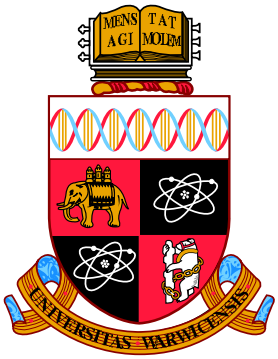

Sampling minimal, adaptive basis sets for multidimensional, nuclear quantum dynamics using simple, semi-classical trajectories.

Maximilian Alexander Christian Saller

Thesis submitted to the University of Warwick towards award of the degree of Doctor of Philosophy.

Department of Chemistry

July 2017

Contents

List of Tables

List of Figures

List of Abbreviations

AI-MCE	*Ab initio* multi-configuration Ehrenfest
AIMC	*Ab initio* multiple cloning
AIMS	*Ab initio* multiple spawning
ATP	Adenosine triphosphate
aTSA	Adaptive trajectory sampling algorithm
BEL-MCG	Basis expansion leaping multi-configuration Ehrenfest
CI	Configuration interaction
DD-vMCG	Direct dynamics variational multi-configuration Gaussian
DOF	Degree of freedom
DVR	Discrete variable representation
EM	Electromagnetic
G-MCTDH	Gaussian multi-configuration time-dependent Hartree
GWP	Gaussian wavepacket
MAE	Mean absolute error
MAPE	Mean absolute percentage error
MCE	Multi-configuration Ehrenfest
MCTDH	Multi-configuration time-dependent Hartree
MD	Molecular dynamics
ML-MCTDH	Multi layer multi-configuration time-dependent Hartree
MP	Matching pursuit
MS	Multiple spawning
NADH	Nicotinamide adenine dinucleotide
OPV	Organic photovoltaic
PES	Potential energy surface
PIMD	Path integral molecular dynamics
QTGB	Quantum trajectories Gaussian basis
SOFT	Split operator Fourier transform
SPF	Single particle function
TDH	Time-dependent Hartree
TDSCFT	Time-dependent self-consistent field theory
TDSE	Time-dependent Schrödinger equation
TISE	Time-independent Schrödinger equation
TSA	Trajectory sampling algorithm
UV	Ultraviolet
vMCG	Variational multi-configuration Gaussian

Acknowledgements

First and foremost I must thank Scott Habershon, for his continued support and guidance over the last four years. I am truly grateful to him for giving me the opportunity to work in a field both exciting and interesting to me. Over the years he has been a truly reliable port of call for questions and advice, both guiding my research efforts and helping me become more independent as a scientist. He has been monumentally patient in putting up with my oddities and I truly feel that his mentoring has given me the chance to seriously pursue a career in science.

Throughout most of my years at Warwick I was joined by a co-conspirator of sorts, to whom I owe a non-trivial debt of gratitude. Lewis Baker has, amongst other things, been a true friend, a brother in arms and, on more than one occasion, the voice of reason to my insanity. Whether it be on, mostly voluntary, outreach assignments or debating intently over what some might argue are trivial or even moot points, his company has always served to keep me sane. I seriously doubt I will meet many people throughout my career so willing to lend an ear or give an opinion on any matter, scientific or not.

I have been extremely lucky with respect to the working environment and my colleagues at Warwick. For four years I have felt that the office was not just a place of work but also a second home, full of friends who were always willing to help and listen. So in no particular order, I would like to extend my deepest thanks to David, Sam, Kirsten, Gareth, Anna-Laura, Jasmine and Christopher.

Finally I would like to thank the Engineering and Physical Sciences Research Council for funding my research project, as well as the Centre for Scientific Computing and the Scientific Computing Research Technology Platform for providing computing resources and infrastructure.

An Wilhelm, Sieglinde, Harald und Elisabeth: Ohne euch gäbe es nichts! Von ganzem Herzen danke ich euch für 26 Jahre Unterstützung, Förderung und Glauben an mich!

Declaration

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree. The work presented (including data generated and data analysis) was carried out by me except in the cases outlined below.

Maximilian Saller, July 2017

Data from the following publications was used for comparative purposes, appropriate citations indicating the origin in every such instance.

- A. Raab, G. A. Worth, H.-C. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 1999, **110**, 936-946
  The author acknowledges the use of multi-configuration time-dependent Hartree results for the quantum dynamics of the 4 and 24 dimensional vibronic pyrazine Hamiltonian.

- S. Habershon, *J. Chem. Phys.*, 2012, **135**, 054109
  The author acknowledges the use of configuration interaction results for the quantum dynamics of the synthetic double well tunnelling Hamiltonian.

Parts of this thesis have been published by the author:

- M. A. C. Saller and S. Habershon, *J. Chem. Theo. Comput.*, 2015, **11**, 8-16

- M. A. C. Saller and S. Habershon, *J. Chem. Theo. Comput.*, 2017, **13**, 3085-3096

Abstract

Methods for the study of nuclear quantum dynamics can be categorised by the nature of the basis set expansion they employ. The wavefunction can be expanded in a static set of time-independent basis functions, the time evolution being described solely *via* the expansion coefficients. Alternatively, basis functions can be propagated in time, along with the coefficients, *via* equations of motion for their parameters. Time-independent basis sets are plagued by exponential scaling, while the equations of motion for time-dependent basis functions are challenging to integrate and, if not derived variationally, can violate energy conservation laws. This work presents a novel basis set sampling method which represents a compromise between these two categories. A set of sampling trajectories, evolving on the potential energy surface of the system, are used to place basis functions in regions of phase space, relevant to wavefunction propagation. These functions then act as a time-independent basis set, the wavefunction being evolved *via* exact, variational equations of motion for the expansion coefficients. This approach is applied to a challenging quantum dynamics benchmark, namely the relaxation dynamics of photoexcited pyrazine, and yields highly encouraging results. In order to address divergence from exact dynamics at longer timescales, which is attributed to the classical sampling trajectories being a sound approximation to quantum propagation of the wavefunction only in the short-time limit, a modification of this method is proposed. Shorter iterations of trajectory sampling and wavefunction propagation are used, linked by a minimisation algorithm that continuously optimises the basis set, preventing unfavourable scaling. This adaptive sampling approach is again applied to the pyrazine benchmark with a significant increase in performance and accuracy. Highly encouraging results are also obtained for a quantum tunnelling benchmark system, which are improved upon even further, and at little extra cost, by the use of path integral sampling trajectories.

Chapter 1

Introduction

This chapter provides an introductory overview of the field in which the work presented herein resides. Light-matter interactions, which are the ultimate target of this research, are briefly discussed and their importance in science and technology highlighted by means of three examples. In the context of existing methods for the study of photophysics and photochemistry, some common drawbacks and limitations are highlighted, which this work seeks to address. Finally, the presentation of the remainder of this work is briefly discussed, giving an overview of the chapters to follow.

1.1 Interactions of light and matter

Given the continued discovery of quantum phenomena, the potential applications of the work presented herein are numerous. Quantum effects have been found to either constitute the driving force of, or having significant impact on, physical processes covering virtually all fields of science.[1–4] However, the one field of particular interest that this work is aimed towards is photophysics and by extension photochemistry. As suggested by their names, deriving from the Greek phós ($\varphi\hat{\omega}\varsigma$) for light, these disciplines are concerned with treating the interaction of light and matter. Of the many physical processes known to science, there are few which have as profound and widespread an impact as this interaction.[5] Furthermore, as, at least on a timescale reasonable to human civilisation, the sun can be considered an infinite source of energy in the form of light, the potential for technological development, harnessing this power, certainly dwarfs that of the heavily relied upon fossil fuels of the current era, but likely also outshines that of other renewable energy sources.

To support these statements, the harnessing of light, both in nature, *via* the process of photosynthesis, and in modern technology, which aims to mimic the former through solar cells, will be briefly discussed.

1.1.1 Photosynthesis

The definition of the term "photosynthesis" varies significantly depending on the scientific context in which it is used, and has evolved over time, as scientific discoveries have continuously broadened the array of processes it might reasonably be applied to.[6] For the purpose of this brief discussion, the interpretation adopted here will be that of a mechanism, by which an organism captures energy in the form of light and uses it to drive chemical reactions.

Of the many approaches that may be taken to introduce the concept of photosynthesis, the one possibly most indicative of its impact on all life on earth, and indeed the one chosen here, is to investigate its emergence during the early stages of earth's planetary history and the very beginnings of life. While their origin is one of the most highly debated topics in science, the earliest life forms, for which no fossil record exists, are thought to have relied on the chemical environment in which they existed, in order to survive and drive the processes of metabolism, information storage and replication, the presence of which constitutes the most common biological definition of life.[7]

Some of the earliest confirmed evidence of life are stromatolites, which are thought to be the result of photosynthetic processes. The oldest examples of these structures date back around 3.5 billion years, suggesting that during this

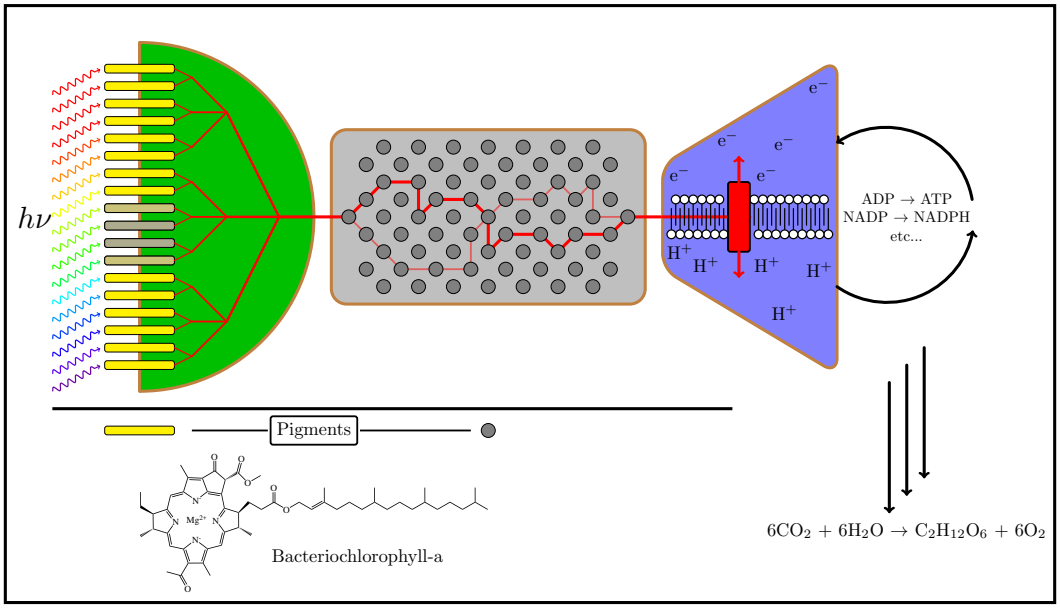

Figure 1.1: The general mechanism of photosynthesis involves the capture of light *via* pigments in the chlorosome antenna complex, from which the electronic energy is then transported *via* a network of more pigment molecules to the reaction centre. There, this energy is used to drive chemical reactions such as the recycling of the coenzymes adenosine triphosphate (ATP) and nicotinamide adenine dinucleotide (NADH).

early stage of life, photosynthesis had, in some form or another, already evolved.[8] While a number of theories for the evolution of the complex metabolic pathways and chemical structures involved with photosynthesis exist,[9] the importance of photosynthetic processes in driving evolution is widely accepted. Harnessing the sun as their major source of energy allowed organisms to become more independent from their local chemical environment and eventually allowed them to develop more complex metabolic pathways, which in turn paved the way for more complex life to evolve.[10]

The number of complex chemical and physical processes involved in photosynthesis is, as for most biological phenomena that are the result of millions of years of evolution, numerous. A general, diagrammatic outline of some of the stages of photosynthesis is shown in Figure 1.1. For many of these processes, theoretical and computational methods can contribute significantly towards the elucidation of the complex underlying mechanisms. To name just one example, a key aspect of the photosynthetic pathway involves the transfer of energy from the regions of the organism, responsible for harvesting it, to the areas which undertake the chemical reactions that are powered by it. This excitation energy transfer has been linked to quantum effects[11] and recent work has shown that a network-based quantum

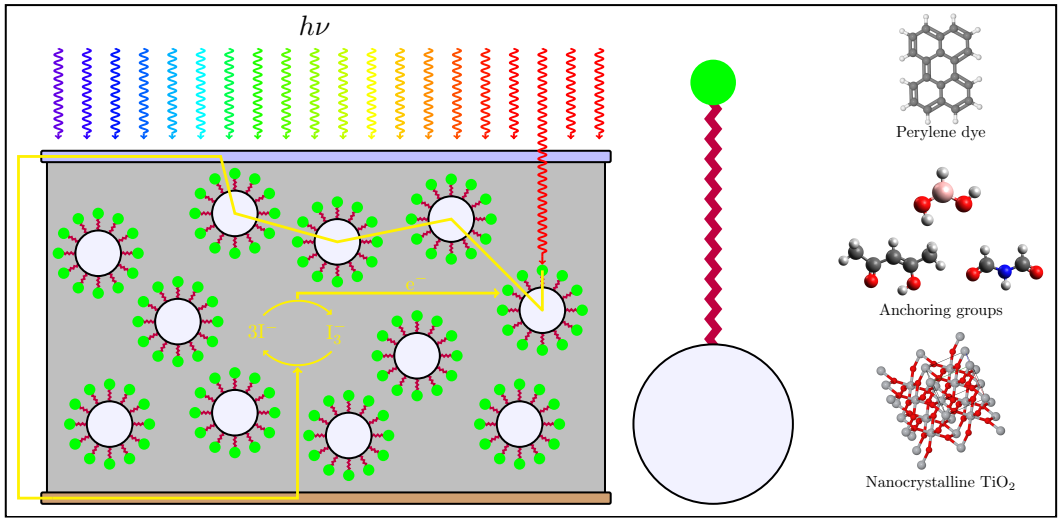

Figure 1.2: Dye sensitised solar cells typically consist of TiO_2 nanoparticles to which light-absorbing, organic dye molecules, like perylene, are attached *via* small organic anchoring groups, which are surrounded by an electrolyte medium. Light is absorbed by the dye molecules, which are oxidised and the resulting electron is transferred, *via* the conduction band of the TiO_2 particles, to the electrode. The ionised dye molecules accept an electron from the electrolyte solution, which in turn, after diffusion to the opposing electrode, is recycled by the electrons re-entering the circuit.

---

dynamics model significantly helps in accounting for and expanding further on experimental observations.[12]

The evolution of the ability to capitalise, *via* photosynthesis, on interactions of light and matter should be considered one of the major paradigm shifting developments that allowed life to flourish, from simple, single-celled organisms, to the vast complexity with which it is observed today. As with most biological processes that have been optimised to such an extent *via* evolution, this therefore has sufficient value to modern technology, as the next section shall discuss.

1.1.2 Solar cells

As mentioned above, the current reliance of human civilisation on a dwindling supply of fossil fuels is almost universally accepted as unsustainable.[13] Thus, the search for an alternative, cost effective renewable energy source is the focus of major international research efforts. Given that energy in the form of sunlight is available, essentially independent of location and more or less unlimited, photovoltaic solar cells thus theoretically constitute a potential solution to this problem. Furthermore, as theoretical and experimental research has elucidated many of the

basic concepts underlying natural light harvesting,[14,15] mimicking them in the context of modern technology and on a scale leading beyond pure research, towards industrial applicability, is increasingly being considered within reach.[16]

The materials used in the construction of photovoltaic cells vary widely, and significant progress has been made in increasing the efficiency for most materials and cell architectures.[17] The most common type of commercial photovoltaic cells are based on silicon which, while consistently leading other materials in terms of efficiency,[18] is also associated with non-trivial costs. Organic photovoltaics (OPVs) constitute a potentially far more cost-effective alternative to the aforementioned silicon cells, and significant research efforts are devoted to improving the efficiency of such devices.[19]

Charge transfer states, one of the key features of OPVs, constitute an excellent example of the benefit in the design of such devices that may be gained *via* theoretical study. Significant progress has been made in predicting the mechanism by which such charge transfer states are formed using theoretical, quantum mechanics based models[20,21] and the importance of quantum effects involving nuclear motion has been highlighted in the context of this process.[21]

Dye-sensitised solar cells are another example of OPVs where theoretical methods have been very successful in paving the way for more efficient and less time-consuming design of new technology. The key feature of such devices is the organic chromophore dye, often developed by mimicking the structural and chemical properties of photosynthetical light-harvesting complexes, which is bound to nanocrystalline TiO_2. Figure 1.2 illustrates the mechanism by which dye sensitised solar cells harvest light and convert it to electric power. There have been a number of very successful theoretical models proposed for these systems, including one which allows for rapid screening of potential new dye molecules,[22,23] as well as the anchoring group used to secure it to the TiO_2.

Furthermore, charge recombination, that is the migration of the electron from the TiO_2 semiconductor back to the chromophore that injected it, is the major source of efficiency loss for dye-sensitized solar cells. A theoretical model predicting the rate of this process for a given dye molecule has been developed,[24] again relying on quantum mechanical methods to screen potential dyes. The ability to predict the rate of charge recombination for a large number of dye molecules, without the need to experimental study, will both help reduce the costs and significantly speed up the development of more advance dye-sensitised solar cells.

1.1.3 Ultrafast spectroscopy and the role of theory

The design of OPVs is only one example of technological processes which require in-depth understanding of the photophysics of small organic molecules they feature. Recent advances in spectroscopy now allow increasingly detailed insights into the photodynamics of such molecules to be gained experimentally. Modern spectroscopic techniques allow the dynamics of atoms and molecules to be probed on the femtosecond scale,[25,26] which allows the finer details of many light-matter interactions to be investigated. While the level of detail obtainable from such ultrafast experiments is significant and deductions can often be made with regards to the underlying mechanisms, they are not able to provide insight into the dynamics on the atomistic and sub-atomic particle scale. This is why theoretical methods are often used in conjunction with experimental results, in order to gain a more complete understanding of the photophysical and photochemical interactions governing the behaviour of the system. This often involves multiple levels of one informing the other, such as theoretical levels which are built from spectroscopic data for a given molecule, which can then be used to further investigate and make predictions for systems of a similar nature.

An excellent, recent example of the benefits that may be reaped from combining theoretical simulations with ultrafast spectroscopic methods is the in depth investigation of the relaxation dynamics in ethylene.[27,28] Although comparatively simple, this system possesses two conical intersections between the ground and first excited, $\pi \rightarrow \pi^*$ electronic states, leading to non-trivial dynamics upon photoexcitation of the molecule. The first avenue of investigation concerned the lifetime of the excited state, which had previously been both measured experimentally, at ~ 50 fs,[29] and predicted using simulations, the latter however yielding a far longer lived excited state, at 89 fs.[30] This discrepancy between experimental results and theoretical predictions was attributed to two factors affecting the former. Firstly, the energy of the pump pulse, that is the radiation used to excited the molecule from its ground to its excited state, had previously been linked to the initial excited state geometry being close to one of the aforementioned conical intersections, thus artificially speeding up relaxation back to the ground state. Secondly, the probe pulse, which is used to ionise the molecule on the excited state to allow its detection, had previously involved multiphoton excitation, which significantly increases the complexity of the theoretical model required to account for the resulting dynamics.

In order to address these issues, a new set of spectroscopic measurements, using femtosecond vacuum ultraviolet pump probe spectroscopy, employing both a higher energy pump as well as a single photon excitation probe pulse, and

theoretical simulations using on-the-fly *ab initio* multiple spawning (see Chapter 2 for more detail), were carried out.[27] While resulting in far better agreement between the predicted excited state lifetimes, this investigation highlighted that the photoion yield, being the measured observable in the experiment, still decays faster than the excited state in simulations. This was linked to an energetic condition whereby the excited molecule is not universally ionizable, even with the higher energy probe pulse used.[27]

A theoretical study can, as mentioned above, also give significant insights at the atomic and molecular scale. The investigation described above also made use of this fact, by measuring the ratio of detected photoions. The higher energy probe pulse employed allowed breaking of the C-C bond in the excited ethylene molecule, which allows for more in depth conclusions to be drawn about the pathways *via* which the molecule decays, following $\pi \rightarrow \pi^*$ excitation, back to the ground state state. Obtaining excellent agreement between experiment and simulation for these ion ratios, experimental evidence for two, theoretically predicted, non-radiative decay pathways, one involving pyramidalization the other hydrogen migration across the C-C bond, were observed for the first time.[28]

The availability of quantum dynamics methods which can model light-matter interactions for chemically interesting systems, at reasonable computational cost, thus allows for more detailed conclusions about the detailed dynamics involved to be drawn. This is especially relevant given the rapid advances in the field of photoelectron spectroscopy,[25,26] which allow for combined experimental and theoretical studies like the one described above. This work seeks to address this need by introducing a novel quantum dynamics approach, aimed at avoiding existing issues of both limited accuracy and unfavourable computational scaling.

1.2 Theoretical methods in photophysics

Amongst the early efforts in the theoretical study of dynamic quantum effects, a significant proportion were focused on describing the interactions of molecules with light, particularly in the context of nuclear dynamics.[31–33] Given the wide scientific interest in light-matter interactions and the potential of theoretical methods to complement experimental studies of systems involving the former, significant progress has since been made in the field of nuclear quantum dynamics, which will be discussed in more detail in Chapter 2.

When studying the interactions of light and matter on the microscopic scale, the main challenge lies in solving the time-dependent Schrödinger equation (TDSE), which will be introduced again in more detail later. While classical mechanics, such as the widely used molecular dynamics method, is often sufficient to simu-

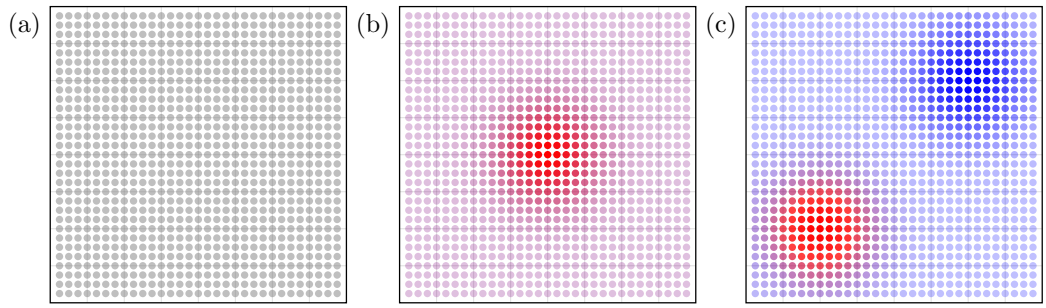

Figure 1.3: (a) Time independent quantum dynamics methods expand the wavefunction statically, (b) the expansion coefficients associated with each basis function result in the correct amplitudes in coordinate space and (c) time evolution of the wavefunction occurs *via* changing of these coefficients.

---

lated molecular processes, the nature of the electromagnetic field and its interaction with molecules, which is discussed in Chapter 2 in more detail, means that for photophysical processes, quantum mechanical effects must be taken into account. The TDSE states that

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = \hat{H}\Psi(\mathbf{r},t), \tag{1.1}$$

where $\Psi(\mathbf{r},t)$ is the wavefunction, defined in terms of the system (nuclear and electronic) coordinates, $\mathbf{r}$, at time t, $\hat{H}$ is the Hamiltonian operator and $\hbar$ is the reduced Planck constant. The natural complexity of the wavefunction requires it to be approximated *via* expansion in a set of mathematically simple functions, in order to solve the TDSE. Quantum dynamics methods can thus be categorised by the type of wavefunction expansion, or *ansatz*, they employ and how the resulting approximate wavefunction is then propagated in time.

1.2.1 Time-independent basis sets

As illustrated in Figure 1.3, the specific example given being the ground and excited state of the harmonic oscillator, one approach is to expand the wavefunction in terms of a large set of basis functions, the positions of which are and remain fixed, often on a grid, throughout the time evolution. The expansion coefficients associated with each function are then propagated in time, usually *via* equations of motion derived from application of a time-dependent variational principle[34–37] to the expanded wavefunction. The wavefunction *ansatz* for such methods may thus be written as

$$\Psi(\mathbf{r},t) \approx \psi(\mathbf{r};t) = \sum_j c_j(t)\phi_j(\mathbf{r})\,, \tag{1.2}$$

where $\phi(\mathbf{r})$ are the time-independent basis functions and $\mathbf{c}(t)$ are the associated, time-dependent expansion coefficients. A number of successful quantum dynamics

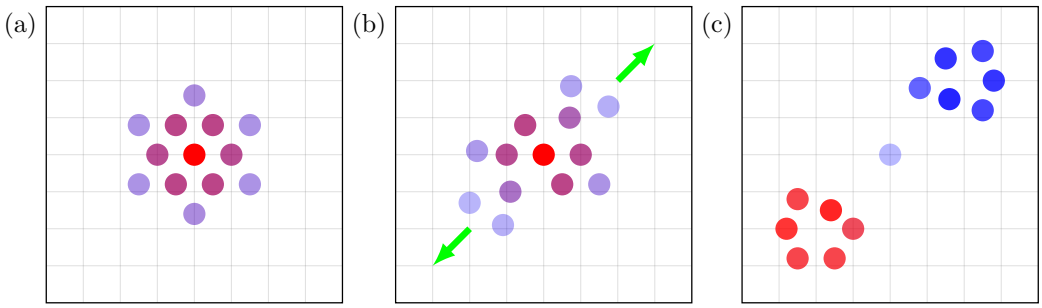

Figure 1.4: (a) The wavefunction is expanded as a small number of time-dependent basis functions which (b) move in phase space over time. (c) As a result, basis functions continuously adapt to the changing shape of the wavefunction.

approaches have been developed based on this rather simple scheme,[38–41] some of which will later be discussed in more detail.

While the aforementioned, relative simplicity of this expansion and propagation of the wavefunction is undeniable, it comes at great expense with respect to the size of the system that may be treated. As basis functions remain fixed in phase space and it usually is not possible to predict, without actually calculating dynamics, which regions of phase space will be relevant to the propagation of the wavefunction, the entirety of accessible phase space must be populated with the basis set. This, quite straightforwardly, leads to extremely unfavourable exponential scaling of the basis set size with respect to the dimensionality of the system. Without a scheme for adaptively changing which parts of the basis set need to be treated and which may be ignored at every step of wavefunction propagation, which do exist and are in part discussed in Chapter 2, methods based on this time-independent *ansatz* are thus limited to calculations of systems consisting of only a few atoms.[41]

1.2.2 Time-dependent basis sets

Figure 1.4 shows the second common wavefunction expansion and propagation scheme shared by many quantum dynamics approaches.[42–46] The basis functions are now propagated in time, *via* equations of motion derived for their parameters and expansion coefficients. The result of this is a dynamic basis set that moves in phase space alongside the wavefunction, which may thus be written as

$$\Psi(\mathbf{r},t) \approx \psi(\mathbf{r};t) = \sum_j c_j(t)\phi_j(\mathbf{r};t)\,, \tag{1.3}$$

where the basis functions, $\phi_j(\mathbf{r};t)$, are now time-dependent.

In comparison to time-independent basis sets, the advantage of this *ansatz*

originates from the continuously changing distribution of basis functions in phase space. As the basis set adapts to the shape of the wavefunction, basis functions naturally occupy the areas of phase space where the amplitude of the wavefunction is largest, so far fewer functions are required to describe it at any given time. Time-dependent methods not only avoid the exponential scaling discussed above, but in fact scale rather favourably with system size, especially when making use of multidimensional basis functions.

There are, however, some disadvantages associated with the time propagation of the basis functions. More specifically, the limiting factor for methods, relying on the *ansatz* above, are the equations of motion for the basis function parameters. If, similarly to the expansion coefficients, these equations are derived *via* application of a time-dependent variational principle to the TDSE featuring the expanded wavefunction, while yielding, within the finite basis limit, formally exact results, practically, their integration is numerically challenging, as they tend to be ill-conditioned. This issue may be avoided by instead deriving non-variational and thus classical or semi-classical equations of motion, however, these, in addition to introducing yet another approximation, have been shown to violate the energy conservation laws, implicit in the TDSE.[47]

This work presents an alternative method for wavefunction expansion and propagation for the purposes of quantum dynamics. While not strictly belonging to either of the above categories, it is inspired by methods from both and is best thought of as a compromise between their associated advantages and limitations.

1.3 Overview of the work presented hereafter

The remainder of this work is structured as follows. Chapter 2 contains a brief introduction to the theoretical background underlying the work presented in later chapters. A number of key historical and current quantum dynamics methods, either related to or in some way inspiring the work presented herein, are also discussed. Chapter 3 introduces a novel scheme for sampling quantum dynamics basis sets *via* simple trajectories and demonstrates their effectiveness for a challenging benchmark problem. Chapter 4 addresses one of the key assumptions of this new method, found to be the key factor limiting accuracy. A modification to the algorithm, designed to overcome said limitation, is presented and applied to two quantum dynamics benchmarks. Chapter 5 investigates the use of a different type of sampling trajectories, specifically aimed at improving the description of strong quantum effects such as tunnelling, within the novel sampling method by applying it to a benchmark system, designed specifically with tunnelling in mind. Finally, Chapter 6 summarises the methods and results presented in preceding chapters and outlines possible avenues for future work.

Chapter 2

Theory

This chapter lays out the essential theoretical background of the work presented thereafter, beginning with a discussion of the laws and equations of quantum mechanics. As the goal of this work is to investigate dynamic, time-dependent phenomena, an overview of the field of quantum dynamics is presented. Focus here is on current and historical methods, either similar in approach or somehow inspiring the strategies presented herein. Finally, given that the phenomena targeted by the work presented herein concern light-matter interactions, this chapter concludes with a brief section on the theoretical framework underlying photophysics.

2.1 Quantum mechanics

A number of experimental observations throughout the late nineteenth and early twentieth century highlighted the fact that the well established mathematical framework of classical mechanics was inadequate when describing the behaviour of elemental particles, atoms and molecules. The first notion of the *quantised* nature of such microscopic systems followed from experiments investigating black-body radiation. Only discrete, not continuous, excitations of the electromagnetic field, could account for the spectral densities, that is the relative intensity with which individual frequencies are observed in black-body radiation at a given temperature. A similar observation was made when heat capacities were found to be overestimated by a model representing each atom as a classical harmonic oscillator. Again the notion of quantisation could rationalise the lower than predicted heat capacities at low temperatures, if each harmonic oscillator required a specific, minimum amount of energy in order to contribute to heat transfer.

Further evidence for the quantisation of the electromagnetic field was the observation of the *photoelectric effect*, that is the emission of an electron from a charge surface, if exposed to ultraviolet radiation. Again this effect was only found to occur if the radiation applied exceeded a certain frequency, however, no such threshold was found for the intensity of radiation.

In order to explain the increasingly obvious link between electromagnetic radiation and the behaviour of matter on the microscopic scale, Louis de Broglie suggested that all moving particles were associated with a wave, the momentum, p, of the former being linked to the wavelength, λ, of the latter *via*

$$\lambda = \frac{h}{p}\,, \tag{2.1}$$

where h is the Planck constant. This proposed wave-like nature of matter was soon confirmed, when electrons were observed to display diffraction (see below), when projected at atomic crystal lattices.

Unifying observations from many experiments it was concluded that electrons, and, as became clear over time, all of the microscopic systems alluded to above, behave sometimes according to the familiar laws of classical mechanics, associated with macroscopic particles, but other times akin to waves. Quantum mechanics accounts for this duality in the behaviour of microscopic systems and links the mathematical framework governing it to macroscopic observations and phenomena.

2.1.1 Heisenberg's uncertainty principle

The wave-like nature of microscopic particles led Werner Heisenberg to propose that if, for example, an electron with linear momentum shares some characteristics with a moving wave, which is defined infinitely in space, then one should not be able to know exactly both said linear momentum and the position of the electron in space at any given time. This, eventually led to the celebrated Heisenberg uncertainty principle which states that,

$$\Delta x \Delta p_x \geq \frac{1}{2}\hbar\,, \tag{2.2}$$

where $\hbar$ is the reduced Planck constant and Δx is the root mean squared deviation of the position[a], that is to say

$$\Delta x = \sqrt{\langle x^2 \rangle - \langle x \rangle^2}\,,$$

where the notation $\langle x \rangle$ denotes the average of x. This may in fact be generalised for any two observables of the two non-commuting operators, $\hat{A}$ and $\hat{B}$, such that

$$\Delta A \Delta B \geq \frac{1}{2}\left|\left\langle \left[\hat{A}, \hat{B}\right] \right\rangle\right|\,, \tag{2.3}$$

$[\hat{A}, \hat{B}]$ denoting the commutator of $\hat{A}$ and $\hat{B}$, defined as

$$[\hat{A}, \hat{B}] = \hat{A}\hat{B} - \hat{B}\hat{A}\,,$$

an operator being generally defined as a function mapping states, or a space of states, to others.

2.1.2 Schrödinger's wavefunction approach

Erwin Schrödinger formulated a mathematical foundation of quantum mechanics, based on the concept of wavefunctions.[48–53] While the behaviour of these wavefunctions does derive from the well-established field of classical wave mechanics, their physical interpretation in quantum mechanics is somewhat more complex. That is to say for a wavefunction, $\Psi(\mathbf{r}, t)$, where $\mathbf{r}$ are the spatial coordinates of the system and t is time, only the square of this often complex-valued function, $|\Psi(\mathbf{r}, t)|^2$ has physical meaning. The wavefunction does contain all information about the system, which can be extracted using operators (for more detail see below).

Expressing quantum mechanics based on wave mechanics allowed for the observations described above to be rationalised in terms of superposition effects. The

[a] And *vice versa* of the momentum.

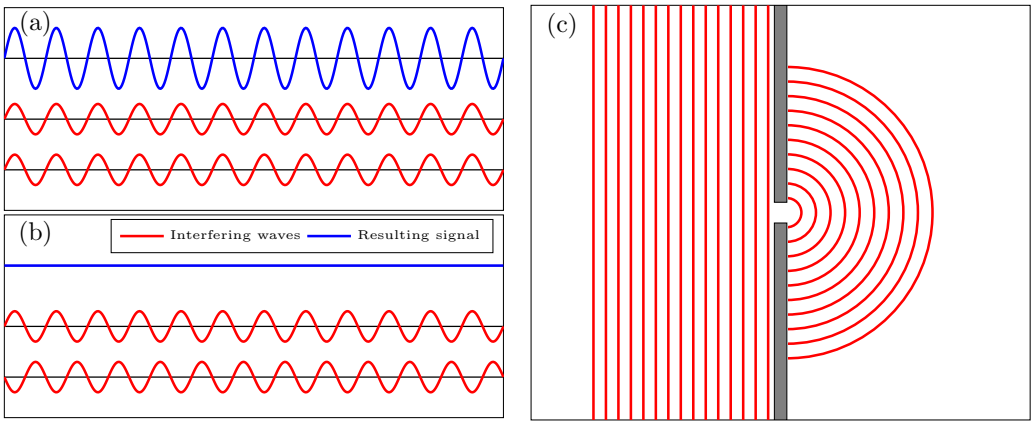

Figure 2.1: Superposition effects of waves: (a) Constructive, in phase, and (b) destructive, out of phase, interference of sine waves based on their phase; (c) diffraction in the context of the classic single slit experiment.

latter are a direct result of the superposition principle,

$$f(x+y) = f(x) + f(y)\,, \tag{2.4}$$

$$f(cx) = cf(x)\,, \tag{2.5}$$

where $f(x)$ is, by definition, a linear function and c is a scalar. In the context of waves, superposition results in constructive or destructive interference and diffraction,[b] as illustrated in Figure 2.1.

2.1.3 Postulates of quantum mechanics

In the context of the wavefunction, the postulates of quantum mechanics are as follows[c]:

1. Any given state of a system is defined by the corresponding wavefunction $\Psi(\mathbf{r}, t)$. Ψ and $\lambda\Psi$, where λ is a scalar, correspond to the same state and thus the wavefunction is usually chosen to be normalised, such that

$$\int \Psi^*(\mathbf{r}, t)\Psi(\mathbf{r}, t)\mathrm{d}\tau = \int |\Psi(\mathbf{r}, t)|^2 \mathrm{d}\tau = 1\,, \tag{2.6}$$

[b]Depending on the field, interference and diffraction are used for superposition phenomena, a common distinction being based on the number of sources. That is, superposition effects due to a small number of sources are often referred to as interference, while those due to a large number of sources are called diffraction. These terms are a matter of preference only.

[c]The number, order and phrasing of these postulates vary significantly in literature, but are all expressions of the underlying fundamental nature of quantum mechanics, simply approached in different contexts.

where $\int \mathrm{d}\tau$ refers to integration over the entire coordinate space of the wavefunction.

2. The probability dP of the particle (or system) being found in the volume element $d\tau$ at the point (or configuration) $\mathbf{r}$ and at time t, can be defined as being proportional to the square of the wavefunction as

$$dP \propto |\Psi(\mathbf{r}, t)|^2 d\tau \,. \tag{2.7}$$

This is the most common interpretation of the wavefunction, as the complex square will be a real, positive number while Ψ itself can be complex and negative.

3. For every observable, there exists a Hermitian operator, $\hat{A}$, satisfying

$$\int a^* \hat{A} b \,\mathrm{d}\tau = \left[\int b^* \hat{A} a \,\mathrm{d}\tau \right]^* , \tag{2.8}$$

where a and b are functions in the domain of $\hat{A}$.[d]

4. The only experimentally measurable values of the observable associated with operator, $\hat{A}$, will be the eigenvalues, a, of $\hat{A}$ arising from

$$\hat{A}\Psi = a\Psi \,, \tag{2.9}$$

assuming the system is in an eigenstate of $\hat{A}$, described by Ψ. In the case where Ψ is not an eigenstate of $\hat{A}$, it may, because the set $\{ \Psi_j \}$ is complete, be expanded as

$$\Psi = \sum_j c_j \Psi_j \,, \tag{2.10}$$

where Ψ_j are eigenfunctions of $\hat{A}$ with eigenvalues a_j and c_j are complex expansion coefficients, determining the weight of each eigenstate. In such a case, the average value of repeated experimental measurements may be calculated from the expectation value of the operator,

$$\langle \hat{A} \rangle = \frac{\int \Psi^* \hat{A} \Psi \mathrm{d}\tau}{\int \Psi^* \Psi \mathrm{d}\tau} \,, \tag{2.11}$$

in general, which can be simplified to $\int \Psi^* \hat{A} \Psi d\tau$, for any normalised wavefunction, where $\int |\Psi|^2 \mathrm{d}\tau = 1$.

[d] Eq. 2.8 formally defines a symmetric operator. While the two are often used interchangeably in literature, the definition of a true Hermitian, or self-adjoint operator is somewhat more complex.

5. The wavefunction evolves in time according to the time-dependent Schrödinger equation,

$$i\hbar\frac{\partial}{\partial t}\Psi = \hat{H}\Psi\,, \tag{2.12}$$

where $\hbar$ is the reduced Planck constant, $\hbar = \frac{h}{2\pi}$, and $\hat{H}$ is the Hamiltonian operator, corresponding to the total energy of the system.

In combination with Schrödinger's wavefunction based interpretation of quantum mechanics, these postulates formally allow the behaviour of microscopic systems to be described. Given knowledge of the wavefunction of any system, any experimental observables may be calculated using the corresponding Hermitian operators. Dynamic information may be extracted by solving the partial differential equation above and extracting the value of operators as the wavefunction evolves in time. However, owing to the fact that Ψ fully describes the state of the system, including any observable quantities that may be measured, wavefunctions of even the most simple molecular systems are mathematically too complex to allow Eq. 2.12 to be solved analytically. Thus a number of approximations are required.

2.1.4 Variable separation and wavepackets

In general, the wavefunction, $\Psi(\mathbf{r}, t)$ is a function of the system's spatial coordinates, $\mathbf{r}$, and time, t. A common approach in simplifying this function is to assume that solutions of the TDSE may be written as a product of functions of only spatial coordinates and time,

$$\Psi(\mathbf{r}, t) = \psi(\mathbf{r})\chi(t)\,. \tag{2.13}$$

Substituting (2.13) into (2.12) and separating the resulting equation by dependence on coordinates and time yields two equations for a time-independent Hamiltonian,

$$\hat{H}\psi(\mathbf{r}) = E\psi(\mathbf{r})\,, \tag{2.14}$$

$$i\hbar\frac{d}{dt}\chi(t) = E\chi(t)\,, \tag{2.15}$$

where E is the eigenvalue of $\hat{H}$, representing the total energy of the system, for $\psi(\mathbf{r})$, which are eigenfunctions of $\hat{H}$. Eq. 2.14 is the time-independent Schrödinger equation (TISE), which, for static systems, may be solved instead of the TDSE to obtain time-independent properties of the system *via* the wavefunction.

Solving Eq. 2.15 and substituting back into (2.13) yields,

$$\Psi(\mathbf{r}, t) = \psi(\mathbf{r})\chi_0 \exp\left(-\frac{i}{\hbar}Et\right)\,. \tag{2.16}$$

It is therefore possible to obtain dynamical information by solving the TISE and then solving Eq. 2.16. However, the spatial coordinates of most systems still

present too much of a challenge for the TISE to be solved analytically. As noted above, the wavefunction, and by extension ψ, refer to the same state, independent of a scalar multiplier. χ_0 is such a scalar the equation above, thus allowing it to be absorbed into the normalisation factor of ψ such that $\Psi(\mathbf{r},t) = \psi(\mathbf{r})\exp(-\frac{i}{\hbar}Et)$.

Determining the probability density, $|\Psi(\mathbf{r},t)|^2$, of this definition of the wavefunction however yields a seemingly paradoxical result.

$$|\Psi(\mathbf{r},t)|^2 = \left(\psi(\mathbf{r})\exp(-\frac{i}{\hbar}Et)\right)^* \psi(\mathbf{r})\exp(-\frac{i}{\hbar}Et) = |\psi(\mathbf{r})|^2\,. \tag{2.17}$$

The apparent lack of time-dependence in the probability density originates from the fact that the TISE is an eigenvalue problem, that is to say that there are a number of $\psi_j(\mathbf{r})$ eigenstates with corresponding eigenvalues, E_j, for any given Ψ. Thus, instead of using a single such eigenstate, a linear combination may be used,[e] such that

$$\Psi(\mathbf{r},t) = \sum_j \left[c_j\psi_j(\mathbf{r})\exp(-\frac{i}{\hbar}E_jt)\right]\,. \tag{2.18}$$

Now the probability density will, in addition to terms originating from each spatial wavefunction, $\psi_j(\mathbf{r})$, contain interference terms which can be thought of as being caused by the superposition of the wave-like solutions to the TDSE. For a two eigenstate expansion for example, the wavefunction may be written as

$$\Psi(\mathbf{r},t) = c_1\psi_1(\mathbf{r})\exp(-\frac{i}{\hbar}E_1t) + c_2\psi_2(\mathbf{r})\exp(-\frac{i}{\hbar}E_2t)\,, \tag{2.19}$$

resulting in a probability density of

$$|\Psi(\mathbf{r},t)|^2 = |c_1|^2|\psi_1(\mathbf{r})|^2 + |c_2|^2|\psi_2(\mathbf{r})|^2 + 2\mathrm{Re}\left[c_1^*c_2\psi_1^*(\mathbf{r})\psi_2(\mathbf{r})\exp(-\frac{i}{\hbar}(E_2-E_1)t)\right]\,, \tag{2.20}$$

where now the final interference term contains the time dependence of the probability density. This distribution of states with differing energies and physical properties is termed a wavepacket.

2.1.5 Born-Oppenheimer approximation

While the time-space variable separation assumption allows for the wavefunction to be simplified, the Born-Oppenheimer approximation is concerned with the Hamiltonian operator, $\hat{H}$. In general, for a molecular system in the absence of any perturbations, the Hamiltonian may be written as

$$\begin{aligned}\hat{H} = &-\sum_j \frac{\hbar^2}{2m_e}\nabla^2_{e,j} + \sum_{k>j}\frac{e^2}{|\mathbf{r}_j-\mathbf{r}_k|} \\ &-\sum_j \frac{\hbar^2}{2M_j}\nabla^2_{N,j} + \sum_{k>j}\frac{Z_jZ_ke^2}{|\mathbf{R}_j-\mathbf{R}_k|} - \sum_{jk}\frac{Z_ke^2}{|\mathbf{r}_j-\mathbf{R}_k|}\,,\end{aligned} \tag{2.21}$$

[e] The eigenstates are per definition, orthogonal and normalised.

where $\mathbf{r}_j$ and $\mathbf{R}_k$ are the positions of electron j and nucleus k respectively, $\nabla^2_{e,j}$ and $\nabla^2_{N,k}$ are the Laplacians with respect to the coordinates of electron j and nucleus k respectively, M_j and Z_j are the mass and charge of nucleus j, and m_e and e are the mass and charge of the electron respectively. The terms in this equation are often assigned specific symbols, based on the interactions, the contribution of which to the total energy they represent, such that

$$\hat{H} = \hat{T}_e + \hat{V}_e + \hat{T}_N + \hat{V}_N + \hat{V}_{eN} \,, \tag{2.22}$$

where $\hat{T}_e$ and $\hat{V}_e$ are the electronic kinetic and potential energy operators, $\hat{T}_N$ and $\hat{V}_N$ the nuclear kinetic and potential energy operators and $\hat{V}_{eN}$ the operator describing the interactions of nuclei and electrons.

Formally the full wavefunction $\Psi(\mathbf{r}, \mathbf{R}, t)$ depends explicitly on all electronic $\mathbf{r}$ and nuclear $\mathbf{R}$ degrees of freedom (DOFs) and time t. As shown in the section above, the spatial and temporal parts of the wavefunction are often assumed to be separable. If the time-independent wavefunction $\Psi(\mathbf{r}, \mathbf{R})$ is similarly separable between nuclear and electronic coordinates, then

$$\Psi(\mathbf{r}, \mathbf{R}) = \psi(\mathbf{r})\chi(\mathbf{R}) \,, \tag{2.23}$$

$\psi(\mathbf{r})$ being the electronic wavefunction at a given nuclear geometry, and thus depending parametrically on $\mathbf{R}$, sometimes also written as $\psi(\mathbf{r}; \mathbf{R})$ and $\chi(\mathbf{R})$ representing the nuclear wavefunction.[f] Substituting this into the TISE yields,

$$\hat{H}\left[\psi(\mathbf{r})\chi(\mathbf{R})\right] = E\left[\psi(\mathbf{r})\chi(\mathbf{R})\right] \,, \tag{2.24}$$

for any $\Psi(\mathbf{r}, \mathbf{R})$ which is an eigenstate of the Hamiltonian. The total Hamiltonian may be separated in a number of ways, including grouping terms which depend on the electronic coordinates $\mathbf{r}$, such that

$$\hat{H} = \left[\hat{T}_e + \hat{V}_e + \hat{V}_{eN}\right] + \left[\hat{T}_N + \hat{V}_N\right] = \hat{H}_e + \hat{H}_N \,, \tag{2.25}$$

where $\hat{H}_e$ and $\hat{H}_N$ are termed the electronic and nuclear Hamiltonian respectively. Application of $\hat{H}_e$ to the expanded wavefunction will, due to the separation of electronic and nuclear coordinates, yield instead of a single energy value, an electronic energy function, depending parametrically on nuclear coordinates,

$$\hat{H}_e\left[\psi(\mathbf{r})\right] = E(\mathbf{R})\left[\psi(\mathbf{r})\right] \,. \tag{2.26}$$

[f] Whilst sharing the same notation as the wavefunctions in the previous section, it should be noted that these are in no way related.

Extending this to the full Hamiltonian, application of the nuclear kinetic energy operator, $\hat{T}_N$, to the expanded wavefunction, results in three gradient-like terms

$$\hat{H}\left[\psi(\mathbf{r})\chi(\mathbf{R})\right] = \left[E(\mathbf{R}) + V_N\right]\left[\psi(\mathbf{r})\chi(\mathbf{R})\right] - \sum_j \frac{\hbar^2}{2M_j}\left[\psi(\mathbf{r})\nabla^2_{N,j}\chi(\mathbf{R}) + 2\nabla_{N,j}\psi(\mathbf{r})\cdot\nabla_{N,j}\chi(\mathbf{R}) + \chi(\mathbf{R})\nabla^2_{N,j}\psi(\mathbf{r})\right], \quad (2.27)$$

where V_N is the nuclear potential energy, which in combination with $E(\mathbf{R})$ forms the potential energy surface of nuclear quantum mechanics.

The terms which involve differentiation of the electronic wavefunction with respect to nuclear coordinates disappear for one state. For multiple states the fact that they are in general dependent on the ratio of nuclear and electronic masses means that they are usually comparatively small and can be approximated as zero.[g] Thus, by integrating out $\psi(\mathbf{r})$, Eq. 2.27 simplifies to

$$\hat{H}\chi(\mathbf{R}) = \left[T_N + E(\mathbf{R}) + V_N\right]\chi(\mathbf{R}), \quad (2.28)$$

where T_N is the nuclear kinetic energy. The elimination of the last two terms of Eq. 2.27 is the Born-Oppenheimer approximation in the context of nuclear quantum mechanics.[h]

2.1.6 Adiabatic and diabatic representations

The quantum dynamics, that is solving the TDSE to obtain the properties of the system as a function of time, of photophysical systems, like those mentioned in Chapter 1, often involve the breakdown of the Born-Oppenheimer approximation. To visualise this, electronic states may be represented in two fundamentally different ways.

As mentioned above, the electronic wavefunction in the Born-Oppenheimer approximation, $\psi(\mathbf{r})$, depends parametrically on the nuclear coordinates and may thus be written as $\psi(\mathbf{r};\mathbf{R})$. Thus, breakdowns of the Born-Oppenheimer approximation may be accounted for exactly, by, instead of simply separating the wavefunction, $\Psi(\mathbf{r},\mathbf{R})$, expanding it in what is called the *adiabatic basis*, such that

$$\Psi(\mathbf{r},\mathbf{R}) = \sum_j^\infty \psi_j(\mathbf{r};\mathbf{R})\chi_j(\mathbf{R}). \quad (2.29)$$

[g] Another rationalisation is as follows: As electrons are very light and fast in comparison to nuclei, any change in nuclear geometry will induce an instantaneous rearrangement of the electrons. Thus, on the nuclear coordinate scale, the gradient of the electronic wavefunction with respect to nuclear coordinates vanishes and terms depending on it may be ignored.

[h] In the case of electronic structure theory, the same mass-ratio argument is used to justify forgoing determination of the nuclear kinetic energy.

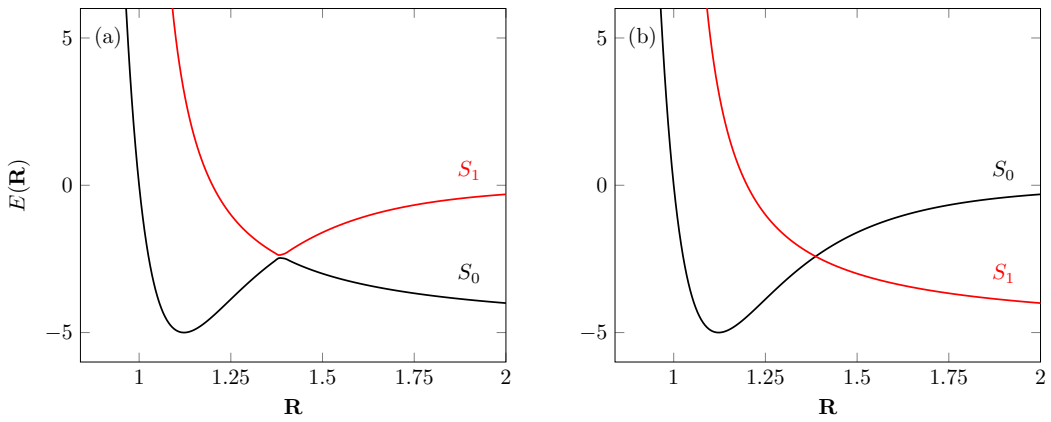

Figure 2.2: (a) The adiabatic representation of a bound (S_0) and dissociative (S_1) electronic state of a diatomic, as a function of internuclear separation, $\mathbf{R}$ and (b) the corresponding diabatic picture.

Adiabatic electronic states are defined in such a way that they preserve the order of energy levels. As a consequence, adiabatic PESs do not cross, as shown in Figure 2.2(a). Around nuclear geometries where electronic states approach each other in energy, this causes discontinuities in the energy of the states, referred to as avoided crossings. The, as per the Born-Oppenheimer approximation, usually negligible contributions to the energy arising from gradients in the electronic wavefunction with respect to the nuclear coordinates, $\nabla_N \psi(\mathbf{r};\mathbf{R})$, make a significant contribution in such areas.

This so-called *non-adiabatic* coupling term is given by

$$\langle \psi_j(\mathbf{r};\mathbf{R}) \,|\, \nabla_N \psi_k(\mathbf{r};\mathbf{R}) \rangle = \frac{\left\langle \psi_j(\mathbf{r};\mathbf{R}) \,\middle|\, \nabla_N \hat{H} \,\middle|\, \psi_k(\mathbf{r};\mathbf{R}) \right\rangle}{E_k - E_j}\,, \tag{2.30}$$

for two states $\psi_j(\mathbf{r};\mathbf{R})$ and $\psi_k(\mathbf{r};\mathbf{R})$ with associated energies E_j and E_k. Thus, as the energies of the two states approach each other, the coupling term increases and at points of intersections, where $E_j = E_k$, a singularity appears.

Alternatively, the electronic basis may be chosen in such a way that it is truly independent of nuclear coordinates, being defined at a fixed nuclear geometry, $\mathbf{R}_0$. This *crude adiabatic basis* may thus be written as $\psi(\mathbf{r};\mathbf{R}_0)$,[i] resulting in a expansion of the wavefunction in the form of

$$\Psi(\mathbf{r},\mathbf{R}) = \sum_j^{\infty} \psi_j(\mathbf{r};\mathbf{R}_0)\chi_j(\mathbf{R})\,. \tag{2.31}$$

While at first glance seemingly rather similar to the adiabatic expansion of the wavefunction, they key difference is that due to the full independence of the elec-

[i] In this case the notation $\psi(\mathbf{r})$ would also be truly appropriate.

tronic basis from the nuclear coordinates, terms involving $\nabla_N \psi(\mathbf{r}; \mathbf{R}_0)$[j] vanish. The result of this, as shown in Figure 2.2(b), is that electronic states now readily cross, their order thus no longer being conserved.

One of the most important uses of the diabatic representation is that in regions of high non-adiabatic coupling, by switching to it, the singularity in the coupling may be avoided, being replaced by smooth, potential-like terms. However, as both these representations are only strictly valid in the infinite basis limit, non-adiabatic couplings can usually not be eliminated completely. Nevertheless, the diabatic representation is rather useful in avoiding singularities in the non-adiabatic coupling.

Both expansions of the wavefunction have benefits and drawbacks, the adiabatic expansion typically requiring fewer states to accurately describe the wavefunction, however in regions of avoided crossings, the ill-behaved gradients in the electronic part with respect to nuclear coordinates often result in the calculation of matrix elements being rather challenging. In such areas, the diabatic expansion is far more well behaved, however, in general it results in significantly more states being required.

2.1.7 Basis set expansion of the nuclear wavefunction

While the section above has outlined strategies for expansion of the electronic part of the total wavefunction, the nuclear part also suffers from the inherent complexity of the total wavefunction alluded to in Section 2.1.3. In practical applications the nuclear wavefunction is thus also usually expanded in a basis set of simple, mathematically well behaved functions. In addition to simplifying solution of the TDSE by reducing it from a single challenging multivariable problem to one of solving a number of linear equations, decomposing the nuclear wavefunction into a number of functions, the time dependence of which is treated individually, results in the time-evolving wavepacket needed to extract time-dependent information from the wavefunction, as shown in Section 2.1.4.

Section 2.2 will investigate in more detail the various strategies for expanding the nuclear wavefunction and propagating the resulting approximate wavefunction in time, alluded to in Chapter 1.

2.1.8 Dirac notation

In the interest of clarity and accessibility of equations, integrals in this Chapter have so far been presented in traditional integral notation, $\int A \mathrm{d}\tau$, representing the

[j] This naturally includes higher powers of the differential operator, such that $(\nabla_N)^n \psi(\mathbf{r}; \mathbf{R}_0) = 0$ for all $n > 0$.

integral of A over all relevant coordinates. Quantum mechanics commonly makes use of the alternative Dirac notation, both in the interest of brevity, as many quantum mechanical equations involve multiple, sometimes nested, integrals and to reflect the fact that often, quantum mechanical symbols only take on their true meaning once represented in terms of an operator, the most common of which is position, $\hat{x}$.

For the remainder of this work, Dirac notation will be used for either of these purposes, the specific integral definition being,

$$A = |A\rangle \ , \tag{2.32}$$

$$A^* = \langle A| \ , \tag{2.33}$$

$$\int A^* B \mathrm{d}\tau = \langle A|B\rangle \ , \tag{2.34}$$

$$\int A^* \hat{O} B \mathrm{d}\tau = \langle A|\hat{O}|B\rangle \ , \tag{2.35}$$

where $\hat{O}$ is an operator and $\mathrm{d}\tau$ refers to all relevant coordinates, unless otherwise specified.

2.2 Quantum dynamics

2.2.1 Early methods and "frozen" Gaussians

Some of the earliest work in quantum dynamics involved very simple expansion, or rather approximation of the nuclear wavefunction in terms of Gaussian functions.[33,54,55] In these early schemes, the wavefunction is expressed as a single Gaussian, the parameters of which follow equations of motion derived directly from the TDSE. A number of trajectories, each associated with a single such Gaussian are then run, in order to model the wavepacket evolving in time.

The general form of a one-dimensional Gaussian function, as commonly used in quantum dynamics is

$$g(r;t) = \exp\left[\frac{i}{\hbar}\alpha(t)\left(r - r(t)\right)^2 + \frac{i}{\hbar}p(t)\left(r - r(t)\right) + \frac{i}{\hbar}\gamma(t)\right] \ , \tag{2.36}$$

where $r(t)$ and $p(t)$ are the real valued position and momentum, the former corresponding to the centre of the Gaussian, of the trajectory the Gaussian is associated with, while $\alpha(t)$ and $\gamma(t)$ are complex numbers, usually chosen in such a way as to satisfy $\langle g|g\rangle = 1$.

There are a number of benefits to be gained from using Gaussian basis functions, which include but are not limited to:

- They are analytic, as
$$\lim_{r\to\pm\infty} g(r;t) = 0\,.$$
- Their integrals can be determined analytically (see Appendix I).
- The product of two Gaussians yields another Gaussian, as
$$g_a(r;t)\cdot g_b(r;t) = g_c(r;t)\,.$$
- As a direct consequence, extension of the one-dimensional form shown above to a multi-dimensional one is trivial.
- Similarly, the Fourier transform of a Gaussian is another Gaussian.

On further, key advantage of Gaussian basis functions is that they are naturally associated with a point in phase space, $(r(t), p(t))$, which allows them to readily be associated with trajectories, by "surrounding" a time evolving trajectory with a Gaussian.[k] Initially Gaussians with all parameters explicitly depending on time were used,[54] however, it was soon found that restricting the width parameter $\alpha(t)$ to a constant, time-independent value, $\alpha(t) \approx \alpha_0$, significantly simplified the problem, avoiding for example the potential for negative widths, $\mathrm{Re}(\alpha(t)) < 0$ which disrupts the Gaussian functional form, while still yielding comparatively accurate results.[33]

Multi-dimensional "frozen" Gaussian functions have thus been used extensively as basis functions in quantum dynamics,[41,42,44–46,56–59] however, other choices, such as discrete variable representations (DVRs) have also been popular.[60] The latter generally represent a basis of continuous functions which is transformed so as to have each function be localised at a point on a grid in coordinate space.

2.2.2 Time-dependent self-consistent field methods

One alternative approach in expanding the wavefunction is the use of Hartree products,[l] as in the time-dependent self-consistent field (TDSCF) method,[38] where the nuclear wavefunction is defined, in a time-dependent basis, as

$$\Psi(\mathbf{r};t) = \prod_j^f \psi_j(r_j;t)\,, \tag{2.37}$$

[k] This is often also though of as evolving the phase space point $(r(t), p(t))$ corresponding to the centre of the Gaussian, and thus the entire function, in time, along the path of a trajectory.

[l] Owing to the use of Hartree products in the wavefunction expansion, the term time-dependent Hartree (TDH) is also used for methods of this kind.

where f is the number of spatial DOFs contained in $\mathbf{r}$. The basis functions are defined as

$$\phi_j(r_j;t) = \exp\left[i\gamma(t)\right]\psi_j(r_j;t) = \sum_k C_{jk}(t)\chi_k(r_j)\exp\left[-\frac{i}{\hbar}E_k^j t\right] , \tag{2.38}$$

each being constructed from a linear combination of the $\chi(r_j)$ eigenstates of the corresponding one-dimensional Hamiltonian operator, $\hat{H}_j$, and its eigenvalues, the energies, E_j^k, obtained from solving

$$\hat{H}_j\chi(r_j) = E_j\chi(r_j)\,. \tag{2.39}$$

Insertion of this expansion into the TDSE yields

$$i\hbar\frac{\partial}{\partial t}\phi_j(r_j;t) = \hat{H}_j^{SCF}(t)\phi_j(r_j;t)\,, \tag{2.40}$$

the method owing its name to the self-consistent definition of the Hamiltonian, $\hat{H}_j^{SCF}(t)$,

$$\hat{H}_j^{SCF}(t) = \hat{T}_j + \hat{V}_j + \sum_{k\neq j}\left\langle \phi_k \left| \hat{V}_{jk} \right| \phi_k \right\rangle . \tag{2.41}$$

While this approach was successfully applied to the dissociation of linear van der Waals molecules,[38] the use of one-dimensional Hamiltonians, $\hat{H}_j$, limits its use as a basis constructed in this way cannot properly treat coupling between DOFs.

2.2.3 "Standard" multi-configuration methods

To improve on the performance of the TDSCF method and others based on a similar expansion of the wavefunction in terms of a Hartree product, a multi-configurational approach may instead be taken. Instead of using a single Hartree product of f functions, the wavefunction is expanded in terms of all possible configurations of f one-dimensional functions, such that

$$\Psi(r_1,...,r_f;t) = \sum_{j_1=1}^{N_1}...\sum_{j_f=1}^{N_f} C_{j_1\ldots j_f}(t)\prod_{\kappa=1}^{f}\chi_{j_\kappa}^{(\kappa)}(r_\kappa)\,. \tag{2.42}$$

The sum of sums of products notation here indicates that for a given number of DOFs, f, and a number of one-dimensional functions, N_κ, in each DOF, κ, all possible combinations (often referred to as configurations) that contain a function for each degree of freedom are formed and associated with time-dependent coefficients, $C_{j_1\ldots j_f}(t)$. This results in a total number of coefficients and thus a rough computational scaling of

$$N_C = \prod_{j=1}^{f} N_j\,, \tag{2.43}$$

however, the scaling is often simplified as N^f, and thus said to scale exponentially with f. Methods relying on this expansion typically use DVR basis functions and have been shown to perform well for small systems.[40,60] This expansion is often considered the "gold standard" of quantum dynamics, as it involves very few approximation and theoretically, similarly to configuration interaction based electronic structure theory, captures all possible interactions within a given basis. However, as already illustrated by Eq. 2.43, this method scales incredibly unfavourably with the size of the system investigated, and in practice is computationally limited to treating only a few DOFs.

2.2.4 Multi-configuration time-dependent Hartree

In order to overcome the unfavourable scaling of the "standard" method, the multi-configuration time-dependent Hartree (MCTDH) approach[39,61–65] reduces the number of one-dimensional functions per DOF by generating configurations from so-called *single-particle* functions (SPFs). The latter, while still one-dimensional, are now time-dependent and as a result, far fewer are needed per DOF, in order to describe the wavefunction to within a given accuracy compared to the "standard" method. This is a direct consequence of the basis set of SPFs moving in coordination space as the wavefunction evolves in time, on paths ideal for the description of the wavefunction at any given point in time, as illustrated in Figure 1.4. The MCTDH wavefunction *ansatz* may be written as

$$\Psi(r_1, ..., r_f; t) = \sum_{j_1=1}^{n_1} ... \sum_{j_f=1}^{n_f} A_{j_1 ... j_f}(t) \prod_{\kappa=1}^{f} \phi_{j_\kappa}^{(\kappa)}(r_\kappa; t)\,. \tag{2.44}$$

The SPFs are, as per their definition above, time-dependent, owing to the fact that they in turn are linearly expanded in a basis of simpler, time-independent DVR functions, with the time-dependence originating from the associated expansion coefficients. A given SPF is defined as

$$\phi_{j_\kappa}^{(\kappa)}(r_\kappa; t) = \sum_{l_\kappa=1}^{N_\kappa} c_{l_\kappa j_\kappa}^{(\kappa)}(t) \chi_{l_\kappa}(r_\kappa)\,, \tag{2.45}$$

where N_κ is the number of primitive basis functions per DOF κ. There are thus two levels of time-dependence in the wavefunction ansatz (2.44), one at the level of the f-dimensional configurations of SPFs and one associated with the coefficients of the primitive basis functions defining a SPF at any given time. This causes redundancy in the wavefunction expansion, thus, in order to derive uniquely defined equations of motion, a number of conditions must thus be imposed on the SPFs.

In general these are

$$\left\langle \phi_j^{(\kappa)}(r_\kappa; t=0) \,\middle|\, \phi_l^{(\kappa)}(r_\kappa; t=0) \right\rangle = \delta_{jl}\,, \tag{2.46}$$

$$\left\langle \phi_j^{(\kappa)}(r_\kappa; t) \,\middle|\, \frac{\partial}{\partial t}\phi_l^{(\kappa)}(r_\kappa; t) \right\rangle = -i \left\langle \phi_j^{(\kappa)}(r_\kappa; t) \,\middle|\, \hat{O}^{(\kappa)} \,\middle|\, \phi_l^{(\kappa)}(r_\kappa; t) \right\rangle\,, \tag{2.47}$$

where $\hat{O}^{(\kappa)}$ may generally be any Hermitian operator, but in practice is often chosen to be $\hat{O}^{(\kappa)} = 0$. In addition to simplifying the problem, the constraint (2.47) ensures that expression of the wavefunction's time-dependence occurs as much as possible on the level of the coefficients, $A_{j_1 \ldots j_f}(t)$, the SPF's variation in time only accounting for a small remainder.

The key benefit of MCTDH lies in the fact that, while, similar to the "standard" method, still scaling exponentially, as per Eq. 2.43, because the number of SPFs required is much lower, this scaling is less unfavourable. The MCTDH method formally scales as $fnN + n^f$ with respect to the memory required, compared to the N^f scaling of the "standard" method. While the MCTDH approach has been shown to perform excellently for a variety of systems,[61,62,65,66] there have also been significant efforts to further increase the accuracy and efficiency of this method.

One such strategy of improvement involves allowing SPFs to depend on more than just one DOF. The so-called mode combination approach re-expresses the SPFs in combined (also known as logical) coordinates, representing multiple DOFs. This significantly reduces the number of $A_{j_1 \ldots j_f}(t)$ coefficients[m] however the SPFs are now no longer simple linear combinations, but rather multi-dimensional products of the primitive basis functions. The latter results in the more complex SPF time-propagation, mode combination thus essentially allowing a balance to be struck between the time-evolution of SPF and the $A_{j_1 \ldots j_f}(t)$ coefficients.

The added challenge of propagating the multi-dimensional primitive basis arising from mode combination may be addressed by adding additional layers of MCTDH. Multi-layer MCTDH[67–69] (ML-MCTDH) involves using equations of motion very similar to those for the SPF coefficients in regular MCTDH, to propagate in time the multi-dimensional functions making up lower layers. One key system that will also be discussed later in the context of this work that serves to highlight the benefit of the ML-MCTDH approach is the relaxation dynamics following photoexcitation of pyrazine. Comparing MCTDH[65] and ML-MCTDH[70] calculations for the full 24-dimensional system, the former required 11 million coefficients and approximately 630 CPU hours the latter only 22 thousand coefficients and 7 minutes.[64]

[m] Although mode combination results in notation changes in Eq. 2.44, these are forgone here, in the interest of brevity.

One final modification that warrants discussion in this section concerns an alternative approach to mode combination. Instead of having multi-dimensional SPFs, certain degrees of freedom, such as environmental or bath modes may not require configurational treatment at all. Instead, these DOFs can be expressed in terms of multi-dimensional Gaussian functions, which evolve in time *via* variational equations of motion, much like in the previously discussed work of Heller.[54] The result of this is the so-called Gaussian MCTDH[71–73] (G-MCTDH) wavefunction *ansatz*

$$\Psi(r_1,...,r_{f-1},\mathbf{r}_f;t)=\sum_{j_1=1}^{n_1}...\sum_{j_{f-1}=1}^{n_{f-1}}\sum_{j=1}^{n_f}A_{j_1...j_{f-1}}(t)\times\left(\prod_{\kappa=1}^{f-1}\phi_{j_\kappa}^{(\kappa)}(r_\kappa;t)\right)g_j^f(\mathbf{r}_f)\,, \tag{2.48}$$

where $g_j^f(\mathbf{r}_f)$ is a multi-dimensional Gaussian function, describing any "secondary", bath DOFs. While this is written in terms of a single set of secondary modes, multiple groupings are also possible and extension to such a case does not significantly complicate the above *ansatz*. Comparing the performance of G-MCTDH[73] to full MCTDH,[65] again for the 24-dimensional pyrazine benchmark problem, the former was able to achieve comparable accuracy at significantly lower costs, further suggesting that, in particular due to the lower memory requirements, G-MCTDH calculations should in theory be scalable closer to the convergence limit than their full MCTDH counterparts. Furthermore in comparison to MCTDH, G-MCTDH can be applied to overall larger systems, especially if a significant number of DOFs act as bath modes.

One key drawback of both the methods based on the MCTDH *ansatz* and those employing some other form of Hartree products is that in order to determine potential energy matrix elements, the PES must be expressible in a sum of product form. Thus the PES must be fitted to the grid of the lowest level of basis functions, which, for many realistic problems, requires non-trivial amounts of computational time and memory. Finally it is worth noting, especially in the context of the photophysical systems the work presented herein is aimed at, that the methods presented so far can, quite easily be used to calculate dynamics on more than one electronic state.

2.2.5 Variational multi-configuration Gaussian method

Taking the theme set by mode combination in MCTDH and inclusion of multi-dimensional Gaussian functions for secondary, environmental DOFs in G-MCTDH to its natural conclusion, the wavefunction may be expanded exclusively in multi-

dimensional Gaussian functions,[44] such that

$$\Psi(\mathbf{r};t) = \sum_{j}^{N} A_j(t) g_j(\mathbf{r};t)\,, \tag{2.49}$$

where $g_j(\mathbf{r};t)$ are multi-dimensional[n] Gaussians, the time-evolution of which is achieved by deriving variational[34–37] equations of motion for their parameters.

The quantum dynamics method resulting from this *ansatz*, termed the variational multiconfiguration Gaussian (vMCG) method has been successfully applied to a number of benchmark problems[44,73,74] and, in the case of the pyrazine model[65] discussed above, has been found to exhibit exceptional scaling with respect to the number of basis functions needed to achieve accurate dynamics.[73,75]

One particular benefit of v-MCG over MCTDH based methods is that, due to the known functional form of the Gaussians, integrals over the potential energy operator can easily be determined analytically, similarly to the overlap, as shown in Appendix I.

This approach indeed requires so few basis functions that it is generally very suitable for direct dynamics applications, that is where instead of the potential energy part of Hamiltonian being obtained analytically or from a fitted surface, the corresponding terms are calculated on-the-fly, using electronic structure methods. Direct dynamics vMCG[76] (DD-vMCG) has been applied to a variety of challenging problems,[76–78] yielding promising results. It finally is worth noting that, while the basis set in v-MCG can be thought of as evolving on trajectories, these are not independent of one another, as in the case of the methods discussed in Section 2.2.8, as the variational[34–37] derivation of the equations of motion for the basis function parameters means that the basis set may only evolve in time as one cohesive unit.

The main drawback of v-MCG, as is discussed in more detail below, relates to the nature of the equations of motion for the basis function parameters. They often tend to be numerically challenging to solve, the matrices which must be calculated and inverted during their solution possessing a strong tendency to be singular, especially as the basis set scales towards being complete.

2.2.6 Recent Gaussian based methods

Given that the work presented herein also employs a "frozen" multi-dimensional Gaussian basis set, there are a number of recent quantum dynamics methods that warrant mentioning in this context.

[n] These functions correspond to the multi-dimensional version of Eq. 2.36. Both "frozen" [33] and thawed versions may be used.

The basis expansion leaping multi-configuration Gaussian (BEL-MCG) method relies on the overcompleteness of Gaussian basis sets.[41] More specifically the wavefunction is expanded in a basis of static Gaussians, the parameters of which are entirely time-independent. The evolution of the wavefunction is expressed solely *via* the expansion coefficients, as per Eq. 1.2, the equations of motion being derived variationally.[34–37] The basis functions are distributed in phase space so as to describe the wavefunction well, however the span of the basis set is limited to regions of non-vanishing amplitudes of the wavefunction at that time. As the wavefunction evolves, *via* the equations of motion for the expansion coefficients, the basis set is regularly "expanded", selecting a new distribution which is again limited to regions of space with wavefunction amplitudes above a threshold. The result of this strategy is a static basis set which, in principle, contains far fewer basis functions than would be required if all relevant system space was occupied with basis functions.

So far this approach has only been tested for up to 3-dimensional problems, however the notion of an adaptive, but static, basis set is certainly promising, and the work presented herein aims to achieve something rather similar, albeit in a non-regular manner, not relying on grids of basis functions.

A recent application of trajectory-guided, and thus at least somewhat related to the work presented herein, Gaussian basis sets is the quantum trajectory Gaussian basis (QTGB) method,[58] which employs Bohmian mechanics to sample a time-dependent Gaussian basis set. While this strategy has shown some promise for 2-dimensional benchmark problems, its applicability to higher-dimensional and otherwise more complex systems remains to be seen.

A final interesting approach is the recently introduced pseudospectral sampling, which uses Dirac delta functions, $\chi_i = \delta(\mathbf{r} - \mathbf{r}_i)$, to simplify the calculation of matrix elements.[59,79] These functions are positioned at the centre of Gaussian functions, which allows the potential matrix element, normally $\langle g_i(\mathbf{r})|\hat{V}(\mathbf{r})|g_j(\mathbf{r})\rangle$ which thus involves the integral over Gaussian basis functions, to instead be determined as $V(\mathbf{r}_i)\,|g_j(\mathbf{r}_i)\rangle$, that is the potential at the coordinates of the delta function, times the corresponding Gaussian. Overall this reduces the number of potential energy calculations necessary. This approach has been successfully applied to a number of low-dimensional benchmark problems[79] and more recently been extended to be able to treat multi-state non-adiabatic systems.[59]

2.2.7 Matching pursuit split-operator Fourier transform

When expanding the Hamiltonian operator, $\hat{H}$, in the TDSE in terms of the kinetic and potential energy operators, $\hat{T}$ and $\hat{V}$, one may write Eq. 2.16 as

$$\Psi(\mathbf{r}, t+\delta t) = \exp\left[-\frac{i}{\hbar}\delta t\left(\hat{T}+\hat{V}\right)\right]\Psi(\mathbf{r}, t)\,, \tag{2.50}$$

for the wavefunction at some small time, δt after a known state, $\Psi(\mathbf{r}, t)$. Recognising that for non-commuting operators $\hat{A}$ and $\hat{B}$, like $\hat{T}$ and $\hat{V}$,

$$\exp\left[\hat{A}+\hat{B}\right] \sim \exp\left[\frac{1}{2}\hat{A}\right]\exp\left[\hat{B}\right]\exp\left[\frac{1}{2}\hat{A}\right]\,. \tag{2.51}$$

The Fourier transform of the exponential term containing the kinetic energy operator can be shown to equal a simple multiplication by an exponential function, thus the matrix elements associated with the former can be relatively easily calculated. This approach, which involves continuously flipping between the position and momentum representation of the wavefunction in order to carry out propagation, is the termed the split operator Fourier transform (SOFT) method.[80]

Some more recent work has combined the SOFT strategy with matching pursuit (MP) method[81] for selecting from an overcomplete basis set the minimal subset that represents the wavefunction to within a desired threshold.[82,83] A more detailed example of the MP algorithm can be found in Chapter 4. The use of the MP algorithm in favour of the Monte-Carlo based importance sampling, normally associated with the SOFT method, was found to significantly improve the scaling of this method, both for a synthetic tunnelling benchmark,[82] which will be discussed later, in the context of the work presented herein, as well as an analytic 2-dimensional Hamiltonian model for intramolecular proton transfer.[83]

2.2.8 Quantum dynamics using independent trajectories

Many systems of real-word interest, especially those involving interactions with light or the electromagnetic field in general, involve multiple electronic states. Furthermore, breakdown of the Born-Oppenheimer approximation is rather common, that is the dynamics of the system occur around geometries where electronic states can rapidly change with respect to the nuclear coordinates. Such *non-adiabatic* dynamics often involve high frequency oscillation of the wavefunction between electronic states, which presents a challenge for quantum dynamics methods. There are a number of quantum dynamics approaches which are rely on independent trajectories to account for dynamics across multiple electronic states, and given the ultimate goal of the work presented herein being light-matter interactions, it is worthwhile briefly discussing some common approaches. Given that the work

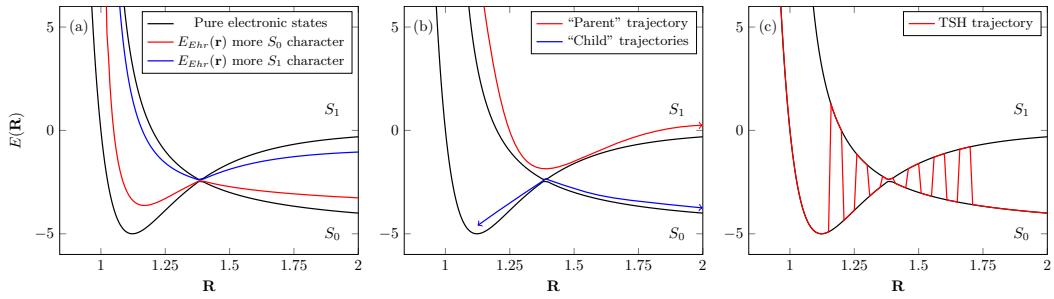

Figure 2.3: Multi-state quantum dynamics methods: (a) MCE trajectories, evolve on a state averaged PES, (b) AIMS trajectories evolve classically, spawning copies of themselves to account for non-adiabatic transitions and (c) TSH trajectories "hop" from state to state with a probability proportional to the non-adiabatic coupling.

presented herein is only loosely related to these methods, in depth discussion is forgone in the interest of brevity here, especially since the scope of methods occupying this area of quantum dynamics is relatively regularly surveyed by far more experienced authors.[75]

The Ehrenfest approximation, which is discussed in more detail in Chapter 3, has been used to derive a set of equations of motion for a "frozen" Gaussian basis set, resulting in the multi-configuration Ehrenfest (MCE) method.[45] While initially developed and applied as a way of treating systems with many environmental modes, the extension to multiple electronic states was quickly achieved,[84] and very successfully applied to the pyrazine benchmark discussed above. One of the key benefits of this approach is its simplicity: the trajectories propagating the basis set evolve on a state-averaged, mean-field PES, the extent of each function on the states determined by a set of time-dependent coefficients, the equations of motion are determined variationally. Figure 2.3 (a) illustrates the nature of the MCE trajectories. The MCE approach has recently been interfaced with *ab initio* electronic structure routines (AI-MCE), thus allowing dynamics to be calculated on the fly, the application of AI-MCE to excited state dynamics of ethylene yielding encouraging results.[85]

Another approach to multi-state dynamics relies on classical trajectories to propagate its basis functions, the equations of motion for the expansion coefficients being determined variationally as usual. In order to account for non-adiabatic transitions, the effective non-adiabatic coupling is calculated and if this value reaches a certain threshold, each trajectory has a chance to spawn identical copies of itself on the state it is coupling with. This approach is referred to as the multiple spawning (MS) method.[43,86] The direct dynamics version of this approach is known

as *ab initio* multiple spawning[56] (AIMS), which has been applied very successfully to a number of challenging systems.[27,28,87,88] The MS strategy is visualised in Figure 2.3(b).

Recently a hybrid method, combining the state-averaged potentials of the MCE method and the basis set expansion strategy from AIMS has been introduced, referred to as *ab initio* multiple cloning[46,89] (AIMC). More specifically, this method addresses the shortcomings of MCE in regions of low non-adiabatic coupling where the wavepacket is spread significantly across multiple states. In such a case, the state-averaged potential does not reproduce the branching of the wavepacket that should occur due to the different shapes of the PESs. The multiple spawning like cloning spawning algorithm allows for this to be corrected however: In AIMC, after the wavepacket has passed through a region of strong non-adiabatic coupling, the wavepacket bifurcates into two initially identical copies, each predominantly on one of the two electronic states involved in the coupling event. Recently, this approach has been applied to the photodissociation of pyrrole with some encouraging initial results.[89]

As an alternative to spawning or cloning in order to account for non-adiabatic transitions, the wavepacket may be expanded in a set of trajectories moving on a single surface, which in regions of strong non-adiabatic coupling have the capability to "hop" and thereby change the electronic state they evolve on. This constitutes one of the earliest approaches to non-adiabatic dynamics and has come to be referred to as trajectory surface hopping (TSH).[42,90–93] Again, the propagation of the time-dependent basis functions occurs via classical mechanics, although in the case of TSH the basis functions are purely electronic, while the equations of motion for the coefficients are derived variationally. At any point, the probability of a "hop" to any given surface may be calculated as a function of the non-adiabatic coupling. Given its relatively long developmental history, a large number of modifications and improvements have been made to TSH,[94] it having been applied to numerous systems. Recently, some rather complex photodynamics problems have been tackled with a direct dynamics version of TSH with encouraging results.[95,96]

In conclusion then, MCE,AIMS and AIMC all share a similar wavefunction *ansatz*,

$$\Psi(\mathbf{r};t) = \sum_j A_j(t)\phi_j(\mathbf{r};t)\,,$$

where the basis functions $\phi(\mathbf{r};t)$ are multi-dimensional "frozen" Gaussians, the time dependence of which is achieved by evolving the centre with classical, or in the case of MCE, semi-classical equations of motion. TSH employs a similar expansion, however the basis is made up of orthonormal electronic basis functions. Time evolution of the coefficients, $A_j(t)$, for all methods occurs *via* variational

equations of motion.[34–37] In order to account for non-adiabatic coupling, each method employs a different strategy to approximate transitions of the wavepacket between electronic states, which are illustrated in Figure 2.3.

2.3 Photophysics

Given that, as outlined in Chapter 1, the ultimate goal of the work presented in the following chapters is to facilitate further study of light-matter interactions, it seems appropriate to begin by discussing the theoretical framework governing such phenomena. As however, the novel methods introduced in this work are just that, new and thus only tested against a limited number of benchmark problems, the following is only a general and by no means in-depth overview, intended mostly to provide perspective on the context of the advances outlined herein.

2.3.1 Electromagnetic field

In the context of his theory of gravity, Isaac Newton postulated the concept of *action at a distance*, whereby physical interactions may occur in the absence of contact between objects. His observations suggested that such interactions occur instantaneously and are unaffected by the medium separating the objects. However, Faraday found over the course of several experiments that in the case of electromagnetic interactions, this framework was not sufficient to describe his observations. He thus suggested that interactions travelled, originating from charged sources, through the intervening medium at finite speed and could thus be affected by the nature of said medium.

Out of this new perspective of a "field" of interactions, arose a number of physical laws which were finally unified by Maxwell in his equations of the electromagnetic (EM) field,

$$\nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0 \tag{2.52}$$

$$\nabla \cdot \mathbf{D}(\mathbf{r}, t) = \rho(\mathbf{r}, t) \tag{2.53}$$

$$\nabla \times \mathbf{E}(\mathbf{r}, t) = -\frac{\partial}{\partial t}\mathbf{B}(\mathbf{r}, t) \tag{2.54}$$

$$\nabla \times \mathbf{H}(\mathbf{r}, t) = \mathbf{J}(\mathbf{r}, t) + \frac{\partial}{\partial t}\mathbf{D}(\mathbf{r}, t)\,, \tag{2.55}$$

where $\mathbf{E}$, $\mathbf{B}$, $\mathbf{D}$, $\mathbf{H}$ and $\mathbf{J}$ are vector fields and ρ is a scalar field, all depending on spatial coordinates $\mathbf{r}$ and time t. The terminology used to refer to the individual fields in $(\mathbf{E}, \mathbf{B}, \mathbf{D}, \mathbf{H})$, making up the *mediating field* varies, however the constituents of the *source field*, $(\mathbf{J}, \rho)$, are usually referred to as the charge distribution, or charge density, ρ, and the current density, $\mathbf{J}$. Often written in other forms, Eq.

2.52 is known as the *magnetic Gauss's law*, Eq. 2.53 as *Gauss's law*, Eq. 2.54 as *Faraday's law* and Eq. 2.55 as *Ampere's law*. These equations, along with supplementary constitutive relationships form a complete, mathematically rigorous description of the behaviour of the electromagnetic field.

Although these equations can be used to formally derive, albeit *via* some rather challenging mathematics, essentially all properties of interacting electromagnetic fields and matter, when the latter is on the microscopic scale, thankfully some approximations can be made. Chief among these is the electric dipole approximation. It states that the electric field, $\mathbf{E}(\mathbf{r}, t)$, is uniform over the extent of each individual molecule. In addition to this, intermolecular electron exchange is neglected. The dipole moment, defined as

$$\mathbf{p}(\mathbf{r}) = \int \rho(\mathbf{r}_0)(\mathbf{r}_0 - \mathbf{r})\mathrm{d}V\,, \tag{2.56}$$

for a continuous charge density, which, in the case of point charges, may be estimated as

$$\rho(\mathbf{r}) = \sum_{j}^{N} q_j \delta(\mathbf{r}_i - \mathbf{r})\,, \tag{2.57}$$

of N discrete charges confined to volume V, is the only molecular multipole coupled to the electromagnetic field. A further result of the electric dipole approximation is the absence of any magnetic interactions.

This approximation is known to be sound when the wavelength of radiation is significantly larger than molecular dimensions. The shortest wavelength of sunlight, even in the absence of atmospheric absorption and filtering is still only around 100 nm,[97] thus the electric dipole approximation is bound to be reasonable for all but the largest bio- or macrocmolecular systems.

2.3.2 Spectroscopic transitions

There are essentially three interactions that can occur between a molecular system and the electromagnetic field.

1. The EM field can be **absorbed** by the molecule, increasing its internal energy levels.

2. A molecule in an excited state may **spontaneously** lose energy *via* **emission** of radiation, the direction of which depends on the spatial orientation of the molecule.

3. Similarly, the EM field may **induce** relaxation of the molecule followed by **emission**, the direction of the emitted radiation being aligned with the inducing radiation.

In order to calculated the probability and rate of a transition between initial state i and final state f, each of these possibilities may be considered in turn. In the case of absorption, the probability of a transition occurring is proportional to the energy density of the EM field, $\rho(\nu)$ at the required frequency (energy), ν. Spontaneous emission on the other hand only depends on the number of molecules in the excited state, n_i. Finally, the probability of induced emission is proportional to both $\rho(\nu)$ and n_i.

The net rate of absorption from the lower i to the higher f state, $\xi_{f\leftarrow i}$ may thus be written in terms of a number of linear coefficients,

$$\xi_{f\leftarrow i} = A_{if} n_i \rho(\nu) - S_{fi} n_f \rho(\nu) - E_{fi} n_f \,, \tag{2.58}$$

where A_{if}, S_{fi} and E_{fi} are the probability coefficients of absorption, stimulated emission and spontaneous emission respectively. It logically follows that the rate of emission, in the absence of any non-radiative transitions, is $\xi_{i\leftarrow f} - \xi_{f\leftarrow i}$.

It can be shown that

$$A_{if} = S_{fi} = \left(\frac{c^3}{8\pi h \nu^3}\right) E_{fi} \,, \tag{2.59}$$

where c is the speed of light and h is the Planck constant.

Absorption probability coefficients, A_{if}, represent the interaction of the EM field with the molecule and may thus be expressed in terms of the transition moment, μ_{if}, such that

$$A_{if} = \left(\frac{2\pi}{3\hbar^2}\right) \mu_{if}^2 \,, \tag{2.60}$$

where the transition moment is defined as the expectation value of the dipole moment operator,

$$\mu_{if} = \langle \Psi_i \,|\, \hat{\mathbf{p}} \,|\, \Psi_f \rangle \,, \tag{2.61}$$

where Ψ_i and Ψ_f are the wavefunctions for states i and f respectively.

This final definition allows the rates and probabilities of interactions between the EM field and a molecular system to be extracted from the wavefunctions of the states involved. Within the scope of this work, which focuses on the propagation of the wavefunction in time, this level of theory is considered to be sufficient, in order to predict or model any light-matter interactions.

Chapter 3

Trajectory-guided Sampling

This chapter introduces a method of sampling basis sets for quantum dynamics using simple trajectories, propagated on the potential energy surface of the system. Initially, the motivation underlying the development of this new approach is outlined, followed by a detailed description of the algorithm which is responsible for the practical implementation of the former. To demonstrate the efficacy of the trajectory sampling method, results for the challenging pyrazine quantum dynamics benchmark are presented, followed finally by an investigation into the individual effects of the specific parameters and overall performance of the algorithm.

The contents of this chapter have, in part, been published:
M. A. C. Saller and S. Habershon, *J. Chem. Theo. Comput.*, **11**, 8-16 (2015)

3.1 Introduction

As already touched on in Chapters 1 and 2, quantum dynamics approaches may be subdivided in a number of ways, one of which concerns the nature with which the basis functions are treated. Some methods evolve the basis functions in time, as the wavefunction is propagated,[42–46]while others choose the basis functions parameters at the outset, effectively generating a grid-like basis set, and propagate only the expansion coefficients in time.[38–41] Both strategies bring with them certain advantages, however, as the novel method proposed here aims to avoid the disadvantages associated with traditional approaches, it will be discussed within the context of the latter.

Relying solely on the expansion coefficients for propagation of the wavefunction, expanded in a static, time-independent basis set, requires the latter to describe the time evolution in phase space over the entire time domain of interest. This leads to extremely unfavourable exponential scaling with respect to basis set size, as without a way of choosing *a priori* where in phase space to place basis functions, usually a grid, covering the entirety of relevant space, is used. Thus these methods are, without some modification to the grid over time,[41] severely limited with regards to the number of DOFs that may be treated explicitly.

In the case of dynamic basis sets, that is expansion of the wavefunction in terms of basis functions that move in phase space over time, this scaling problem is potentially avoided, as the basis moves in phase space alongside the wavefunction, thus requiring only enough basis functions to describe it at one moment in time.[44] There arises however the question of how to propagate the basis functions in time. Usually, equations of motion for their parameters, often positions and momenta, may be readily obtained using variational or approximate treatment of the TDSE. Should these equations be derived using variational principles,[34–37] while the resulting dynamics of the basis functions are formally exact, the solution often proves practically challenging due to numerical issues of ill-conditioning.[74] While this may be avoided by employing non-variational equations of motion, the latter have been shown to violate the energy conservation condition, imposed by the TDSE.[47]

The trajectory sampling method proposed here represents a compromise between the two aforementioned categories. The aim of this approach is to avoid both the issue of exponential scaling as well as those associated with basis functions propagation. The latter is achieved by avoiding moving the basis functions during the solution of the TDSE altogether, thus employing a purely static basis set for wavefunction propagation. In order to improve on the scaling of methods using this expansion of the wavefunction, instead of using a grid of basis functions, covering

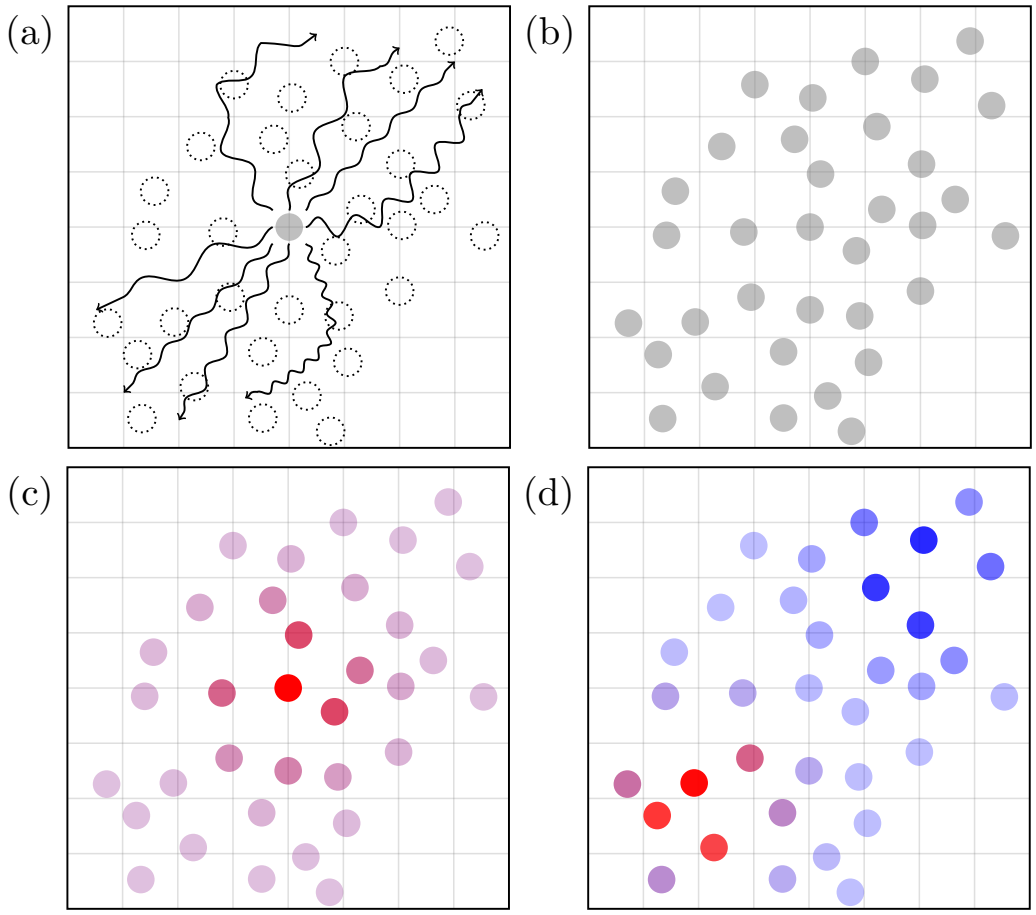

Figure 3.1: Trajectory sampling: (a) Based on the initial wavefunction, a set of trajectories sample phase space, storing basis functions along their path. (b) The resulting time-independent basis set is propagated in time by (c) setting expansion coefficient to represent the initial wavefunction and (d) evolving these coefficients variationally.

phase space, a set of simple, classical-like sampling trajectories are evolved on the PES of the system, in order to sample regions of phase space, relevant to wavefunction propagation, and populate them with basis functions.

This strategy for sampling the basis set, which is illustrated in Figure 3.1, is expected to significantly improve scaling at little to no extra cost, as the dynamics driving the trajectories are simple and in essence governed by classical mechanics. The reason this is possible is that, while classical mechanics would constitute a poor choice to evolve the wavefunction in time, being inherently incapable of describing quantum behaviour, during the sampling stage, the trajectories are not required to follow the exact quantum solution, only to, at some point during their propagation, visit the regions of phase space relevant to the latter. Furthermore, the chance of this may be arbitrarily increased by sampling a larger number of

trajectories, with varying initial conditions, representative of the wavefunction at the start of propagation.

The remainder of this chapter focuses on introducing the algorithm responsible for implementing the trajectory sampling method described above, followed by testing the validity of the assumptions underlying it and investigating its performance for a challenging benchmark problem.

3.1.1 Gaussian wavepacket basis functions

As discussed in Chapter 2, there are numerous advantages to be gained from using multi-dimensional Gaussian wavepacket basis functions as the basis of quantum dynamics.

The Gaussian wavepacket (GWP) basis functions used in this work are defined as

$$|g(\mathbf{r};t)\rangle = \prod_{j=1}^{f}\left[\left(\frac{2\alpha_j}{\pi\hbar}\right)^{\frac{1}{4}}\exp\left[-\alpha_j\frac{1}{\hbar}\left(r_j - r_j(t)\right)^2 + p_j(t)\frac{i}{\hbar}\left(r_j - r_j(t)\right)\right]\right], \quad (3.1)$$

where f is the number of degrees of freedom, α_j is the width of the GWP in DOF j and $\mathbf{r}(t)$ and $\mathbf{p}(t)$ are f-dimensional vectors, representing the position and momentum associated with the centre of the GWP, referred to from here on out as its phase space coordinates. It is important to note that, while the basis functions, according to the notation above, are parametrically time dependent, that is only meant to reflect the sampling stage, the propagation stage treating each basis function as time-independent.

Throughout this work, the width of the GWPs, α_j, remains unchanged as a function of time, thus these basis functions can be referred to as being "frozen".[54] Additionally, $\boldsymbol{\alpha}$, takes the same value across all degrees of freedom f. Although it is possible to vary the individual widths, given that they remain fixed over time, the lack of any intuitive way of choosing what to set the width for a given DOF to, eliminates most benefits that may be reaped by doing so. Given the key role the width of GWPs plays in their behaviour with respect to quantum dynamics,[74] allowing them to change as a function of time (thus "thawing" the Gaussian), may add flexibility to the method presented herein, though this would come at additional computational cost and complexity.

Appendix I outlines the expression used to calculate the integral over these GWPs, as well as the formula determining the overlap of two such functions.

In the case of multiple electronic states, the electronic wavefunction is expanded in a set of orthonormal basis functions, which for the purposes of this method, only act as labels in determining which PES contributes to the Hamiltonian (see Section 3.3). These electronic basis functions are omitted whenever they do not impact

the expression at hand. Nevertheless, a single (nuclear) basis function, $|g(\mathbf{r};t)\rangle$, is thus technically always associated with a corresponding electronic basis function, $|\alpha\rangle$.

3.2 Implementation of trajectory sampling

Firstly, a distinction must be made between the algorithm that implements the sampling strategy described above, and that which, immediately following the former, handles the propagation of the sampled basis set. The latter, henceforth referred to as the propagation algorithm, is the well established solution of applying the time-dependent variational principle[34–37] to the TDSE, expanded in a time-independent basis, and has been used in a number of other quantum dynamics approaches.[38,40,79] Both algorithms are discussed in turn here, however given the extensive literature available on the propagation algorithm, the focus is primarily on the implementation of the basis set sampling.

3.2.1 Sampling algorithm

Starting from the initial, nuclear wavefunction, $\psi_0(\mathbf{r}) \equiv \psi(\mathbf{r}, t = 0)$, the first challenge lies in sampling from this function of coordinate space, $\mathbf{r}$, positions and momenta, (q, p), to allow classical time propagation of sampling trajectories. The quasiprobability distribution resulting from application of the Wigner function, $f_W(q,p)$, to ψ_0, allows phase space variables q and p to be sampled with correct quantum probabilities,[98] where

$$f_W(q,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{+\infty} \psi^*(\mathbf{r}+s)\,\psi(\mathbf{r}-s)\,e^{\frac{i}{\hbar}ps}\mathrm{d}s\,. \tag{3.2}$$

Given the choice of Gaussian basis functions, the wavefunction can be represented as a linear combination of GWPs, $|g(\mathbf{r})\rangle$, with complex expansion coefficients, $\mathbf{c}$, such that

$$|\psi_0(\mathbf{r})\rangle = \sum_j^{N_0} c_j\,|g_j(\mathbf{r})\rangle\ . \tag{3.3}$$

Note that, with reference to Section 3.1.1, in the majority of cases, the GWPs defining the initial wavefunction are assigned only positions, resulting in zero valued momenta in Eq. 3.1. The label of the GWPs as $|g_j(\mathbf{r})\rangle$ is meant to reflect this and highlight the fact that the nature by which the positions for these initial GWPs are chosen, results in them effectively being defined in coordinate space. In cases where the nature of the system requires multiple Gaussians to describe the initial wavefunction ($N_0 > 1$), before application of $f_W(q,p)$, a single basis function, $|g\rangle$, is chosen first. A variety of selection criteria may be employed here,

a simple and common strategy being to sample $|g\rangle$ according to its overlap with ψ_0, specifically $\langle g \,|\, \psi_0 \rangle$, a more detailed discussion of such a selection process can be found in Chapter 4.

A set of m trajectories are generated by sampling the initial conditions as described above and initialising new Gaussian basis functions at the corresponding coordinates, which are then independently propagated on the PES. Setting aside systems with multiple electronic states for now, the propagation occurs completely classically, by integrating Newtons equation of motion in the well established molecular dynamics (MD) method, using the popular Velocity-Verlet algorithm,[99] outlined in Appendix II. All that is required at this stage is that the force due to the PES of the system may be obtained at the coordinates of the sampling trajectory GWP. For a discussion of alternative methods of sampling trajectory propagation, see Chapter 5.

The sampling trajectories are propagated for n_t timesteps. While setting n_t to be equal to the number of timesteps the TDSE will be solved for, n_p, may be the most immediately obvious choice, the trajectory sampling method in no way limits n_t to a specific set of values, with respect to n_p or otherwise. Similar to the number of timesteps to propagate for, the timestep duration of the sampling trajectories Δt_t may be, but is not required to be equal to that with which the TDSE is solved, Δt_p. The impact the magnitudes of n_t with respect to n_p and Δt_t with respect to Δt_p have on the accuracy of TDSE propagation and the concept of 'oversampling' are discussed in Section 3.4. The total trajectory sampling time is thus defined as $t_t = n_t \Delta t_t$, while the duration of propagation, and thus the total time for which dynamics are calculated, is given by $t_{max} = n_p \Delta t_p$.

As the sampling trajectories are propagated, the guiding GWPs are stored with a probability of $1/n_s$ at every timestep, where n_s may be interpreted as the average sampling frequency. Stored GWPs retain the position and momentum of the guiding trajectory they originate from, as well as inheriting any other parameters, which are not propagated during the sampling stage, such as width (in the case of "frozen" Gaussians). The size of the basis set resulting from an iteration of trajectory sampling, N_s, may thus be approximated as

$$N_t \approx m \frac{n_t}{n_s} \,. \tag{3.4}$$

3.2.2 Propagation algorithm

Upon completion of the sampling algorithm outlined above, the basis set for propagation of the wavefunction is formed by including first the N_0 functions defining the initial wavefunction, with coefficients as per Eq. 3.3, followed by the approximately N_t trajectory sampled GWPs, $|g_j\rangle$, which are initially assigned zero

valued coefficients, $c_j(t)$, in order to reflect that at time $t = 0$, the wavefunction remains unchanged from its initial conditions. The wavefunction at the start of the propagation algorithm, expanded in the full basis set, may thus be written as

$$|\psi(t=0)\rangle = c_0 \,|\psi_0\rangle + \sum_{j}^{N\approx N_t} c_j(t=0)\,|g_j\rangle \;. \tag{3.5}$$

To propagate the wavefunction in time, this basis set expansion is inserted into the TDSE, which yields

$$i\hbar\frac{\partial}{\partial t}\left[\sum_{j}^{N_{total}} c_j(t)\,|g_j\rangle\right] = \hat{H}\left[\sum_{j}^{N_{total}} c_j(t)\,|g_j\rangle\right] , \tag{3.6}$$

where the indices of summation for ψ_0 and the trajectory sampled basis functions have been combined for convenience and the total basis set size, $N_{total} \approx N_0 + N_t$. The wavefunction only depends explicitly on time through the expansion coefficients, $c_j(t)$, the basis functions remaining fixed in phase space now that the sampling algorithm has been completed.

The Dirac-Frenkel variational principle[34,35] may then be applied in order to derive equations of motion for the coefficients, however note that alternative versions do exist,[36,37] the derivation of equations of motion from them following closely the one given below. This variational principle minimises

$$\left\langle \delta\Psi \,\middle|\, \hat{H} - i\hbar\frac{\partial}{\partial t} \,\middle|\, \Psi \right\rangle = 0 \tag{3.7}$$

In the time-independent basis set expansion of the wavefunction,

$$|\delta\Psi\rangle = \sum_{j} \delta c_j(t)\,|g_j(\mathbf{r})\rangle \;. \tag{3.8}$$

thus

$$\begin{aligned} &\left\langle \sum_j \delta c_j(t) g_j(\mathbf{r}) \,\middle|\, \hat{H} \,\middle|\, \sum_k c_k(t) g_k(\mathbf{r}) \right\rangle \\ -i\hbar &\left\langle \sum_j \delta c_j(t) g_j(\mathbf{r}) \,\middle|\, \frac{\partial}{\partial t} \,\middle|\, \sum_k c_k(t) g_k(\mathbf{r}) \right\rangle = 0 \,, \end{aligned} \tag{3.9}$$

$$\sum_{j,k}\left[\delta c_j^* c_k \left\langle g_j \,\middle|\, \hat{H} \,\middle|\, g_k \right\rangle\right] - i\hbar\sum_{j,k}\left[\delta c_j^* \dot{c}_k \left\langle g_j \,|\, g_k \right\rangle\right] = 0 \,, \tag{3.10}$$

where $\dot{c}_k = \frac{\mathrm{d}}{\mathrm{d}t}c_k$. Therefore, for the arbitrary variation δc_j,

$$\sum_j \delta c_j^* \left(\sum_k \left[c_k \left\langle g_j \,\middle|\, \hat{H} \,\middle|\, g_k \right\rangle\right] - i\hbar\sum_k \dot{c}_k \left\langle g_j \,|\, g_k \right\rangle\right) = 0 \,, \tag{3.11}$$

and therefore this simplifies to the matrix equation

$$\mathbf{Hc} = i\hbar\mathbf{S}\dot{\mathbf{c}}\,, \tag{3.12}$$

This may be rearranged to yield

$$\boxed{\dot{\mathbf{c}} = -\frac{i}{\hbar}\mathbf{S}^{-1}\mathbf{Hc}\,,} \tag{3.13}$$

where $\mathbf{H}$ is the Hamiltonian matrix, the elements of which are defined as

$$H_{ij} = \langle\alpha_i| \langle g_i|\hat{H}|g_j\rangle |\alpha_j\rangle \ , \tag{3.14}$$

and $\mathbf{S}$ is the overlap matrix, the elements of which are

$$S_{ij} = \langle\alpha_j| \langle g_j|g_i\rangle |\alpha_i\rangle \ , \tag{3.15}$$

where α_j denotes the electronic state occupied by $|g_j\rangle$ and due to the orthonormality of electronic states in the vibronic pyrazine Hamiltonian, $\langle\alpha_i|\alpha_j\rangle = \delta_{i,j}$.

The first computationally expensive step of the propagation algorithm involves the calculation of the Hamiltonian and overlap matrices, $\mathbf{H}$ and $\mathbf{S}$, followed by inversion of $\mathbf{S}$. In the case of analytic Hamiltonians such as the vibronic pyrazine Hamiltonian discussed in Section 3.3, this step is the most computationally expensive, the sampling trajectories being relatively cheap by comparison. Should the determination of Hamiltonian matrix elements occur on-the-fly however, the electronic structure calculations required to determine the potential energy and non-adiabatic coupling coefficients (in case of a multi-state system) are highly likely to constitute the rate limiting step of the algorithm.

The set of first order differential equations, defined by Eq. 3.13, is then solved iteratively using the common Runge-Kutta method,[100] for n_p timesteps of duration, Δt_p, resulting a total time t_{max} of dynamics. More specifically the 4th order Runge-Kutta method is used, the details of which are outlined in Appendix III.

3.3 Pyrazine dynamics benchmark

3.3.1 Photophysics of pyrazine

Pyrazine, $C_4N_2H_4$, shown in Figure 3.2(a), forms part of the azaaromatic and more specifically the azabenzene class of molecules, which have continued to be the subject of strong interest in the field of spectroscopy for a number of years.[101–105] Structurally similar to benzene, C_6H_6, azabenzenes possess a rather diverse number of transitions, both radiative and radiationless, between their ground and low-lying excited states, as highlighted in Figure 3.2(b). Of particular interest in

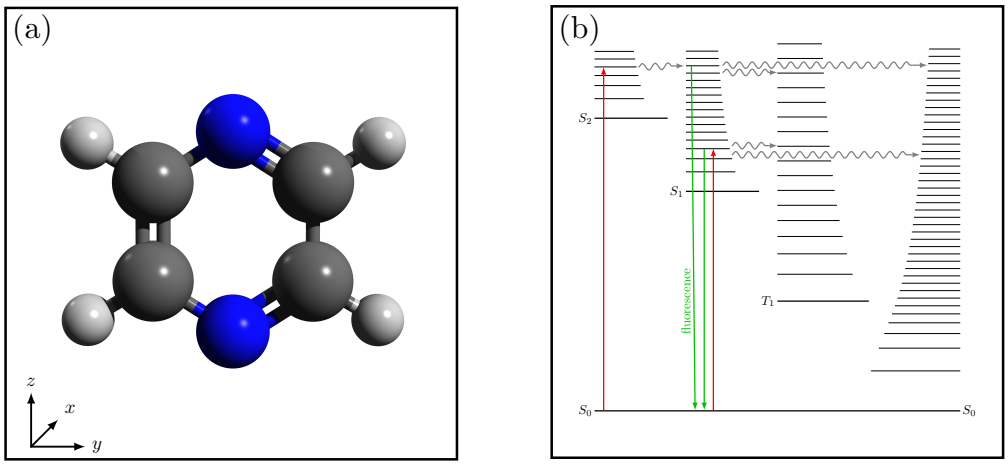

Figure 3.2: (a) Molecular structure of pyrazine, $C_4N_2H_4$ and (b) diagram of low-lying electronic states and transitions of azabenzenes.

this case is the radiationless decay from the second excited singlet state, S_2, to the first excited singlet state, S_1, following excitation from the ground state, S_0.

In pyrazine the ultraviolet (UV) absorption spectrum hints at the rapid nature of this transition by the distinct structure observed in the region of the S_1 state, contrasted with the relatively broad, featureless signal corresponding to the S_2 state.[101,105] Recent time resolved photoelectron spectroscopy results have confirmed the extremely fast timescale of this relaxation, the lifetime of the S_2 state having been measured at 22 ± 3 fs in a gaseous molecular beam of pyrazine seeded in He.[104] Such fast transitions can often be attributed to the presence of a conical intersection in the PES of the system, due to strong coupling of the two states involved. In this case, presence of a conical intersection has indeed been confirmed computationally.[106]

It is well known that the symmetries of the S_2 and S_1, being $^1B_{2u}(\pi\pi^*)$ and $^1B_{3u}(n\pi^*)$ respectively, allow them to be linearly coupled by modes exclusively of B_{1g} symmetry. The only deformation of the pyrazine molecule that satisfies this condition is the asymmetric out-of-plane bending mode of the hydrogen atoms, commonly termed ν_{10a}, shown in Figure 3.3(h). The remaining modes may however indirectly contribute to the coupling between the two states, especially those with identical A_g symmetry, as they may affect the energy separation of S_2 and S_1, without affecting the symmetry of the B_{1g} mode coupling the two states.[106]

3.3.2 Vibronic Hamiltonian

Computational study of the relaxation dynamics of pyrazine involves solving the TDSE, which in turn requires a Hamiltonian to describe the PES arising from

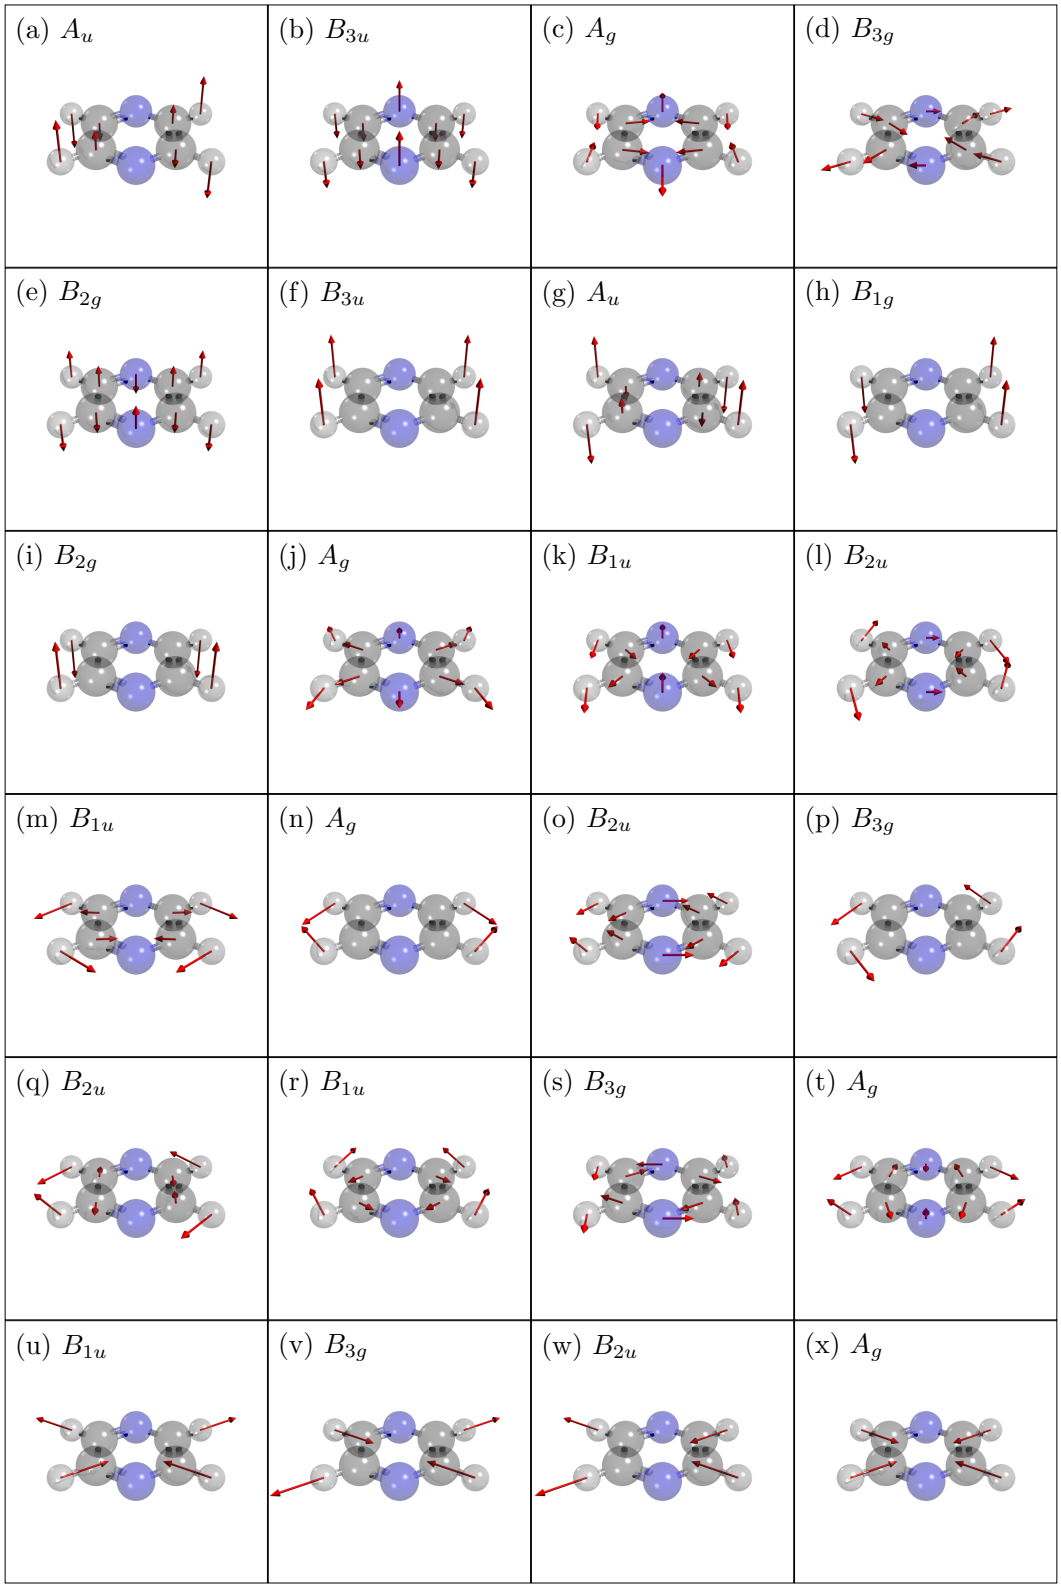

Figure 3.3: Vibrational modes, including symmetry point groups, of pyrazine in the S_0 ground state, obtained using Møller-Plesset 2nd order perturbation theory,[107] with a D95** basis set[108] in Gaussian03.[109]

these dynamics. A number of different models have been proposed,[62,103,106,110,111] however by far the most complete and effective to this date is the second order (bilinear), vibronic Hamiltonian, capable of incorporating all 24 degrees of freedom, parameterised using *ab initio* electronic structure calculations,[65] given by

$$\begin{aligned}\hat{\mathrm{H}} = &\sum_{i=1}^{f}\left[-\frac{\omega_i}{2}\frac{\partial^2}{\partial q_i^2}+\frac{\omega_i}{2}q_i^2\right]+\begin{pmatrix}-\Delta & 0\\ 0 & \Delta\end{pmatrix}+\sum_{i\in G_1}\begin{pmatrix}a_i & 0\\ 0 & b_i\end{pmatrix}q_i\\ &+\sum_{(i,j)\in G_2}\begin{pmatrix}a_{ij} & 0\\ 0 & b_{ij}\end{pmatrix}q_iq_j+\sum_{i\in G_3}\begin{pmatrix}0 & c_i\\ c_i & 0\end{pmatrix}q_i+\sum_{(i,j)\in G_4}\begin{pmatrix}0 & c_{ij}\\ c_{ij} & 0\end{pmatrix}q_iq_j\,,\end{aligned} \tag{3.16}$$

where q_i is the normal-mode coordinate of the ith vibrational mode, ω_i corresponding to the associated frequency and 2Δ is the energy splitting between S_1 and S_2 at the origin of nuclear coordinate-space. The PES is described in terms of the normal modes by linear, a_i and b_i, and bilinear, a_{ij} and b_{ij}, expansion terms, while the coupling between states is expressed linearly and biliearly by c_i and c_{ij} respectively. The subdivision of the 24 modes in Eq. 3.16 reflects the symmetry of the system: The set G_1 thus contains all modes with identical, A_g, symmetry, G_2 is comprised of pairs of modes resulting in A_g symmetry, G_3 corresponds to the single, ν_{10a}, mode with B_{1g} symmetry and G_4 is the set of all pairs of modes, the product of which has B_{1g} symmetry.

One stand-out feature of this particular Hamiltonian is the fact that it can incorporate a varying number of degrees of freedom as follows: Representing an improvement on a previously developed system,[110] a 4-dimensional ($f = 4$) version has been shown to qualitatively reproduce experimental spectra.[65] In this version of the above Hamiltonian, aside from the B_{1g} coupling mode, ν_{10a}, the three modes of A_g symmetry with the largest linear coupling, ν_{6a}, ν_1 and ν_{9a}, act as tuning modes. Coupling coefficients for the remaining two modes with A_g symmetry, ν_2 and ν_{8a}, are small enough to indicate little involvement of the corresponding vibrations in coupling the two electronic states.[110,112] A second version of this Hamiltonian incorporates 12 vibrational modes, including in addition to those mentioned above, (ν_{10a}, ν_{6a}, ν_1, ν_{9a}), all remaining modes with g symmetry, (ν_2, ν_3, ν_4, ν_5, ν_{6b}, ν_{7b}, ν_{8a}, ν_{8b}), to reflect the fact that S_1 and S_2 are both of u symmetry and thus the coupling between them must overall be of g symmetry. Finally, all 24 degrees of freedom can be incorporated, resulting in the most accurate version of the above Hamiltonian, that treats all vibrations explicitly.

This Hamiltonian has been used in conjunction with MCTDH calculations to calculate the pyrazine S_2 absorption spectrum, employing the three versions outlined above, by Cederbaum *et al.*[65] It was found to yield excellent results for the dynamics of pyrazine, in the case of the full 24-dimensional Hamiltonian,

approaching numerical accuracy with respect to experimental spectra.

Given the excellent agreement with experiment and flexibility of the model, the vibronic pyrazine Hamiltonian has become a standard benchmark for quantum dynamics methods. It has been studied extensively using a variety of different approaches[73,75,113–115] and thus presents an excellent opportunity to gauge the performance and rigour of the trajectory sampling method set forth in this chapter.

3.3.3 Initial conditions

In order to simulate the dynamics following photoexcitation from the S_0 ground state to the S_2 excited state, the initial wavefunction, $|\psi_0\rangle$, is set to a single GWP on the S_2 state, with width $\alpha = 0.5$ in every DOF, such that

$$|\psi_0\rangle = |g_0\rangle\,|2\rangle\ , \tag{3.17}$$

where $|g_0\rangle$, placed at the origin of nuclear coordinate space and assigned zero valued momenta, is given by

$$|g_0(\mathbf{r})\rangle = \prod_j^f \left[\left(\frac{1}{\pi\hbar}\right)^{\frac{1}{4}} \exp\left(-\frac{1}{2\hbar}r_k^2\right)\right]\ . \tag{3.18}$$

These initial conditions are intended to reflect a vibrational wavepacket on the ground state, projected onto the S_2 state, which in turn constitutes a simple, but effective model for photoexcitation. The wavepacket formed on the upper state could also be seen as the result of f excitations from the ground state to each individual vibrational state of the S_2 electronic state.

3.3.4 Ehrenfest dynamics

Given that the object of the benchmark described above is to calculate the relaxation dynamics from S_2 to S_1 after excitation from the ground state, it is clear that the Velocity-Verlet driven MD propagation of the sampling trajectories, alluded to in Section 3.2.1, must be adapted to suit a system with multiple electronic states. While there exists a variety of strategies to accomplish this, some of which have been discussed in Chapter 2, the Ehrenfest approach[45] was chosen here, for the relative simplicity with which it can be adapted from purely classical MD.

The core of Ehrenfest's theorem lies in linking the expectation values of operators in quantum formalism, more specifically position, $\hat{q}$, and momentum, $\hat{p}$, with their variable counterparts in classical mechanics, q and p. It is relatively simple to show that

$$i\hbar\frac{d}{dt}\left\langle\hat{O}\right\rangle = \left\langle\Psi\,\middle|\,\left[\hat{O},\hat{H}\right]\,\middle|\,\Psi\right\rangle\ , \tag{3.19}$$

for any operator $\hat{O}$ without explicit time-dependence, where $\hat{H}$ is the Hamiltonian operator and Ψ is the wavefunction. In the case of $\hat{q}$ and $\hat{p}$ this yields equations of motion, which may be classically integrated.

Applying this theorem to the case of multiple electronic states, one may derive a set of equations, which allow the position and momentum to evolve classically on an artificial PES, which represents the quantum potential, averaged across the electronic states, such that

$$\frac{\partial q_\kappa}{\partial t} = \frac{p_\kappa}{m_\kappa} \tag{3.20}$$

$$\frac{\partial p_\kappa}{\partial t} = -\frac{\partial V_{Ehr}}{\partial q_\kappa}\,, \tag{3.21}$$

where m_κ is the mass of degree of freedom κ and V_{ehr} is the aforementioned state-averaged potential, which, in this particular case of the two states in pyrazine, is defined as

$$V_{Ehr} = \frac{|a_1|^2 V_{11} + |a_2|^2 V_{22} + 2Re\left(a_1^* a_2 V_{12}\right)}{|a_1|^2 + |a_2|^2}\,, \tag{3.22}$$

where V_{ij} is the ijth element of the potential energy matrix, defined by $\langle\Psi|\hat{V}|\Psi\rangle$, and a_i are expansion coefficients, which, in combination with the GWP basis functions described in Section 3.1, form a wavepacket spanning both states,

$$|\phi(t)\rangle = \left[a_1(t)\,|1\rangle + a_2(t)\,|2\rangle\right]|g(\mathbf{q},\mathbf{p};t)\rangle\ . \tag{3.23}$$

Equations of motion for these coefficients may be derived similarly to those of the GWP basis functions. The major difference lying in the nature of the basis set expansion

$$|\Psi(\mathbf{r},t)\rangle = \sum_j \left[a_1(t)\,|1\rangle + a_2(t)\,|2\rangle\right] c_j(t)\,|g_j(\mathbf{q},\mathbf{p};t)\rangle\ , \tag{3.24}$$

where now, as the basis functions $g_j(\mathbf{q},\mathbf{p},t)$ are evolving in time as part of the sampling trajectories and are thus time-dependent. Furthermore, the electronic basis functions α_j are orthonormal, $\langle\alpha_i|\alpha_j\rangle = \delta_{ij}$. In the interest of brevity these are absorbed into the basis functions ϕ below, which thus also become orthonormal.

This introduces another term in the expansion of the Dirac-Frenkel variational principle,[34,35] such that

$$\begin{aligned}&\left\langle \sum_j \delta a_j(t)\phi_j(\mathbf{r};t) \,\middle|\, \hat{H} \,\middle|\, \sum_k a_k(t)\phi_k(\mathbf{r};t) \right\rangle \\ -i\hbar &\left\langle \sum_j \delta a_j(t)\phi_j(\mathbf{r};t) \,\middle|\, \frac{\partial}{\partial t} \,\middle|\, \sum_k a_k(t)\phi_k(\mathbf{r};t) \right\rangle = 0\,,\end{aligned} \tag{3.25}$$

$$\sum_{j,k}\left[\delta a_j^* a_k \left\langle \phi_j \middle| \hat{H} \middle| \phi_k \right\rangle\right] - i\hbar \sum_{j,k}\left[\delta a_j^* \dot{a}_k \left\langle \phi_j \mid \phi_k \right\rangle + \delta a_j^* a_k \left\langle \phi_j \middle| \dot{\phi}_k \right\rangle\right] = 0\,, \tag{3.26}$$

$$\sum_j \delta a_j^* \left(\sum_k \left[a_k \left\langle \phi_j \middle| \hat{H} \middle| \phi_k \right\rangle\right] - i\hbar \sum_k \left[\dot{a}_k \left\langle \phi_j \mid \phi_k \right\rangle + a_k \left\langle \phi_j \middle| \dot{\phi}_k \right\rangle\right]\right) = 0\,, \tag{3.27}$$

$$\mathbf{H}\mathbf{a} = i\hbar\left(\dot{\mathbf{a}} + \dot{\mathbf{S}}\mathbf{a}\right)\,, \tag{3.28}$$

$$i\hbar\dot{\mathbf{a}} = \left[\mathbf{H} - i\hbar\dot{\mathbf{S}}\right]\mathbf{a}\,, \tag{3.29}$$

$$\boxed{\dot{\mathbf{a}} = -\frac{i}{\hbar}\left[\mathbf{H} - i\hbar\dot{\mathbf{S}}\right]\mathbf{a}\,,} \tag{3.30}$$

where $\mathbf{H}$ is the Hamiltonian matrix in the electronic basis, such that $H_{ij} = \langle g| \langle i|\hat{H}|j\rangle |g\rangle$, and $\dot{\mathbf{S}}$ is the time-dependent overlap matrix, with elements

$$\dot{S}_{ij} = \delta_{ij}\left[\sum_{\kappa=1}^{f}\left\langle g \middle| \frac{\partial g}{\partial q_\kappa}\right\rangle \frac{\partial q_\kappa}{\partial t} + \left\langle g \middle| \frac{\partial g}{\partial p_\kappa}\right\rangle \frac{\partial p_\kappa}{\partial t}\right]\,. \tag{3.31}$$

Again, akin to the treatment of the expansion coefficients of the GWP basis, described in Section 3.2.2, these equations of motion are integrated using the RK4 method.

3.3.5 Calculated quantities

In order to compare the performance and accuracy of the quantum dynamics obtained from the trajectory sampling algorithm with MCTDH results, which are, as outlined above, extremely close to experimental data, the approach of Cederbaum *et al.* for calculating the S_2 absorption spectrum[65] was replicated here. The photoabsorption spectrum can be obtained by the Fourier transform

$$I(\omega) \propto \omega \int_{-\infty}^{+\infty} C(t) e^{i\omega t} \mathrm{d}t\,, \tag{3.32}$$

where the wavefunction autocorrelation function, $C(t)$, is given by

$$C(t) = \left\langle \Psi(\mathbf{r},\mathbf{R};t=0) \middle| e^{-\frac{i}{\hbar}\hat{H}t} \middle| \Psi(\mathbf{r},\mathbf{R};t=0)\right\rangle = \langle \Psi(\mathbf{r},\mathbf{R};t=0) \mid \Psi(\mathbf{r},\mathbf{R};t)\rangle\,. \tag{3.33}$$

To account for the finite resolution of experimental spectrometers, the spectral lines resulting from Eq. 3.32 can be convoluted, which, given that the autocorrelation function is directly calculated, amounts to a dampening of $C(t)$. The function

chosen here, again following the literature approach,[65] was

$$d(t) = \exp\left[-\frac{|t|}{t_d}\right], \tag{3.34}$$

where t_d is a user selected parameter. Furthermore, since the Fourier transform range is finite, the resulting spectrum is prone to numerical anomalies, which can however be minimised, by introducing a further function, $s(t)$, which smoothly reduces $C(t)$ to zero at t_{max}, defined by

$$s(t) = \cos\left[\frac{\pi t}{2t_{max}}\right]. \tag{3.35}$$

Here, t_{max} is the total time, the TDSE is solved for. Bringing all these functions together, the S_2 absorption spectrum, $I(\omega)$ was calculated as

$$I(\omega) = \omega \int_{-\infty}^{+\infty} \left[s(t)d(t)C(t)\right] e^{i\omega t} \mathrm{d}t\,. \tag{3.36}$$

In addition to comparison to MCTDH spectra,[65] the diabatic electronic state populations can provide detailed information on the nature of the dynamics occurring as a function of time. In the two-state pyrazine model, the population of state α, $P_\alpha(t)$ can be calculated as

$$P_\alpha(t) = \sum_i^N \sum_j^N \left[c_i^* c_j \left\langle g_i \,|\, g_j \right\rangle \delta_{\lambda_i,\alpha}\delta_{\lambda_j,\alpha}\right], \tag{3.37}$$

where λ_j denotes the electronic state which is occupied by basis function j. In the analysis below, the population of the lower excited state, S_1, was chosen, however this choice is completely arbitrary, as the population of S_2 will give identical information as to the nature of the dynamics. For the Hamiltonian containing only two electronic states, outlined above, the two populations are exactly additive, such that $P_1(t) + P_2(t) = 1$, for any given time t. This condition is a direct result of the normalisation of the wavefunction, $|\Psi|^2 = 1$.

3.3.6 4-dimensional results

The first step in testing the validity and accuracy of the trajectory sampling algorithm, set out above, was to benchmark its performance against MCTDH results[65] for the 4-dimensional vibronic pyrazine Hamiltonian. This was chosen for its relative accessibility in terms of computational resources required, while still providing a challenging benchmark problem for quantum dynamics methods. Additionally, this version of the Hamiltonian has been studied using a number of different approaches,[73,75,113–115] thus allowing the results of the strategy discussed

here to be put into the broader context of quantum dynamics methods commonly found in literature.

A set of 4 calculations was carried out using 1500 timesteps, lasting 0.1 fs each, of trajectory sampling, followed by the same number of wavefunction propagation. The basis function sampling frequency was $1/n_s = 1/50$ while the GWP width was chosen to be $\alpha = 0.5$ for all DOFs. The four calculations differed only in the number of sampling trajectories employed, which were $m = 33, 200, 400$ and 800, resulting in approximately 1000, 6000, 12000 and 24000 GWPs respectively in the final basis set (N_{total}) upon completion of the sampling algorithm. These input parameters along with the exact basis set sizes are summarised in Table 3.1.

As outlined in Section 3.3.5, the wavefunction autocorrelation function, $C(t)$, was calculated, and from this the S_2 photoabsorption spectrum, $I(\omega)$, was obtained *via* Fourier transform, using a dampening constant of $t_d = 30$ fs, as this was the value chosen for the original MCTDH benchmark results.[65] Figure 3.4 shows $C(t)$ and the corresponding $I(\omega)$ for the four calculations, compared to MCTDH data. Inspecting the accuracy of $C(t)$ from trajectory sampling with respect to MCTDH, it is clear that even at relatively small basis set sizes, the short time behaviour up to around 50 fs is captured relatively well. At longer timescales however, the method presented here fails to capture both the frequency and amplitude of the oscillations in $C(t)$. As the basis set size increases the error in the autocorrelation function does decrease, to the point where, in the case of 24000 basis functions, the MCTDH results are reproduced qualitatively, across the entire time domain. Given the relative simplicity of the sampling algorithm outlined above, as well as its low computational cost[a], this is highly encouraging.

The S_2 photoabsorption spectra, also shown in Figure 3.4, interestingly do not reflect the same trend as the autocorrelation functions they were obtained from. Even in the case of the smallest basis sets, the spectra from trajectory sampling

[a]Section 3.4.4 contains a more detailed analysis of computational expenses

Table 3.1: Input parameters and final basis set sizes for trajectory sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figures 3.4 and 3.6.

n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	m	N_{total}
1500	0.1	50	1500	0.1	33	1019
1500	0.1	50	1500	0.1	200	5850
1500	0.1	50	1500	0.1	400	12140
1500	0.1	50	1500	0.1	800	24144

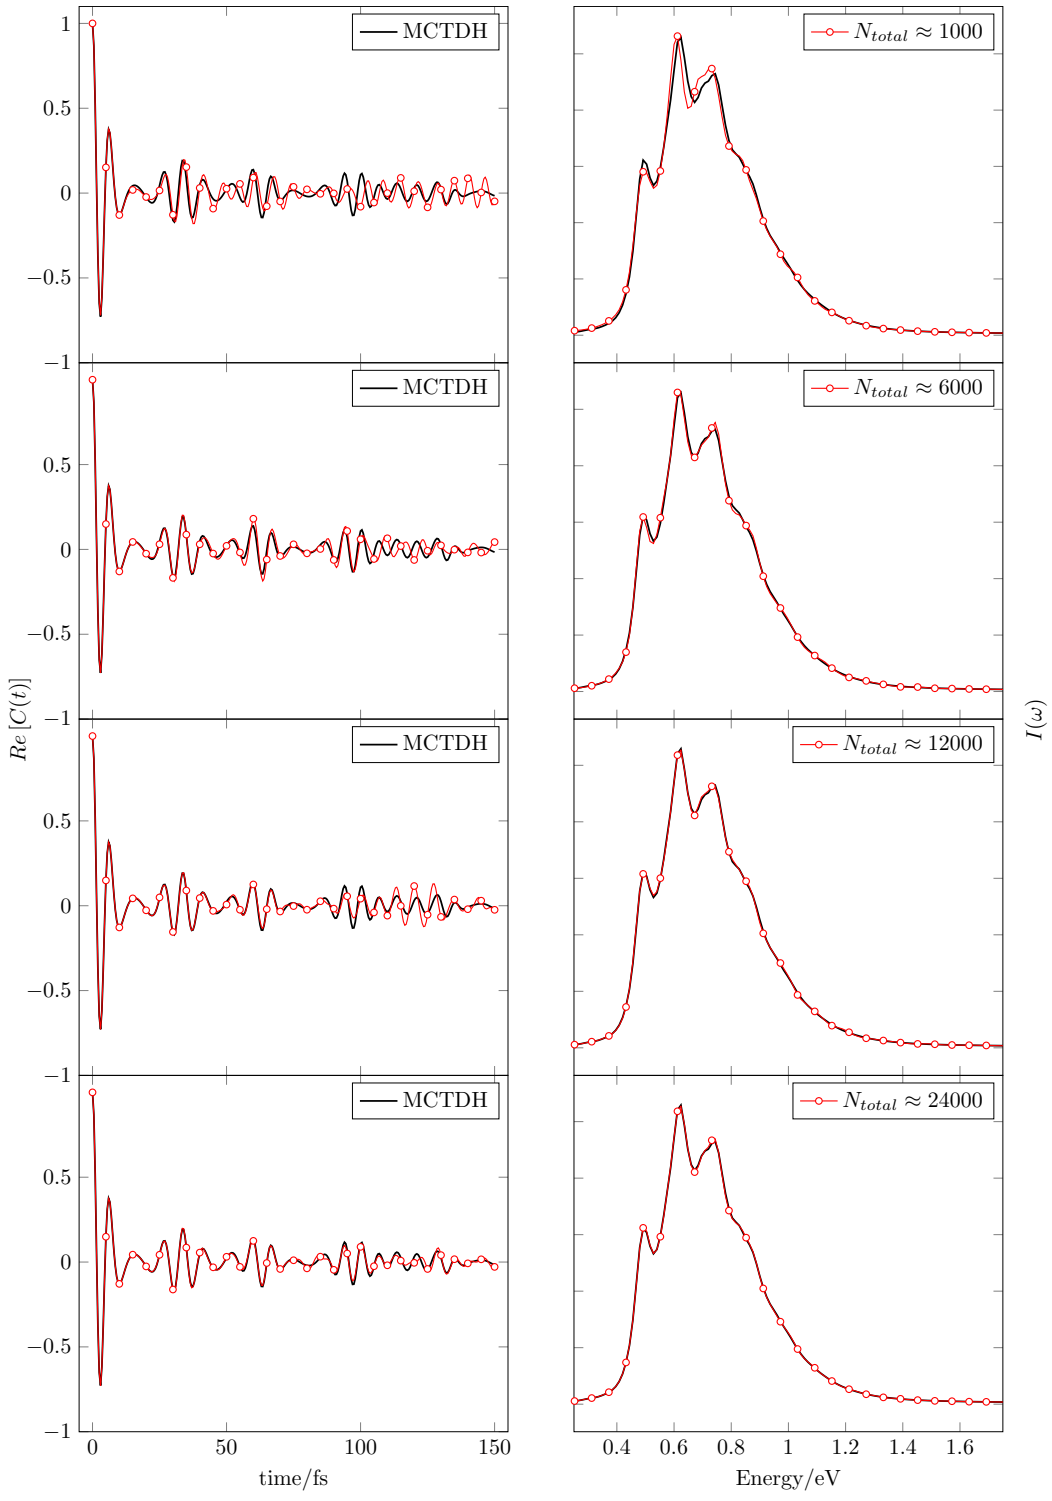

Figure 3.4: Wavefunction autocorrelation functions, $C(t)$, and corresponding S_2 photoabsorption spectra, $I(\omega)$, for 4 trajectory sampling calculations of the 4D pyrazine Hamiltonian at varying basis set size N_{total}, compared to MCTDH results.[65]

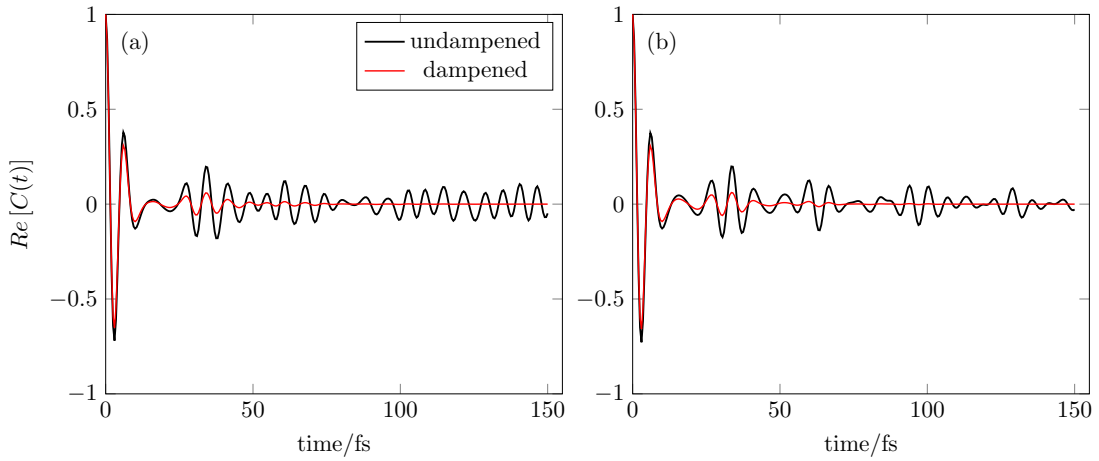

Figure 3.5: Effects of the dampening and sampling functions, $d(t)$ and $s(t)$, on the wavefunction autocorrelation function, $C(t)$, for the 4D pyrazine Hamiltonian, obtained using basis set sizes of (a) $N_{total} \approx 1000$ and (b) $N_{total} \approx 24000$.

are in excellent agreement with the MCTDH benchmark. Furthermore, increasing the number of basis functions does not significantly improve the results beyond $N_{total} \approx 6000$. Closer inspection of the dampening and sampling functions, $d(t)$ and $s(t)$, discussed in Section 3.3.5, gives some insight into the origin of this discrepancy. Figure 3.5 shows the dampened and undampened autocorrelation functions for the two most extreme basis set sizes, $N_{total} \approx 1000$ and $N_{total} \approx 24000$. It is clear that the long-time oscillations of $C(t)$ are effectively completely dampened out, thus making the absorption spectrum only truly sensitive to the short time dynamics up to about $t = 75$ fs.

While comparison to experimentally observable properties of the system is always useful in assessing the accuracy of any given computational method, the extent of detail lost during the dampening, in order to replicate finite spectrometer resolution, suggests that there is significantly more detailed information about the dynamics of the system accessible to computational methods. The state populations, being a more direct measure of the real time dynamics across the entirety of the system, should give more insight into the relaxation dynamics of pyrazine, following photoexcitation. Thus the populations of the lower excited S_1 state, $P_1(t)$, were calculated as per Eq. 3.37 for the four calculations discussed above, and are shown in Figure 3.6.

It becomes immediately clear that the populations give a much more detailed description of the dynamics, as the basis set size dependence of the accuracy with respect to MCTDH, observed in the autocorrelation functions of Figure 3.4, is far more pronounced in Figure 3.6. Smaller basis set sizes result in the reabsorbance feature occurring between 75 fs and 100 fs being inadequately described and the

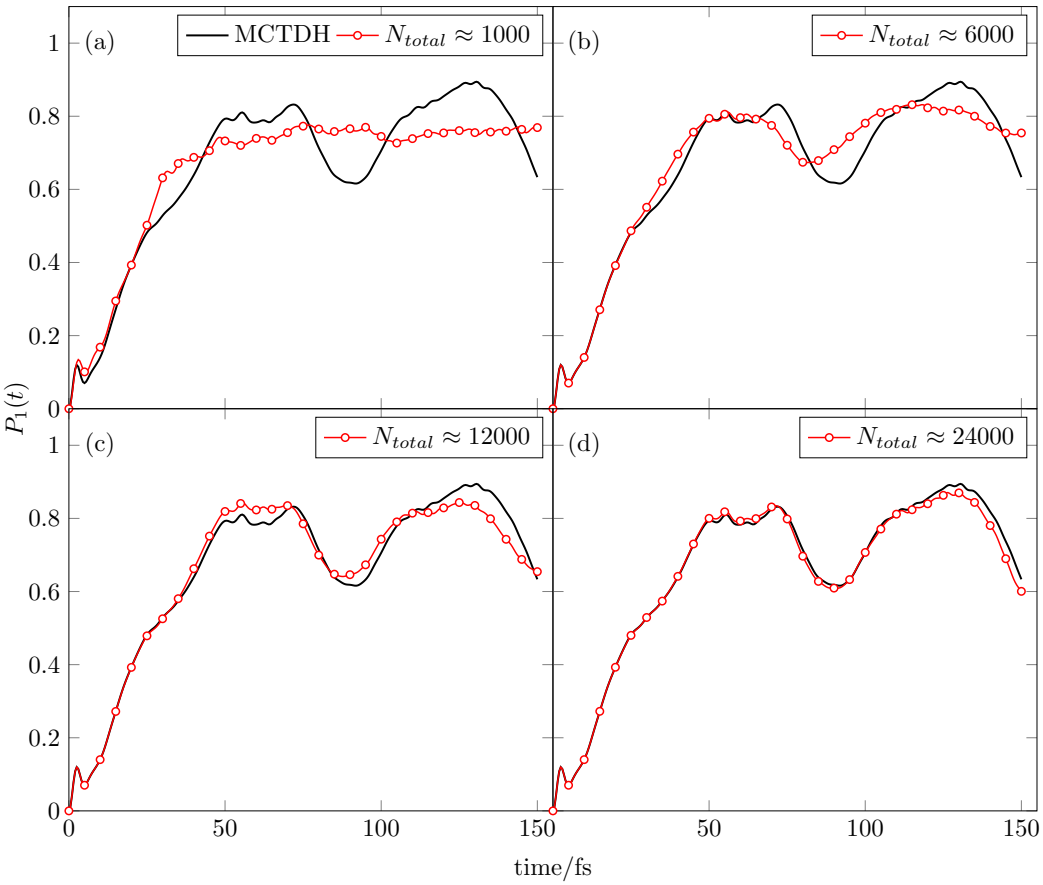

Figure 3.6: Populations of the lower S_1 excited state of pyrazine, $P_1(t)$, for the 4D pyrazine Hamiltonian, using varying basis set size, N_{total}, compared to exact MCTDH data.[65]

population spuriously levelling off upon reaching a value of $P_1(t) \approx 0.75$. Some similarity with the scaling observed for the autocorrelation functions can however be found on examination of the short time accuracy of the populations. Even the smallest basis set employed ($N_{total} \approx 1000$) is able to accurately reproduce the MCTDH dynamics up to the shoulder-like feature occurring at approximately $t = 25$ fs. The cause of this short time accuracy, or more correctly, the long-time drop off in accuracy is likely related to the nature of the sampling trajectories. As the Ehrenfest dynamics employed here, while taking into account multiple electronic states, include no quantum effects, the longer they are propagated for, the more the divergence of quantum and classical trajectory paths is likely to affect the basis function placement. Thus, as the basis set is required to describe the evolution of the wavefunction in time for the entire timescale of propagation, at longer times, the regions of phase space, sampled by the classical trajectories, result in inaccurate dynamics.

As the basis set size is increased, the overall accuracy with respect to the

MCTDH benchmark scales favourably, the largest calculation ($N_{total} \approx 24000$) yielding results that are numerically accurate within the variation introduced by the stochastic nature of the trajectory sampling, which is addressed in more detail in Section 3.4. Encouragingly, the calculation employing only half as many basis functions ($N_{total} \approx 12000$) still captures all essential features of the relaxation dynamics with qualitative accuracy. The improvements that are gained from increasing the basis set size beyond this point are marginal, suggesting that the method has converged.

Overall these results for the 4-dimensional pyrazine Hamiltonian are extremely encouraging. The relatively simple trajectory sampling method is clearly able to provide qualitatively accurate quantum dynamics results for this challenging benchmark problem, even using relatively small basis sets and thus incurring little computational cost. Furthermore, the method converges rather quickly with respect to basis set size, approaching quantitative accuracy for the largest basis sets, which are, due to the cheap nature of the sampling algorithm, still very computationally feasible. A more detailed investigation of the scaling and convergence of this method is presented in Section 3.4.

3.3.7 24-dimensional results

In order to test the performance of the trajectory sampling algorithm for systems of higher-dimensionality, the full 24-dimensional Hamiltonian was also considered. This constitutes an extremely challenging benchmark for any quantum dynamics method, given the high number of DOFs which, to obtain correct dynamics, must be treated explicitly across both electronic states. The availability of benchmark MCTDH results, known to be numerically exact with respect to experiment,[65] facilitates a direct measure of the accuracy of the algorithm set forth herein.

In keeping with the approach chosen for the 4-dimensional version of this Hamiltonian, a set of 4 calculations were run, employing similar input parameters, the details of which are shown in Table 3.2.

Figure 3.7 shows the wavefunction autocorrelations functions, $C(t)$, compared to MCTDH results,[65] for varying basis set size. Given the limited insight into the detailed dynamics of the system that can be gained from the photoabsorption spectra due to the necessity of dampening and sampling functions, as outlined in Section 3.3.6, the aforementioned are omitted here. Inspecting these results, it becomes immediately clear that the 24-dimensional Hamiltonian constitutes a far more challenging problem, as the accuracy with which the dynamics are reproduced scales far more slowly with basis set size.

The smaller sized basis sets, that in the case of the 4-dimensional model, were

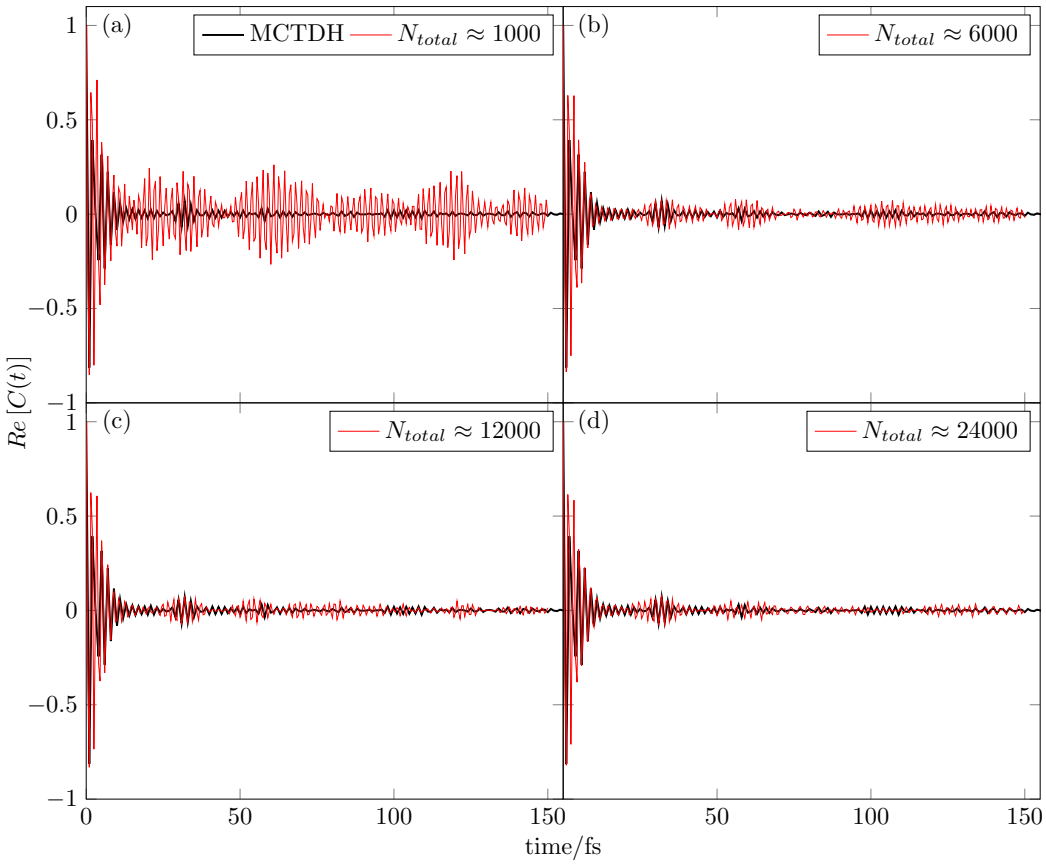

Figure 3.7: Wavefunction autocorrelation functions, $C(t)$, for the 24-dimensional pyrazine Hamiltonian, using varying basis set sizes, N_{total}, compared to MCTDH benchmark results.[65]

able to provide at least qualitative accuracy, now produce decidedly incorrect dynamics, although the short time accuracy, alluded to in the section above, is retained to a certain extent. The oscillations in $C(t)$ are clearly overestimated and in the case of the smaller basis set sizes, $C(t)$ looks rather similar to the 4-dimensional data shown in Figure 3.4. As basis set size is increased, the correct shape of the autocorrelation function is recovered to the point of qualitative accuracy, however the amplitude of oscillation remains too high.

Figure 3.8 shows the corresponding populations of the lower excited S_1 state, $P_1(t)$. Again it is clear that the smaller basis sets are incapable of properly describing the dynamics even in the very short-time limit, as the small reabsorbance feature at around $t \approx 5$ fs is not described at all, and result in spurious dynamics at longer time-scales. For the larger basis sets employed, the dynamics are relatively accurate in the short-time limit, up to approximately $t = 40$ fs, and at longer timescales, while not introducing spurious fine structure like the smaller basis sets, fail to capture the significant reabsorbance feature at $t \approx 75$ fs.

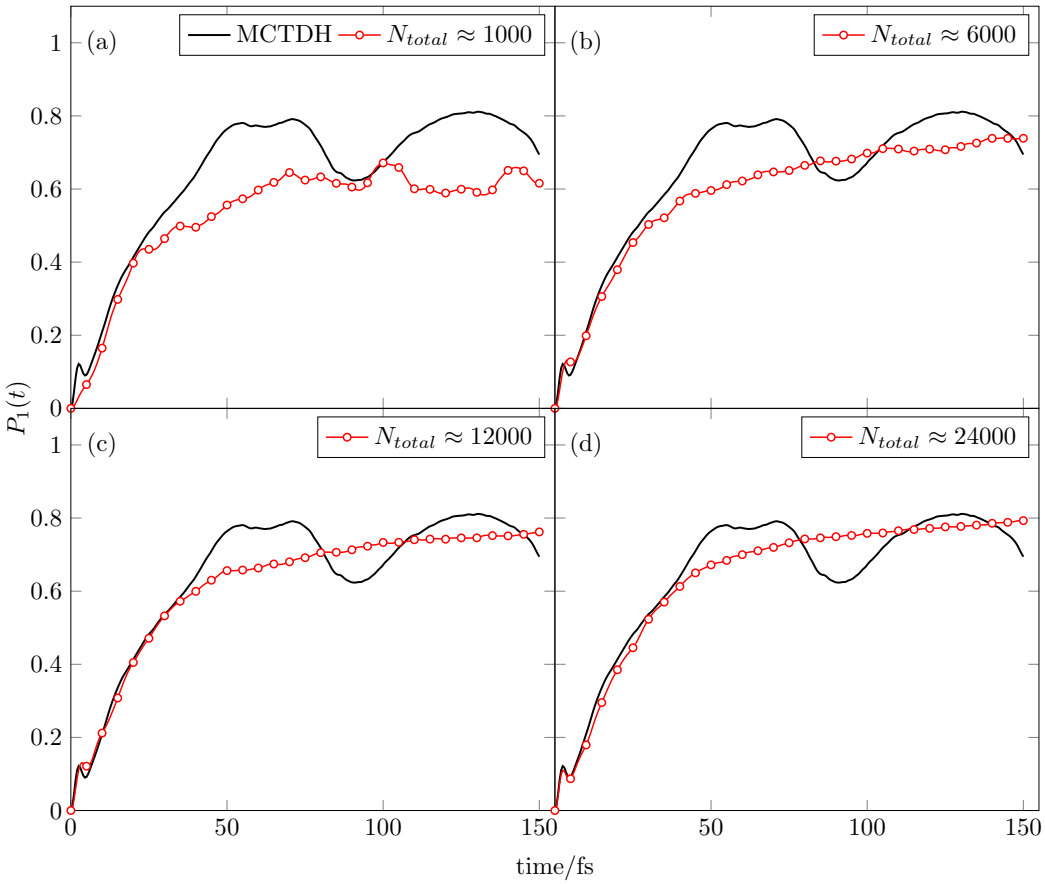

Figure 3.8: Populations of the lower S_1 excited state, $P_1(t)$, for the 24-dimensional pyrazine Hamiltonian, using varying basis set sizes, N_{total}, shown in comparison with exact MCTDH data.[65]

Table 3.2: Input parameters and final basis set sizes for trajectory sampling calculations of the 24D pyrazine Hamiltonian, the results of which are shown in Figures 3.7 and 3.8.

n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	m	N_{total}
1500	0.1	50	1500	0.1	33	940
1500	0.1	50	1500	0.1	200	5904
1500	0.1	50	1500	0.1	400	12059
1500	0.1	50	1500	0.1	800	23673

Overall, whilst not performing as well, treating the full 24-dimensional Hamiltonian, as in the 4D case, it is worth noting that qualitative accuracy could still be achieved, given that enough basis functions were sampled. This is rather encouraging, as it suggests that even for large, multi-dimensional problems, the trajectory sampling method can in principle generate basis sets, that allow accurate dynamics to be calculated, albeit in the short-time limit. As already mentioned in Section 3.3.6, the inaccurate dynamics observed in the long-time limit for both versions of the pyrazine Hamiltonian investigated here, are likely due to the divergence of the semi-classical Ehrenfest sampling trajectories from the correct quantum behaviour. This results in a shift of the basis set in phase space, away from regions relevant to wavefunction propagation, thus reducing the accuracy of the dynamics resulting from the propagation algorithm. Chapter 4 addresses this aspect of the trajectory sampling algorithm with a simple but effective modification.

3.4 Algorithm parameters and performance

The pyrazine benchmark discussed above clearly demonstrates that the trajectory sampling algorithm outlined in Section 3.2 may be used to obtain qualitatively accurate quantum dynamics without suffering from some of the drawbacks of traditional approaches, as mentioned in Section 3.1. This however raises questions with respect to the detailed performance and scaling of this strategy, as well as to how its computational performance compares to other popular quantum dynamics methods, all of which will be addressed below. Furthermore, the algorithm is driven by a number of complementary input parameters, the effects of which are investigated in detail, as significant insights into the details of the algorithm's performance may be gained from this.

3.4.1 Scaling and convergence with respect to basis set size

As already discussed in Section 3.3, the accuracy of the dynamics produced by the trajectory sampling algorithm scale expectedly favourably with increasing basis set size. As this approach is, at its core, a finite basis set method, this is not at all surprising, however the numerical nature of this scaling, rate and limits of convergence will be of great use in comparing to alternative quantum dynamics strategies, as well as allowing a detailed measure of the improvements introduced by the modification, proposed in Chapter 4. The extent to which the accuracy of the algorithm is affected by basis set size has already been observed to be greatest in the case of the 4-dimensional version of the pyrazine Hamiltonian, the results for which are presented in Section 3.3.6. This particular benchmark was thus

chosen for the detailed investigation into basis set scaling, set forth below.

In order to ensure that stochastic variations are kept to a minimum, 4 independent iterations of the algorithm were run and the results averaged. Following closely the parameters of those presented in Section 3.3.6, all calculations involved 1500 timesteps of sampling followed by 1500 of wavefunction propagation ($n_t = n_p = 1500$ fs), both employing timesteps of 0.1 fs ($\Delta t_t = \Delta t_p = 0.1$). The sampling frequency was set to $1/50$ ($n_s = 50$), while the number of sampling trajectories was varied from $m = 33$ to $m = 800$, resulting in total basis set sizes from $N_{total} \approx 1000$ to $N_{total} \approx 24000$. The number of sampling trajectories, m, as well as the resulting average basis set size, $\overline{N}_{total} = \frac{1}{4}\sum N_{total}$, are shown in Table A.1 of Appendix IV .

Then, to assess the overall accuracy of these calculations, the mean absolute error (MAE) and mean absolute percentage error (MAPE) with respect to MCTDH populations, $P_1^M(t)$ were calculated for each data set using

$$\mathrm{MAE} = \frac{1}{N_o}\sum_j^{N_o} |P_1(j) - P_1^M(j)| \qquad \mathrm{MAPE} = 100\frac{1}{N_o}\sum_j^{N_o} \left|\frac{P_1(j) - P_1^M(j)}{P_1^M(t)}\right| , \tag{3.38}$$

where N_o is the number of data points available for the population, $P_1(t)$. The resulting error values were then averaged for each set of 4 calculations ($\overline{\mathrm{MAE}}$ and $\overline{\mathrm{MAPE}}$) and standard deviations, σ, for were calculated for both, as a measure of the extent of stochastic variation, using the general formula

$$\sigma = \frac{1}{4}\sum_j^4 \sqrt{\left(E_j - \overline{E}\right)^2} , \tag{3.39}$$

where E_j is the error (MAE or MAPE) for the jth data set (out of 4) and $\overline{E}$ is the average error calculated as $\overline{E} = \frac{1}{4}\sum_j^4 E_j$. The values for the averaged errors and the corresponding standard deviations are also shown in Table A.1

Figure 3.9 shows the scaling of both errors in the population of the lower excited S_1 state. The fast rate of convergence for smaller basis sets ($\overline{N}_{total} < 10000$), is highly encouraging, approaching exponential scaling. This suggests that the trajectory sampling method scales excellently when aiming to obtain qualitatively accurate dynamics. Convergence at larger basis set sizes ($\overline{N}_{total} > 10000$), even more encouragingly, does not slow down as much as might be expected in the case of purely exponential scaling of the accuracy, instead the rate of convergence appears to depend roughly linearly on the number of basis functions. Finally, the low values of the percentage errors across the entire range of basis set sizes, shown in Figure 3.9, indicate that for this particular benchmark, the algorithm presented herein can provide results within 1% of the exact dynamics at relatively low computational cost.

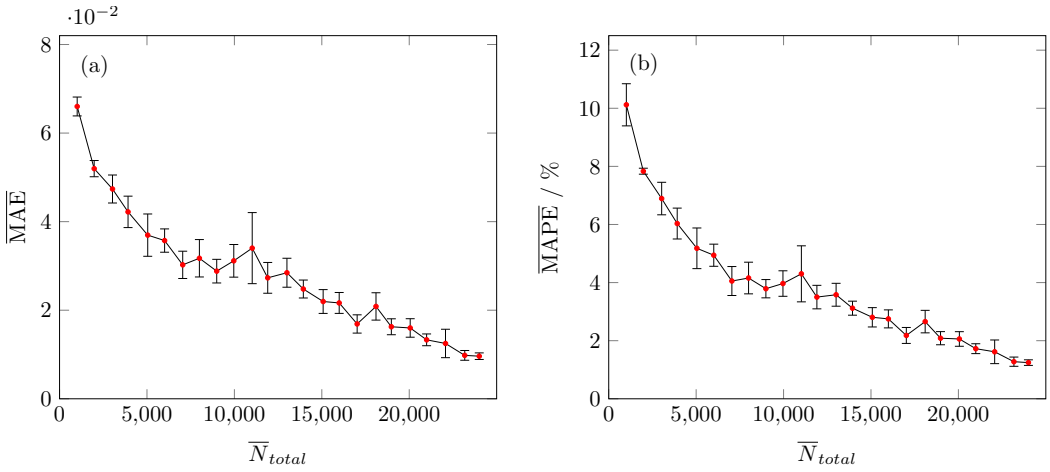

Figure 3.9: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian as a function of total basis set size, $\overline{N}_{total}$, and (b) corresponding mean absolute percentage errors.

3.4.2 Basis function sampling frequency

As mentioned in Section 3.2.1, the size of the basis set sampled during the trajectory algorithm is influenced by three parameters. The number of timesteps the trajectories are propagated for, n_t, the total number of trajectories employed, m, and the frequency with which basis functions are, on average, stored for each trajectory, $1/n_s$. Of these three, m and n_s, are not linked in any way to the propagation algorithm, while n_t is often chosen to be equal to the number of timesteps the TDSE will we solved for. The effects of changing the latter will be addressed in Section 3.4.3, however it stands to reason that the sampling algorithm will not perform equally well for all values of m and n_s, thus warranting methodical investigation.

Considering for a moment the two most extreme choices possible for these two parameters, these constitute either only running a single trajectory which samples a very dense but narrow basis set across phase space, or a large number of trajectories, which each only yield a single basis function. Both these cases are subject to significant limitations on the accuracy that may be achieved using the basis sets they produce. Using only a single trajectory involves only a single set of initial conditions being sampled from the initial wavefunction, which is unlikely to yield a large enough spread of the basis set across phase space. A large number of trajectories, each sampling only a single basis function on the other hand, is likely to result in too random a distribution of the basis set, as the sampling across any given trajectory is stochastic, every timestep storing a basis function with equal probability, $1/n_s$.

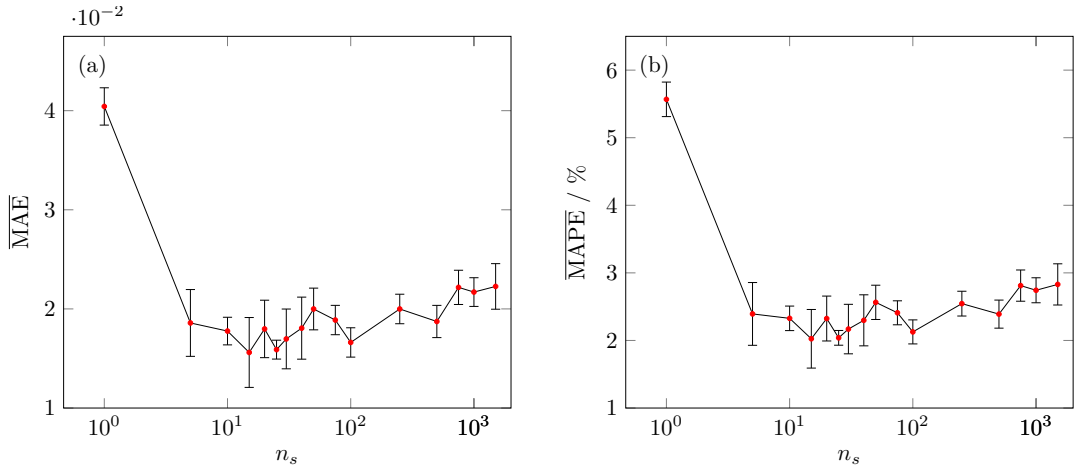

Figure 3.10: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian as a function of the number of the sampling frequency parameter, n_s; the sampling trajectories employed, m, is determined by Eq. 3.4 and (b) corresponding mean absolute percentage errors, using an average basis set size of $N_{total} \approx 18000$.

To investigate the effect the ratio of the number of sampling trajectories to the sampling frequency has on the accuracy of the trajectory sampling algorithm, a set of calculations for the 4-dimensional pyrazine Hamiltonian were run at a constant basis set size, $N_{total} \approx 18000$, varying the values of m and n_s, as outlined in Table A.2. The remaining parameters follow those employed for the calculations in Section 3.3.6 and are summarised in Table 3.3. Similarly to the investigation presented in Section 3.4.1, in order to keep variation due to the stochastic nature of the sampling algorithm to a minimum, 4 identical calculations were run for each of the 16 $m:n_s$ ratios investigated, and the results averaged. Again in keeping with the methodology employed above, the average MAE, MAPE and corresponding standard deviations were calculated using Eqs. 3.38 and 3.39.

Figure 3.10 shows the result of the 4×16 calculations run for the 4D pyrazine Hamiltonian. The MAE and MAPE suggest that the hypothesis with respect to the accuracy dependence on the ratio of $m:n_s$ was indeed correct. Both extremes of this ratio result in clearly less accurate dynamics, due to the basis set not cover-

Table 3.3: Input parameters for trajectory sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figure 3.10.

n_t	Δt_t/fs	n_p	Δt_p/fs	N_{total}
1500	0.1	1500	0.1	≈ 18000

ing the areas of phase space, relevant to wavefunction propagation. Interestingly, using few trajectories with very frequent basis set sampling seems to result in far less accurate dynamics, highlighting the importance of properly sampling the initial conditions of the wavefunction. Taking into account the standard deviations shown in Figure 3.10, it can be seen that, as long as the limit of extremely few trajectories is avoided, the ratio of $m:n_s$ otherwise results in relatively comparable accuracy, with only a small decrease due to a large number of sparsely sampling trajectories. Finally it is worth pointing out the the number of sampling trajectory timesteps, n_t, acts as a soft limiter for the sampling frequency parameter, n_s, as cases where $n_s > n_t$ allow for sampling trajectories which store no basis functions whatsoever. While this does not hinder the trajectory sampling algorithm, the relatively precise control over the total basis set size, illustrated in Eq. 3.4 is forfeited in such a calculation.

3.4.3 Timestep ratios and "oversampling"

Another set of parameters modifying the behaviour of the trajectory sampling method are the timestep durations for both the sampling and propagation algorithm, Δt_t and Δt_p respectively, introduced in Section 3.2.2. The latter is often dictated by the system under investigation, however Δt_t may essentially be chosen freely, although should the value differ significantly from that of Δt_p, a negative impact on performance might be expected. Assuming similar basis set sizes, N_{total}, a drastically smaller timestep duration Δt_t can be interpreted as sampling phase space in a very dense manner, similar to a high sampling frequency, $1/n_s$. Conversely, significantly longer sampling timesteps represent sparser sampling, as the distance covered in phase space, between potential basis function sampling points, is larger.

In order to investigate this hypothesis, following the approach outlined above, a set of 4×11 calculations was run, varying the values of Δt_t, as well as n_t and m, in order to maintain a total basis set size of $N_{total} \approx 18000$. The details of the input parameters used can be found in Table 3.4. Again the average MAE and MAPE error as well as the corresponding standard deviation were calculated,

Table 3.4: Input parameters for trajectory sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figure 3.11.

n_s	n_p	Δt_p/fs	N_{total}
50	1500	0.1	≈ 18000

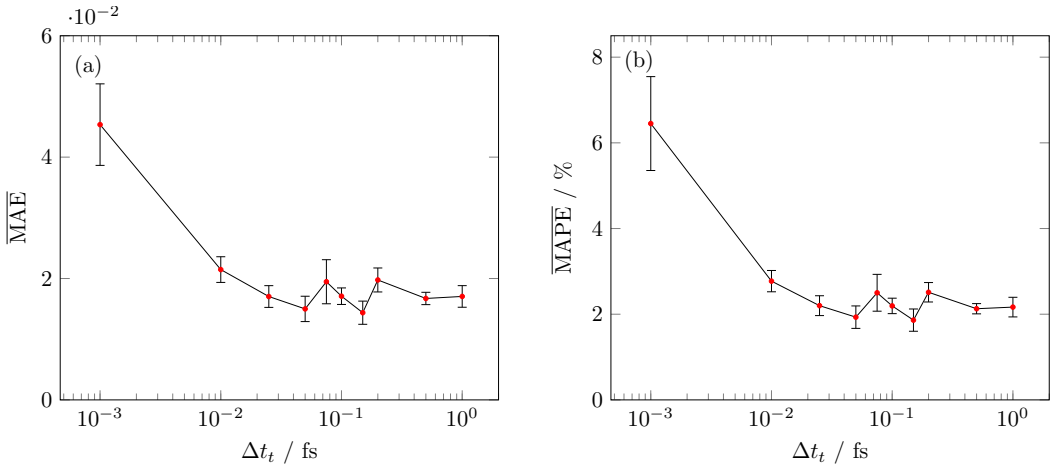

Figure 3.11: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian using varying timestep durations, Δt_t, and (b) corresponding mean absolute percentage errors, with an average basis set size of $N_{total} \approx 18000$.

using Eqs. 3.38 and 3.39, the results of which are shown in Figure 3.11 and can in detail be found in Table A.3. Inspecting first of all the absolute values of the error, the above hypothesis appears to hold, although the aside from the very extreme cases where $\Delta t_t \gg \Delta t_p$ and $\Delta t_t \ll \Delta t_p$, the effect of changing Δt_t seems to be minimal. Further confirmation of this can be found *via* inspection of the standard deviations in the errors. The latter are comparatively large, suggesting that a significant proportion of the variation in the $\overline{\text{MAE}}$ and $\overline{\text{MAPE}}$ is due to the stochastic nature of the sampling trajectories.

Overall this is encouraging, as the choice of the propagation timestep for some systems may not be obvious, but will clearly significantly affect the nature of the dynamics. Given that, within reasonable limits, the choice of the sampling timestep, Δt_t, is unlikely to significantly change the accuracy of the trajectory sampling algorithm, it can, in the case described above, be kept constant as Δt_p is varied in order to find the optimal value.

One final aspect of the sampling algorithm that warrants investigation relates to the duration for which the sampling trajectories are propagated, with respect to the timescale of the TDSE solution. Again, while it may seem natural to sample phase space for the same amount of time as the wavefunction will be propagated for, the trajectory sampling algorithm does not strictly require this to be the case. Given the classical nature of the sampling trajectories, it is likely that the rate at which they explore phase space will vary from that of the actual dynamics of the wavefunction, the latter being quantum in nature and thus inherently different.

There are two possible cases here, the first, in which the sampling trajectories are propagated for a longer period of time than the TDSE solution, is termed "over-

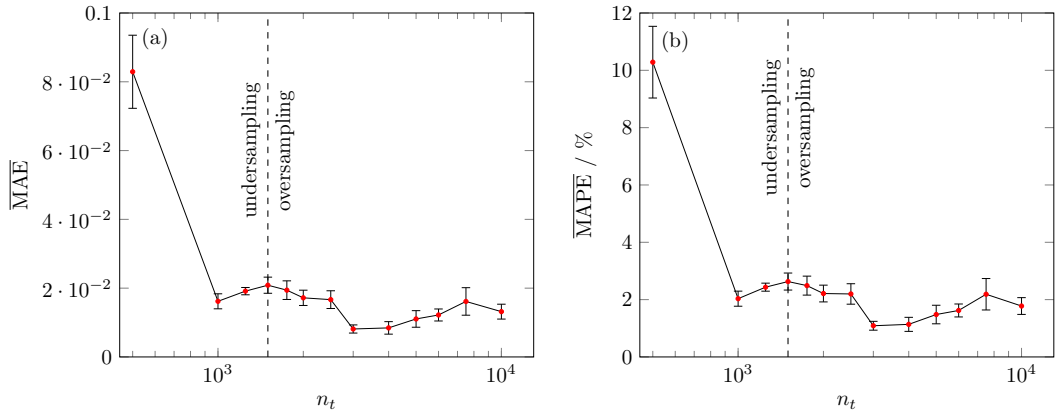

Figure 3.12: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian varying the duration of the sampling trajectories via the total number of sampling timesteps, n_t as well as the number of trajectories, m, and (b) corresponding mean absolute percentage errors, with an average basis set size of $N_{total} \approx 18000$.

sampling", the converse is, by logical extension, referred to as "undersampling". In order to explore these concepts, a set of calculations was run, following closely the well documented approach taken for the other investigations presented in this Section. Thus, a set of 4×14 calculations was run, varying the number of timesteps of trajectory sampling, n_t and the number of trajectories run, m, in order to maintain a total basis set size of $N_{total} \approx 18000$, the remaining parameters being shown in Table 3.5. In keeping with the methodology used above, the average MAE and MAPE error were calculated using Eq. 3.38, while the extent of stochastic variation was determined via the standard deviation from Eq. 3.39. Average basis set sizes and resulting errors are shown in Table A.4.

Figure 3.12 shows clearly that varying the duration of the sampling trajectories has a significant impact on the accuracy with which the resulting basis set is able to represent the wavefunction moving in phase space. Unsurprisingly, "undersampling", that is sampling for drastically less time than the wavefunction will be propagated for considerably increases the error, as, the trajectories simply do not reach certain areas of phase space, which are however visited by the wavefunc-

Table 3.5: Input parameters for trajectory sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figure 3.12.

n_s	n_p	Δt_p/fs	Δt_t/fs	N_{total}
50	1500	0.1	0.1	≈ 18000

tion. The absence of any basis functions in the aforementioned areas reduces the accuracy with which the wavefunction can be described, as it approaches them, thus reducing the overall accuracy of the dynamics.

Conversely, "oversampling" by allowing the sampling trajectories to evolve for longer than the duration for which the TDSE will be solved, does improve the accuracy with respect to exact MCTDH results,[65] however this increase in the quality of the dynamics is marginal and evolving the trajectories for too long negates this benefit. The latter can be rationalised by considering for one the fact that due to the increased number of timesteps, n_t, in order to maintain a constant basis set size, N_{total}, overall fewer trajectories are run, which has, in Section 3.4.2, been shown to negatively effect the quality of results, due to insufficient initial conditioning sampling. Furthermore, as trajectories start exploring more and more of phase space, the proportion of basis functions in any given area decreases, thus, given that using classical-like trajectories is likely to result in sampling of at least some irrelevant areas, will decrease the number of basis functions in areas that are relevant to wavefunction propagation.

The practices of "over-" and "undersampling" were found to indeed affect the accuracy of the trajectory sampling algorithm, the former providing a slight improvement unless $n_t \gg n_p$, while the latter consistently negatively impacted the results, for reasons discussed above.

In conclusion, the trajectory sampling algorithm, while conceptually extremely simple, incorporates a number of parameters which may be used to tune its performance. With regards to applying this method to a variety of different problems, this is considered an advantage as it provides flexibility and allows the performance to be tuned to specifically suit the system at hand. There are undoubtedly more correlations that could be investigated, both between different parameters as well as the input and the accuracy with which the algorithm performs for any given problem, however the basic relationships described above should allow this method to be applied to most quantum dynamical systems, without too much further benchmarking being necessary.

3.4.4 Computational performance and cost

A significant motivation for the introduction of the trajectory sampling method discussed above is to minimise computational costs and improve performance over other, time-independent quantum dynamics methods,,[41,58] which in turn should allow larger systems to be treated. Thus, in order to assess both the aforementioned computational factors, the four calculations discussed in detail in Section 3.3.6 are investigated here in terms of both the required processing time, as well

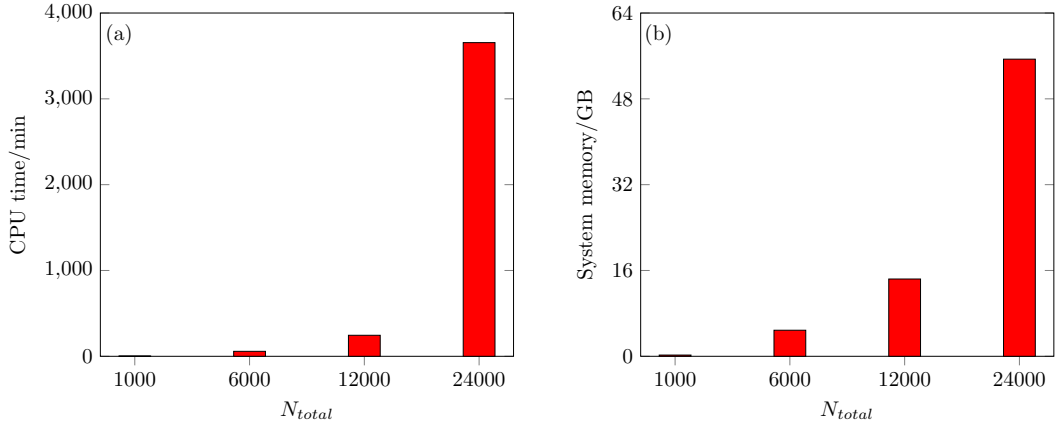

Figure 3.13: Total CPU time and (b) maximum system memory required by trajectory sampling calculations of the 4D pyrazine Hamiltonian, the input parameters of which are shown in Table 3.1.

as the maximum system memory used during runtime.

All calculations were performed on single compute nodes, containing either two Intel Xeon E5-2630 v3 2.4 GHz 8-core processors, for a total of 16 cores per calculation, or two Intel Xeon E7-4809 v3 2.0 GHz 8-core processors, for a total of 16 cores (or 32 threads) per calculation. 16 core nodes were equipped with a maximum of 64GB of DDR4 RAM, or 4GB per core, while 32 thread nodes supported up to 1024GB of DDR4 RAM, or 32GB per core. The adaptive sampling algorithm is implemented using the FORTRAN programming language and accelerated using OpenMP and LAPACK.

Table 3.1 summarises the input parameters and total basis set size of the four calculations investigated, while Figure 3.13 shows the total CPU time and maximum requested memory per calculation.

3.5 Conclusions

This chapter has introduced a novel algorithm that allows basis sets for quantum dynamics simulations to be sampled using simple classical-like trajectories, evolving on the PES of the system. The proposed benefits of this approach are twofold. The paths of the sampling trajectories, while not fully quantum, are guided by the shape of the PES and thus the resulting basis set is not subject to the exponential scaling that limits the approach of methods relying on grids of static basis functions. Furthermore, as the GWPs remain fixed in phase space once the sampling stage of this method is complete, the time evolution of the wavefunction being expressed solely *via* propagation of the expansion coefficients, the numer-

ical issues, encountered when solving variational equations of motion for basis set parameters, are avoided entirely. This method also conserves energy, which while implicit in the TDSE, is not necessarily the case for methods, propagating basis functions with non-variational equations of motion.[47]

In order to test the validity of the assumptions laid out above and measure the performance of the algorithm presented here, the challenging vibronic pyrazine Hamiltonian, modelling the relaxation dynamics following excitation to the second excited state,[65] was chosen as a first benchmark problem. Due to availability of numerically exact MCTDH reference results,[65] both the reduced, 4-dimensional model and the full Hamiltonian, incorporating all 24 normal modes of the molecule, were treated with the new trajectory sampling approach. Encouragingly, in the case of the 4-dimensional version, even comparatively small basis sets were able to capture the dynamics of the system with qualitative accuracy. The method furthermore scaled excellently with respect to basis set size, readily converging towards nearly quantitative results. These two observations suggests that this approach for sampling basis sets is indeed subject to more favourable scaling than comparable time-independent methods.[39,41,58] The full 24-dimensional model of the pyrazine Hamiltonian proved a more challenging problem, however the trajectory sampling algorithm did yield qualitatively accurate dynamics. While this required larger basis sets than in the case of the 4-dimensional Hamiltonian, the computational cost of these calculations was still comparatively low, considering the large size of the system and complexity of the model.

The input parameters for the trajectory sampling algorithm were systematically investigated, revealing favourable scaling and convergence, as well as a number of insights into the nature of the sampling of phase space achieved with it. The balance between the number of independent sampling trajectories and the frequency of basis function storage was found to have a key impact on the accuracy of the algorithm. Very few trajectories, sampling GWPs very regularly, were found to significantly limit the accuracy of the dynamics resulting from the basis set sampled in this way, which was linked to insufficient sampling of the initial conditions of wavefunction. Using, in the other extreme, a large number of trajectories, each infrequently sampling phase space, was also found to degrade the quality of results, albeit not by as much. Overall this suggests that, as long as the initial conditions are sampled in a rigorous manner, the algorithm is rather flexible with respect to the number of trajectories (and thereby the sampling frequency) employed.

The duration of each sampling timestep was found to not significantly affect the accuracy of the method, as long as its value is not extremely different from the duration of the timestep with which the TDSE is solved. In the case of sys-

tems where the ideal propagation timestep is not known *a priori*, this allows the sampling timestep to be kept constant while varying the former, thus reducing the difficulty of the optimisation procedure. Allowing the duration of the sampling trajectories vary with respect to the timescale of TDSE solution was found to significantly influence the accuracy of the trajectory sampling algorithm. Unsurprisingly, if the sampling trajectories are propagated for a significantly shorter time than the wavefunction evolves for, the resulting basis set is not very well suited to describing the dynamics. Increasing the sampling time beyond the propagation limit did yield slightly more accurate basis sets, however this increase was only marginal.

There are however a number of assumptions associated with the trajectory sampling method, which are to an extent, highlighted by the results for the pyrazine benchmark. For both the diabatic state populations and especially for the wavefunction autocorrelation functions calculated, the algorithm was found to perform excellently in the short-time limit. Deviation from the exact solutions was, in the majority of cases, limited to longer timescales. This is easily explained, as the key assumption this method is based on, expects the classical-like trajectories to sample, based on the PES of the system, regions of phase space relevant to quantum propagation of the wavefunction. In the short time limit, classical mechanics are well known to constitute a fair approximation for quantum behaviour, however at longer timescales, the two diverge, resulting in the placement of essentially useless basis functions, in phase space completely irrelevant to the solution of the TDSE.

There exist a variety of potential causes for this divergence, one or more of which, depending on the nature of the system, may be influencing the effectiveness of the sampling trajectories at longer times. There is however a rather straightforward modification that can be made to the trajectory sampling method, which addresses these issues. Chapter 4 introduces this improved version of the algorithm discussed here and investigates the extent to which it improves performance for a number of challenging quantum dynamics benchmarks.

Chapter 4

Adaptive Trajectory Sampling

This chapter introduces a modified version of the trajectory sampling method presented in Chapter 3, capitalising on the validity of classical mechanics as an approximation to the quantum solution in the short time limit. After a brief introduction to the underlying ideas, the algorithm implementing this adaptive sampling strategy is presented. The new method is then applied to two challenging quantum dynamics benchmarks and finally, its performance and the effects of input parameters are systematically investigated.

The contents of this chapter have, in part, been published:
M. A. C. Saller and S. Habershon, *J. Chem. Theo. Comput.*, **13**, 3085–3096 (2017)

4.1 Introduction

Chapter 3 introduced the trajectory sampling method, employing classical-like trajectories evolving on the PES of the system, to sample basis sets for quantum dynamics simulations. Although this rather simple approach was shown to perform encouragingly well for the challenging benchmark modelling relaxation dynamics of pyrazine, a number of shortcomings were identified. While the short-time accuracy of this method was found to be excellent, even for relatively small basis sets, in the long-time limit performance dropped off significantly, which, in higher-dimensional cases, could not even be counteracted by drastically increasing the basis set size.

This apparent limitation to short timescales was linked to the assumption, underlying this particular approach, that classical-like trajectories, evolving on the PES of the system, will sample regions of phase space, relevant to the quantum dynamics of the wavefunction. Classical mechanics are well known to constitute a relatively good approximation to quantum mechanics, in the short-time limit,[42,43] however at longer times, there are a number of causes for the divergence of these two.

Classical mechanics are, for example, inherently unable to properly treat tunnelling through an energy barrier, when the barrier height exceeds the energy of the classical trajectory. Even in cases where enough energy is available to overcome such a barrier, the path described by a classical trajectory is fundamentally different from the true quantum tunnelling path. Furthermore the inability of classical mechanics to properly treat zero point energy has been well documented.[116,117] The unrestricted flow of zero point energy between degrees of freedom, allowed by classical mechanics, can, in extreme cases, lead to violation of the uncertainty principle and readily introduces significant differences between the motion of quantum and classical trajectories in phase space.

This chapter recognises the limitations of the trajectory sampling method, caused by this essentially unavoidable classical-quantum divergence in the long-time limit. A modification of the algorithm outlined in Chapter 3 is proposed, the aim of which is to, instead of using a single set of trajectories to sample the basis set across the entire time domain of desired dynamics, split each calculation into a number of shorter instances of trajectory sampling followed by wavefunction propagation. This limits the sampling stage to relatively short times, thus hopefully minimising the error due to classical-quantum divergence and improving the overall accuracy of the method.

In order to facilitate seamless dynamics, a novel algorithm for minimising and optimising the basis set at the end of a single instance of trajectory sampling and

propagation is introduced, based on the matching pursuit scheme, used previously by other quantum dynamics methods.[82,83] The aim of this step is to select, from the basis set used to propagate the wavefunction up to the current moment in time, only those basis functions which are essential to describe the wavefunction at the current point. These functions, which should constitute only a small proportion of the old basis set, then become the initial wavefunction of the next instance of the method, from which the initial conditions of the sampling trajectories are sampled.

For the remainder of this chapter, the algorithms responsible for the implementation of multiple instances of the trajectory sampling method, strung together, and the minimisation and optimisation of the inherited basis set are first discussed in detail. Where convenient, the trajectory sampling and adaptive sampling algorithms will be referred to as TSA and aTSA respectively, the latter having been chosen as the it represents a modification and still incorporates the idea of using simple classical like trajectories as the driver for basis set sampling. Then, the new, adaptive sampling method is applied to two challenging quantum dynamics benchmarks. First, in order to allow the performance and accuracy increase over the old method to be quantified, the vibronic Hamiltonian describing the relaxation dynamics following photoexcitation of pyrazine is revisited. Secondly, as minimising the negative impact of classical-quantum divergence should improve the ability of this method to treat systems with strong quantum effects, a synthetic benchmark for quantum tunnelling is investigated. Finally, the performance dependence on input parameters of this new adaptive sampling algorithm is systematically probed, giving further insight into the nature of this new sampling strategy.

4.2 Adaptive sampling

The detailed description of the trajectory sampling method found in Section 3.2 is still relevant here, as the adaptive algorithm is mainly concerned with facilitating the seamless transition between instances of the former. Thus, the details of these single "bursts", as they shall be referred to from now on, are not discussed in detail here, although the relevant sections of the previous chapter will be highlighted where ever it is deemed necessary. This section is split into two parts, which are however intimately linked. The first introduces the overall algorithm, implementing the adaptive sampling method, while the second focuses on the details of the matching pursuit step, bridging the gap between two bursts.

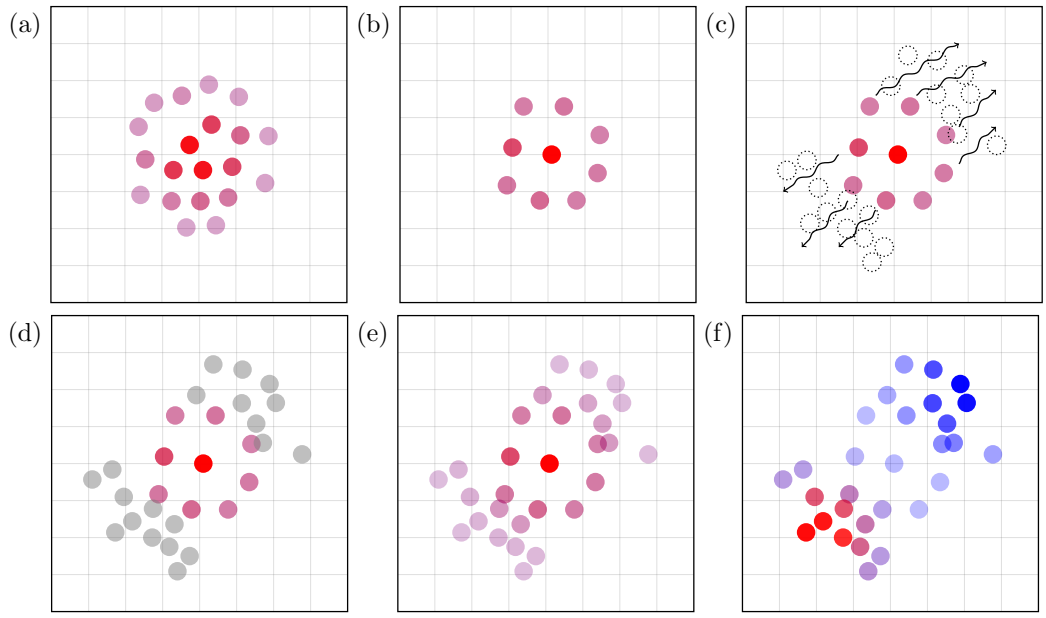

Figure 4.1: The adaptive trajectory sampling algorithm: (a) Starting with a wavefunction expanded in a basis, the number of basis functions is minimised and their parameters optimised, resulting in (b) a more compact wavefunction expansion. (c) Sampling trajectories are initialised from the basis functions making up the wavefunction, which sample phase space and store basis functions along their path for a short amount of time, resulting in (d) a set of additional basis functions, which then may be used in conjunction with the minimised basis to propagate the wavefunction in time by (e) assigning all functions coefficients and (f) propagating them in time variationally. This procedure is then iterated to give long-time dynamics.

4.2.1 Trajectory burst algorithm

In comparison to the trajectory sampling method, where basis set sampling and wavefunction propagation is carried out only once, the adaptive sampling algorithm splits any given calculation into several instances of the former, as shown in Figure 4.1. In the interest of generality, the order of steps discussed below will be applicable to one such burst, which has been preceded by at least one completed iteration. The initial burst differs slightly from this, however in this special case, algorithm is identical to a short calculation, purely following the trajectory sampling scheme, discussed in detail in Chapter 3.

Initially, the basis set from the previous burst, that is the static basis functions and coefficients after wavefunction propagation has been completed, is read in, representing the exact initial wavefunction for this iteration. As this set is usually over complete with respect to the number of basis functions required to describe

the wavefunction at the current point in time, the number of basis functions is reduced to N_{MP}, using the matching pursuit minimisation and optimisation algorithm, outlined in detail in Section 4.2.2.

From the resulting, minimal basis set, initial conditions for a set of m trajectories are then sampled, using the Wigner quasi-probability distribution, covered in Section 3.3.3. These trajectories are then propagated on the PES of the system for n_t timesteps of length Δt_t, resulting in a total sampling time of $t_t = n_t \Delta t_t$. Note that t_t in the case of the adaptive sampling algorithm is drastically shorter than in Chapter 3. As before, basis functions are stored with a probability of $1/n_s$ as trajectories evolve, resulting, at the end of the sampling stage, in a static basis set of size $N_t \approx m\frac{n_t}{n_s}$.

To allow a seamless transition between bursts, the wavefunction at the time of the restart, should be represented as accurately as possible, thus the set of basis functions from the MP algorithm are added to the trajectory sampled ones, resulting in a total basis set size of $N_{total} = N_{MP} + N_t$. The expansion coefficients of the optimised set of N_{MP} GWPs are retained, in order to reflect the initial wavefunction at time zero (for this burst), and those of the N_t trajectory sampled basis functions are set to zero.

Following the algorithm presented in Section 3.2.2, the wavefunction is then propagated in time, using only the expansion coefficients, for n_p timesteps lasting Δt_p each, resulting in a total propagation time of $t_p = n_p \Delta t_p$. The wavefunction at the end of this propagation is then saved, in order to act as the input for the following iteration of the adaptive sampling method.

4.2.2 Matching pursuit minimisation and optimisation

Capitalising on the overcomplete nature of the basis set inherited from the previous burst, alluded to above, the MP algorithm is used to select the minimal number of GWPs required to represent the wavefunction to within a user-determined accuracy. This is further sped up by optimising the selected function's parameters, thus allowing each function to contribute to the wavefunction description as much as possible. Matching pursuit, originally a method for signal decomposition,[81] has been used as part of quantum dynamics calculations in the past, MP-SOFT constituting the most notable example, where the MP algorithm forms a key part of the method.[82,83]

The MP minimisation and optimisation algorithm gets initialised at the beginning of an aTSA calculation, assuming at least one iteration has already been completed. The input thus consists of the basis set describing the wavefunction at the end of the propagation stage of the previous iteration. This initial set of

basis functions and coefficients defines the wavefunction to within the highest accuracy the aTSA is capable. The goal of the MP minimisation and optimisation algorithm is to select from this large set, the lowest number of basis functions that determine an approximate wavefunction to within an accuracy threshold and determine the optimal expansion coefficients for this minimal basis set. To begin, the full inherited basis set defining the wavefunction is,

$$\Psi = \sum_{j}^{N^{(I)}} c_j^{(I)} \phi_j^{(I)} \,, \tag{4.1}$$

where $\phi_j^{(I)}$ are GWP basis functions and $c_j^{(I)}$ the corresponding expansion coefficients at the end of wavefunction propagation of the previous burst.

The initial basis set is then expanded in order to introduce a wider variety of phase space parameters and allow the MP algorithm to operate a higher efficiency, as well as reducing the number of optimisation steps necessary, as will become clear later. This is achieved, similarly to initial condition sampling for trajectories, using the Wigner quasi-probability distribution. From Ψ basis functions are selected at random and from the Wigner distribution corresponding to the chosen function, positions and momenta are sampled for a new GWP, which is added to the initial set with a zero valued coefficient. While in this case, the selection of basis functions to sample new parameters from is random, it is also possible for the former to occur in a biased manner, based for example on the overlap of the particular chosen function, $\phi^{\dagger}$, with the initial wavefunction, $\langle \phi^{\dagger} | \Psi \rangle$. This process is continued until the initial set has grown by a user determined factor, γ, such that $N^{(I)} \to \gamma N^{I}$. The parameter γ takes a value of 3.0 for all calculations presented in this chapter, with the exception of those in Section 4.5.5, where its effects are investigated in detail. At this point it is helpful to introduce the set $\{ \phi^{(I)} \}$ containing all basis functions defining Ψ.

Let the initial expanded wavefunction therefore be defined by the initial basis set

$$\Psi^{(I)} = \sum_{j}^{N^{(I)}} c_j^{(I)} \phi_j^{(I)} \,.$$

This will remain unchanged throughout the algorithm, as it forms the overlap target for which the algorithm measures its convergence. Thus define two more basis sets, the active and minimal sets, indicated by superscripts (A) and (M) respectively. The former represents the remainder of the wavefunction after subtraction of the latter, which is the approximation of the full wavefunction.

$$\Psi^{(M)} = \sum_{j}^{N^{(M)}} c_j^{(M)} \phi_j^{(M)} \,,$$

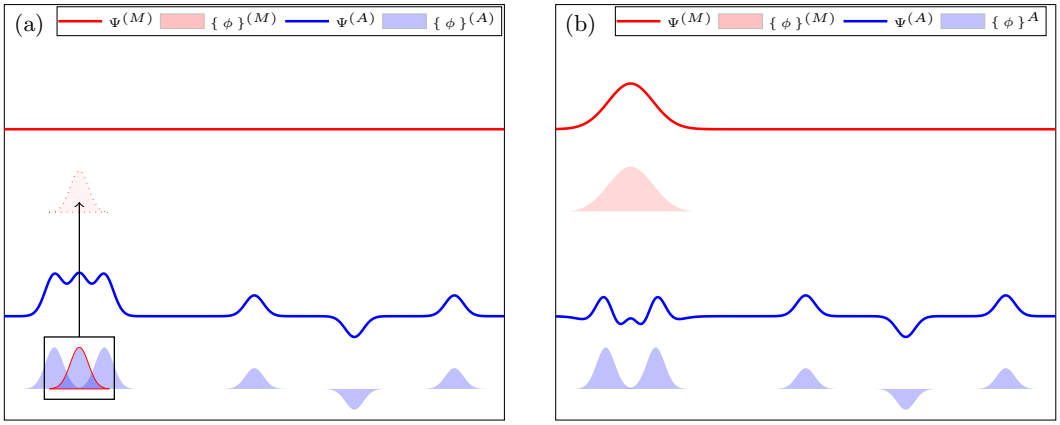

Figure 4.2: The two major stages of the MP minimisation and optimisation algorithm: (a) From the active basis set, $\{\phi\}^{(A)}$, the basis function possess the largest overlap with $\Psi^{(A)}$ is selected, removed from $\{\phi\}^{(A)}$ and (b) optimised with respect to overlap with $\Psi^{(A)}$ and added to $\{\phi\}^{(M)}$.

$$\Psi^{(A)} = \sum_{j}^{N^{(I)}} c_j^{(I)} \phi_j^{(I)} - \sum_{j}^{N^{(M)}} c_j^{(M)} \phi_j^{(M)} ,$$

$$\{\phi\}^{(A)} = \{\phi\}^{(I)} - \{\phi\}^{(M)} .$$

Note however that the last line actually reflects the removal of functions from $\{\phi\}^{(I)}$ before optimisation, which will be discussed in more detail later.
Figure 4.2 illustrates the two essential stages of the algorithm: MP basis function selection and optimisation. Note that at the start of the algorithm, $N^{(A)} = N^{(I)}$ and $N^{(M)} = 0$ and in general, at the start of each iteration of the algorithm $N^{(A)} = N^{(I)} - N^{(M)}$. The algorithm proceeds as follows

1. Select from $\{\phi\}^{(A)}$ the basis function, ${\phi^{(A)}}^{\dagger}$ that maximises the overlap with the remainder wavefunction

$$ {\phi^{(A)}}^{\dagger} = \underset{\phi \in \{\phi\}^{(A)}}{\operatorname{argmax}} \left\langle \phi \,\middle|\, \Psi^{(A)} \right\rangle .$$

2. Optimise the parameters, $(\mathbf{r}^{\dagger}, \mathbf{p}^{\dagger})$ of ${\phi^{(A)}}^{\dagger}(\mathbf{r}^{\dagger}, \mathbf{p}^{\dagger})$ to yield $\phi^{\ddagger}(\mathbf{r}^{\ddagger}, \mathbf{p}^{\ddagger})$, which maximises the overlap with the remainder wavefunction, such that

$$\phi^{\ddagger}(\mathbf{r}^{\ddagger}, \mathbf{p}^{\ddagger}) = \operatorname{argmax} \left\langle {\phi^{(A)}}^{\dagger}(\delta\mathbf{r}^{\dagger}, \delta\mathbf{p}^{\dagger}) \,\middle|\, \Psi^{(A)} \right\rangle .$$

3. Remove ${\phi^{(A)}}^{\dagger}$ from the active set,

$$\{\phi\}^{(A)} \rightarrow \{\phi\}^{(A)} - {\phi^{(A)}}^{\dagger} ,$$

$$N^{(A)} \rightarrow N^{(A)} - 1 ,$$

add $\phi^{\ddagger}$ to the minimal basis set

$$\{\,\phi\,\}^{(M)} \rightarrow \{\,\phi\,\}^{(M)} + \phi^{\ddagger}\,,$$

$$N^{(M)} \rightarrow N^{(M)} + 1\,,$$

and determine the new expansion coefficients of the minimal wavefunction via projection. If

$$\Psi^{(M)} = \Psi^{(I)}\,,$$

then

$$\sum_{j}^{N^{(M)}} c_j^{(M)} \phi_j^{(M)} = \sum_{j}^{N^{(I)}} c_j^{(I)} \phi_j^{(I)}\,.$$

Projecting from the left with $\{\,\phi\,\}^{(M)}$ yields

$$\sum_{k}^{N^{(M)}} \sum_{j}^{N^{(M)}} c_j^{(M)} \left\langle \phi_k^{(M)} \,\middle|\, \phi_j^{(M)} \right\rangle = \sum_{k}^{N^{(M)}} \sum_{j}^{N^{(I)}} c_j^{(I)} \left\langle \phi_k^{(M)} \,\middle|\, \phi_j^{(I)} \right\rangle\,,$$

which may be simplified to matrix notation and rearranged to yield

$$\mathbf{c}^{(M)} = \left(\mathbf{S}^{(M)}\right)^{-1} \mathbf{w}\,,$$

where $\mathbf{S}^{(M)}$ is the minimal basis overlap matrix and $\mathbf{w}$ is a vector of length $N^{(M)}$ with elements $w_k = \langle \phi_k^{(M)} | \Psi^{(I)} \rangle$.

4. Check for convergence of the algorithm. If

$$\left| \left\langle \Psi^{(M)} \,\middle|\, \Psi^{(I)} \right\rangle \right| \geq \zeta\,,$$

where ζ is the user-defined convergence criterion, the algorithm has converged and $\Psi^{(M)}$ should be accepted as the minimal representation of $\Psi^{(I)}$, else go to step 1.

4.3 Pyrazine benchmark

In order to allow the adaptive sampling algorithm to be compared to the trajectory sampling method, the first benchmark problem chosen is the vibronic pyrazine Hamiltonian, studied extensively with the latter. A detailed discussion of this model can be found in Section 3.3. Given that the lower, 4-dimensional version was used during the systematic investigation of the trajectory sampling algorithm's performance and the impact of parameters, presented in Section 3.4, the focus will remain on this version of the Hamiltonian, however the more challenging 24-dimensional version will be briefly discussed as well.

4.3.1 4-dimensional results

In order to assess the improvement over the trajectory sampling algorithm, afforded by the adaptive sampling method, a set of calculations was run, following closely those presented in Section 3.3.6. Timestep durations for both the sampling and propagation stages were $\Delta t_t = \Delta t_p = 0.1$ fs and the total time for which dynamics were obtained was 150 fs. A set of $N_b = 15$ bursts of equal length, $t_b = 10$ fs was employed, resulting in a total of 14 instances of the MP minimisation and optimisation algorithm, each with a convergence factor of $\zeta = 0.95$.

Consequently, each individual burst consisted of $n_t = 100$ timesteps of trajectory sampling followed by $n_p = 100$ steps of wavefunction propagation. The number of trajectories employed, m, and the basis function sampling frequency, $1/n_s = 1/10$, were kept constant across all bursts as well. The number of sampling trajectories was varied, taking values of $m = 100$, 300, 900 and 1800. A distinction must be made between the number of basis functions resulting from trajectory sampling and those inherited from any previous iterations of the algorithm as a result of the MP step. To this end, the input parameters for these calculations are shown in Table 4.1, along with the average (across all 15 bursts) number of total, trajectory sampled and inherited basis functions, N_{total}, N_t and N_{MP} respectively.

Figure 4.3 shows the populations of the lower excited S_1 state of pyrazine, compared to exact MCTDH results.[65] Discussion of the wavefunction autocorrelation functions and corresponding S_2 photoabsorption spectra is foregone here, as it was shown in Section 3.3.6 that the population provides the most in-depth measure of the accuracy with which the dynamics are described, across all of phase space.

It is immediately apparent that the adaptive sampling algorithm significantly outperforms the trajectory sampling method. Even the smallest basis set size employed captures many of the essential features of the dynamics in a qualitative manner. Furthermore, the adaptive method also appears to scale much more

Table 4.1: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figures 4.3 and 4.4.

N_b	n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ	m	N_{total}	N_t	N_{MP}
15	100	0.1	10	100	0.1	0.95	100	1749	996	752
15	100	0.1	10	100	0.1	0.95	300	3192	2972	220
15	100	0.1	10	100	0.1	0.95	900	9135	9023	112
15	100	0.1	10	100	0.1	0.95	1800	18111	18009	101

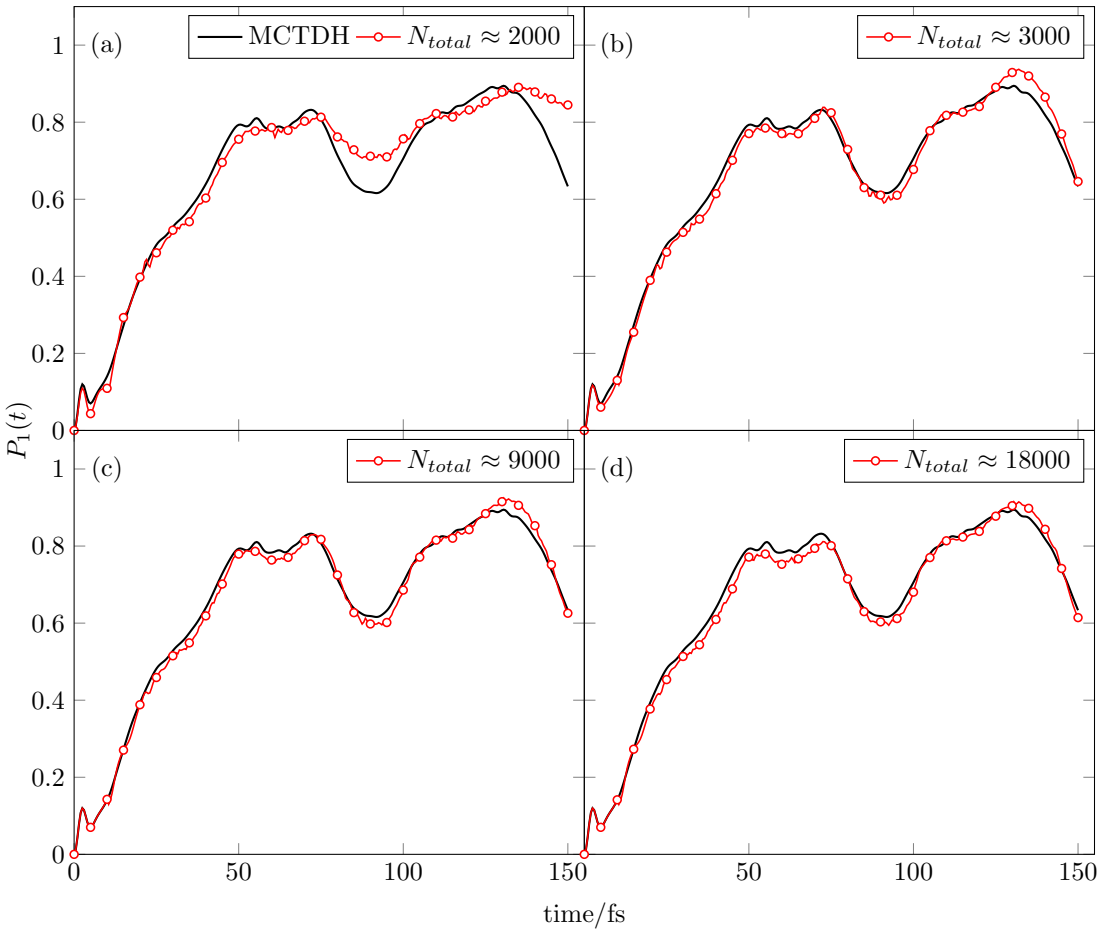

Figure 4.3: Populations of the lower S_1 excited state of pyrazine, resulting from adaptive sampling calculations of the 4-dimensional model Hamiltonian with varying basis set sizes, compared to exact MCTDH data.[65]

favourably, with an increase to only 3000 basis functions resulting in almost quantitative accuracy. This is a strong indication that this algorithm converges more quickly towards the finite basis set limit, shared by both these approaches. Further evidence of this is provided by the relatively small increase in accuracy, compared to $N_{total} \approx 3000$, associated with the significantly larger basis sets shown in Figure 4.3(c) and (d).

For a more direct comparison between the two methods, Figure 4.4 shows results for both the trajectory sampling and adaptive sampling method, at comparable total basis set sizes, N_{total}, with respect to MCTDH data. Here the advantages of the relatively frequent minimisation and optimisation of the basis set describing the wavefunction are even more apparent, as adaptive sampling results consistently outperforms the corresponding trajectory sampled data. In order to quantify this increase in accuracy, the MAE and MAPE with respect to MCTDH of both sets of data were calculated at each basis set size, using Eqs. 3.38, the

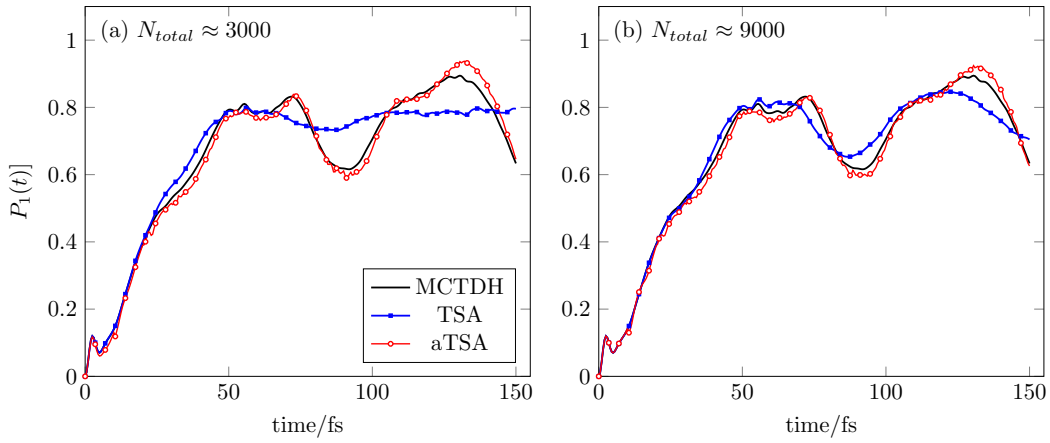

Figure 4.4: Comparison of the accuracy in the populations of the S_1 state of pyrazine for the 4-dimensional Hamiltonian for the trajectory sampling and adaptive sampling algorithm, with respect to MCTDH reference results.[65]

results being shown in Table 4.2. Again this data supports the observations made above, as both errors are significantly lower using the adaptive sampling algorithm over the trajectory sampling method.

Close inspection of Figure 4.3 reveals a phenomenon that warrants further discussion. Careful examination reveals small discontinuities in the population which, in this case, occur every 15 fs. This feature is closely linked with the MP minimisation and optimisation algorithm. As the former results in a minimal basis set which only reproduces the wavefunction to within the user set parameter ζ, which in this case was chosen to be 0.95, the populations resulting from this minimal set will differ from those due to the full set. This essentially yields two populations at the same point in time. One results from the full basis set at the end of the previous burst and thus is due to the more accurate description of the wavefunction. Conversely, the population at the start of the next iteration is due to the MP minimised and optimised basis set, which is inherently less accurate in

Table 4.2: Mean absolute and mean absolute percentage errors for TSA and aTSA calculations of the 4D vibronic pyrazine Hamiltonian at varying total basis set sizes with respect to exact MCTDH results.[65]

		MAE		MAPE / %	
m	N_{total}	TSA	aTSA	TSA	aTSA
300	≈ 3000	0.045	0.019	6.3	3.5
900	≈ 9000	0.025	0.015	3.4	2.4

describing the wavefunction. At the time of the transition between iterations of the algorithm, the former is thus chosen as the reported value.

As mentioned in Section 4.2.1, after completion of the MP algorithm, the expansion coefficients resulting from the latter are retained, while the trajectory sampled basis functions are initially assigned a value of zero. This, while reflecting the fact that, at the time of the restart, the wavefunction should be described only by the basis set resulting from MP, introduces a spurious shift away from the exact population, which often takes a number of timesteps of wavefunction propagation to correct itself. A good exampled of this occurring can be found in Figure 4.3(a), at 20 fs, which corresponds to the transition between the second and third iteration of the algorithm. Here a visible discontinuity occurs, resulting in a sharp peak in the otherwise smooth shape of the population in this area, however the general shape of and trend in $P_1(t)$ is recovered after approximately 1-2 fs.

This aspect of the adaptive sampling algorithm is unavoidable, as the basis set resulting from MP will always describe the wavefunction with limited accuracy. By increasing the convergence parameter of the minimisation and optimisation stage, ζ, the magnitude of this discontinuity may however be decreased, and naturally it is also affected by the number of iterations of the algorithm, N_b, and thus the number of MP iterations necessary. The effects of both these parameters are investigated in more detail in Section 4.5.

4.3.2 24-dimensional results

As already outlined in Section 3.3.7, the full 24-dimensional version of the vibronic pyrazine Hamiltonian constitutes an extremely difficult benchmark problem for quantum dynamics methods. The increased number of degrees of freedom significantly limited the accuracy of the trajectory sampling algorithm. Nevertheless, using a relatively larger basis set, the dynamics of the system could be reproduced qualitatively with respect to MCTDH.[65]

As the number of dimensions increases, so does the computational expense of all quantum dynamics methods, however the adaptive sampling algorithm, presented here, is expected to be impacted more significantly than most. The reasons for this are twofold but both relate to the MP minimisation and optimisation step.

Firstly, while, as shown above, drastically improving the performance of the algorithm for any given basis set size, this stage involves repeated calculation of overlap matrices over parts of, or indeed the entire basis set. A number of these calculations occur during every iteration of the MP algorithm. The higher-dimensionality of the 24D Hamiltonian was found to, in general, require significantly more basis functions to achieve qualitative accuracy in Section 3.3.7. Thus,

every iteration of the MP algorithm involves repeated calculation of inherently large matrices, the elements of which are high-dimensional integrals over the GWP basis functions. The computational cost of this is non-trivial and, in the case of the 24-dimensional Hamiltonian, the MP minimisation and optimisation algorithm was indeed found to be the computational bottleneck.

The second reason, while closely related to the first and, at its core, also relating to repeated matrix element evaluation, concerns the optimisation algorithm. As alluded to in Section 4.2.2, the specific method chosen, to optimise GWP positions and momenta with respect to the overlap of the function with the total wavefunction, was steepest descent. As, opposed to more advanced algorithms, such as the conjugate gradient method, this does not involve finding optimal parameter directions, in which to optimise the function in question, but rather considers each degree of freedom and within that each parameter in turn. In addition to directly scaling with respect to the number of degrees of freedom, this again results in each iteration involving calculation of many integrals over the basis functions.

Although the performance of the adaptive sampling algorithm is expected to, for these reasons, be significantly limited by computational cost in the case of the full 24-dimensional Hamiltonian, a successful calculation was run; however the input parameters had to be adjusted for optimal computational performance. A total of $N_b = 15$ iterations of the adaptive sampling algorithm, each consisting of $n_t = n_p = 100$ timesteps, lasting $\Delta t_t = \Delta t_p = 0.1$ fs, were carried out. The sampling frequency was $1/n_s = 1/500$ and the MP minimisation and optimisation convergence criterion was $\zeta = 0.75$. The latter was necessary in order to deal with the computational bottleneck associated with the MP algorithm, as discussed above, although it is recognised that this, in comparison to the 4D calculations, lower value of ζ will limit the accuracy of the aTSA. Finally the number of sampling trajectories was $m = 1000$. These input parameters and the resulting average basis set sizes are summarised in Table 4.3.

Figure 4.5 shows a comparison of the TSA and aTSA results with an average of $N_{total} = 9000$ basis functions. As discussed above, the lower MP minimisation and optimisation convergence criterion that had to be employed in order to counteract

Table 4.3: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of the 24D pyrazine Hamiltonian, the results of which are shown in Figure 4.5.

N_b	n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ	m	N_{total}	N_t	N_{MP}
15	100	0.1	10	100	0.1	0.75	1000	9030	8774	256

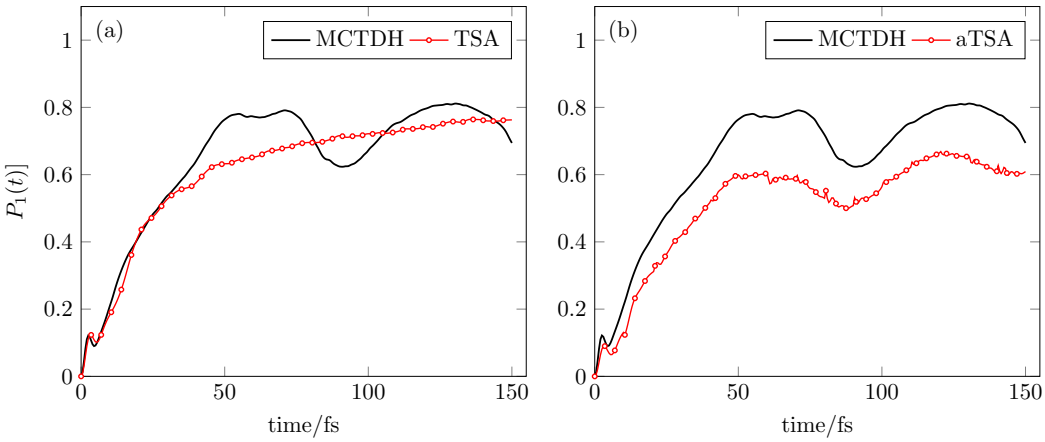

Figure 4.5: Comparison of the accuracy in the populations of the S_1 state of pyrazine for the 24-dimensional Hamiltonian for the (a) trajectory sampling and (b) adaptive sampling algorithm, with respect to MCTDH reference results.[65]

the computational bottleneck of the MP algorithm will limit the performance of the aTSA results. Nevertheless, comparing the TSA and aTSA S_1 populations, the benefits of the aTSA are still relatively clear. While in the case of the TSA, after approximately 40 fs, the population diverges from the exact MCTDH result,[65] the aTSA qualitatively reproduces the shape of the population, although the extent of relaxation to S_1 from S_2 is underestimated. The key features of the population dynamics are captured by the aTSA, as the reabsorbance features around 5 fs and 75 fs are clearly observed.

Given the computational challenge associated with the MP minimisation and optimisation algorithm, especially for systems with many degrees of freedom, these results are considered very encouraging. It should also be noted that adapting existing MP algorithm to either allow restarts *via* checkpoint files stored on disk or to make use of more highly parallel computing environments would significantly reduce the bottleneck introduced during this stage of the aTSA.

4.4 Tunnelling benchmark

As mentioned in Section 4.1, the short nature of the sampling trajectories in the case of the adaptive sampling algorithm should minimise the error due to classical-quantum divergence at longer times. Thus, this method should be able to treat systems which are subject to strong quantum effects much more effectively than the trajectory sampling method, presented in Chapter 3. To investigate this hypothesis, a synthetic benchmark system, modelling the tunnelling of a quantum particle through a barrier, was chosen as the second benchmark for the adaptive

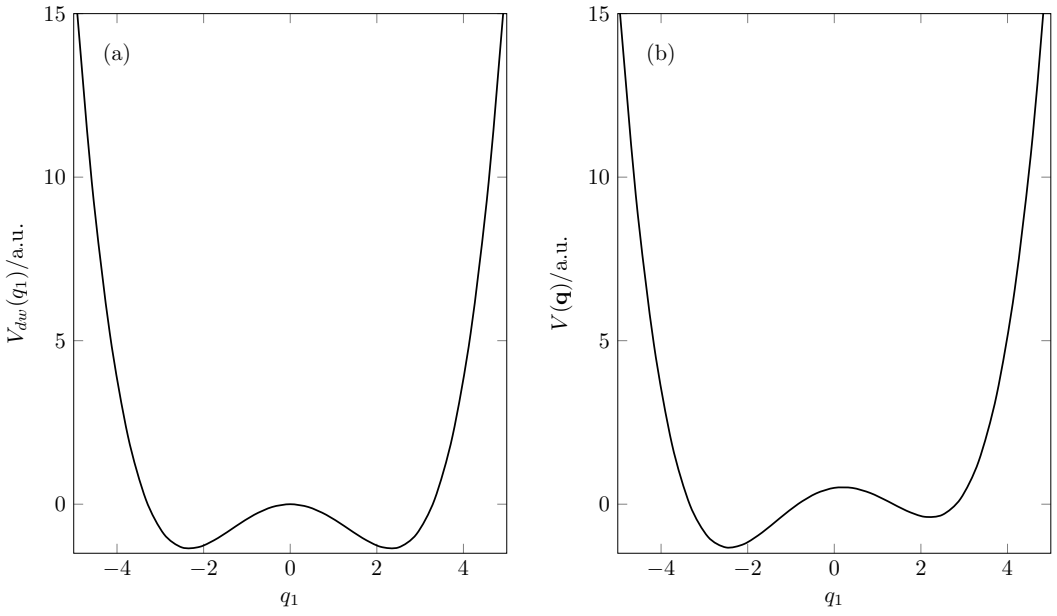

Figure 4.6: (a) Potential energy, $V_{dw}(q_1)$, as a function of the tunnelling coordinate of the double well benchmark and (b) $V(\mathbf{q})$ for Model I, with $q_2 = 1$ in order to demonstrate the asymmetric effect of the harmonic bath.

sampling method. One further reason the performance of this sampling strategy for tunnelling problems is the fact that most other trajectory-based quantum dynamics methods are not capable of capturing tunnelling effects.[42,43,45,46,56,118]

This model, which has been studied using other quantum dynamics approaches in the past,[47,82,119] consists of a one-dimensional energy barrier, through which the particle tunnels, coupled to a number of harmonic oscillators, representing the environment, such that

$$V(\mathbf{q}) = V_{dw}(q_1) + V_{env}(\mathbf{q})\,. \tag{4.2}$$

The potential energy barrier in the tunnelling coordinate is a double well, given by

$$V_{dw}(q_1) = \frac{1}{16\eta} q_1^4 - \frac{1}{2} q_1^2\,. \tag{4.3}$$

where $\eta = 1.3544$, following previous investigations of this model.[47,82,119] The shape of this potential is shown in Figure 4.6(a). Coupling to the remaining DoFs introduces an asymmetry, raising the energy of the tunnelling target well, as shown in Figure 4.6(b). While inherently flexible with respect to the number of degrees of freedom that make up the harmonic oscillator environment, the nature of the coupling between the former and tunnelling coordinate can be either linear or quadratic.

The initial conditions of this system are chosen in such a way that the energy expectation value of the initial wavefunction, ψ_0 is lower than the height of the

barrier in the tunnelling coordinate. Thus, tunnelling due to classical mechanics will be negligible, instead being driven almost entirely by the quantum nature of the system. This wavefunction is a single, normalised GWP, given by

$$\psi_0(\mathbf{q}) = \prod_{j=1}^{f} \left[\left(\frac{\alpha}{\pi\hbar} \right)^{\frac{1}{4}} \exp\left[-\frac{\alpha}{\hbar} \left(q_j - q_j^{\dagger} \right) \right] \right] , \tag{4.4}$$

where $q_1^{\dagger} = -2.5$, $q_{j>1}^{\dagger} = 0$ and $\alpha = 0.5$. Furthermore, the mass of each degree of freedom was set to $m_j = 1$.

In order to quantify the amount of tunnelling as function of time, the tunnelling autocorrelation function, $C_t(t)$, is calculated as

$$C_t(t) = \left\langle \psi_0'(\mathbf{q}) \,\middle|\, \psi(\mathbf{q}, t) \right\rangle , \tag{4.5}$$

where $\psi_0'(\mathbf{q})$ is the mirror image of the initial wavefunction, given by Eq. 4.4 with values of $q_1^{\dagger} = 2.5$ and $q_{j>1}^{\dagger} = 0$. The value of this function can be interpreted as a direct measure of the extent to which the wavefunction has tunnelled through the barrier at any given moment in time.

As this model does not feature multiple electronic states, instead of the Ehrenfest trajectories employed for the vibronic pyrazine Hamiltonian, classical MD trajectories are used. While the Velocity Verlet scheme is retained, the equations of motion are now

$$\frac{\partial q_\kappa}{\partial t} = \frac{p_\kappa}{m_\kappa} \tag{4.6}$$

$$\frac{\partial p_\kappa}{\partial t} = -\frac{\partial V}{\partial q_\kappa} , \tag{4.7}$$

where V is potential as described by Eq. 4.2. The sampling trajectories for this benchmark are thus purely classical and therefore inherently incapable of describing a strong quantum event such as tunnelling. The reason the adaptive sampling algorithm should nevertheless correctly describe quantum dynamics, given that a sound basis set was sampled, lies with the propagation stage described in detail in Section 3.2.2. Although the basis set in this case will be sampled purely classically, the propagation of the wavefunction, expanded in the former, occurs *via* formally exact, variational equations of motion, derived from the TDSE. The performance of the algorithm for this model will thus give a good indication to what extent, in the absence of long-time classical-quantum deviation, regions of phase space, relevant to wavefunction propagation, may be accessed using classical trajectories.

4.4.1 Linear coupling

The first version of this model investigated involves linear coupling between pairs of modes, such that the environmental potential energy is given by

$$V_{env}(\mathbf{q}) = \sum_{j=2}^{f} \left[\frac{1}{2} q_j^2 + a q_{j-1} q_j \right] , \tag{4.8}$$

where, in accordance with previous work on this problem, $a = 0.2$.[47,82,118,119] Thus, the nature of the coupling is such that the tunnelling coordinate is only directly coupled to one degree of freedom ($j = 2$), while the rest are only indirectly coupled and the strength of this indirect interaction decreases with increasing j. Two separate versions of this benchmark are considered here. The simplest case, where $f = 2$ and there are no indirectly coupled degrees of freedom, is termed Model I from here on. Whilst rather uncomplicated and not fully representative of tunnelling in a real system, the simplicity associated with this low-dimensional model does grant unobstructed insight into the ability of the adaptive sampling algorithm to treat quantum tunnelling effects. Secondly, in order to more closely represent a system which exhibits linear-like coupling of the degrees of freedom, the case of $f = 5$, referred to as Model II, is also investigated.

Calculations for Model I employed sampling and propagation timesteps of $\Delta t_t = 0.002$ a.u. and $\Delta t_p = 0.002$ a.u. A set of $N_b = 12$ iterations of the adaptive sampling algorithm were employed, each consisting of $n_t = 10^4$ timesteps of trajectory sampling and $n_p = 10^4$ timesteps of wavefunction propagation, resulting in a total timespan of 240 a.u. of dynamics calculated. A set of 4 calculations was run, employing a MP convergence factor of $\zeta = 0.99$ and sampling frequency of $1/n_s = 1/500$. The number of sampling trajectories was varied, taking values of $m = 8$, 12, 16 and 20. The resulting average (over all 12 bursts) total, trajectory sampled and inherited basis set sizes, N_{total}, N_t and N_{MP} respectively, are shown, along with critical input parameters, in Table 4.4.

Table 4.4: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of Model I, the results of which are shown in Figure 4.7.

N_b	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	ζ	m	N_{total}	N_t	N_{MP}
12	10000	0.002	500	10000	0.002	0.99	8	171	160	11
12	10000	0.002	500	10000	0.002	0.99	12	252	241	11
12	10000	0.002	500	10000	0.002	0.99	16	331	319	12
12	10000	0.002	500	10000	0.002	0.99	20	418	406	12

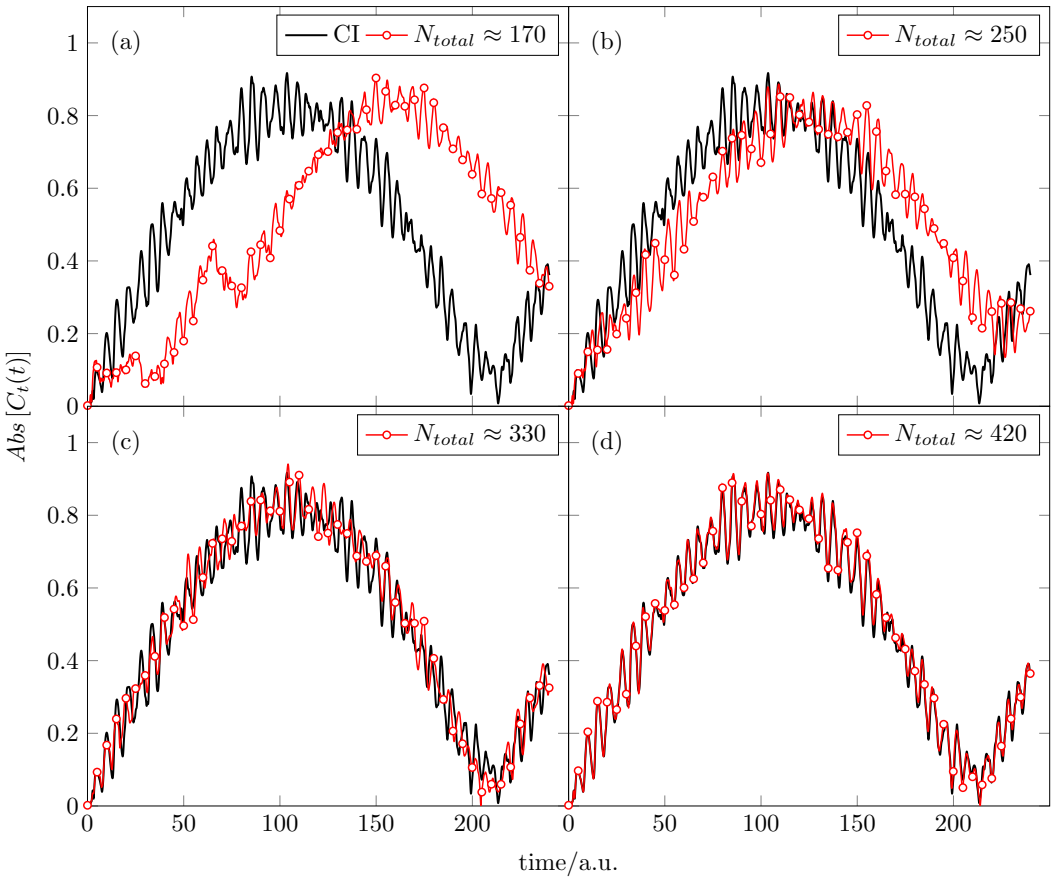

Figure 4.7: Tunnelling autocorrelation functions for Model I, calculated using the adaptive sampling algorithm with varying average total basis set sizes, N_{total}, compared to exact CI results.[119]

Figure 4.7 shows the resulting tunnelling autocorrelation functions, compared to numerically exact configuration interaction (CI) results.[119] For the low-dimensional Model I, the adaptive sampling algorithm performs excellently. All but the smallest basis set, $N_{total} \approx 170$ shown in Figure 4.7(a), clearly capture the qualitative nature of the tunnelling dynamics. Increasing to $N_{total} \approx 250$, while improving on the overall description of the tunnelling, still fails to capture the longer time dynamics. The two largest basis sets employed however, excellently reproduce the behaviour of the wavefunction, correctly showing it having tunnelled almost completely through the barrier after around 100 a.u. and returning to its original position after a further 110 a.u.

Only $N_{total} \approx 420$ basis functions, as shown in Figure 4.7(d), are needed to converge to essentially exact dynamics with respect to the CI reference data.[119] Interestingly, the number of GWPs required to describe the wavefunction to within an accuracy of $\zeta = 0.99$ remains effectively constant for all 4 calculations. This speaks, for one, to the relative simplicity of the wavefunction, however also in-

dicates that the MP minimisation and optimisation algorithm for this particular problem readily converges, allowing for the small total basis set sizes shown in Table 4.4. Overall, the performance of the adaptive sampling algorithm for Model I is very encouraging, as it clearly demonstrates that the purely classical sampling trajectories employed, are capable of generating a basis set which can accurately describe the, inherently quantum, tunnelling dynamics of the wavefunction.

To investigate to what extent this excellent performance is retained as the number of DOFs increases, a set of 4 calculations was run for Model II. In this case, the number of iterations of the adaptive sampling algorithm was $N_b = 24$, with the same sampling and propagation timestep duration of $\Delta t_t = \Delta t_p = 0.002$ a.u., each iteration consisting of $n_t = n_p = 5000$ steps of sampling and propagation. Each such iteration thus consisted of 10 a.u. of sampling and propagation, leading, as above, to a total simulation time of 240 a.u. A sampling frequency of $1/n_s = 1/500$ was used, while the MP convergence factor was $\zeta = 0.99$. Again the number of sampling trajectories was varied as $m = 100$, 200, 400 and 800. The average basis set sizes as well as important input parameters for these calculations are shown in Table 4.5.

Figure 4.8 shows the tunnelling autocorrelation functions, $Abs\left[C_t(t)\right]$ for these 4 calculations, compared to exact CI results.[119] The increased number of degrees of freedom clearly results in an increase of the minimum number of basis functions, required to accurately describe the wavefunction. The calculation shown in Figure 4.8(a), while still comparatively small, captures the essential dynamics of the tunnelling, however both the timescale and the oscillatory fine structure of the autocorrelation function are not accurately reproduced. Increasing the average number of basis functions to $N_{total} \approx 2000$, as shown in Figure 4.8(b), does significantly improve the description of the tunnelling, as the timescale of the latter is now qualitatively reproduced. Notably, conversely to the CI reference,[119] the autocorrelation function does not quite return to $Abs\left[C_t(t)\right] = 0$ around 210 a.u.,

Table 4.5: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of Model II, the results of which are shown in Figure 4.8.

N_b	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	ζ	m	N_{total}	N_t	N_{MP}
24	5000	0.002	500	5000	0.002	0.99	100	1026	995	31
24	5000	0.002	500	5000	0.002	0.99	200	2025	1989	36
24	5000	0.002	500	5000	0.002	0.99	400	4003	3974	29
24	5000	0.002	500	5000	0.002	0.99	800	8016	7990	26

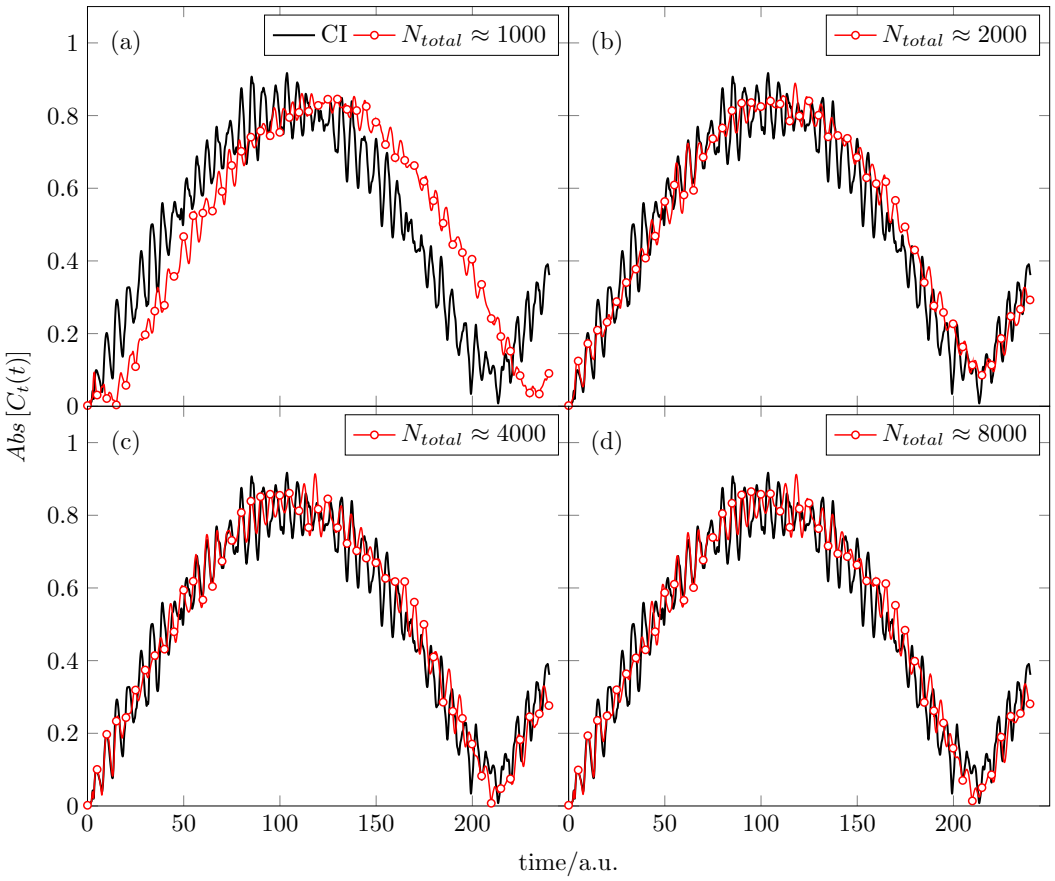

Figure 4.8: Tunnelling autocorrelation functions for Model II, calculated using the adaptive sampling algorithm with varying average total basis set sizes, N_{total}, with respect to an exact CI benchmark.[119]

indicating that there are still limitations to the accuracy of the aTSA at this basis set size.

Further increase of the latter to $N_{total} \approx 4000$, shown in Figure 4.8(c), improves on the qualitative accuracy of the tunnelling dynamics, the wavefunction now fully returning to its original configuration around 210 a.u., although the oscillatory structure is still not captured accurately. Increasing the basis set beyond this point appears to not significantly improve results, as is evident from Figure 4.8(d), thus suggesting that the adaptive sampling algorithm has converged for this particular problem. Again, similarly to Model II, the number of basis functions selected by the MP minimisation and optimisation algorithm to facilitate a smooth transition between iterations of the aTSA, N_{MP} remains relatively constant for varying total basis set sizes, N_{total}.

While exact dynamics, with respect to the CI reference, were not obtained, the tunnelling behaviour was captured with qualitative accuracy, using a still rather small basis set. Given the relatively low computational cost associated with basis

sets smaller than $N_{total} \approx 10000$ and the purely classical sampling trajectories, this is considered rather encouraging.

4.4.2 Quadratic coupling

The second type of coupling considered for the tunnelling benchmark outlined above involves quadratic coupling between the tunnelling coordinate and every environmental oscillator. In this case the potential energy due to the environment is given by

$$V_{env}(\mathbf{q}) = \sum_{j=2}^{f} \left[\frac{1}{2} q_j^2 + a q_1 q_j^2 \right] , \tag{4.9}$$

where, again following previous investigations of this system, $a = 0.05$.[47,82,118,119] From here on referred to as Model III, this constitutes a far more challenging problem than either Model I or II. Now the tunnelling coordinate is directly coupled to every other degree of freedom, the coupling itself is quadratic instead of linear and the strength of the latter is constant across all $f - 1$ coordinates. In order to demonstrate the applicability of the adaptive sampling algorithm beyond systems with only a few degrees of freedom, already represented by Model I and II, the number of DOFs for Model III was chosen to be $f = 20$. While still a synthetic benchmark, Model III is far more representative of real higher-dimensional systems involving quantum tunnelling and consequently a significant challenge for any quantum dynamics method.

As a result of this, a common approach,[118] when sampling the basis set for solving this problem, is to sample basis functions not only from the initial wavefunction, centred around $q_1 = -2.5$ and $q_{j>1} = 0$, but also from the mirrored wavefunction, representing complete tunnelling to the other well, centred around $q_1 = -2.5$ and $q_{j>1} = 0$. This, while resulting in excellent dynamics,[118] somewhat reduces the sampling challenge, as the basis set generated by this approach inherently supports tunnelling, independent of whether the sampling strategy captures it or not. However, Model III still constitutes a challenging benchmark with respect to the accuracy with which a given method captures the dynamics of the wavefunction, irrespective of the basis set's inclusion of tunnelling effects.

To assess both the performance of the aTSA for this benchmark, as well as comparing to the TSA, two calculations, one for each approach, were run. Both calculations were run for a total sampling time of 120 a.u. and employed sampling and propagation timesteps of $\Delta t_t = \Delta_p = 0.002$ a.u., as well as $m = 50$ sampling trajectories. The TSA calculation consisted of $n_t = n_p = 60000$ timesteps of sampling and propagation and a sampling frequency of $1/n_s = 1/4000$. For the aTSA calculation, on the other hand, $N_b = 60$ sampling and propagation iterations of

$n_t = n_p = 1000$ timesteps were run, storing basis functions with a frequency of $1/n_s = 1/200$ and employing a MP minimisation and optimisation convergence criterion of $\zeta = 0.995$. These input parameters are summarised in Table 4.6. Finally, it is worth noting that, in accordance with the sampling strategy set out above and with previous investigations of this benchmark,[118] the initial burst of sampling involved selecting basis functions from both wells in the potential energy.

As with the two models of linear coupling, the tunnelling autocorrelation function, $C_t(t)$, was calculated for both methods, the results of which are shown in Figure 4.9. Comparing again to exact CI data,[119] the challenging nature of Model III is immediately apparent. The TSA completely fails to capture the tunnelling dynamics of the system beyond about 10 a.u., as too large a proportion of the wavefunction tunnels to the opposite well, furthermore failing to return to its original configuration, instead remaining effectively trapped. The latter is indicated by the decreasing amplitudes of the oscillations in $C_t(t)$ in Figure 4.9(b).

The aTSA, on the other hand, does not exhibit either of the two shortcomings above, capturing at the very least, the broad qualitative features of $C_t(t)$. While by no means following the CI results[119] exactly, the extent to which the wavefunction tunnels is reproduced relatively well, with the exception of the last 25 a.u. of time. Given the relatively small basis sets employed for both calculations, as $N_{total} = 994$ and 884, for the TSA and aTSA respectively (trajectory sampled and MP inherited basis set sizes are shown in Table 4.6), this is extremely encouraging.

The underlying conditions of Model III, that is the high number of DOFs and the complex, quadratic coupling of the tunnelling coordinate to all others, makes this a relatively realistic benchmark for systems exhibiting strong quantum tunnelling. The adaptive sampling method, while, to an extent, addressing the classical-quantum divergence, negatively affecting the TSA, still relies on sampling trajectories driven by purely classical mechanics. Thus, even the limited qualitative accuracy achieved here indicates that sufficiently frequent resampling of the

Table 4.6: Input parameters and basis set sizes for calculations of Model III, using the TSA and the aTSA, the results of which are shown in Figure 4.9.

	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	m	N_{total}
TSA	60000	0.002	4000	60000	0.002	50	994
aTSA	1000	0.002	200	1000	0.002	50	884
	N_b	ζ	N_t	N_{MP}			
aTSA	60	0.995	759	125			

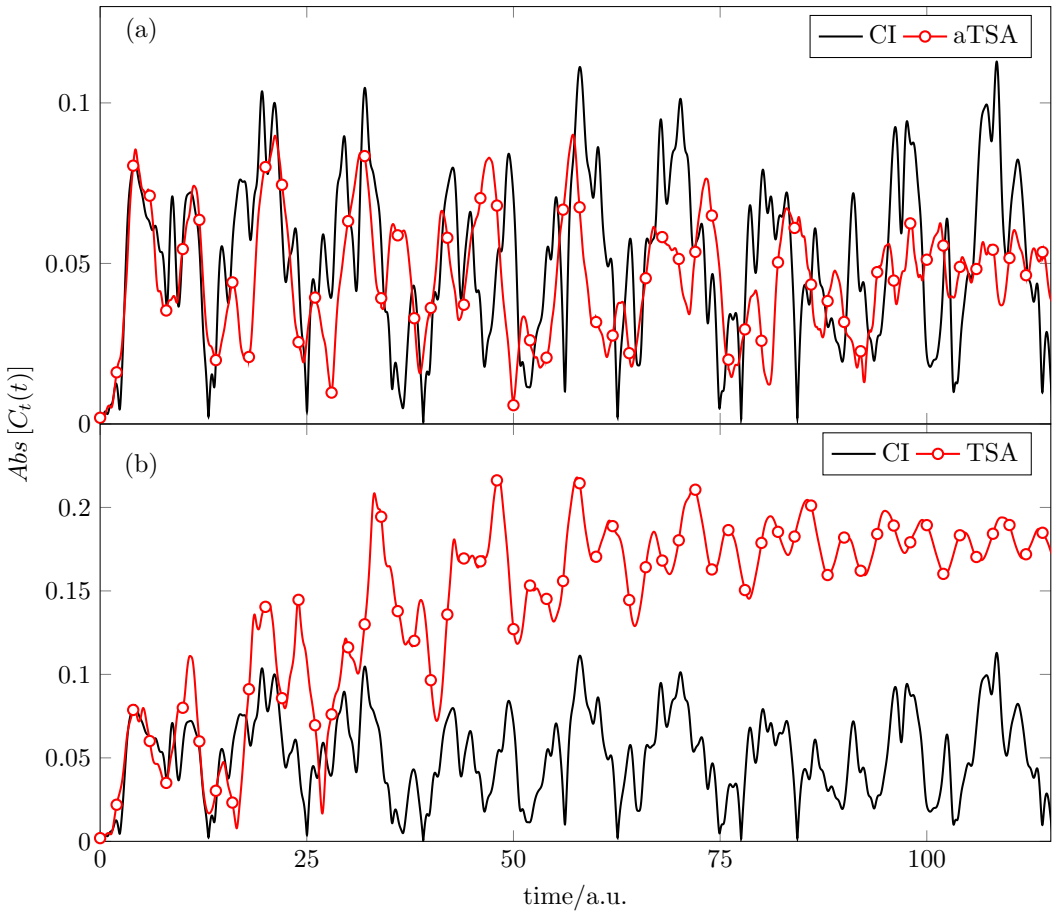

Figure 4.9: Tunnelling autocorrelation function for Model III, calculated using (a) the adaptive sampling algorithm and (b) with the trajectory sampling algorithm, both with respect to an exact CI benchmark.[119]

wavefunction makes strong quantum effects accessible with such relatively simple and computationally inexpensive tools.

4.5 Algorithm parameters and performance

The approach taken in this section follows very closely that used in Section 3.4. The 4-D pyrazine benchmark was again used here, to allow comparison to the results in the previous chapter.

4.5.1 Basis set size consistency

One of the key advantages of the adaptive sampling algorithm is that due to the regular resampling of the basis set using trajectories, the overall number of basis functions required to describe the wavefunction to a given accuracy is much lower than in the case of the trajectory sampling method. The MP minimisation

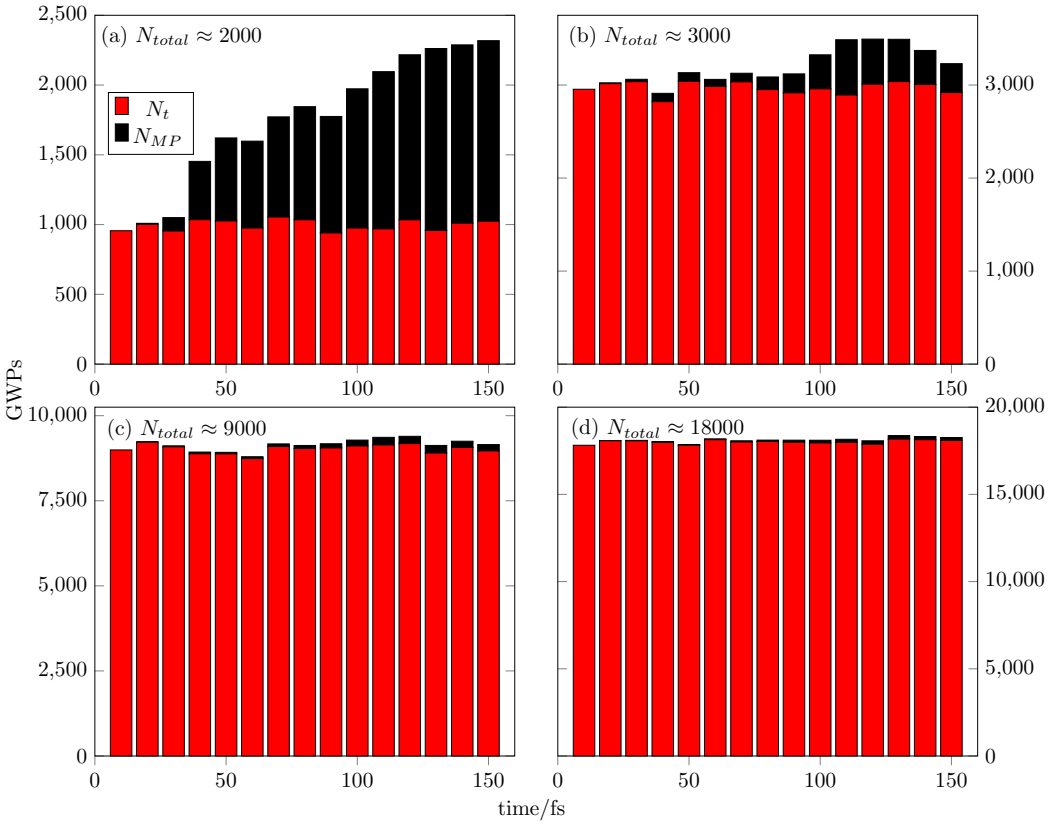

Figure 4.10: Comparison of the numbers of trajectory sampled GWP basis functions and those inherited from the MP minimisation and optimisation algorithm, for the adaptive sampling calculations of the 4D pyrazine Hamiltonian, presented in Section 4.3 inputs for which are shown in Table 4.1 and results in Figure 4.3 and Table 4.2.

and optimisation algorithm, while allowing for seamless transitions between these iterations, does also have the potential to systematically grow the basis set, which would, in the long-time limit, counteract the aforementioned benefit.

More specifically, if the number of inherited basis functions was to grow at every restart, resulting in a steadily increasing total basis set size, this method would clearly be of limited use, especially concerning systems where a large number of iterations are required. On the other hand, if the number of basis functions required to describe the wavefunction at the transition between bursts remains relatively constant, there are no inherent limitations to the number of iterations of the adaptive sampling algorithm that may be used, ignoring computational demands.

To investigate which is the case, the results for the calculations shown in Figure 4.3 and Table 4.2 can be investigated further. Figure 4.10 shows the constitution of the basis set at the end of each iteration of the MP algorithm for 4 different

total basis set sizes. In the case of the three larger basis sets, the number of inherited basis functions does remain relatively constant over all 15 iterations and, more importantly, only constitutes a small percentage of the total basis set. Furthermore, it is interesting to note that the iterations between which the number of MP inherited basis functions did fluctuate, albeit not significantly, corresponding to areas of complex dynamics, when comparing to the populations shown in Figure 4.3. This indicates that when there is significant flux of the wavefunction between the two electronic states, the number of basis functions required to describe the former increases, which is unsurprising, given that each function, while multi-dimensional, is only associated with one of the two diabatic states.

The case of the smallest basis set, corresponding to the calculation employing only $m = 100$ sampling trajectories, does not follow this general trend. Here the number of basis functions inherited during each iteration of the MP minimisation and optimisation algorithm does significantly contribute to the overall basis set size, making up as much as 56% of all basis functions. Furthermore, the number of functions inherited steadily increases for each iteration, thus making this calculation increasingly expensive. Although only 15 iterations of the algorithm were employed in this case, the total basis set at the end of the final sampling stage is more than twice the size of that sampled initially. Clearly, this calculation would not be sustainable in the long-time limit, as the basis set would continue to grow, increasing computational costs.

This non-conformance of the smallest calculation to the otherwise observed trend is however considered rather encouraging, as it serves as a system of easy self-diagnosis. Thus, if calculating dynamics without reference data to compare to, should the number of inherited basis functions continue to grow in a similar manner, it is likely that the number of sampling trajectories employed is insufficient to properly describe the dynamics of the wavefunction. Given that the number of trajectories to use, as well as the balance between this and the sampling frequency, $1/n_s$, which is discussed in more detail below, is not likely to be known *a priori* for most systems, a diagnostic tool of this kind is guaranteed to be useful when applying the trajectory sampling algorithm to a wider variety of problems.

4.5.2 Scaling and MP convergence

Given the increased performance of the aTSA over the TSA observed for the 4D pyrazine Hamiltonian benchmark in Section 4.3, the scaling of this new method with respect to total basis set size, N_{total}, warrants further investigation. In addition, possibly the most critical parameter of the aTSA is the convergence criterion of the MP minimisation and optimisation algorithm, ζ, as it directly controls

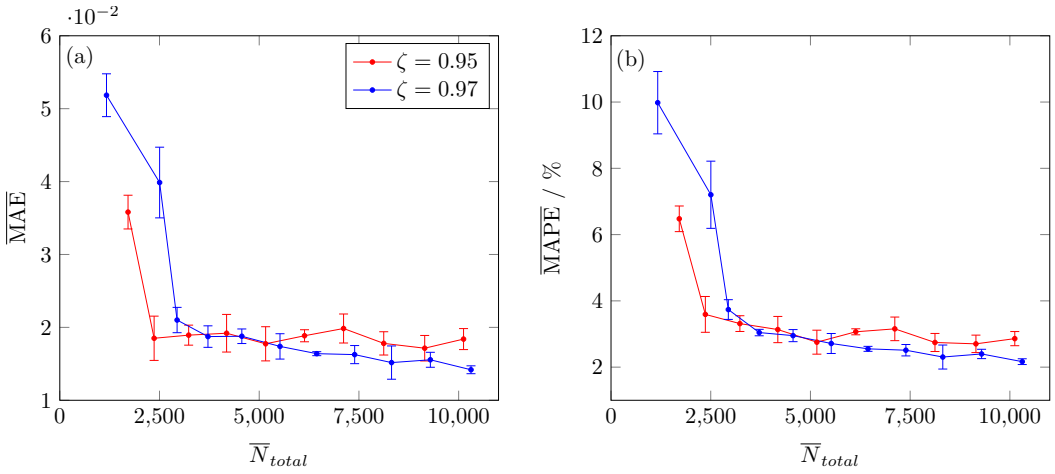

Figure 4.11: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian from adaptive sampling calculations, employing basis sets of varying size and (b) corresponding mean absolute percentage errors, both for MP convergence criteria of $\zeta = 0.95$ and $\zeta = 0.97$.

the accuracy with which the wavefunction is retained between two iterations of sampling and propagation.

To investigate the effects of N_{total} and ζ, two sets of 4×10 calculations were run, ζ taking values of 0.95 and 0.97. All calculations employed $N_b = 15$ iterations of the aTSA, which consisted of $n_t = n_p = 100$ timesteps of trajectory sampling and propagation, each timestep lasting $\Delta t_t = \Delta t_p = 0.1$ fs, thus resulting in a total simulation time of 150 fs. The basis function sampling frequency was $1/n_s = 1/10$ and the number of sampling trajectories was varied from $m = 100$ to $m = 1000$ in steps of 100. Due to the increased number of basis functions inherited with $\zeta - 0.97$ one additional set of calculations was run with $m = 50$. These input parameters are summarised in Table 4.7.

Following the approach used throughout Section 3.4, the average MAE and MAPE error, as well as corresponding standard deviations, were calculated using Eqs. 3.38 and 3.39, all with respect to exact MCTDH results.[65] Again, the total basis set size, averaged across sets of four calculations with identical input

Table 4.7: Input parameters for adaptive sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figures 4.11 and 4.12.

N_b	n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ
15	100	0.1	10	100	0.1	0.95
15	100	0.1	10	100	0.1	0.97

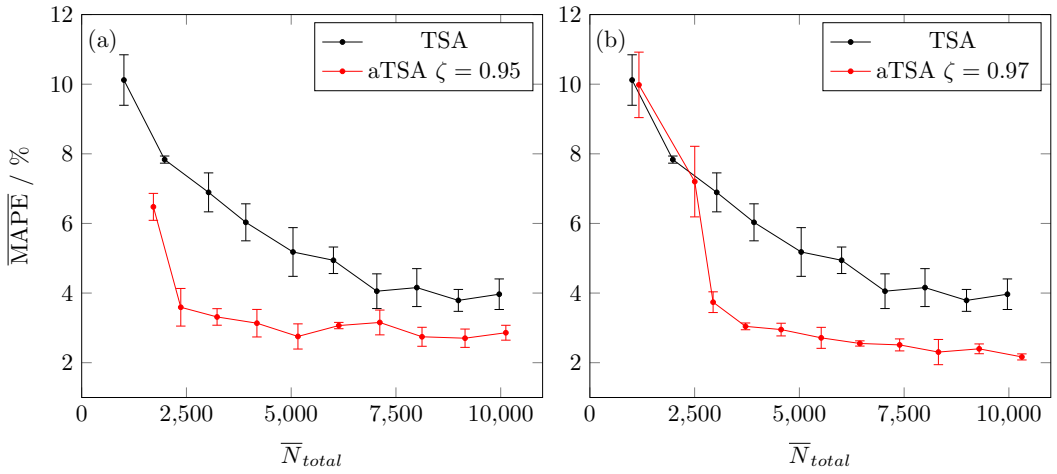

Figure 4.12: Mean absolute percentage errors in the P_1 population of the 4D pyrazine Hamiltonian, comparing the performance of the aTSA and the TSA for MP minimisation and optimisation convergence criteria of (a) $\zeta = 0.97$ and (b) $\zeta = 0.97$.

parameters, was calculated as $\overline{N}_{total} = \frac{1}{4}\sum N_{total}$. It is important to note that this average is different from what has been and continues to be referred to as the average total basis set size, N_{total}, where the latter refers to only one aTSA calculation, the average being taken over the N_b iterations of the algorithm.

Figure 4.11 shows these errors as a function of $\overline{N}_{total}$ for both $\zeta = 0.95$ and $\zeta = 0.97$. Furthermore, Figure 4.12 compares the average MAPE of both sets of calculations to TSA results with comparable total basis set sizes, $\overline{N}_{total}$, previously shown in Figure 3.9. The values of the average MAE, average MAPE, standard deviation and basis set sizes are also shown in Tables A.5 and A.6.

The increased accuracy of the aTSA over the TSA at a given basis set size, previously observed in Section 4.3, is clearly evident from Figure 4.12. Both sets of calculations clearly outperform the TSA results by a significant margin. Even more encouragingly, the rate at which the TSA appears to converge for this problem is much higher, the average MAPE sharply dropping around $\overline{N}_{total} \approx 2500$, as opposed to the more gradual decrease, observed for the TSA. Overall this is highly encouraging, as it quantitatively confirms that the aTSA constitutes a consistent performance and accuracy increase over the TSA for a variety of total basis set sizes.

In addition to this, the MP minimisation and optimisation convergence criterion, ζ, appears to have a significant impact on the performance of the aTSA. Close inspection of Figure 4.11 reveals that the accuracy of the dynamics obtained does not drastically improve as ζ is increased, the rate of convergence of

the aTSA on the other hand does seemingly depend on the choice of ζ. More specifically, while in the case of $\zeta = 0.95$, the algorithm effectively converges around $\overline{N}_{total} \approx 3000$, further increase in the total basis set size not translating into noticeably higher accuracy, with a higher convergence criterion of $\zeta = 0.97$, the average MAE and MAPE continue to decrease, albeit not as rapidly as initially, suggesting that the algorithm has not yet fully converged.

Given that certain systems, such as the tunnelling dynamics discussed in Section 4.4, are bound to be more sensitive to the accuracy with which the wavefunction description is retained between iterations of the aTSA, this is rather encouraging. As the ζ parameter that controls this aspect of the algorithm is increased, the maximum accuracy theoretically achievable will, as shown above, also increase. The aTSA can thus, within limits imposed by available computational resources, be adapted to a variety of systems and its theoretical performance be tuned through the MP minimisation and optimisation convergence criterion.

4.5.3 Resampling frequency

The MP minimisation and optimisation stage of the adaptive sampling algorithm is mainly controlled by two parameters. The number of iterations of the algorithm that are employed, N_b, and the convergence criterion, ζ, which controls the accuracy with which the wavefunction must be reproduced by the minimised basis set, and is discussed in more detail below.

For a given system the most likely parameters to be known *a priori* are the total time for which dynamics need to be calculated and the timestep of wavefunction propagation, which is often linked to the duration of the shortest action, likely to influence the dynamics in any significant way, that the system is capable of. The timestep duration for the sampling trajectories may be chosen to be equal to the latter, but is not limited to this value, as has already been discussed, with respect to the TSA, in Section 3.4.3 and is also investigated for the aTSA below. Once the total number of timesteps to sample and propagate for has been determined, the number of iterations of the adaptive sampling algorithm determined the duration of each individual burst, which can also be though of as the frequency with which the basis set is resampled.

Given that the performance increase due to these resampling events stems from the constraining of the sampling trajectories to the short-time limit, minimising errors due to classical-quantum divergence, the aTSA is expected to be relatively sensitive with respect to the choice of this frequency. Clearly, too small a value of N_b, corresponding to too infrequent resampling, will allow for the aforementioned errors due to divergence to again impact the dynamics. Conversely,

choosing a value which is too large is likely to result in disproportionate computational expense, given that, as outlined above, the MP minimisation and optimisation algorithm often acts as the bottleneck limiting performance of the algorithm. Furthermore, as each iteration of the MP algorithm introduces a systematic error, due to the approximate nature of the inherited basis set, it stands to reason that a large number of bursts will result in the introduction of a significant error, assuming ζ remains unchanged.

In order to investigate the effect of the resampling frequency, a set of 4×6 calculations was run, varying the number of iterations of the aTSA, N_b and thus the number of sampling and propagation timesteps, n_t and n_p respectively. In order to maintain an approximately constant total basis set size of $N_{total} \approx 3000$, the number of sampling trajectories, m, was varied as well. Timestep durations were $\Delta t_t = \Delta t_p = 0.1$ fs, while the basis function storage frequency was $1/n_s = 1/10$ and the MP minimisation and optimisation convergence criterion was $\zeta = 0.95$. These input parameters are summarised in Table 4.8.

Following the approach used above, the average MAE and MAPE, along with corresponding standard deviations were calculated using Eqs. 3.38 and 3.39, all with respect to exact MCTDH results.[65] Figure 4.13 shows these errors as a function of the number of algorithm iterations N_b, the data and corresponding average basis set sizes also being shown in Table A.7. The trend predicted above is clearly reflected by this data, as both a large number of short sampling and propagation bursts as well as very few long ones result in a clear loss of accuracy. Using between $N_b = 30$ and $N_b = 15$ bursts appears to strike the balance between resampling the wavefunction often enough to prevent large errors due to classical-quantum divergence, while at the same time avoiding the significant systematic error due to the finite accuracy of the MP minimisation and optimisation step.

In order to address the computational expense of the MP algorithm, alluded to above, a second balance may be considered. The total basis set size, N_{total} drastically affects the rate at which classical-quantum divergence degrades the accuracy of the TSA, as is discussed in detail in Chapter 3. Thus, if the goal is to

Table 4.8: Input parameters and average total basis set sizes for adaptive sampling calculations of the 4D vibronic pyrazine Hamiltonian, the results of which are shown in Figures 4.13 and 4.14.

Δt_t/fs	n_s	Δt_p/fs	ζ	N_{total}
0.1	10	0.1	0.99	≈ 3000
0.1	10	0.1	0.99	≈ 9000

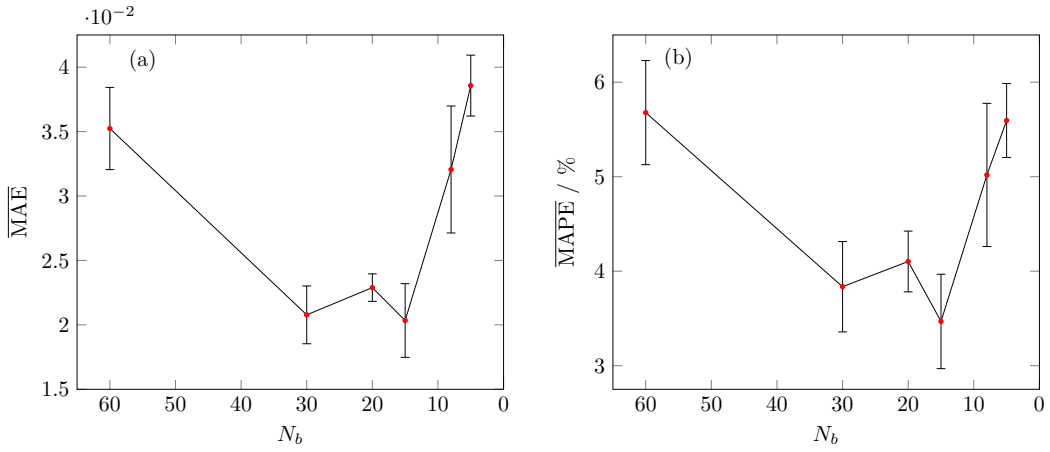

Figure 4.13: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian using varying the number of iterations of the aTSA and (b) corresponding mean absolute percentage errors, with an average total basis set size of $N_{total} \approx 3000$.

minimise the computational cost of the MP minimisation and optimisation step, using a larger total basis set size, N_{total}, should result in more accurate single burst performance, thus requiring fewer restarts of the aTSA algorithm.

To investigate this, a second set of 4×6 calculations was run, using identical parameters to those employed above, with the exception of the number of sampling trajectories, which was scaled, so as to result in an average total basis set size of $N_{total} \approx 9000$. Again, average MAE and MAPE, as well as corresponding standard deviations were calculated, the results being shown in Figure 4.14 and Table A.8.

On inspection of these errors it becomes immediately clear that the hypothesis laid out above does indeed hold. In the case of the larger basis sets employed here, the number of iterations of the aTSA algorithm required to accurately reproduce the dynamics is rather low. To illustrate this, the average MAPE errors shown in Figure 4.14 may be compared to those for TSA calculations in Figure 3.9. Using a total of $N_{total} \approx 9000$ basis functions yielded an average MAPE of approximately 4% for the TSA, while only 5 instances of wavefunction resampling using MP reduced this error to about 2% for the aTSA.

The calculation employing the fewest, $N_b = 4$, bursts did result in increased errors, thus suggesting that the balance between classical-quantum divergence and MP systematic errors still plays a role for larger basis set sizes, albeit a less significant one. The total basis set size, N_{total}, may thus, within reason, be leveraged against the number of MP minimisation and optimisation steps required, in order to maximise computational performance. Given that the MP step constitutes the computational bottleneck for most calculations, as discussed in Section 4.3, this is

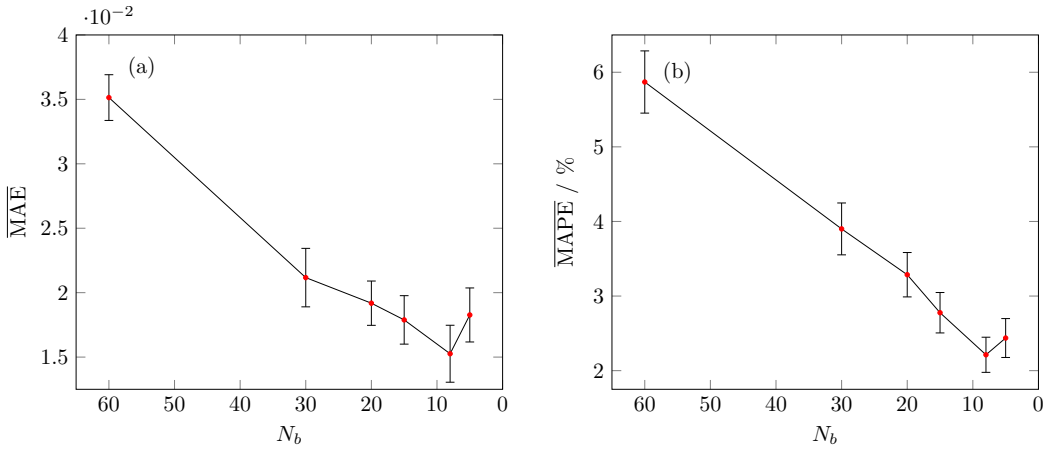

Figure 4.14: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian using varying the number of iterations of the aTSA and (b) corresponding mean absolute percentage errors, with an average total basis set size of $N_{total} \approx 9000$.

highly encouraging as it adds some flexibility to the set up of aTSA calculations.

4.5.4 "Oversampling"

The concepts of "over-" and "undersampling" have previously been introduced in Section 3.4.3. These terms, referring to the cases of $n_t > n_p$ and $n_t < n_p$ respectively, may also be extended to the adaptive sampling algorithm. "Undersampling" has previously been found to result in a significant decrease in dynamics accuracy for the TSA and is thus omitted here. "Oversampling" however, as it has been linked to an increase in accuracy in Section 3.4.3, albeit not a drastic one, is considered for the aTSA as well.

Given that the sampling in the case of the aTSA is limited to much shorter times, the effects of "oversampling" may be more significant, especially considering the MP minimisation and optimisation algorithm, which naturally greatly benefits from a more diverse basis function library. Aside from the classical-quantum divergence inevitably experienced by the sampling trajectories, it stands to reason that while certain quantum-relevant areas of phase space may be accessible to classical mechanics, the time taken for a trajectory, driven by the latter, to reach these may be significantly longer than under exact quantum propagation. Therefore, extending the duration of the sampling stage beyond that of wavefunction propagation, may result in a basis set, which more accurately reflects the distribution of the wavefunction in phase space, at the end of the TDSE solution.

In order to investigate the extent to which "oversampling" can improve the accuracy of the adaptive sampling algorithm, a set of 4×12 calculations of the

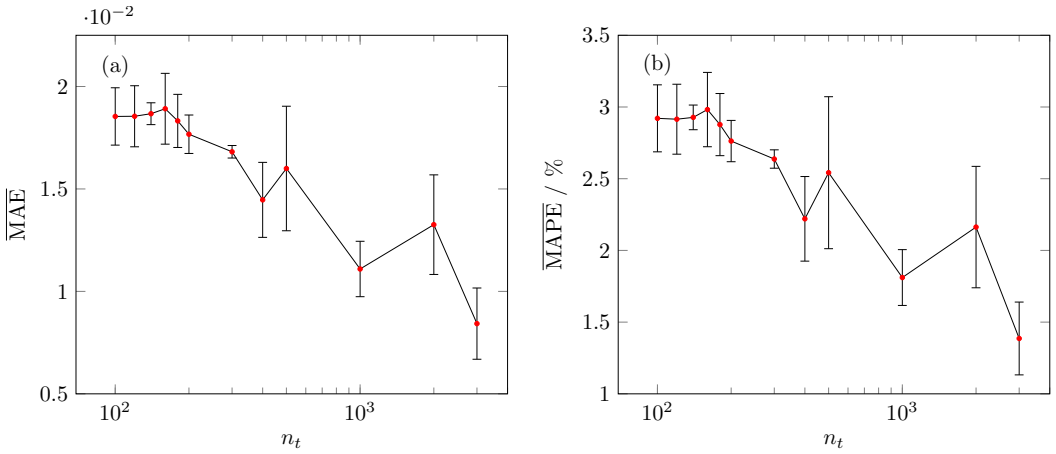

Figure 4.15: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian varying the number of sampling timesteps, n_t, resulting in "oversampling" and (b) corresponding mean absolute percentage errors, with an average total basis set size of $N_{total} \approx 9000$.

4-dimensional pyrazine Hamiltonian benchmark were run, the latter being chosen for the reasons already discussed above. A total of N_b iterations of the aTSA were run for each calculation, made up of $n_p = 100$ timesteps of wavefunction propagation, each lasting $\Delta t_p = 0.1$ fs. The duration of the sampling timesteps was identical, $\Delta t_t = 0.1$ fs, while the number of such steps, n_t, and the number of sampling trajectories employed, m, were varied, in order to maintain an average total basis set size of $N_{total} \approx 9000$. Finally, the basis functions sampling frequency was $1/n_s = 1/500$ and the MP minimisation and optimisation convergence criterion was $\zeta = 0.99$. These input parameters are also summarised in Table 4.9

Following the approach above, the average MAE and MAPE, along with corresponding standard deviations, were calculated using Eqs. 3.38 and 3.39, the results being shown in Figure 4.15 and Table A.9. As suggested above, for the adaptive sampling algorithm, "oversampling" can result in a significant increase in accuracy. Focusing on Figure 4.15(b), using $n_t = 3000$ sampling timesteps in comparison to $n_p = 100$ propagation steps, resulted in a MAPE comparable to the

Table 4.9: Input parameters and average total basis set sizes for adaptive sampling calculations of the 4D vibronic pyrazine Hamiltonian, the results of which are shown in Figure 4.15.

N_b	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ	N_{total}
15	0.1	500	100	0.1	0.99	≈ 9000

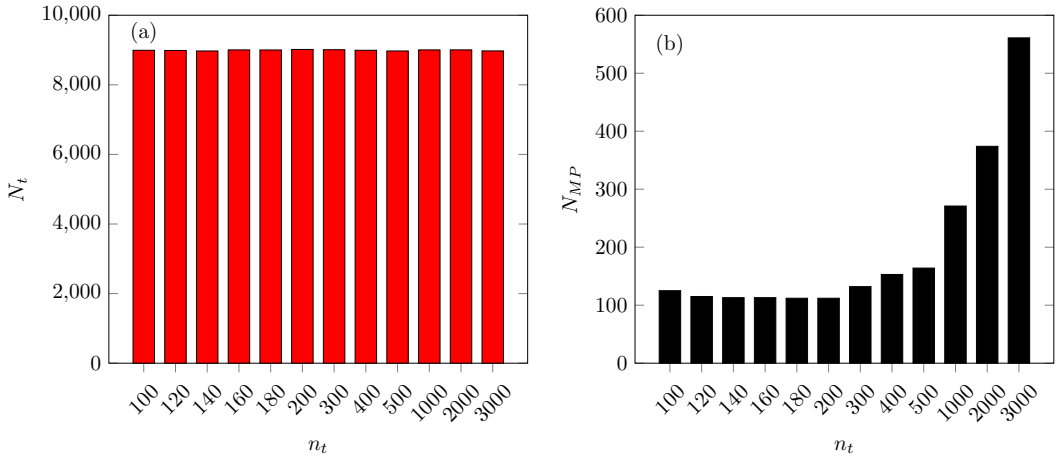

Figure 4.16: Average basis set constitution, comparing average numbers of (a) trajectory sampled, N_t, and (b) inherited basis functions, N_{MP}, for adaptive sampling calculations of the 4D pyrazine Hamiltonian, varying the number of sampling timesteps, the results of which are shown in Figure 4.15.

---

best results obtained using the TSA, shown in Figure 3.9(b). Notably however, the latter involved a total basis set size of $N_{total} \approx 24000$ GWPs, while the average number of functions employed by this specific implementation of the aTSA was only $N_{total} \approx 9000$. Given the relatively low computational cost of the sampling trajectories, "oversampling" could thus constitute a cheap way to increase the performance of any given aTSA calculation, barring any negative impact on other parts of the algorithm.

This caveat warrants further investigation when considering the effect that "oversampling" inevitably has on the distribution of the trajectory sampled basis set in phase space. More specifically, as the duration of the sampling trajectories starts increasing far beyond the duration of wavefunction propagation, $n_t >> n_p$, the resulting basis set will be more dispersed in phase space. As a result, after propagation, the MP minimisation and optimisation algorithm will have fewer functions, which overlap significantly with the wavefunction at the end of propagation, thus increasing the number of iterations required. As already discussed above, for most problems, the computational cost associated with the MP algorithm far outweighs that of any other part of the aTSA, thus constituting the rate-limiting bottleneck. It stands to reason that the increase in accuracy gained by excessive "oversampling" will be overshadowed by unfavourable computational cost, due to the MP algorithm.

To investigate this further, Figure 4.16 shows the average number of trajectory sampled and MP inherited basis functions, N_t and N_{MP} respectively, for each of the "oversampling" cases shown in Figure 4.15. N_t remains relatively constant, which

is to be expected, given that the number of sampling trajectories, m, was scaled so as to roughly maintain a constant basis set size across all calculations. The number of basis functions resulting from the MP minimisation and optimisation algorithm, N_{MP}, and therefore the number of computationally expensive iterations of the latter, does however increase significantly with the extent of oversampling.

While this was not found to drastically limit performance in this lower-dimensional version of the pyrazine benchmark, it is expected that for more complex systems, such as the full 24D Hamiltonian, the approximately 400% increase in N_{MP} for $n_t = 3000$ would result in infeasible computational cost. Therefore, as suggested above, while "oversampling" is certainly more beneficial for the aTSA than the TSA, the increase in accuracy resulting from it comes at the cost of rising computational expense.

4.5.5 Basis function library expansion

As outlined in Section 4.2.1, before the MP minimisation and optimisation algorithm is iterated, the basis set at the end of the previous burst is expanded using a parameter, γ. This generates a library of basis functions, based on and still containing the aforementioned set, but providing a more diverse selection of basis functions to the MP algorithm. This is associated with little additional computational cost, as the functions added to the library are all assigned zero valued coefficients and thus, while adding to the overall size of the basis set handled by the MP algorithm, do not contribute to the overlap matrices, the calculation and inversion of which constitutes the computational bottleneck of this stage.

Nevertheless, the effect of this parameter on the accuracy of the adaptive sampling algorithm is worth investigating, especially as it does not significantly increase computational costs, but has the potential to drastically speed up the most expensive part of the majority of calculations, namely the MP minimisation and optimisation step.

To investigate this, a set of 4×9 calculations of the 4D pyrazine Hamiltonian were run, with values of γ, varying from $\gamma = 1.0$ to $\gamma = 5.0$ in steps of 0.5. Each

Table 4.10: Input parameters and average total basis set sizes for adaptive sampling calculations of the 4D vibronic pyrazine Hamiltonian, the results of which are shown in Figures 4.17 and 4.18.

N_b	n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ	m	N_{total}
15	100	0.1	10	100	0.1	0.95	300	≈ 3000

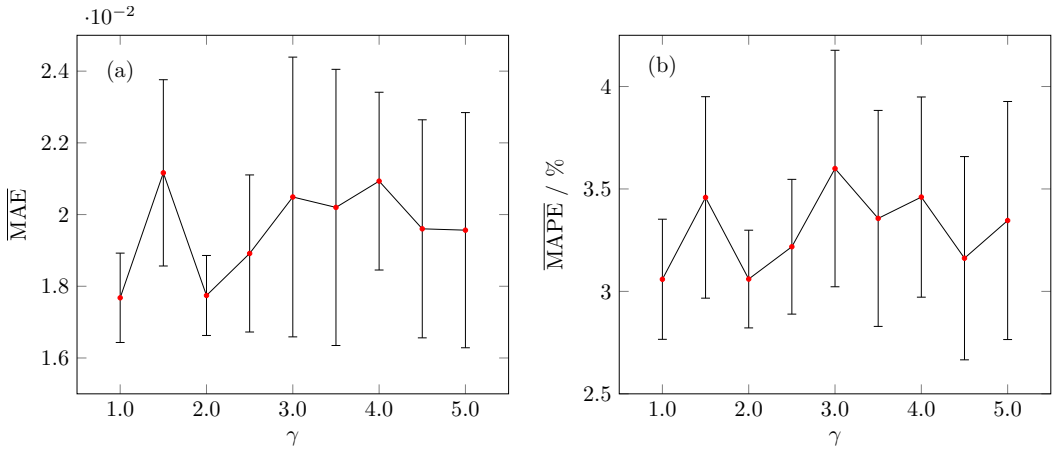

Figure 4.17: (a) Mean absolute error in the P_1 population of the 4D pyrazine Hamiltonian varying the basis library expansion factor, γ, of the MP minimisation and optimisation algorithm and (b) corresponding mean absolute percentage errors, with an average total basis set size of $N_{total} \approx 3000$.

calculation consisted of $N_b = 15$ iterations of the aTSA, made up of $m = 300$ trajectories with $n_t = n_p = 100$ timesteps of sampling and propagation, each lasting $\Delta t_t = \Delta t_p = 0.1$ fs. The basis function sampling frequency was $1/n_s = 1/10$ and a MP minimisation and optimisation convergence criterion of $\zeta = 0.95$. These input parameters are summarised in Table 4.10.

In accordance with the approach taken throughout this section, the average MAE and MAPE were calculated using Eqs. 3.38 and 3.39, the results of which are shown in Figure 4.17 and Table A.10. Inspecting these errors, it becomes immediately clear that the basis library expansion parameter does not significantly affect the accuracy of the adaptive sampling algorithm. Furthermore, the rather large standard deviations suggest that changing γ simply increases the variation in the accuracy between calculations. This, naturally, is not ideal, as it complicates judging the accuracy of any given calculation.

To investigate the effect of this parameter on the MP minimisation and optimisation algorithm, for the calculations above, basis set compositions, specifically the average number of trajectory sampling and MP inherited basis functions, N_t and N_{MP} respectively, were calculated.

Figure 4.18 shows these compositions. Inspecting, especially the number of basis functions yielded by the MP algorithm, N_{MP}, shown in Figure 4.17(b), it becomes clear that there is merit in changing the basis library expansion factor, as values of $\gamma > 1.0$ result in fewer basis functions required to describe the wavefunction between bursts. As outlined above, the MP algorithm is computationally rather expensive and often constitutes the computational bottleneck of a given

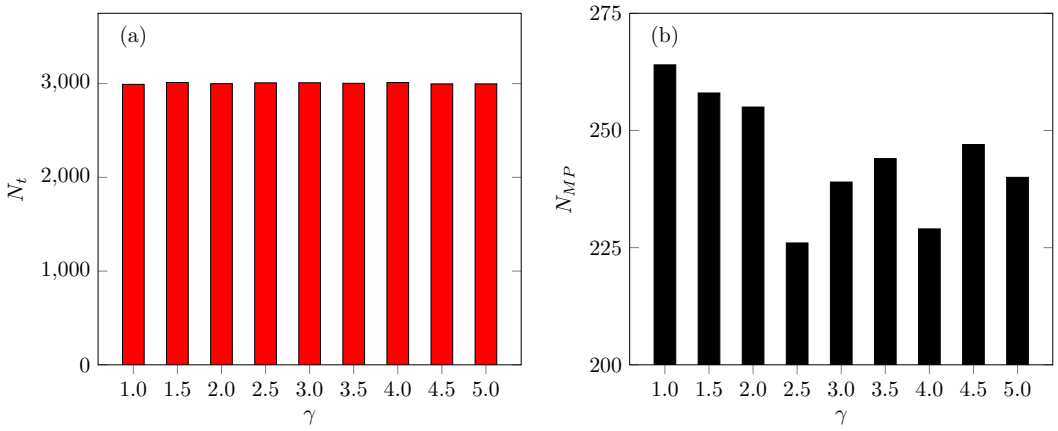

Figure 4.18: Average basis set constitution, comparing average numbers of (a) trajectory sampled, N_t, and (b) inherited basis functions, N_{MP}, for adaptive sampling calculations of the 4D pyrazine Hamiltonian, varying the basis library expansion parameter, γ, the results of which are shown in Figure 4.17.

calculation. Thus, as changing γ does not significantly degrade the accuracy of the aTSA, the latter may be sped up slightly by selecting a value of $\gamma > 1.0$. Given however that drastically increasing this parameter will require storage of very large matrices, values such as $\gamma = 3.0$, which was chosen for the calculations presented throughout this work, are likely to yield the best balance between speeding up the MP algorithm and limiting overall memory requirements.

Overall this parameter is considered rather useful, as, within reasonable limits, it allows the computational cost of the adaptive sampling algorithm to be distributed between the number of operations required per iteration and the necessary system memory.

4.5.6 Computational performance and cost

In order to quantitatively assess the computational cost associated with the adaptive sampling algorithm, a number of previously discussed calculations were analysed for total CPU time taken and maximum system memory required. All calculations were performed on single compute nodes, containing either two Intel Xeon E5-2630 v3 2.4 GHz 8-core processors, for a total of 16 cores per calculation, or two Intel Xeon E7-4809 v3 2.0 GHz 8-core processors, for a total of 16 cores (and 32 threads) per calculation. 16 core nodes were equipped with a maximum of 64GB of DDR4 RAM, or 4GB per core, while 32 thread nodes supported up to 1024GB of DDR4 RAM, or 32GB per core. The adaptive sampling algorithm is implemented using the FORTRAN programming language and accelerated using OpenMP and LAPACK.

In total, a set of 8 calculations was analysed, consisting of 4 different total basis set sizes, each of which was investigated for two values of the MP minimisation and optimisation convergence criterion. These calculations, or in some cases ones with identical input parameters, were already discussed in Section 4.5.2. The specific input parameters and average basis set sizes of the calculations investigated are summarised in Table 4.11.

The resulting total CPU time per calculation and maximum system memory used for each calculation are shown in Figure 4.19. The total computation time provides further evidence that the MP minimisation and optimisation algorithm does indeed constitute the computational bottleneck, even for this low-dimensional 4D version of the pyrazine Hamiltonian. Not only were the calculations resulting in only 1000 trajectory sampled basis functions longer than those resulting in 3000, but increasing the MP convergence criterion from $\zeta = 0.95$ to $\zeta = 0.97$, also resulted in consistent and significant increases in the total CPU time, for comparable basis set sizes. Memory requirements on the other hand seem to be less sensitive to the MP algorithm, only increasing slightly as ζ was altered. This is to be expected, as, while contributing to the number of large matrices that must be stored in memory, the MP step is by no means the only part of the algorithm requiring this.

Comparing the computational requirements of the aTSA to those of the TSA discussed in Chapter 3, while both the maximum memory usage and overall computation time were higher for the aTSA, the benefit from using the latter far outweighs these costs. Overall the the adaptive sampling algorithm performs well for the kind of systems investigated in this chapter. Both memory and computing time requirements are well within the bounds of what is commonly available in terms of computational infrastructure and feasible with respect to the general timescale of research.

4.6 Conclusions

This chapter has introduced a modification to the trajectory sampling method, presented in Chapter 3. The adaptive sampling algorithm recognises the limitations placed on the former by the deviation of the classical-like sampling trajectories, employed to place basis functions in phase space, from the exact quantum solution in the long time limit. By constraining the sampling to short timescales and splitting the desired calculation into a number of shorter iterations of sampling and wavefunction propagation, this challenge can however be overcome. In order to ensure smooth transition between these bursts, a minimisation and optimisation algorithm, based on the popular matching pursuit approach,[81–83] has been

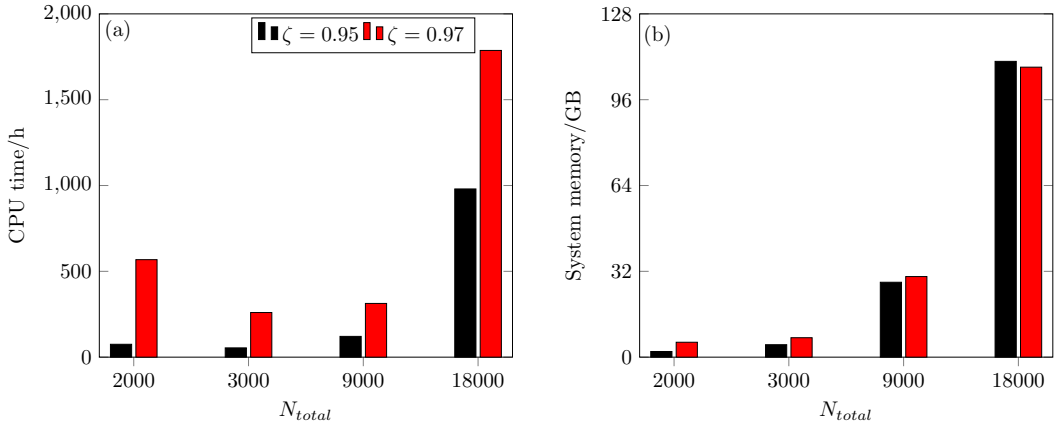

Figure 4.19: (a) Total CPU time and (b) maximum system memory required by adaptive sampling calculations of the 4D pyrazine Hamiltonian, the input parameters of which are shown in Table 4.11.

Table 4.11: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of the 4D pyrazine Hamiltonian, the results of which are shown in Figure 4.19.

N_b	n_t	Δt_t/fs	n_s	n_p	Δt_p/fs	ζ	m	N_{total}	N_t	N_{MP}
15	100	0.1	10	100	0.1	0.95	100	1646	993	653
15	100	0.1	10	100	0.1	0.95	300	3230	3005	225
15	100	0.1	10	100	0.1	0.95	900	9145	9030	115
15	100	0.1	10	100	0.1	0.95	1800	18136	18034	102
15	100	0.1	10	100	0.1	0.97	100	2615	996	1619
15	100	0.1	10	100	0.1	0.97	300	3691	2996	695
15	100	0.1	10	100	0.1	0.97	900	9295	8996	299
15	100	0.1	10	100	0.1	0.97	1800	18208	17979	229

presented. In order to investigate both the validity of the approximation, underlying this modified method, as well as quantify and compare its performance and accuracy to that of the trajectory sampling algorithm, two challenging quantum dynamics benchmarks were chosen.

First, the vibronic Hamiltonian modelling the relaxation dynamics of pyrazine, already extensively studied with the underlying method, was investigated. Comparing performance of the adaptive sampling algorithm for the 4-dimensional version of this model to the trajectory sampling method, the advantages gained from frequent minimisation and optimisation of the wavefunction became immediately evident. Then number of basis functions required to reach near quantitative accuracy with respect to exact MCTDH data was significantly lower, with as few as 3000 GWPs yielding extremely precise dynamics.

Extending this new modification to the full 24-dimensional version of the pyrazine Hamiltonian was found to significantly increase the computational demands of the adaptive sampling algorithm. It was discovered that the matching pursuit minimisation and optimisation stage is especially sensitive to the number of degrees of freedom, as it involves repeated calculation of large matrices, the elements of which are multi-dimensional integrals between pairs of basis functions across all system coordinates.

Although the matching pursuit minimisation and optimisation algorithm scales unfavourably with the number of degrees of freedom, encouragingly a calculation for the 24-dimensional pyrazine Hamiltonian revealed that even with the accuracy of said algorithm being severely limited by computational constraints, a clear improvement over the trajectory sampling algorithm was still achieved. More specifically, the population dynamics produced by the adaptive sampling algorithm qualitatively reproduced most features of the exact MCTDH results[65] using approximately 9000 basis functions, something could previously not be achieved with the trajectory sampling algorithm, even using as many as 24000 basis functions. Given the potential to further optimise and accelerate the matching pursuit algorithm, this is highly encouraging, as it suggests that this strategy should be applicable to systems even larger than the pyrazine benchmark.

The second benchmark was chosen to investigate the hypothesis that, as the adaptive sampling algorithm should reduce errors due to classical-quantum divergence of the sampling trajectories, this method can be used to calculated the dynamics of systems, subject to strong quantum effects. Thus, a synthetic double well benchmark, modelling the tunnelling of a quantum particle through a one-dimensional energy barrier, was treated with this new approach. Two types of coupling, linear and quadratic, as well as a varying number of degrees of freedom were investigated.

For the simplest, linear, 2-dimensional model, near quantitative accuracy could be achieved with a relatively small basis set. Increasing the number of dimensions to 5 without altering the nature of the coupling was found to increase the challenge somewhat, however qualitatively accurate dynamics could still be obtained at low computational cost. The final 20-dimensional model, featuring quadratic coupling to all environmental degrees of freedom, proved a significantly more difficult problem, however the adaptive sampling algorithm was able to capture the basic, qualitative nature of the dynamics of this system, which given the complexity of the model, is considered highly encouraging. Furthermore, the trajectory sampling algorithm completely fails to capture the tunnelling beyond the very short time limit, thus indicating, even using classically driven sampling trajectories, resampling the wavefunction with sufficient frequency can, at least qualitatively, describe strong quantum events such as tunnelling.

Investigating in more detail the various input parameters and overall scaling of the new adaptive sampling algorithm resulted in a number of interesting and favourable discoveries. Encouragingly the basis set size was found to be relatively constant between bursts, the MP algorithm only contributing a small fraction of basis functions every time. Additionally, an insufficient number of GWPs will result in this behaviour to break down, thus allowing for the easy identification of cases where the number of trajectory sampled basis functions must be increased.

The algorithm was found to consistently outperform the trajectory sampling approach, yielding significantly more accurate results at similar total basis set size. Increasing the convergence criterion of the MP algorithm was found to raise the theoretical maximum accuracy of the algorithm, allowing the rate of convergence to be tuned by the user. A balance between the number of resampling events and the accuracy of the algorithm was discovered, where both too frequent and too infrequent sampling degrades the performance.

The concept of "oversampling", introduced for the trajectory sampling algorithm, was found to be significantly more beneficial in the case of adaptive sampling. While increasing the performance of the algorithm, this benefit is, to an extent, leveraged against increased computational cost, as more iterations of the costly MP algorithm are required under strong "oversampling" conditions. The expansion step preceding this MP minimisation and optimisation step was found to, while not significantly affecting the accuracy of the algorithm, be useful in balancing the computational expense between the number of required operations and necessary system memory.

While the general performance increase over the trajectory sampling method resulting from the adaptive sampling algorithm clearly indicates that there is significant merit to limiting the sampling stage to short timescales, there are still a

number of limitations to this strategy. As discussed in Chapter 2 and mentioned in Section 4.1, the inability of classical trajectories to properly treat zero point energy can allow for regions of phase space being visited which are inherently inaccessible to the corresponding quantum trajectories. Furthermore, as investigation of the tunnelling benchmark in Section 4.4 has made abundantly clear, certain quantum effects, such as tunnelling, are simply not replicated by purely classical dynamics, which limits the effectiveness of the basis sets sampled using the latter. As this is a well known fact however, there exist a number of semi-classical approaches to overcoming this limitation some of which may quite easily be adapted as driving mechanics for the sampling trajectories.

One such method, which may be particularly suitable, given its similarity to classical mechanics, is the path integral molecular dynamics (PIMD) framework, based on Feynman path integral theory. Trajectories driven by this approach should, in theory, sample phase space in a way that more accurately follows the exact quantum path. This may be especially relevant in the case of strong quantum tunnelling effects, such as is the case for the benchmark presented in Section 4.4, which is why Chapter 5 is dedicated to the implementation of PIMD trajectories as part of the adaptive sampling algorithm.

Chapter 5

Path Integral Sampling Trajectories

This chapter investigates the use of path integral molecular dynamics as the driving trajectories for the adaptive sampling algorithm. After introducing the underlying theory of such trajectories, the algorithm responsible for implementing them in the framework of the adaptive sampling method is presented. In order to investigate whether they can provide an improvement over classical sampling trajectories, the double well tunnelling benchmark from the previous chapter is investigated.

5.1 Introduction

The adaptive sampling algorithm was shown to constitute a significant improvement over the TSA in Chapter 4. Comparing results for the vibronic pyrazine benchmark, the aTSA required only a fraction of the basis functions needed by the TSA to reach quantitative accuracy with respect to exact MCTDH results.[65] In order to demonstrate that this novel method is not limited in its application to non-adiabatic systems, a benchmark system featuring strong quantum tunnelling was also investigated.

While the aTSA results for this synthetic double well benchmark were encouraging, they also highlighted a key limitation of the aTSA as used in Chapter 4, specifically employing classical trajectories. For the lowest, 2-dimensional version of the double well Hamiltonian, featuring linear coupling of the tunnelling coordinate to the environment, quantitative accuracy with respect to exact CI results[119] could be achieved using the aTSA. However as the number of degrees of freedom was increased to 5, the performance of the algorithm degraded somewhat, the highest accuracy achievable only capturing qualitative dynamics. This loss of accuracy was even more noticeable in the case of the 20-dimensional, quadratically coupled version of the double well Hamiltonian, where the aTSA, while still significantly outperforming the TSA, altogether failed to capture some of the tunnelling features of the system, only retaining rough quantitative accuracy with respect to exact CI results.[119]

The origin of this degradation in performance as the complexity of the system increases can be explained when considering the nature of the sampling trajectories employed. As discussed at length in Chapter 4, the classical trajectories driving the sampling are rather ill-suited for modelling strong quantum effects, such as the tunnelling in the double well benchmark. The reasons for this are manifold, however all stem from the inherent difference between classical and quantum mechanics.

One such limitation of classical trajectories is the possible leakage of the zero point energy from one degree of freedom into another.[116,117] This, in addition to violating the law of conservation of ZPE, implicit in the TDSE, allows for regions of phase space to be sampled which violate the uncertainty principle. Placing basis functions in such areas, which is formally allowed and, depending on the number and length of sampling trajectories, even likely when using classical mechanics as the driver of sampling, has the potential to, on top of introducing an essentially superfluous basis function, partially corrupt the basis set, as its presence allows for nonphysical deformations of the wavefunction.

A further issue with using classical sampling trajectories, albeit related to the

one discussed above, concerns the inherent difference in phase space paths taken by classical and quantum mechanics and the differing rate at which phase space may be sampled. Certain areas of space may be inaccessible to classical trajectories, but readily visited by quantum trajectories. A good example of such a case are systems, like the double well benchmark, where two local minima in the PES are separated by a energy barrier, the height of which exceeds the energy of the wavefunction in one of these minima. The other minimum, while formally inaccessible to classical trajectories, is often readily tunnelled to by quantum trajectories. Even in cases where the classical trajectory has sufficient energy to cross the barrier, that energy may result in the trajectory failing to remain in the other well, returning over the barrier to its original position.

Even in cases where the phase space region of interest is formally accessible by classical mechanics, for the example above this would constitute an energy barrier lower than the energy of the wavefunction in one minimum configuration, the rate at which this area may be sampled by a classical trajectory may be significantly lower than the corresponding quantum rate. Thus, the difference in paths taken and rate of phase space sampling between classical and quantum trajectories is likely to also negatively impact the effectiveness of the former as a choice of driving force for the aTSA.

The immediately obvious alternative of propagating the sampling trajectories using exact quantum derived equations of motion would reduce the aTSA method to one employing time-dependent basis functions. This, on top of rendering the sampling stage obsolete, as dynamics may just as well be collected during the propagation of the basis functions, would be associated with a number of issues, which have, in the context of time-dependent quantum dynamics methods, been discussed in Chapters 1 and 2.

The most desirable solution to this issue thus constitutes replacing the classical sampling trajectories of the aTSA with a semi-classical alternative, which, while capturing quantum effects such as tunnelling, avoids the issues commonly associated with time-dependent quantum dynamics methods. While a number of choices for such a substitution exist, the approach chosen here was to employ path integral molecular dynamics (PIMD),[120,121] amongst other reasons for its relative algorithmic similarity to purely classical molecular dynamics.

The remainder of this chapter focuses on first introducing the concept of PIMD, by putting it in the context of path integrals, from which it derives. The algorithm implementing this new type of trajectories in the aTSA is then outlined, followed by application to the tunnelling double well benchmark discussed in detail in Chapter 4.

5.2 Path integral sampling trajectories

5.2.1 Path integral Hamiltonian

The fundamental underlying principle of semi-classical methodologies lies in, by some means, averaging over a set of classically evolving trajectories in order to approximate quantum dynamical information. The essence of the path integral approach lies in this average being taken over a number of paths, connecting the start and end point of the dynamics under investigation. Ring polymer molecular dynamics capitalises on the isomorphism of the path integral description of quantum mechanics and the classical behaviour of a hypothetical ring polymer.[120,121] In the context of the aTSA, PIMD only serves as a sampling tool and is not used to collect dynamical information, in depth discussion of the finer points of the method is therefore omitted here.

The harmonic PIMD Hamiltonian is given by

$$H_n(\mathbf{q},\mathbf{p}) = \sum_{j=1}^{n}\left[\frac{p_j^2}{2m} + \frac{1}{2}m\omega_n^2\left(q_j - q_{j-1}\right)^2 + V(q_j)\right], \tag{5.1}$$

where (q_j, p_j) are the phase space coordinates of the jth bead, $\omega_n = \frac{1}{\beta_n\hbar}$, which in turn depends on the inverse temperature $\beta_n = \frac{1}{n}\beta = \frac{1}{nk_bT}$ and the index j is bound cyclicly such that $(q_0, p_0) = (q_n, p_n)$. The phase space coordinates of the centroid of this ring may be determined as a simple average over all beads such that

$$q_n = \frac{1}{n}\sum_{j=1}^{n} q_j \qquad\qquad p_n = \frac{1}{n}\sum_{j=1}^{n} p_j\,. \tag{5.2}$$

As this Hamiltonian is purely classical, equations of motion may be readily derived. In the case of the beads these are

$$\frac{\partial q_j}{\partial t} = \frac{p_j}{m}\,, \tag{5.3}$$

$$\frac{\partial p_j}{\partial t} = -m\omega_n^2\left(2q_j - q_{j-1} - q_{j+1}\right) - \frac{\partial V(q_j)}{\partial q_j}\,. \tag{5.4}$$

While for the centroid they are given by

$$\frac{\partial q_n}{\partial t} = \frac{p_n}{m}\,, \tag{5.5}$$

$$\frac{\partial p_n}{\partial t} = -\frac{1}{n}\sum_{j=1}^{n}\frac{\partial V(q_j)}{\partial q_j}\,. \tag{5.6}$$

However, in practice, the time evolution of the centroid is dictated by the motion of the beads, that is to say, rather than propagating the centroid using the Eqs. 5.5 and 5.6, it is recalculated for each new set of bead coordinates, using Eq. 5.2.

5.2.2 Path integral sampling algorithm

Employing PIMD sampling trajectories in the aTSA requires only minor adjustments of the algorithm described in Section 4.2. After the initial conditions of a given sampling trajectory haven been sampled from the Wigner distribution of either the initial wavefunction or of a selected inherited basis function, a ring polymer is generated by placing a number of n identical copies of the resulting GWP at the sampled phase space coordinates.

At this point it becomes necessary to estimate the temperature of the GWP in order to allow evaluation of the Hamiltonian given in Eq. 5.1. While playing a critical role in applications of PIMD where dynamics are calculated directly from the trajectories, in the aTSA, a reasonable approximation of temperature is likely to suffice, given that the trajectories only fulfil a sampling role and are not directly involved in the calculation of dynamics. Thus, the relatively simple equipartition theorem describing the relationship between the average kinetic energy and temperature of an ideal gas was adapted to fit this case, such that

$$\left\langle g \left| \hat{T} \right| g \right\rangle = \frac{3}{2} k_B T \,, \tag{5.7}$$

where $|g\rangle$ is the selected GWP, k_B is the Boltzmann constant and the kinetic energy operator takes the general form $\hat{T} = -\frac{\hbar^2}{2m}\nabla^2$. From this, the inverse temperature, β may be calculated as

$$\beta = \frac{3}{2\left\langle g|\hat{T}|g\right\rangle} \,. \tag{5.8}$$

After initially placing all n beads of the new ring polymer at the phase space coordinates of the GWP representing the desired initial conditions, the next step is to equilibrate the polymer. This is achieved by integrating the equations of motion for each bead, given by Eqs. 5.3 and 5.4, using the Velocity Verlet algorithm, discussed in detail in Appendix II. The parameters used for this equilibration are identical to those used for the sampling trajectories. In order to prevent the polymer moving away from the coordinates, specified by the initial conditions, after each equilibration step, the centroid of the ring polymer is calculated using Eq. 5.2 and the positions of each bead are shifted by the vector corresponding to the difference between the new centroid and the initial positions. This is illustrated in Figure 5.1.

To determine when to terminate equilibration, the radius of gyration, r_g, of the ring polymer is calculated at each step using

$$r_g^2 = \frac{1}{n}\sum_j^n \left[\sum_k^f \left(q_k^j - q_k^n\right)^2\right]^{\frac{1}{2}} , \tag{5.9}$$

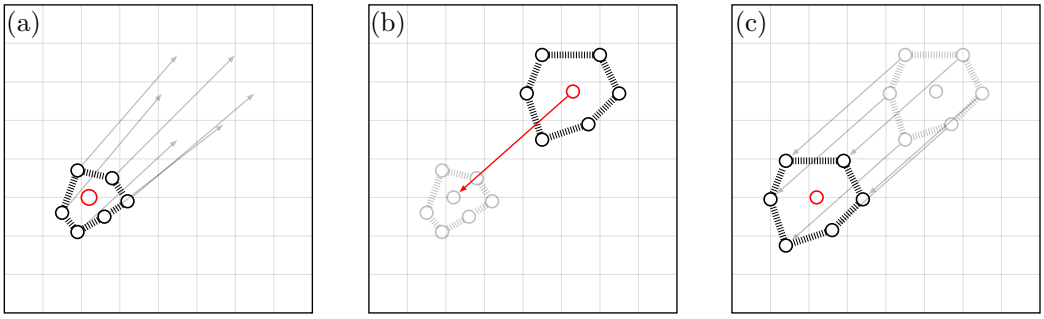

Figure 5.1: Illustration of the PIMD equilibration algorithm: (a) Each bead is evolved using the Velocity Verlet algorithm, (b) the centroid of the ring is calculated based on the new bead positions, (c) each bead is shifted by the vector connecting the new and old centroid, returning the former to the latter.

where q_k^j refers to the position of the jth bead of the polymer in DOF k and similarly, q_k^n refers to the position of the centroid in DOF k. As initially all beads of the ring polymer have identical positions, r_g starts at zero and increases as the ring expands. This is monitored and once the fractional increase in r_g between two steps falls below a user set parameter, ξ, the equilibration algorithm is considered to be complete.

The equilibrated ring polymer is then propagated in time, again using the Velocity Verlet algorithm to integrate the equations of motion for each bead, given by Eqs. 5.3 and 5.3. Other than these equations of motion being of a slightly different form than those discussed in Section 4.4, individual sampling steps proceed identically, using the same input parameters for the number of sampling timesteps and their duration, n_t and Δt_t respectively. When storing basis functions, which is again controlled by the storage frequency parameter, $1/n_s$, the fact that each trajectory now consists of an evolving ring polymer, each bead having different phase space coordinates, allows a choice with regards to which parameters to choose for the stored GWP basis function.

The most immediately obvious choice is to store a basis function with positions and momenta corresponding to the centroid of the ring polymer, which may be calculated using Eq. 5.2. Alternatively, a single bead may be chosen at random and the corresponding phase space coordinates then used as the parameters of the stored basis function. As it is not immediately clear *a prioiri* which of these approaches will yield more accurate results, both will be investigated.

Another parameter that's key to the performance of the path integral trajectories is the number of beads employed in the ring polymer. As again, there is not instinctive way of choosing this for this particular implementation of PIMD, the normal approach of using the vibrations of the system to inform this choice not

applying here, the number of beads was simply scaled until no further increase in accuracy was observed.

5.3 Tunnelling benchmark

In order to assess whether the use of PIMD trajectories to drive the sampling stage of the aTSA can significantly improve the accuracy with which strong quantum effects, specifically tunnelling, can be treated, the synthetic double well benchmark discussed at length in Section 4.4 was investigated again. The parameters defining the potential, as well as the terminology referring to three separate versions of this Hamiltonian as Model I, II and III are retained, thus detailed discussion of the system is omitted here. Also retained is the method of analysis, specifically calculating the tunnelling autocorrelation function, $C_t(t)$, using the approach outlined in Section 4.4, specifically Eq. 4.5.

5.3.1 Linear coupling

Starting with Model I, the use of PIMD sampling trajectories resulted in a noticeable improvement of the accuracy with which exact CI results[119] were reproduced. Two calculations, both employing $N_b = 12$ bursts of trajectory sampling, each consisting of $n_t = n_p = 10000$ steps of trajectory sampling and wavefunction propagation, were run. Timestep durations were $\Delta t_t = \Delta t_p = 0.002$ a.u. with a sampling frequency of $1/n_s = 1/500$ and a MP minimisation and optimisation convergence criterion of $\zeta = 0.99$. Both calculations employed ring polymers constituted of $n = 4$ beads which were equilibrated until the fractional variation in the radius of gyration was less than $\xi = 0.001$. The method of basis function storage was

Table 5.1: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of Model I, using PIMD sampling trajectories, the results of which are shown in Figure 5.2.

N_b	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	ζ	m	n	ξ
12	10000	0.002	500	10000	0.002	0.99	8	4	0.001
12	10000	0.002	500	10000	0.002	0.99	12	4	0.001

N_{total}	N_t	N_{MP}
181	169	12
249	237	12

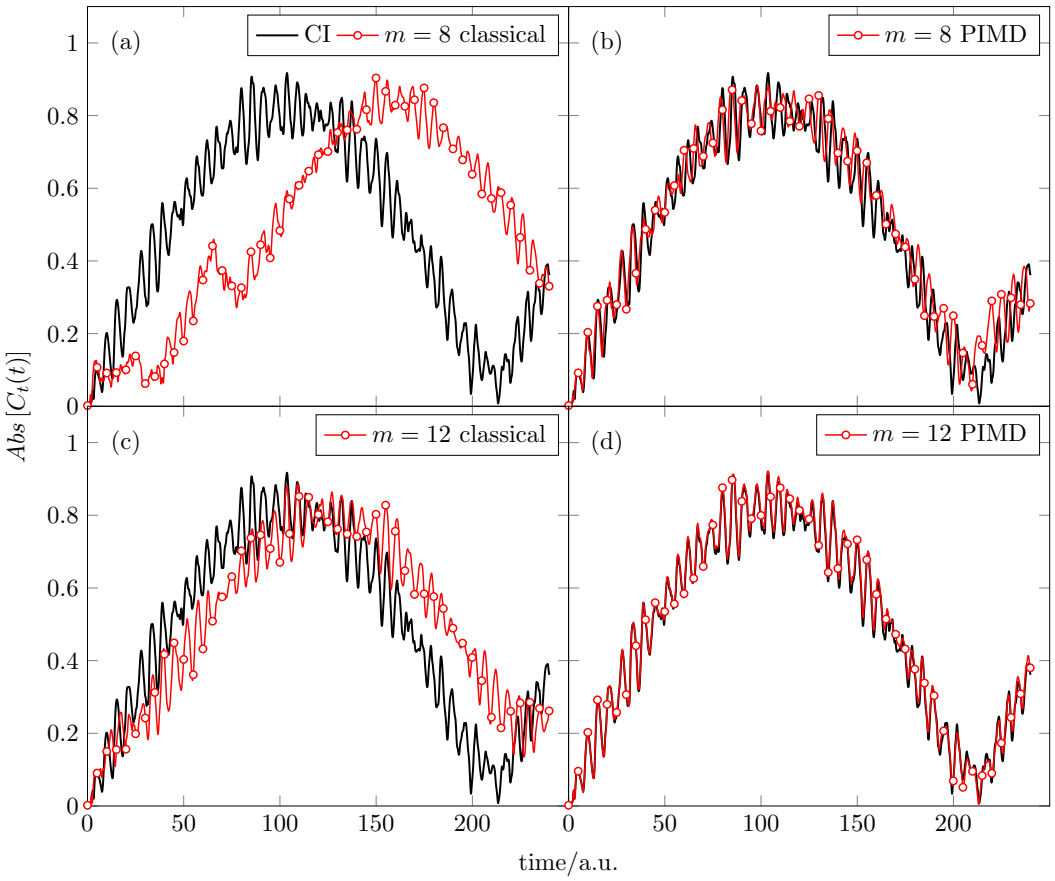

Figure 5.2: Tunnelling autocorrelation functions for Model I, comparing the effect of using classical ((a) and (c)) and PIMD ((b) and (d)) sampling trajectories in the aTSA, varying the number of sampling trajectories, m, compared to exact CI results.[119]

to randomly select a given bead of the ring polymer. These calculations differed only in the number of sampling trajectories employed, which, in order to allow comparison to the data presented in Section 4.4.1, $m = 8$ and $m = 12$. These input parameters and the resulting average basis set sizes are summarised in Table 5.1.

The results of these two calculations and, to allow the improvement afforded by the use of PIMD sampling trajectories to be quantified, results for calculations with the classical aTSA employing identical input parameters, previously shown in Figure 4.7, are shown in Figure 5.2. Comparing the results for classical and PIMD sampling with the same number of sampling trajectories, the benefit of the latter becomes immediately obvious. In the case of $m = 8$ sampling trajectories, PIMD sampling results in a basis set which qualitatively captures the tunnelling dynamics, however fails to correctly describe the short time oscillations of the wavefunction. This does nevertheless represent a significant improvement over

classical sampling, which fails to describe the tunnelling with nearly as much accuracy, introducing both a superfluous peak around 75 a.u. as well as significantly underestimating the rate with which the wavefunction tunnels through the barrier.

This observation is further cemented when considering the results using $m = 12$ sampling trajectories. The quality of results from classical sampling improves slightly, however it is still not comparable to that achieved using $m = 8$ PIMD sampling trajectories. In fact, referring back to Figure 4.7, to reach similar accuracy to that shown in Figure 5.2(b), at least $m = 16$ classical trajectories are necessary. The results from $m = 12$ PIMD trajectories on the other hand now capture both long and short time oscillations of the wavefunction, reproducing the exact CI reference results[119] quantitatively.

Overall these results are highly encouraging, as they clearly indicate that a significant improvement in the accuracy with which the sampled basis set can describe wavefunction evolution may be derived from using semi-classical PIMD sampling trajectories. Given that this is not associated with a drastic increase in computational cost, especially considering the previously discussed cost of the MP minimisation and optimisation algorithm, which remains unaffected by the nature of sampling trajectories, this suggest that PIMD constitutes an effective and worthwhile choice of sampling for the aTSA.

Moving to the more challenging version of the linear coupling Hamiltonian, Model II, again 2 calculations with input parameters similar to those of calculations presented in Section 4.4 were run. In keeping with the latter, both calculations involved $N_b = 24$ bursts of trajectory sampling and wavefunction propagation, each consisting of $n_t = n_p = 5000$ steps lasting $\Delta t_t = \Delta t_p = 0.002$ a.u. Both calculations also employed a basis function storage frequency of $1/n_s = 1/500$. The

Table 5.2: Input parameters and average total, trajectory sampled and inherited basis set sizes for adaptive sampling calculations of Model II, using PIMD sampling trajectories, the results of which are shown in Figure 5.3.

N_b	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	ζ	m	n	ξ
24	5000	0.002	500	5000	0.002	0.99	100	4	0.001
24	5000	0.002	500	5000	0.002	0.995	200	10	0.001

N_{total}	N_t	N_{MP}
1038	1002	36
2116	2007	109

remaining input parameters differed so as to probe two levels of stringency with which the aTSA and PIMD trajectories operated.

The first calculation involved $m = 100$ sampling trajectories, with a MP minimisation and optimisation criterion of $\zeta = 0.99$. It employed a $n = 4$ bead PIMD ring polymer, which was equilibrated until the fractional radius of gyration variation was below $\xi = 0.001$ and used the parameters of a randomly chosen bead for each stored basis function. The second calculation featured $m = 200$ sampling trajectories, with $\zeta = 0.995$, $n = 20$, $\xi = 0.001$ and also used random bead parameters for basis function storage. This second calculation thus involved both an inherently more quantum-like ring polymer and a more accurate implementation of the MP minimisation and optimisation algorithm, as well as overall more sampling trajectories and thus basis functions. The input parameters and the resulting average basis set sizes for both these calculations are summarised in Table 5.2.

Figure 5.3 compares the tunnelling autocorrelation function, calculated for both these calculations using Eq. 4.5, to those obtained with classical sampling trajectories, previously discussed in Section 4.4.1 and shown in Figure 4.7(a) and (b). The observations made above for Model I are very much replicated in the results for Model II. Again, as can be seen when comparing Figure 5.3(a) and (b), PIMD sampling trajectories result in a basis set which, for a given number of basis functions, much more effectively describes the tunnelling behaviour of the wavefunction. The PIMD sampled basis set clearly captures the tunnelling of the wavefunction qualitatively (only the short time oscillations not being accurately described), which classical sampling fails to do, underestimating the rate of tunnelling, which is observed as a shift in $C_t(t)$ of approximately 25 a.u.

In fact, using only $m = 100$ PIMD sampling trajectories and thus a basis set consisting, on average, of only 1000 basis functions, yielded results as, if not more, accurate than those obtained with $m = 800$ classical trajectories and a basis set of 8000 GWPs, which are shown in Figure 4.8(d). More specifically, while similarly to classical sampling, PIMD trajectories only captured the qualitative aspects of tunnelling, the "period" of tunnelling, that is the point in time at which the wavefunction has returned completely to its original configuration, at around 210 a.u., is reproduced almost exactly, with respect to the CI reference,[119] in the case of the latter.

This suggests, as in the case of Model I, significantly quicker convergence of the aTSA with PIMD trajectories, which is further supported by the minimal increase in accuracy observed between Figure 5.3(b) and (d). The latter, as discussed above, is the result of a calculation, both employing overall more basis functions, as well as a larger, that is composed of more beads, ring polymer. Additionally

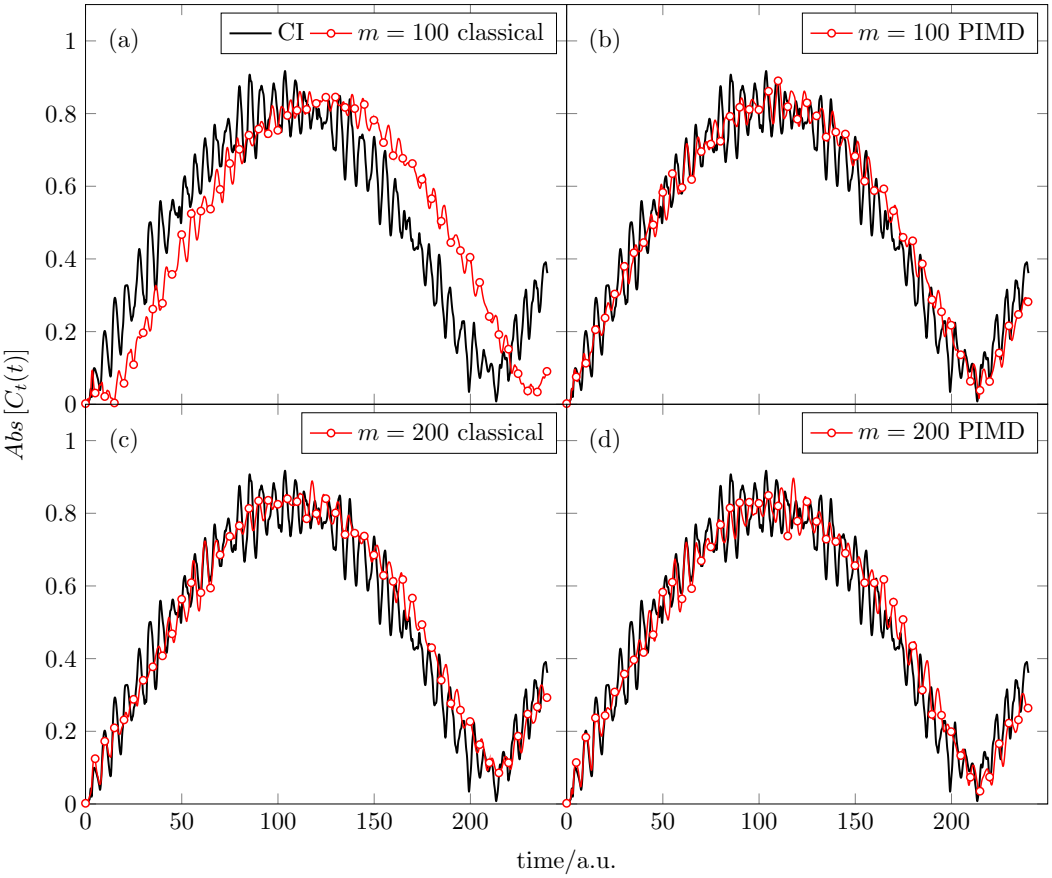

Figure 5.3: Tunnelling autocorrelation functions for Model II, comparing the effect of using classical ((a) and (c)) and PIMD ((b) and (d)) sampling trajectories in the aTSA, varying, amongst other parameters, the number of sampling trajectories, m and beads n (further input parameters are given in Table 5.2), compared to exact CI results.[119]

more stringent control parameters of the equilibration and the MP minimisation and optimisation algorithm were used as well.

These results further cement the observation, made above for Model I, that using PIMD sampling trajectories in combination with the aTSA significantly improves the accuracy with which strong quantum effects such as tunnelling can be described. Furthermore, as the number of beads in the ring polymer required in both cases was comparatively low, the increase in computational cost over purely classical trajectories is minimal, especially in comparison to that associated with the MP minimisation and optimisation algorithm.

5.3.2 Quadratic coupling

In Chapter 4, the 20-dimensional, quadratic coupling version of the double-well Hamiltonian, termed Model III, was found to constitute a significantly more challenging benchmark problem than either Model I or II. This is a result of both the increased number of DOFs and the nature of the coupling between them. The latter plays a particularly interesting role, as in the case of Model III, each DOF is coupled directly to the tunnelling coordinate, which makes this version of the benchmark more representative of a real system, where some coupling, albeit of varying strength, would be expected for most pairs of DOFs.

The intention of introducing the PIMD sampling trajectories in the context of the aTSA is to improve the description of strong quantum effects such as tunnelling. The performance for Model III of the algorithm with this new type of sampling trajectories should be indicative of the their applicability to realistic systems featuring strong quantum effects.

In order to investigate the performance of the aTSA with PIMD sampling trajectories, a calculation, consisting of $N_b = 12$ bursts was run, each consisting of $n_t = n_p = 5000$ timesteps of trajectory sampling and propagation, each lasting $\Delta t_t = \Delta t_p = 0.002$ a.u. Basis functions were stored with a frequency of $1/n_s = 1/50$, while the MP minimisation and optimisation convergence criterion was $\zeta = 0.99$. A total of $m = 100$ PIMD sampling trajectories were run, each consisting of a ring of $n = 4$ beads, which were equilibrated with a radius of gyration partial variation condition of $\xi = 0.001$. For each storage event, a random bead was chosen. These input parameters and the resulting average basis set sizes are summarised in Table 5.3.

Figure 5.4 shows the tunnelling autocorrelation functions, calculated using Eq. 4.5 for calculations employing the TSA, aTSA with classical and PIMD trajectories. As already discussed in Chapter 4, the aTSA significantly improves the

Table 5.3: Input parameters and average total, trajectory sampled and inherited basis set sizes for an adaptive sampling calculation of Model III, using PIMD sampling trajectories, the results of which are shown in Figure 5.4(c).

N_b	n_t	Δt_t/a.u.	n_s	n_p	Δt_p/a.u.	ζ	m	n	ξ
12	5000	0.002	50	5000	0.002	0.99	100	4	0.001

N_{total}	N_t	N_{MP}
10424	9962	462

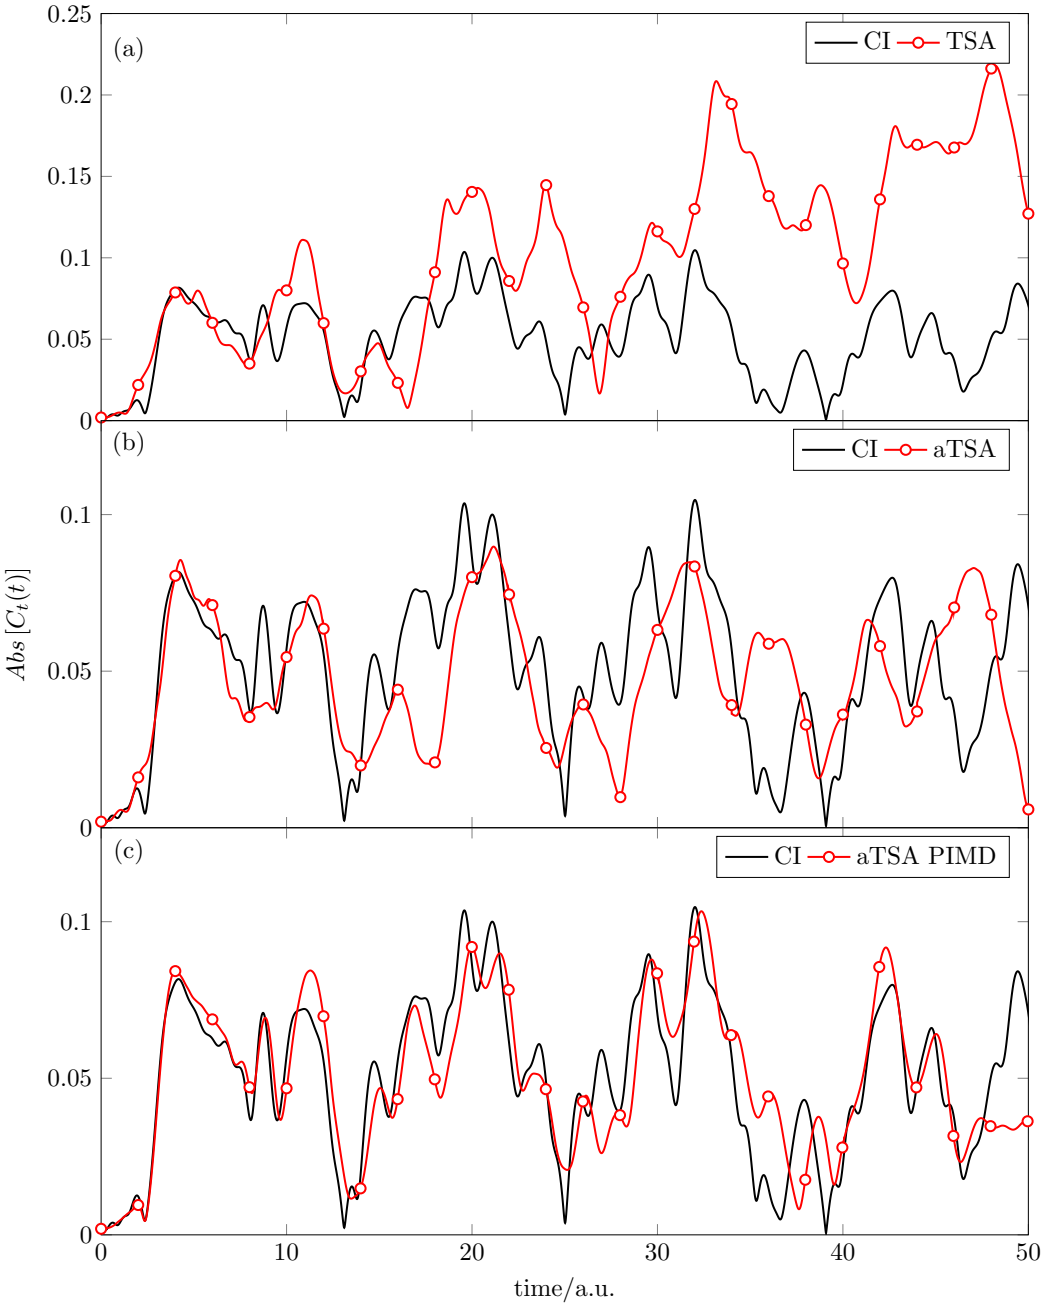

Figure 5.4: Tunnelling autocorrelation functions for Model III, calculated using (a) the trajectory sampling algorithm, (b) the adaptive trajectory sampling algorithm with classical sampling trajectories and (c) the aTSA with PIMD sampling trajectories, all compared to an exact CI benchmark.[119]

description of the tunnelling over the TSA. The latter overestimates the extent of tunnelling and fails to capture the oscillations of the wavefunction between the two wells of the tunnelling coordinate. Encouragingly, the inclusion of PIMD trajectories significantly improves the descriptions of these oscillations. While for classical sampling the aTSA only qualitatively reproduced the tunnelling dynamics, the PIMD sampled basis reached close to quantitative accuracy with respect

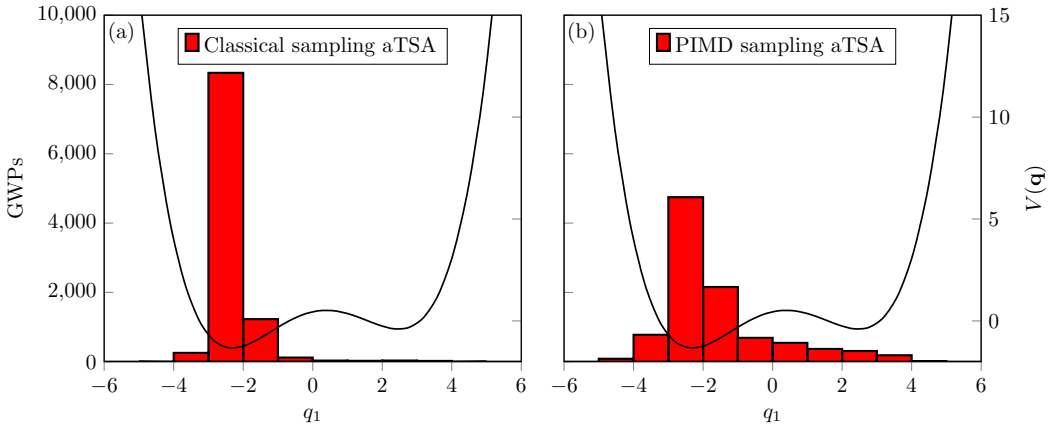

Figure 5.5: Distribution of basis functions across the q_1 tunnelling coordinate of Model III, for a single iteration of the aTSA with (a) classical and (b) PIMD sampling trajectories.

to the CI reference.[119]

The main discrepancy between the PIMD aTSA and CI results lies in the amplitudes of the tunnelling autocorrelation function. Towards the end of a iteration of the aTSA propagation stage, a slight shift with respect to the benchmark results also becomes apparent, however this is often corrected by the MP minimisation and optimisation followed by trajectory sampling during the beginning of the next iteration.

In order to visualise the benefit of the PIMD sampling trajectories over the purely classical approach taken in Chapter 4, Figure 5.5 shows the positions of the basis functions in the tunnelling coordinate, q_1 for the first iteration of the calculations shown in Figure 5.4 (b) and (c). As can clearly be seen from comparing Figure 5.5 (a) to (b), even for the first iteration of the aTSA the difference in basis function placement is significant. Using PIMD sampling trajectories not only distributes the basis set more evenly across the left well, where the initial wavefunction is located, but also populates the other well and areas corresponding to the energy barrier, thus allowing for tunnelling to be described by the resulting basis set with far greater accuracy. Given that the PIMD sampling trajectories scale linearly with the number of beads, the benefits afforded by implementation of the aTSA far outweigh the costs.

Overall these results are highly encouraging, not only due to the significant increase in accuracy over the classic sampling results, but also because, as for Model I and II, only $n = 4$ beads were required per PIMD trajectory to achieve this improvement. This means that even for a relatively high-dimensional problem like Model III, these trajectories will not incur a significant increase in the computational resources required.

5.4 Conclusions

The aim of this Chapter was to test the applicability of the PIMD Hamiltonian as the driving mechanic for the sampling trajectories of the aTSA. PIMD trajectories can capture quantum events using a simple expansion of the particle as a ring polymer of identical beads, connected by harmonic springs, which evolve according to classical equations of motion. Given the well documented[120,121] ability of PIMD trajectories to increase the accuracy with which strong quantum effects such as tunnelling are described, they were expected to provide an excellent choice for the sampling algorithm of the aTSA. One of the main shortcomings of the latter, discovered in Chapter 4, relating to the classical nature of the sampling trajectories.

While the wavefunction expansion of the aTSA is formally exact in the finite basis limit, the time evolution of the wavefunction being expressed solely through the expansion coefficients, which move according to variational equations of motion,[34–37] the classical mechanics driving the basis set sampling severely limit the accuracy with which strong quantum effects such as tunnelling can be described.

In order to address this, a simple implementation of the PIMD formalism was developed for use as the sampling stage of the aTSA. The PIMD trajectories are not used to collect dynamical information, instead they are only required for the placement of basis functions in regions of phase space, relevant to wavefunction propagation. Thus, a number of simplifications could be made, the first of which concerns the estimation of the temperature, required as part of the PIMD Hamiltonian. In this case, the temperature was estimated by simply calculating the kinetic energy of the basis function providing the initial conditions of the sampling trajectory, which was then inserted into a relationship borrowed from the ideal gas.

The second approximation that could be made concerns the equilibration of the ring polymers, previous to propagation of the sampling trajectories. Again a relatively simple strategy was chosen here, whereby the ring is evolved in time, using the Velocity Verlet algorithm, for one timestep, then the centroid is calculated and every bead shifted by the vector connecting the new centroid to the one immediately before movement of the ring. This is repeated until the variation of the ring's radius of gyration fell below a user-determined threshold.

Applying this formalism to the tunnelling double well benchmark studied in Chapter 4, yielded highly encouraging results. For the smallest, linear coupling model, exact dynamics with respect to the CI reference[119] was converged to much faster using PIMD trajectories than their classical counterparts from Chapter 4. The same was found for the higher-dimensional version of the linearly coupled double well Hamiltonian, where qualitatively accurate results were again obtained

using far fewer basis functions with the PIMD aTSA and a small improvement was observed for the timescale of the tunnelling.

Finally, applying the PIMD sampling trajectories to the most challenging, high-dimensional quadratic coupling version of the double well benchmark yielded excellent results. In this case a significant improvement, not only in scaling but also in the overall quality of dynamics achievable, was observed with respect to classical sampling.

These initial results obtained suggest that using PIMD trajectories as the driving mechanics of the sampling stage in the aTSA can, as predicted, significantly improve its ability to treat strong quantum effects such as tunnelling. Given however that this version of the aTSA has only been tested for three version of the same benchmark so far and that the dependence on input parameter has so far not been systematically tested, further work should be carried out before generally adopting PIMD trajectories.

Given however that all results presented in this chapter suggest that only a few beads in the ring polymer are required in order to capture tunnelling, further investigation is considered to be highly worthwhile, especially given the significantly more favourable scaling of the sampling stage of the aTSA than either the propagation of the wavefunction or the MP minimisation and optimisation algorithm, the latter of which was found to somewhat limit performance in Chapter 4. Thus, if the increased sampling accuracy of the PIMD trajectories could be leveraged against the computational expense incurred by the MP algorithm, this could serve to increase the maximum size of systems that may be investigated with the aTSA, as well as improving performance and reducing computational costs.

Chapter 6

Conclusions and Further Work

This chapter states again the goal of this work, putting it into the wider context of the field it resides in. Then follows a summary of the methods and results presented herein and the conclusions that can be drawn from them. Finally a number of avenues for potential further work are highlighted.

6.1 Summary and conclusions

The quantised nature of the electromagnetic field means that for many interactions between light and molecules, quantum effects need to be taken into account when studying such processes theoretically. The benefit of applying theory to light-matter interactions can translate into significant cost and time savings in both industrial and academic research and development. Significant efforts are thus directed towards improvement of the theoretical methods that model the quantum dynamics of photophysical and photochemical systems.

Existing methods for nuclear quantum dynamics can be broadly categorised based on how the wavefunction is expanded. The natural complexity of the latter requires it to be expressed in a basis of mathematically well behaved functions, the ensemble of which can be treated in order to simulate the time-evolution of the wavefunction. The manner in which this is achieved and the nature of this basis set for most common approaches falls into one of two categories, each with its own advantages and limitations.

If the basis functions remain fixed in phase space, the propagation of the wavefunction occurs *via* the expansion coefficients associated with the basis set, for which equations of motion are derived, using the time-dependent variational principle.[34–37] The benefit of such time-independent basis sets[38–41] is twofold. The aforementioned equations of motion are well behaved and easy to integrate and within the limit of a sufficiently large basis set, the dynamics resulting from their propagation are formally exact. There are however two significant limitations of such methods. The number of basis functions typically required are rather large, as, without a way of choosing *a priori* where in phase space to place coordinates, a grid of functions, sufficiently dense to cover all feasibly accessible space, must be employed. Furthermore, and closely related to this, methods of this type exhibit extremely unfavourable exponential scale with the number of degrees of freedom of the system, thus severely limiting the size of molecules that may be treated using them.

Alternatively, the basis functions can be allowed to, in addition to the expansion coefficients, evolve in time, resulting in a time-dependent basis set.[42–46,74] This is achieved by deriving a set of equations of motion for the parameters defining the basis functions. The resulting basis set moves in phase space and can therefore adapt to the changing amplitude of the wavefunction. The benefit of this is that in comparison to time-independent basis sets, far fewer functions are required in order to achieve a given level of accuracy. In addition to this, these methods typically exhibit scaling with respect to system size, which while not necessarily favourable, still allows them to outperform time-independent methods.

This benefit does however come at a price. The equations of motion for the basis function parameters have a tendency to be ill-conditioned, which leads to their integration being numerically challenging.[74] Furthermore, if these equations are not derived using the same variational approach[34–37] that is employed for the expansion coefficients, the resulting trajectories have been shown to violate energy conservation laws.[47]

This work has presented an alternative method for sampling basis sets for nuclear quantum mechanics,[57] which while borrowing aspects from both these archetypes, avoids some of their commonly associated limitations. The key idea is to sample a time-independent basis set, in a manner which is inspired by the wavefunction propagation it is intended for and avoids having to place too many functions. To achieve this, a set of simple trajectories, with initial conditions sampled from the wavefunction, evolve on the potential energy surface of the system, storing basis functions in phase space along their path. The key idea being that, as the trajectories are not used to actually calculate dynamics, they need no follow the exact quantum path. As long as some of these sampling trajectories visit the regions of phase space, relevant to the propagation of the wavefunction, and basis functions are stored there, the subsequent propagation *via* the expansion coefficients will be accurate. This approach maintains all the benefits of time-independent basis sets, while avoiding scaling issues, as the basis set only occupies regions of phase space relevant to propagation of the wavefunction.

In order to test the effectiveness of this approach, it was applied to the challenging quantum dynamics benchmark problem, modelling the relaxation dynamics of photoexcited pyrazine. This small, organic molecule decays, upon excitation to its second excited state, rapidly, via a conical intersection, to the lower lying excited state. In addition to proving both a low and higher-dimensional version of the Hamiltonian to test the novel method proposed herein against, this particular problem has been studied with a number of other quantum dynamics approaches[62,65,70,73,84,106,110,112–114] and exact results are available.[65]

The trajectory sampling method[57] yielded some extremely encouraging results for the lower-dimensional version of the pyrazine benchmark, approaching numerical accuracy. For the higher-dimensional version, only qualitative accuracy could be achieve, however, in combination with the relatively low computational expense associated with the calculations carried out, these results are still highly competitive, especially when comparing to other, purely time-independent methods.

One common feature of these results for the pyrazine benchmark was that the dynamics obtained were, even for the higher-dimensional version, essentially exact in the short-time limit, errors only appearing at longer times. This was linked to the main assumption of this method, namely that the classical sampling

trajectories constitute a sound approximation to the quantum propagation of the wavefunction. While this is well known to be true in the short-time limit, at longer times a number of effects, such as leakage of zero point energy into other degrees of freedom,[116,117] cause the divergence of quantum and classical paths in phase space. This results in inefficient basis function placement at longer times, which translates into sub par dynamics being produced during the wavefunction propagation stage.

In order to address this shortcoming, a modification to the trajectory sampling algorithm was thus presented.[122] The rather simple underlying principle being that the existing method can yield accurate dynamics, as long as it is confined to short periods of time. Therefore, the modified algorithm involves short "bursts" of trajectory sampling and subsequent wavefunction propagation *via* the expansion coefficients, strung together in order to make up the desired simulation time. In order to prevent the basis set from continuously expanding, which would reintroduce the exponential scaling that plagues traditional time-independent methods, a novel algorithm was developed to minimise and optimise the set of functions carried over from one iteration to the next.

Based on the matching pursuit method,[81] which has been used as part of other quantum dynamics methods before,[82,83] this relatively simple scheme results in a basis set which, while remaining time-independent as far as wavefunction propagation is concerned, can as part of the multiple iterations of the algorithm adapt to the change in the wavefunction, not dissimilar to a time-dependent basis set.

Applying this adaptive trajectory sampling method[122] to the pyrazine benchmark yielded extremely encouraging results. For the lower-dimensional version of the Hamiltonian, near quantitative accuracy could be achieve with a fraction of the basis functions required by the trajectory sampling method, consequently also incurring significantly lower computational expense. For the higher-dimensional version of the benchmark a marked increase in the accuracy of the dynamics obtained was observed, with essentially all qualitative features being reproduced.

In order to further test the performance of this method, a double-well benchmark modelling the tunnelling of a quantum particle through a one-dimensional energy barrier was also considered. Again this system was chosen not only for the significant challenge it poses for most quantum dynamics methods, but also because it has previously been studied with a number of approaches[47,82,119] and exact results are again available.[119]

The adaptive sampling method[122] preformed well for two versions of this benchmark, featuring linear coupling between the tunnelling coordinate and its environment. In the lower-dimensional case, exact results were reproduced using small

basis sets, while for the higher-dimensional version, highly qualitatively accurate dynamics were calculated, albeit using larger basis sets.

For the significantly more challenging, quadratic ally coupled version of the double-well benchmark however, the adaptive sampling algorithm did not perform as well. While still significantly outperforming the simpler trajectory sampling method,[57] the results obtained only roughly reproduced the qualitative features of the exact dynamics. This was attributed to the fact that the combination of the significantly higher number of degrees of freedom and more complex coupling make this version of the benchmark much more representable of a realistic system. While the short bursts of sampling and propagation, interspaced with minimisation and optimisation of the basis set, can address some classical-quantum divergence effects, their ability so correctly sample strong quantum events, such as the tunnelling in this case, will always be limited.

In order to address this shortcoming, path integral molecular dynamics[117,120] were employed to drive the sampling trajectories. In this relatively simple scheme, each particle is expanded into a ring of identical beads, connected by harmonic springs. Evolving this ring classically has been shown to result in sampling of quantum effects at essentially classical costs. As the trajectories only act as a sampling mechanism and do not actually calculate dynamics, which is normally the role of path integrals, a number of approximations could be introduced. Firstly, the temperature, from which the stiffness of the ring is derived, was approximated based on the kinetic energy of the basis function following the trajectory. Secondly, in order to equilibrate the polymer, it was simply evolved normally for a single step, then its centroid was calculated and the entire ring shifted by the vector connecting this new centroid to the old one. This was repeated until the variation in the rings radius of gyration fell below a threshold.

This rather simple implementation of path integral dynamics was found to result in significant improvements for the description of the tunnelling dynamics of the double-well benchmark. For the lower-dimensional version of linear coupling, the trajectory sampling algorithm converged to the exact result significantly quicker than with classical sampling trajectories. For the higher-dimensional, linear coupling benchmark, in addition to similarly improved convergence, a small increase in the accuracy of the dynamics was also observed.

Finally and most encouragingly, for the quadratically coupled, high-dimensional tunnelling problem, the quality of the dynamics was significantly increased, approaching quantitative agreement with exact results. This more drastic increase in the performance of the adaptive sampling algorithm with path integral trajectories may be linked to the more realistic nature of this particular version of the benchmark. The benefit of using these semi-classical trajectories may in fact

only truly come into play once the complexity of the problem reaches a certain threshold, the path integral otherwise not significantly affecting the sampling with respect to purely classical sampling trajectories.

In addition to the improvements in quality, a further encouraging aspect of the results for the double-well benchmark was that the number of path integral beads required in order to reap the benefits outlined above was extremely small, which meant that the use of the path integral sampling trajectories came at little additional cost. Given that strong quantum effects are a common occurrence in real systems of interest, this points towards the use of this type of trajectories as a relatively cheap solution to the challenge of modelling the behaviour in such environments.

Overall the trajectory based scheme for sampling quantum dynamics basis sets presented and systematically improved upon herein[57,122] has been demonstrated to be quite effective for the types of problems often encountered in the study of photophysical and photochemical processes. In contrast to traditional time-independent[38–41] and time-dependent[42–46,74] methods, this novel approach to the issue of expanding and evolving the wavefunction has demonstrated that a compromise between the advantages and limitations of these two categories is indeed possible and, as outlined above, certainly warrants further investigation.

6.2 Further work

There are a number of potential avenues for further work regarding the trajectory sampling method presented herein. There are of course a number of benchmark problems, highlighting different challenges to quantum dynamics methods, which should be investigated. In addition, while systematic investigations for both the trajectory sampling and adaptive trajectory sampling algorithm with classical trajectories have been carried out, probing the effects of the various input parameters, no such work has yet been undertaken for path integral sampling trajectories. Given the insights into the workings of these methods gained from these studies, this is certainly considered worthwhile.

Another direction to pursue would be the investigation of the effect of different methods as the driving mechanics of the sampling trajectories. Trajectory surface hopping[42,90–93] and *ab initio* multiple spawning[43,56,86–88] could for example replace the Ehrenfest dynamics employed in order to account for the multiple electronic states of the pyrazine benchmark.

A third area to expand into is direct dynamics.[85,95,96] This would involve determining Hamiltonian matrix elements, and in the case of multiple electronic

states, the coupling elements between them, from electronic structure calculations during each timestep of trajectory sampling. The choice of Gaussian basis functions is helpful in achieving this, not only are these functions associated with a classical centre, thus allowing the required integrals to be determined around this point, these integrals can also usually be solved analytically. This avenue of further development is considered particularly worthwhile as it would allow for the departure from analytical Hamiltonians and extend the range of systems that can be treated using this method to all those which are computationally feasible.

Bibliography

[1] J. P. Dowling and G. J. Milburn, *Philos. Trans. Roy. Soc. A*, 2003, **361**, 1655–1674.

[2] J. L. O'Brien, A. Furusawa and J. Vuckovic, *Nat. Photon.*, 2009, **3**, 687–695.

[3] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe and J. L. O'Brien, *Nature*, 2010, **464**, 45–53.

[4] N. Lambert, Y.-N. Chen, Y.-C. Cheng, C.-M. Li, G.-Y. Chen and F. Nori, *Nat. Phys.*, 2013, **9**, 10–18.

[5] J. F. Rabek, *Progress in PhotochemPhoto & Photophysics*, CRC Press, Boca Raton, 1990-1992, vol. I-VI.

[6] H. Gest, *Photosynth. Res.*, 2002, **73**, 7–10.

[7] D. E. Koshland, *Science*, 2002, **295**, 2215–2216.

[8] J. W. Schopf, *Science*, 1993, **260**, 640–646.

[9] R. E. Blankenship, *Molecular Mechanisms of Photosynthesis*, Blackwell Science, Oxford, 2002.

[10] G. M. Cooper, *The Cell: A Molecular Approach*, Sinauer Associates, Sunderland (MA), 2nd edn., 2000.

[11] Y.-C. Cheng and G. R. Fleming, *Annu. Rev. Phys. Chem.*, 2009, **60**, 241–262.

[12] L. A. Baker and S. Habershon, *J. Chem. Phys.*, 2015, **143**, 105101.

[13] A. Y. Hoekstra and T. O. Wiedmann, *Science*, 2014, **344**, 1114–1117.

[14] G. D. Scholes, G. R. Fleming, A. Olaya-Castro and R. van Grondelle, *Nat. Chem.*, 2011, **3**, 763–774.

[15] W.-D. Quan, A. Pitto-Barry, L. A. Baker, E. Stulz, R. Napier, R. K. O'Reilly and V. G. Stavros, *Chem. Commun.*, 2016, **52**, 1938–1941.

[16] N. M. Haegel, R. Margolis, T. Buonassisi, D. Feldman, A. Froitzheim, R. Garabedian, M. Green, S. Glunz, H.-M. Henning, B. Holder, I. Kaizuka, B. Kroposki, K. Matsubara, S. Niki, K. Sakurai, R. A. Schindler, W. Tumas, E. R. Weber, G. Wilson, M. Woodhouse and S. Kurtz, *Science*, 2017, **356**, 141–143.

[17] M. A. Green, Y. Hishikawa, W. Warta, E. D. Dunlop, D. H. Levi, J. Hohl-Ebinger and A. W. Ho-Baillie, *Prog. Photov.*, 2017, **25**, 668–676.

[18] *Best Research-Cell Efficiencies*, `https://www.nrel.gov/pv/assets/images/efficiency-chart.png`, 2017.

[19] Y. Lin, Y. Li and X. Zhan, *Chem. Soc. Rev.*, 2012, **41**, 4245–4272.

[20] H. Ma and A. Troisi, *Adv. Mater.*, 2014, **26**, 6163–6167.

[21] M. H. Lee, J. Arag and A. Troisi, *J. Phys. Chem. C*, 2015, **119**, 14989–14998.

[22] N. Martsinovich and A. Troisi, *J. Phys. Chem. C*, 2011, **115**, 11781–11792.

[23] F. Ambrosio, N. Martsinovich and A. Troisi, *J. Phys. Chem. Lett.*, 2012, **3**, 1531–1535.

[24] E. Maggio, N. Martsinovich and A. Troisi, *J. Phys. Chem. C*, 2012, **116**, 7638–7649.

[25] V. G. Stavros and J. R. Verlet, *Annu. Rev. Phys. Chem.*, 2016, **67**, 211–232.

[26] L. A. Baker, B. Marchetti, T. N. V. Karsili, V. G. Stavros and M. N. R. Ashfold, *Chem. Soc. Rev.*, 2017, **46**, 3770–3791.

[27] H. Tao, T. K. Allison, T. W. Wright, A. M. Stooke, C. Khurmi, J. van Tilborg, Y. Liu, R. W. Falcone, A. Belkacem and T. J. Martínez, *J. Chem. Phys.*, 2011, **134**, 244306–.

[28] T. K. Allison, H. Tao, W. J. Glover, T. W. Wright, A. M. Stooke, C. Khurmi, J. van Tilborg, Y. Liu, R. W. Falcone, T. J. Martínez and A. Belkacem, *J. Chem. Phys.*, 2012, **136**, 124317–.

[29] K. Kosma, S. A. Trushin, W. Fuss and W. E. Schmid, *J. Phys. Chem. A*, 2008, **112**, 7514–7529.

[30] H. Tao, B. G. Levine and T. J. Martínez, *J. Phys. Chem. A*, 2009, **113**, 13656–13662.

[31] G. Orlandi and W. Siebrand, *J. Chem. Phys.*, 1973, **58**, 4513–4523.

[32] F. Kaspar, W. Domcke and L. Cederbaum, *Chem. Phys.*, 1979, **44**, 33–44.

[33] E. J. Heller, *J. Chem. Phys.*, 1981, **75**, 2923–2931.

[34] P. A. M. Dirac, *Proc. Cambridge Phil. Soc.*, 1930, **26**, 376.

[35] J. Frenkel, *Wave Mechanics*, Oxford University Press, Oxford, 1934.

[36] A. D. McLachlan, *Mol. Phys.*, 1964, **8**, 39.

[37] J. Broeckhove, L. Lathouwers, E. Kesteloot and P. Van Leuven, *Chem. Phys. Letters*, 1988, **149**, 547–550.

[38] R. B. Gerber, V. Buch and M. A. Ratner, *J. Chem. Phys.*, 1982, **77**, 3022–3030.

[39] H.-D. Meyer, U. Manthe and L. Cederbaum, *Chem. Phys. Lett.*, 1990, **165**, 73–78.

[40] J. Light, in *Time-Dependent Quantum Molecular Dynamics*, ed. J. Broeckhove and L. Lathouwers, New York:Plenum, 1992, pp. 185–199.

[41] W. Koch and T. J. Frankcombe, *Phys. Rev. Lett.*, 2013, **110**, 263202.

[42] J. C. Tully and R. K. Preston, *J. Chem. Phys.*, 1971, **55**, 562–572.

[43] T. J. Martínez, M. Ben-Nun and R. D. Levine, *J. Phys. Chem.*, 1996, **100**, 7884.

[44] G. A. Worth and I. Burghardt, *Chem. Phys. Letters*, 2003, **368**, 502.

[45] D. V. Shalashilin, *J. Chem. Phys.*, 2009, **130**, 244101.

[46] D. V. Makhov, W. J. Glover, T. J. Martínez and D. V. Shalashilin, *J. Chem. Phys.*, 2014, **141**, 054110.

[47] S. Habershon, *J. Chem. Phys.*, 2012, **136**, 014109.

[48] E. Schrödinger, *Phys. Rev.*, 1926, **28**, 1049–1070.

[49] E. Schrödinger, *Ann. d. Physik*, 1926, **79**, 361.

[50] E. Schrödinger, *Ann. d. Physik*, 1926, **79**, 489.

[51] E. Schrödinger, *Ann. d. Physik*, 1926, **79**, 734.

[52] E. Schrödinger, *Ann. d. Physik*, 1926, **80**, 437.

[53] E. Schrödinger, *Ann. d. Physik*, 1926, **81**, 109.

[54] E. J. Heller, *J. Chem. Phys.*, 1975, **62**, 1544–1555.

[55] M. J. Davis and E. J. Heller, *J. Chem. Phys.*, 1981, **75**, 246–254.

[56] M. Ben-Nun and T. J. Martínez, *Adv. Chem. Phys.*, 2002, **121**, 439.

[57] M. A. C. Saller and S. Habershon, *The Journal of Chemical Theory and Computation*, 2015, **11**, 8–16.

[58] B. Gu and S. Garashchuk, *J. Phys. Chem. A*, 2016, **120**, 3023–3031.

[59] C. W. Heaps and D. A. Mazziotti, *J. Chem. Phys.*, 2016, **145**, 064101.

[60] J. C. Light and T. Carrington, in *Advances in Chemical Physics*, John Wiley & Sons, Inc., 2007, pp. 263–310.

[61] U. Manthe, H.-D. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 1992, **97**, 3199–3213.

[62] G. A. Worth, H. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 1996, **105**, 4412–4426.

[63] *Multidimensional Quantum Dynamics: MCTDH Theory and Applications*, ed. H.-D. Meyer, F. Gatti and G. A. Worth, Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, 2009, pp. 1–7.

[64] H.-D. Meyer, *WIREs Comput. Mol. Sci.*, 2012, **2**, 351–374.

[65] A. Raab, G. A. Worth, H.-D. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 1999, **110**, 936 – 946.

[66] C. Cattarius, G. A. Worth, H.-D. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 2001, **115**, 2088–2100.

[67] H. Wang and M. Thoss, *J. Chem. Phys.*, 2003, **119**, 1289–1299.

[68] U. Manthe, *J. Chem. Phys.*, 2008, **128**, 164116.

[69] U. Manthe, *J. Chem. Phys.*, 2009, **130**, 054109.

[70] O. Vendrell and H.-D. Meyer, *J. Chem. Phys.*, 2011, **134**, 044135.

[71] I. Burghardt, H.-D. Meyer and L. S. Cederbaum, *J. Chem. Phys.*, 1999, **111**, 2927–2939.

[72] I. Burghardt, M. Nest and G. A. Worth, *J. Chem. Phys.*, 2003, **119**, 5364–5378.

[73] I. Burghardt, K. Giri and G. A. Worth, *J. Chem. Phys.*, 2008, **129**, 174104.

[74] G. Richings, I. Polyak, K. Spinlove, G. Worth, I. Burghardt and B. Lasorne, *Int. Rev. Phys. Chem.*, 2015, **34**, 269–308.

[75] G. A. Worth, M. A. Robb and B. Lasorne, *Mol. Phys.*, 2008, **106**, 2077.

[76] G. A. Worth, M. A. Robb and I. Burghardt, *Faraday Discuss.*, 2004, **127**, 307–323.

[77] B. Lasorne, M. J. Bearpark, M. A. Robb and G. A. Worth, *Chem. Phys. Lett.*, 2006, **432**, 604–609.

[78] B. Lasorne, M. A. Robb and G. A. Worth, *Phys. Chem. Chem. Phys.*, 2007, **9**, 3210–3227.

[79] C. W. Heaps and D. A. Mazziotti, *J. Chem. Phys.*, 2016, **144**, 164108.

[80] B. M. Garraway and K. A. Suominen, *Rep. Prog. Phys.*, 1995, **58**, 365–.

[81] S. Mallat and Z. Zhang, *IEEE Trans. Signal Process.*, 1993, **41**, 3397 –3415.

[82] Y. Wu and V. S. Batista, *J. Chem. Phys.*, 2003, **118**, 6720–6724.

[83] Y. Wu and V. S. Batista, *J. Chem. Phys.*, 2006, **124**, 224305.

[84] D. V. Shalashilin, *J. Chem. Phys.*, 2010, **132**, 244111.

[85] K. Saita and D. V. Shalashilin, *J. Chem. Phys.*, 2012, **137**, 22A506.

[86] M. Ben-Nun and T. J. Martínez, *J. Chem. Phys.*, 1998, **108**, 7244.

[87] L. Liu, J. Liu and T. J. Martínez, *J. Phys. Chem. B*, 2016, **120**, 1940–1949.

[88] B. F. E. Curchod, A. Sisto and T. J. Martínez, *J. Phys. Chem. A*, 2017, **121**, 265–276.

[89] D. V. Makhov, T. J. Martínez and D. V. Shalashilin, *Faraday Discuss.*, 2016, **194**, 81–94.

[90] J. C. Tully, *J. Chem. Phys.*, 1990, **93**, 1061–1071.

[91] D. F. Coker and L. Xiao, *J. Chem. Phys.*, 1995, **102**, 496–510.

[92] M. S. Topaler, M. D. Hack, T. C. Allison, Y.-P. Liu, S. L. Mielke, D. W. Schwenke and D. G. Truhlar, *J. Chem. Phys.*, 1997, **106**, 8699–8709.

[93] M. D. Hack and D. G. Truhlar, *J. Phys. Chem. A*, 2000, **104**, 7917–7926.

[94] L. Wang, A. Akimov and O. V. Prezhdo, *J. Phys. Chem. Lett.*, 2016, **7**, 2100–2112.

[95] Y. Gan, L. Yue, X. Guo, C. Zhu and Z. Cao, *Phys. Chem. Chem. Phys.*, 2017, **19**, 12094–12106.

[96] X. Zheng, G. Zhai, W. Gao, Y. Lei, L. Yu and C. Zhu, *Phys. Chem. Chem. Phys.*, 2016, **18**, 8971–8979.

[97] L. Colina, R. C. Bohlin and F. Castelli, *Astron. J.*, 1996, **112**, 307.

[98] E. Wigner, *Phys. Rev.*, 1932, **40**, 749–759.

[99] W. C. Swope, H. C. Andersen, P. H. Berens and K. R. Wilson, *J. Chem. Phys.*, 1982, **76**, 637–649.

[100] J. Butcher, *Appl. Numer. Math.*, 1996, **20**, 247–260.

[101] I. Yamazaki, T. Murao, T. Yamanaka and K. Yoshihara, *Faraday Discuss. Chem. Soc.*, 1983, **75**, 395–405.

[102] K. Innes, I. Ross and W. R. Moomaw, *J. Mol. Spectrosc.*, 1988, **132**, 492 – 544.

[103] G. Stock, C. Woywod, W. Domcke, T. Swinney and B. S. Hudson, *J. Chem. Phys.*, 1995, **103**, 6851–6860.

[104] Y.-I. Suzuki, T. Fuji, T. Horio and T. Suzuki, *J. Chem. Phys.*, 2010, **132**, 174302.

[105] M. Stener, P. Decleva, D. M. P. Holland and D. A. Shaw, *J. Phys. B*, 2011, **44**, 075203.

[106] R. Schneider and W. Domcke, *Chem. Phys. Lett.*, 1988, **150**, 235 – 242.

[107] C. Møller and M. S. Plesset, *Phys. Rev.*, 1934, **48**, 618–622.

[108] T. H. Dunning and P. J. Hay, in *Modern Theoretical Chemistry*, ed. H. F. Schaefer, Plenum, New York, 1977, vol. Vol. 3.

[109] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, J. A. Montgomery, Jr., T. Vreven, K. N. Kudin, J. C. Burant, J. M. Millam, S. S. Iyengar, J. Tomasi, V. Barone, B. Mennucci, M. Cossi, G. Scalmani, N. Rega, G. A. Petersson, H. Nakatsuji, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima,

Y. Honda, O. Kitao, H. Nakai, M. Klene, X. Li, J. E. Knox, H. P. Hratchian, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, P. Y. Ayala, K. Morokuma, G. A. Voth, P. Salvador, J. J. Dannenberg, V. G. Zakrzewski, S. Dapprich, A. D. Daniels, M. C. Strain, O. Farkas, D. K. Malick, A. D. Rabuck, K. Raghavachari, J. B. Foresman, J. V. Ortiz, Q. Cui, A. G. Baboul, S. Clifford, J. Cioslowski, B. B. Stefanov, G. Liu, A. Liashenko, P. Piskorz, I. Komaromi, R. L. Martin, D. J. Fox, T. Keith, M. A. Al-Laham, C. Y. Peng, A. Nanayakkara, M. Challacombe, P. M. W. Gill, B. Johnson, W. Chen, M. W. Wong, C. Gonzalez and J. A. Pople, *Gaussian 03, Revision D.02*, Gaussian, Inc., Wallingford, CT, 2004.

[110] C. Woywod, W. Domcke, A. L. Sobolewski and H. Werner, *J. Chem. Phys.*, 1994, **100**, 1400–1413.

[111] S. Krempl, M. Winterstetter, H. Plöhn and W. Domcke, *J. Chem. Phys.*, 1994, **100**, 926–937.

[112] L. Seidner, G. Stock, A. L. Sobolewski and W. Domcke, *J. Chem. Phys.*, 1992, **96**, 5298–5309.

[113] X. Chen and V. S. Batista, *J. Chem. Phys.*, 2006, **125**, 124313.

[114] P. Puzari, B. Sarkar and S. Adhikari, *J. Chem. Phys.*, 2006, **125**, 194316.

[115] D. V. Shalashilin, *J. Chem. Phys.*, 2010, **132**, 244111.

[116] U. Müller and G. Stock, *J. Chem. Phys.*, 1999, **111**, 77 – 88.

[117] S. Habershon and D. E. Manolopoulos, *J. Chem. Phys.*, 2009, **131**, 244518.

[118] P. A. Sherratt, D. V. Shalashillin and M. S. Child, *Chem. Phys.*, 2006, **322**, 127.

[119] S. Habershon, *J. Chem. Phys.*, 2012, **136**, 054109.

[120] I. R. Craig and D. E. Manolopoulos, *J. Chem. Phys.*, 2004, **121**, 3368–3373.

[121] S. Habershon, D. E. Manolopoulos, T. E. Markland and T. F. Miller, 3rd, *Annu. Rev. Phys. Chem.*, 2013, **64**, 387–413.

[122] M. A. C. Saller and S. Habershon, *J. Chem. Theory Comput.*, 2017, **13**, 3085–3096.

Appendices

I Gaussian wavepacket integration and overlap

For a general Gaussian basis function,

$$|g(\mathbf{r};t)\rangle = \prod_{j=1}^{f}\left[\left(\frac{2\alpha_j}{\pi\hbar}\right)^{\frac{1}{4}}\exp\left[-\alpha_j\frac{1}{\hbar}\left(r_j - r_j(t)\right)^2 + p_j(t)\frac{i}{\hbar}\left(r_j - r_j(t)\right)\right]\right] ,$$

the overlap

$$\langle g_i(\mathbf{r};t)\,|\,g_j(\mathbf{r};t)\rangle \ ,$$

may be determined as a product of one dimensional overlaps, such that

$$\langle g_i(\mathbf{r};t)\,|\,g_j(\mathbf{r};t)\rangle = \prod_{k=1}^{f}\langle g_i(r_k;t)\,|\,g_j(r_k;t)\rangle_k \ .$$

This in turn may be expanded as

$$\langle g_i(r_k;t)\,|\,g_j(r_k;t)\rangle_k = \left(\frac{2\alpha_i}{\pi\hbar}\right)^{\frac{1}{4}}\left(\frac{2\alpha_j}{\pi\hbar}\right)^{\frac{1}{4}} e^c \int e^{-ar_k^2+br_k}\,\mathrm{d}r_k \,,$$

where

$$a = \frac{1}{\hbar}\left(\alpha_i + \alpha_j\right) \ ,$$

$$b = \frac{1}{\hbar}\left(2\alpha_i r_k^{(i)} + 2\alpha_j r_k^{(j)} + i(p_k^{(j)} - p_k^{(i)})\right) \ ,$$

$$c = \frac{1}{\hbar}\left(-\alpha_i {r_k^{(i)}}^2 - \alpha_j {r_k^{(j)}}^2 + ip_k^{(i)} r_k^{(i)} - ip_k^{(j)} r_k^{(j)}\right) \ .$$

This may be analytically solved, as it constitutes the general Gaussian integral, such that

$$\langle g_i(r_k;t)\,|\,g_j(r_k;t)\rangle_k = \left(\frac{2\alpha_i}{\pi\hbar}\right)^{\frac{1}{4}}\left(\frac{2\alpha_j}{\pi\hbar}\right)^{\frac{1}{4}} e^c\sqrt{\frac{\pi}{a}}e^{\frac{b^2}{4a}} \ .$$

II Velocity Verlet algorithm

The Velocity Verlet algorithm integrates Newton's equations of motion

$$\frac{\mathrm{d}r_k}{\mathrm{d}t} = \frac{p_k}{m_k}\,,$$

$$\frac{\mathrm{d}p_k}{\mathrm{d}t} = -\frac{\mathrm{d}V(\mathbf{r})}{\mathrm{d}r_k}\,,$$

where r_k is the position, p_k the momentum and m_k the mass in DOF k and $V(\mathbf{r})$ is the potential energy.

The Velocity Verlet algorithm proceeds as follows

$$r_k(t+\delta t) = r_k + \delta t \frac{p_k(t)}{m_k} - \frac{1}{2}\delta t^2 \frac{\mathrm{d}V(\mathbf{r}(t))}{\mathrm{d}r_k}\,,$$

$$\frac{p_k(t+\frac{1}{2}\delta t)}{m_k} = \frac{p_k(t)}{m_k} - \frac{1}{2}\delta t \frac{\mathrm{d}V(\mathbf{r}(t))}{\mathrm{d}r_k}\,,$$

$$\frac{p_k(t+\delta t)}{m_k} = \frac{p_k(t+\frac{1}{2}\delta t)}{m_k} - \frac{1}{2}\delta t \frac{\mathrm{d}V(\mathbf{r}(t+\delta t))}{\mathrm{d}r_k}\,.$$

III Iteration using the 4^{th} order Runge-Kutta method

The iteration of the equations of motion both for the basis set expansion coefficients (see above) and the Ehrenfest coefficients (see below) are solved numerically using the 4^{th}-order Runge-Kutta algorithm.
For a problem where

$$\begin{aligned} \dot{y} &= f(y,t)\,, \\ y(t_0) &= y_0\,. \end{aligned}$$

For a given timestep δt,

$$\begin{aligned} y(t_{n+1}) \approx y_{n+1} &= y_n + \frac{1}{6}\delta t\left(k_1 + 2k_2 + 2k_3 + k_4\right)\,, \\ t_{n+1} &= t_n + \delta t\,, \end{aligned}$$

where

$$\begin{aligned} k_1 &= f\left(y_n, t_n\right)\,, \\ k_2 &= f\left(y_n + \frac{1}{2}\delta t k_1, t + \frac{1}{2}\delta t\right)\,, \\ k_3 &= f\left(y_n + \frac{1}{2}\delta t k_2, t + \frac{1}{2}\delta t\right)\,, \\ k_4 &= f\left(y_n + \delta t k_3, t + \delta t\right)\,. \end{aligned}$$

IV Scaling of the trajectory sampling algorithm with basis set size

Table A.1: Number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 3.9, the resulting average basis set size, $\overline{N}_{total}$ and average errors with corresponding standard deviations.

m	$\overline{N}_{total}$	$\overline{\mathrm{MAE}}$	$\sigma_{\overline{\mathrm{MAE}}}$	$\overline{\mathrm{MAPE}}\,(\%)$	$\sigma_{\overline{\mathrm{MAPE}}}\,(\%)$
33	1008	0.0660	0.0021	10.1	0.7
67	1980	0.0520	0.0018	7.83	0.10
100	3028	0.0474	0.0032	6.89	0.56
133	3917	0.0422	0.0035	6.03	0.53
167	5043	0.0370	0.0048	5.18	0.70
200	6008	0.0357	0.0026	4.94	0.38
233	7044	0.0302	0.0031	4.05	0.50
267	7998	0.0317	0.0042	4.16	0.54
300	8989	0.0288	0.0027	3.79	0.31
333	9965	0.0312	0.0037	3.97	0.44
367	11019	0.0340	0.0080	4.30	0.96
400	11906	0.0273	0.0035	3.50	0.40
433	13005	0.0285	0.0033	3.58	0.39
467	13943	0.0248	0.0020	3.12	0.24
500	15076	0.0220	0.0027	2.81	0.33
533	15994	0.0216	0.0024	2.75	0.31
567	17021	0.0169	0.0021	2.18	0.27
600	18099	0.0209	0.0031	2.66	0.39
633	18977	0.0163	0.0018	2.09	0.23
667	20055	0.0160	0.0021	2.06	0.25
700	20994	0.0133	0.0013	1.72	0.17
733	22074	0.0125	0.0032	1.62	0.41
767	23175	0.00980	0.00110	1.28	0.16
800	24005	0.00962	0.00074	1.25	0.10

V Scaling of the trajectory sampling algorithm with basis function sampling frequency

Table A.2: Number of sampling trajectories, m, and sampling frequency parameters, n_s, used in sets of 4 calculations to obtain the data shown in Figure 3.10, the resulting average basis set size, $\overline{N}_{total}$ and average errors with corresponding standard deviations.

m	n_s	$\overline{N}_{total}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
12	1	18028	0.0404	0.0019	5.57	0.26
60	5	18104	0.0186	0.0034	2.39	0.46
120	10	17998	0.0178	0.0014	2.33	0.18
180	15	17939	0.0156	0.0035	2.02	0.43
240	20	18061	0.0180	0.0029	2.32	0.33
300	25	18099	0.0159	0.0010	2.04	0.11
360	30	17922	0.0170	0.0030	2.17	0.37
480	40	17938	0.0181	0.0031	2.30	0.38
600	50	18015	0.0200	0.0021	2.56	0.25
900	75	18150	0.0189	0.0015	2.41	0.18
1200	100	18111	0.0166	0.0015	2.13	0.18
3000	250	17900	0.0200	0.0015	2.54	0.18
6000	500	18064	0.0187	0.0016	2.39	0.21
9000	750	17914	0.0222	0.0017	2.81	0.23
12000	1000	17948	0.0217	0.0014	2.74	0.18
18000	1500	17996	0.0223	0.0023	2.83	0.30

VI Scaling of the trajectory sampling algorithm with timestep duration and "over-" and "undersampling"

Table A.3: Sampling timestep duration, Δt_t, number of sampling timesteps ,n_t, and number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 3.11, the resulting average basis set size, $\overline{N}_{total}$ and average errors with corresponding standard deviations.

Δt_t/fs	n_t	m	$\overline{N}_{total}$	$\overline{\mathrm{MAE}}$	$\sigma_{\overline{\mathrm{MAE}}}$	$\overline{\mathrm{MAPE}}\,(\%)$	$\sigma_{\overline{\mathrm{MAPE}}}\,(\%)$
0.001	150000	6	17939	0.0454	0.0067	6.45	1.10
0.010	15000	60	18113	0.0215	0.0021	2.77	0.25
0.025	6000	150	17947	0.0170	0.0018	2.20	0.23
0.050	3000	300	17988	0.0150	0.0021	1.93	0.26
0.075	2000	450	18046	0.0195	0.0036	2.50	0.43
0.100	1500	600	17996	0.0171	0.0014	2.19	0.18
0.150	900	1000	17927	0.0144	0.0019	1.86	0.26
0.200	750	1200	18130	0.0198	0.0020	2.51	0.23
0.500	300	3000	18048	0.0167	0.0010	2.13	0.12
1.000	150	6000	17951	0.0170	0.0018	2.17	0.23

Table A.4: Number of sampling timesteps ,n_t, and number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 3.12, the resulting average basis set size, $\overline{N}_{total}$ and average errors with corresponding standard deviations.

n_t	m	$\overline{N}_{total}$	$\overline{\mathrm{MAE}}$	$\sigma_{\overline{\mathrm{MAE}}}$	$\overline{\mathrm{MAPE}}\,(\%)$	$\sigma_{\overline{\mathrm{MAPE}}}\,(\%)$
500	1800	17936	0.0829	0.0106	10.3	1.3
1000	900	17943	0.0162	0.0022	2.03	0.26
1250	720	18092	0.0191	0.0011	2.43	0.14
1500	600	18077	0.0209	0.0023	2.63	0.30
1750	515	17970	0.0194	0.0027	2.49	0.33
2000	450	18040	0.0172	0.0022	2.21	0.29
2500	360	17890	0.0167	0.0026	2.20	0.36
3000	300	17970	0.0081	0.0012	1.09	0.15
4000	225	17936	0.0084	0.0018	1.14	0.25
5000	180	18073	0.0110	0.0024	1.48	0.32
6000	150	18020	0.0122	0.0017	1.62	0.23
7500	120	17834	0.0161	0.0040	2.19	0.55
10000	90	18024	0.0132	0.0022	1.78	0.29

VII Scaling of the adaptive sampling algorithm with basis set size and MP convergence criterion, ζ

Table A.5: Number of sampling trajectories, m, used in sets of 4 calculations with $\zeta = 0.95$ to obtain the data shown in Figure 4.11, the resulting average basis set sizes and average errors with corresponding standard deviations.

m	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
100	1711	995	716	0.0358	0.0023	6.48	0.39
200	2365	1999	366	0.0185	0.0030	3.59	0.54
300	3231	2998	233	0.0189	0.0014	3.31	0.24
400	4182	3993	189	0.0192	0.0026	3.13	0.40
500	5161	4992	169	0.0178	0.0023	2.75	0.36
600	6136	5995	141	0.0189	0.0008	3.07	0.09
700	7115	6986	129	0.0199	0.0020	3.16	0.36
800	8122	8000	122	0.0178	0.0016	2.74	0.27
900	9149	9030	119	0.0171	0.0018	2.70	0.26
1000	10123	10002	121	0.0184	0.0014	2.86	0.21

Table A.6: Number of sampling trajectories, m, used in sets of 4 calculations with $\zeta = 0.97$ to obtain the data shown in Figure 4.11, the resulting average basis set sizes and average errors with corresponding standard deviations.

m	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\mathrm{MAE}}$	$\sigma_{\overline{\mathrm{MAE}}}$	$\overline{\mathrm{MAPE}}\,(\%)$	$\sigma_{\overline{\mathrm{MAPE}}}\,(\%)$
50	1171	501	670	0.0519	0.0029	9.98	0.94
100	2503	996	1507	0.0399	0.0049	7.20	1.01
200	2941	1994	947	0.0210	0.0017	3.74	0.30
300	3715	3005	710	0.0187	0.0015	3.04	0.10
400	4564	3989	575	0.0188	0.0010	2.95	0.18
500	5521	5007	514	0.0174	0.0017	2.71	0.30
600	6443	5991	452	0.0164	0.0003	2.55	0.08
700	7393	7013	380	0.0163	0.0012	2.51	0.17
800	8318	7983	335	0.0152	0.0023	2.31	0.36
900	9291	8968	323	0.0156	0.0010	2.40	0.14
1000	10310	10009	301	0.0142	0.0005	2.17	0.09

VIII Scaling of the adaptive sampling algorithm with restart frequency

Table A.7: Number instances of the aTSA, N_b, number of sampling and propagation timesteps, n_t and n_p respectively, and number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 4.13, the resulting average basis set sizes and average errors with corresponding standard deviations.

N_b	n_t, n_p	m	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
60	25	1200	3005	2997	8	0.0352	0.0032	5.68	0.55
30	50	600	3046	3010	36	0.0208	0.0022	3.84	0.48
20	75	400	3090	2999	91	0.0229	0.0011	4.10	0.32
15	100	300	3235	3007	228	0.0203	0.0029	3.47	0.50
8	200	150	4770	3426	1344	0.0321	0.0049	5.02	0.76
5	300	100	4222	2996	1226	0.0386	0.0024	5.59	0.39
4	400	75	5559	4012	1547	0.0336	0.0017	4.91	0.21
3	500	60	3876	3009	867	0.0271	0.0035	4.02	0.52

Table A.8: Number instances of the aTSA, N_b, number of sampling and propagation timesteps, n_t and n_p respectively, and number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 4.14, the resulting average basis set sizes and average errors with corresponding standard deviations.

N_b	n_t, n_p	m	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
60	25	3600	9011	9003	8	0.0351	0.0018	5.87	0.42
30	50	1800	9013	8981	32	0.0212	0.0023	3.90	0.35
20	75	1200	9078	9005	73	0.0192	0.0017	3.29	0.30
15	100	900	9120	9001	119	0.0179	0.0019	2.78	0.27
8	200	450	10913	10327	586	0.0153	0.0022	2.21	0.24
5	300	300	10066	9022	1044	0.0183	0.0021	2.44	0.26
4	400	225	13212	11959	1253	0.0098	0.0005	1.29	0.07
3	500	180	9740	9028	712	0.0088	0.0005	1.25	0.06

IX Scaling of the adaptive sampling algorithm with respect to "oversampling"

Table A.9: Number of sampling timesteps, n_t, and number of sampling trajectories, m, used in sets of 4 calculations to obtain the data shown in Figure 4.15, the resulting average basis set sizes and average errors with corresponding standard deviations.

n_t	m	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
100	900	9119	8993	125	0.0185	0.0014	2.92	0.23
120	750	9103	8988	115	0.0186	0.0015	2.91	0.24
140	643	9086	8973	113	0.0187	0.0005	2.93	0.09
160	563	9118	9004	113	0.0189	0.0017	2.98	0.26
180	500	9114	9002	112	0.0183	0.0013	2.88	0.22
200	450	9129	9017	112	0.0177	0.0009	2.76	0.14
300	300	9142	9009	132	0.0168	0.0003	2.64	0.06
400	225	9147	8993	153	0.0145	0.0018	2.22	0.29
500	180	9138	8973	164	0.0160	0.0030	2.54	0.53
1000	90	9276	9004	271	0.0111	0.0014	1.81	0.19
2000	45	9379	9005	374	0.0133	0.0024	2.16	0.42
3000	30	9536	8975	561	0.0084	0.0017	1.39	0.25

X Scaling of the adaptive sampling algorithm with the basis library expansion parameter, γ

Table A.10: Basis set expansion parameter, γ, used in sets of 4 calculations to obtain the data shown in Figure 4.17, the resulting average basis set sizes and average errors with corresponding standard deviations.

γ	$\overline{N}_{total}$	$\overline{N}_t$	$\overline{N}_{MP}$	$\overline{\text{MAE}}$	$\sigma_{\overline{\text{MAE}}}$	$\overline{\text{MAPE}}\,(\%)$	$\sigma_{\overline{\text{MAPE}}}\,(\%)$
1.0	3257	2992	264	0.0177	0.0012	3.06	0.29
1.5	3271	3012	258	0.0212	0.0026	3.46	0.49
2.0	3254	2999	255	0.0177	0.0011	3.06	0.24
2.5	3234	3008	226	0.0189	0.0022	3.22	0.33
3.0	3248	3009	239	0.0205	0.0039	3.60	0.58
3.5	3247	3003	244	0.0202	0.0039	3.36	0.53
4.0	3240	3011	229	0.0209	0.0025	3.46	0.49
4.5	3244	2997	247	0.0196	0.0030	3.16	0.50
5.0	3238	2997	240	0.0196	0.0033	3.35	0.58