy Transfer Dynamics in Biomaterial Systems*; Springer: 2009.

(143) Runeson, J. E.; Richardson, J. O. *The Journal of Chemical Physics* **2020**, *152*, 084110.

(144) Runeson, J. E.; Mannouch, J. R.; Amati, G.; Fiechter, M. R.; Richardson, J. O. *CHIMIA* **2022**, *76*, 582.

(145) Egorov, S. A.; Rabani, E.; Berne, B. J. *The Journal of Chemical Physics* **1999**, *110*, 5238–5248.

(146) Shi, Q.; Geva, E. *The Journal of Chemical Physics* **2005**, *122*, 064506.

(147) Lieberherr, A. Z.; Kelly, J. E.; Runeson, J. E.; Markland, T. E.; Manolopoulos, D. E., `https://github.com/alieberherr/2des-ehrenfest-mash`, currently a private repository.

(148) Hein, B.; Kreisbeck, C.; Kramer, T.; Rodríguez, M. *New Journal of Physics* **2012**, *14*, 023018.

(149) Berkelbach, T.; Fetherolf, J.; Shih, P.; Iansdunn, `https://github.com/berkelbach-group/pyrho`, 2020.

(150) Runeson, J. E.; Richardson, J. O. *The Journal of Chemical Physics* **2019**, *151*, 044119.

(151) Lawrence, J. E.; Mannouch, J. R.; Richardson, J. O. *The Journal of Chemical Physics* **2024**, *160*, 244112.

(152) Kimura, G. *Journal of the Physical Society of Japan* **2003**, *72*, 185–188.

(153) Kimura, G.; Kossakowski, A. *Open Systems & Information Dynamics* **2005**, *12*, 207–229.

(154) London, N.; Limbu, D. K.; Faruque, M. O.; Shakib, F. A.; Momeni, M. R. *The Journal of Physical Chemistry A* **2025**, *129*, 4015–4028.

(155) Limbu, D. K.; London, N.; Faruque, M. O.; Momeni, M. R. *The Journal of Chemical Physics* **2025**, *162*, 014111.

(156) Nguyen, H.; Mandal, A.; Mahajan, A.; Reichman, D. R. Mixed Quantum-Classical Methods for Polaron Spectral Functions, 2025.

Appendix

A Ring polymer normal modes

In the absence of a molecular potential, the ring polymer Hamiltonian is diagonalised by its normal modes $\tilde{\boldsymbol{P}}, \tilde{\boldsymbol{Q}}$, which are related to the spatial coordinates $\boldsymbol{P}, \boldsymbol{Q}$ via

$$\begin{aligned} \tilde{\boldsymbol{Q}} &= C^T \boldsymbol{Q} \\ \tilde{\boldsymbol{P}} &= C^T \boldsymbol{P}. \end{aligned} \tag{5.1}$$

The free ring polymer Hamiltonian expressed in terms of normal modes is

$$\mathcal{H}_P(\tilde{\boldsymbol{P}}, \tilde{\boldsymbol{Q}}) = \sum_{k=0}^{P-1} \sum_{i=1}^{N} \left(\frac{|\tilde{\boldsymbol{P}}_i^{(k)}|^2}{2m_i} + \frac{1}{2} m_i \omega_k |\tilde{\boldsymbol{Q}}_i^{(k)}|^2 \right), \tag{5.2}$$

where $\omega_k = 2\omega_P \sin \frac{\pi k}{P}$. If P is even, the orthogonal transformation matrix is

$$C_{jk} = \begin{cases} \sqrt{\frac{1}{P}} & k = 0 \\ \sqrt{\frac{2}{P}} \cos \frac{2\pi jk}{P} & 1 \leq k \leq \frac{P}{2} - 1 \\ \sqrt{\frac{1}{P}} (-1)^j & k = \frac{P}{2} \\ \sqrt{\frac{2}{P}} \sin \frac{2\pi jk}{P} & \frac{P}{2} + 1 \leq k \leq P - 1 \end{cases}, \tag{5.3}$$

and if P is odd, it is

$$C_{jk} = \begin{cases} \sqrt{\frac{1}{P}} & k = 0 \\ \sqrt{\frac{2}{P}} \cos \frac{2\pi jk}{P} & 1 \leq k \leq \frac{P-1}{2} \\ \sqrt{\frac{2}{P}} \cos \frac{2\pi jk}{P} & \frac{P-1}{2} \leq k \leq P - 1 \end{cases} \tag{5.4}$$

$1 \leq j \leq P$ enumerates the spatial coordinates, and $0 \leq k \leq P-1$ enumerates the normal modes. Since C is an orthogonal matrix, the back-transformation from normal modes to spatial coordinates is

$$\begin{aligned} \boldsymbol{Q} &= C\tilde{\boldsymbol{Q}}, \\ \boldsymbol{P} &= C\tilde{\boldsymbol{P}}. \end{aligned} \tag{5.5}$$

One normal mode ($k = 0$), for which $\omega_k = 0$, is special. It is related to the centroid of the ring polymer, $\overline{\boldsymbol{Q}} = \frac{1}{P}\sum_{j=1}^{N} \boldsymbol{Q}^{(j)}$ through

$$\tilde{\boldsymbol{Q}}^{(0)} = \sum_{j=1}^{P} C_{j0}\boldsymbol{Q}^{(j)} = \sqrt{\frac{1}{P}}\sum_{j=1}^{P} \boldsymbol{Q}^{(j)} = \sqrt{P}\,\overline{\boldsymbol{Q}}. \tag{5.6}$$

The derivative with respect to the centroid mode is

$$\frac{\partial}{\partial \tilde{\boldsymbol{Q}}^{(0)}} = \sum_{j=1}^{P} \underbrace{\left(\frac{\partial \boldsymbol{Q}^{(j)}}{\partial \tilde{\boldsymbol{Q}}^{(0)}}\right)}_{=\sqrt{\frac{1}{P}}\ \text{(Eqs. (5.3), (5.4))}} \frac{\partial}{\partial \boldsymbol{Q}^{(j)}} = \sqrt{\frac{1}{P}}\sum_{j=1}^{P} \frac{\partial}{\partial \boldsymbol{Q}^{(j)}}, \tag{5.7}$$

and the derivative with respect to the centroid coordinate is

$$\frac{\partial}{\partial \overline{\boldsymbol{Q}}} = \sqrt{P}\frac{\partial}{\partial \tilde{\boldsymbol{Q}}^{(0)}} = \sum_{j=1}^{P} \frac{\partial}{\partial \boldsymbol{Q}^{(j)}}. \tag{5.8}$$

B Note on spin mapping initial conditions

Section 4.7, which describes the spin-mapping simulation of two-dimensional electronic spectra, describes an initialisation of the electronic mapping variables with focused initial conditions. In practice, a different scheme is used with the aim of improving the convergence of spin mapping results. This scheme, which we have termed 'doubly focused sampling', is described here and was used to compute the results presented in the thesis.

Doubly focused sampling

At the beginning of the t_2 propagation, the density matrix of interest is $|\phi\rangle\langle\phi|$, where $|\phi\rangle$ is in the singly excited manifold and has been obtained from the equatorial pure state decom-

position. (The case where $|\phi\rangle$ is proportional to the ground state is not considered here, since the spin mapping and Ehrenfest treatments are identical in that case.) First, we normalise the state to obtain $|\tilde{\phi}\rangle\langle\tilde{\phi}|$, where $|\tilde{\phi}\rangle = \langle\phi|\phi\rangle^{-1/2}|\phi\rangle$. The goal is now to sample a set of N vectors $|c_n\rangle$, such that

$$\frac{1}{N}\sum_{n=1}^{N}\left[\sqrt{N+1}|c_n\rangle\langle c_n| - \frac{\gamma_{\mathrm{W}}}{2}\hat{I}_N\right] = |\tilde{\phi}\rangle\langle\tilde{\phi}|, \tag{5.9}$$

with $\gamma_{\mathrm{W}} = \frac{2}{N}(\sqrt{N+1}-1)$ such that the initial density matrix is reproduced exactly with the N samples.

It is easiest to do all this in a basis $|k\rangle$, $k = 1,\ldots,N$, where $|\tilde{\phi}\rangle$ is the eigenvector with $k = 1$. This can be done via a Gram-Schmidt orthonormalisation of $N-1$ randomly generated $SU(M)$ coherent states. In this basis, the expansion coefficients of the states $|c_n\rangle$ are $c_{n,k} = \langle k|c_n\rangle$, and the coefficients of the target state $|\tilde{\phi}\rangle$ are $\langle k|\tilde{\phi}\rangle = \delta_{1,k}$. For these coefficients, Equation (5.9) becomes

$$\frac{1}{N}\sum_{n=1}^{N}\left[\sqrt{N+1}c_{n,k}c^{*}_{n,k'} - \frac{\gamma_{\mathrm{W}}}{2}\delta_{k,k'}\right] = \delta_{1,k}\delta_{1,k'}. \tag{5.10}$$

The diagonal entries ($k = k'$) constrain the square moduli of the coefficients,

$$\frac{1}{N}\sum_{n=1}^{N}[\sqrt{N+1}|c_{n,k}|^2 - \frac{\gamma_{\mathrm{W}}}{2}] = \delta_{1,k}. \tag{5.11}$$

We assume that the moduli do not depend on n. Equation (5.11) then gives

$$|c_{n,k=1}|^2 = \alpha \qquad\qquad |c_{n,k\neq 1}|^2 = \beta, \tag{5.12}$$

with $\alpha = \frac{2+\gamma_{\mathrm{W}}}{2\sqrt{N+1}}$ and $\beta = \frac{\gamma_{\mathrm{W}}}{2\sqrt{N+1}}$. Note that these coefficients give normalised states, since

$$\alpha + (N-1)\beta = \frac{2+N\gamma_{\mathrm{W}}}{2\sqrt{N+1}} = 1. \tag{5.13}$$

Let us now turn to the conditions on the off-diagonal ($k \neq k'$) contributions. Having deter-

mined the magnitude of the coefficients $c_{n,k}$, we can make the general ansatz

$$c_{n,k=1} = \sqrt{\alpha} e^{i\phi_{n,k=1}} \qquad\qquad c_{n,k>1} = \sqrt{\beta} e^{i\phi_{n,k>1}}. \tag{5.14}$$

Inserting it into Equation (5.10) gives

$$\frac{1}{N} \sum_{n=1}^{N} e^{i(\phi_{n,k} - \phi_{n,k'})} = 0, \tag{5.15}$$

where we have omitted constant prefactors as they are irrelevant to the result. A general set of phases that satisfies Equation (5.15) is

$$\phi_{n,k} = \phi_k + (n-1)(k-1)\frac{2\pi}{N}, \tag{5.16}$$

where the ϕ_k are sampled randomly between 0 and 2π. After sampling all electronic coefficients, the states are transformed from the Gram-Schmidt basis back into the site basis.

This sampling strategy increases the computational effort by a factor N, but we expect that it reduces the number of total samples necessary for convergence. Futher, we call it doubly focused sampling because it is an extension of focused sampling, which constrains only the moduli of the electronic coefficients. In fact, their magnitudes are identical to the ones used here (Equation (5.12)). Focused sampling is usually used to efficiently calculate population dynamics. However, since the third order response functions depend on the full density matrix, additional constraints on the off-diagonal elements should improve their convergence.