

Low-dimensional systems: a quantum Monte Carlo study



David Martin Thomas

Department of Physics

Lancaster University

This thesis is submitted for the degree of

Doctor of Philosophy

September 2021

Abstract

We study low-dimensional materials and devices through use of the variational and diffusion quantum Monte Carlo methods. Firstly, we use models of nanostructures in semiconductor heterostructures that confine charge-carriers in one (or more) dimensions to investigate the energetics of the charge-carrier complexes that form in such structures. For type-II quantum rings and superlattices, we present energy data to aid in experimental identification of these complexes and show that these energies are relatively insensitive to the geometrical dimensions of the devices.

Secondly, we study similar models of charge-carrier complexes but this time where the confinement is provided by the two-dimensional nature of the material, rather than by artificial construction. Application of an in-plane electric field shifts the binding energies of complexes in monolayer transition metal dichalcogenides such that charged complexes can be identified from neutral ones. The truly two-dimensional character of these materials results in a Keldysh interaction between charge-carriers, rather than a screened Coulomb interaction. In such materials, modelling the two-dimensional electron gas using a more realistic Keldysh interaction acts to lower the Wigner crystallisation density, when compared to using a Coulomb interaction.

Thirdly, and finally, we perform *ab-initio* calculations of the defect formation energy for mono-vacancies in graphene, with the aim of benchmarking the accuracy of the widely-used density functional theory method in these types of calculation. The mono-vacancy defect formation energy is shown to be significantly underestimated by density functional theory.

Declaration

The content of this thesis describes work carried out in the Department of Physics at Lancaster University between April 2018 and September 2021, as a member of the Graphene NOWNANO Centre for Doctoral Training, and under the supervision of Dr. N. D. Drummond. The following sections of this thesis are included in work that has been published, or is in preparation to be submitted for publication:

Chapter 2.2 D. M. Thomas, R. J. Hunt, N. D. Drummond, and M. Hayne, “Binding energies of excitonic complexes in type-II quantum rings from diffusion quantum Monte Carlo calculations,” *Phys. Rev. B*, vol. 99, p. 115306, 2019.

Chapter 3.3 E. Mostaani, R. J. Hunt, D. M. Thomas, M. Szyniszewski, A. R. P. Montblanch, M. Barbone, M. Atatüre, N. D. Drummond, and A. C. Ferrari, “Charge-carrier complexes in monolayer semiconductors,” in preparation.

Chapter 3.4 R. J. Hunt, D. M. Thomas, and N. D. Drummond, “Quantum Monte Carlo study of Wigner crystallization in weakly doped monolayer MoSe₂,” in preparation.

Chapter 4 D. M. Thomas, Y. Asiri, and N. D. Drummond, “Defect formation energies in graphene from diffusion quantum Monte Carlo and density functional theory,” in preparation.

I declare that this thesis is my own work, and has not been submitted in substantially the same form for the award of a higher degree elsewhere. This thesis does not exceed the maximum word count of 80,000 words.

David Martin Thomas

Lancaster, September 2021

Acknowledgements

I would like to thank my supervisor, Neil Drummond, for his support and guidance throughout my time at Lancaster. His knowledge, willingness to have discussions, and his timely responses have been invaluable to me, and for which I am particularly grateful.

My thanks also goes to Ryan Hunt, Yasmin Asiri, Marcin Szyniszewski, Manus Hayne, and the rest of the Condensed Matter theory group for their collaborations and discussions. Special thanks are owed to Ryan Hunt for his good humour and discussions beyond our common scientific interests.

Last, but by no means least, I thank Rachel Hendrikse for the good times we have shared in the last few years, and for listening to my endless ramblings regarding the progress of this project.

This work has been financially supported by the Engineering and Physical Sciences Research Council (EPSRC), under grant number EP/L01548X/1 (Graphene NOWNANO CDT). Computational facilities have been provided by the Lancaster High End Computing (HEC) cluster, and by the Cambridge Service for Data Driven Discovery (CSD3).

Table of Contents

Abstract	ii
Declaration	iii
Acknowledgements	v
List of Figures	x
List of Tables	xv
Introduction	1
1 Introduction to quantum Monte Carlo	3
1.1 The quantum many-body problem	3
1.2 Statistical foundations	8
1.2.1 Monte Carlo integration	8
1.2.2 Metropolis sampling algorithm	10
1.2.3 Serial correlation	12
1.3 Variational Monte Carlo	13
1.3.1 Expectation values	13
1.3.2 Trial wave function	16
1.3.3 Optimisation of the trial wave function	23
1.4 Diffusion Monte Carlo	25
1.4.1 Imaginary-time Schrödinger equation	25
1.4.2 Imaginary-time evolution	26

1.4.3	Importance sampling	28
1.4.4	Fixed-phase approximation	31
1.4.5	Expectation values	33
1.4.6	Time-step and population-control biases	34
2	Excitonic complexes in semiconductor nanostructures	37
2.1	Models of excitonic complexes in QMC	37
2.1.1	Effective mass approximation	37
2.1.2	Excitonic complexes	40
2.1.3	Excitonic units	42
2.2	Type-II quantum rings	43
2.2.1	Quantum rings	43
2.2.2	The model	44
2.2.3	Computational details	46
2.2.4	Results and discussion	50
2.2.5	Conclusion	57
2.3	Type-II superlattices	58
2.3.1	Superlattices	58
2.3.2	Computational details	59
2.3.3	Results and discussion	62
2.3.4	Conclusion	65
3	Monolayer transition metal dichalcogenides	66
3.1	Material properties	66
3.2	The Keldysh interaction	68
3.3	Electric fields	69
3.4	The two-dimensional electron gas	75
3.4.1	Ewald interaction	76
3.4.2	Wigner crystallisation	78

4	Mono-vacancies in graphene	84
4.1	Point defects in graphene	84
4.2	Computational methodology	86
4.2.1	Defect formation energy	86
4.2.2	Migration energy barrier	87
4.2.3	Twisted periodic boundary conditions	87
4.2.4	DFT calculations	89
4.2.4.1	Total energy, geometry optimisation, and phonon calculations	89
4.2.4.2	QMC orbital generation	90
4.2.4.3	Transition state calculations	91
4.2.5	QMC calculations	91
4.2.5.1	Trial wave functions	91
4.2.5.2	DMC calculations	92
4.2.6	Finite-concentration and finite-size effects	93
4.2.6.1	Periodic supercells	93
4.2.6.2	Twist averaging	95
4.2.6.3	Long-range effects	98
4.2.7	Backflow	100
4.3	Results and discussion	100
4.3.1	Atomic structures	100
4.3.1.1	Pristine graphene	100
4.3.1.2	MV	101
4.3.1.3	MV diffusion transition state	103
4.3.2	Defect formation energy	104
4.3.3	Migration energy barrier	107
4.4	Conclusion	108
	Conclusion	110

Bibliography	113
A Importance sampling transformation of the imaginary-time Schrödinger equation	136

List of Figures

2.1	Model of the quantum ring shown in (a) cross-section and (b) plan view. The parameters R_z , r_i , and r_o represent the half height, the inner radius, and the outer radius of the ring, respectively.	45
2.2	The nine different regions of space (1–9) used to construct an electron orbital with a continuous first derivative around a quantum ring with a rectangular cross section. The parameters R_z , r_i , and r_o represent the half height, the inner radius, and the outer radius of the ring, respectively.	47
2.3	Ground-state total energies per hole of a single hole (h^+), two holes ($2h^+$), an exciton (X), a negative trion (X^-), a positive trion (X^+), and a biexciton (XX) in a quantum ring plotted against the aspect ratio $2R_z/(r_o - r_i)$ of the ring's cross section. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The exciton Rydberg R_y^* is 4.45 meV for the experimentally relevant geometry.	51
2.4	Expected peak positions for the excitonic complexes in a photoluminescence spectrum relative to the exciton peak, for a model of the quantum rings reported in Ref. [5]. The peak heights represent the relative stability of the complexes.	53

2.5	De-excitonisation energies against the aspect ratio $2R_z/(r_o - r_i)$ of a quantum ring's cross-section for different charge-carrier complexes. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The dashed line shows the experimentally relevant aspect ratio [5].	54
2.6	Binding energies against the aspect ratio $2R_z/(r_o - r_i)$ of a quantum ring's cross-section for different charge-carrier complexes. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The dashed line shows the experimentally relevant aspect ratio [5].	55
2.7	Electronic charge density ρ for (a) an exciton, (b) a negative trion, (c) a positive trion, and (d) a biexciton in the experimentally relevant quantum-ring geometry [5]. The shaded regions represent the ring and ρ_m is the maximum density across all four plots. The free exciton Bohr radius is $a_0^* = 12.5$ nm.	56
2.8	Diagrams showing (a) the superlattice (SL) and (b) the multiple quantum well (MQW) heterostructure models. The green coloured layers are the InAs electron-bearing layers and the red coloured layers are the InAsSb hole-bearing layers. The black and white circles represent electrons and holes, respectively. Hence, (a) shows the negative quinton $XX^-(e h e h e)$ and (b) shows the biexciton $XX(e hh e)$	60
3.1	Diagram illustrating the conduction band minimum and valence band maximum at the \mathbf{K} (left) and \mathbf{K}' (right) valleys in (a) MoX_2 and (b) WX_2 materials. Spin-up bands are red, and spin-down bands are blue. At room temperature the lower spin-split valence band is not expected to contain any holes.	67

3.2	DMC binding energy shift for (a) X, (b) XX, and (c) D ⁰ as a function of F^2 for different monolayer TMDCs in vacuum and encapsulated by hBN. Error bars in (a) and (c) are smaller than the symbols. The solid and dashed lines show the binding energies determined by the polarisabilities in Table 3.2, for monolayer TMDCs in vacuum and encapsulated by hBN. The vertical dotted lines correspond to $F = 50 \text{ mV nm}^{-1}$, beyond which VMC energy minimisation does not result in bound-state wave functions.	72
3.3	DMC binding energy shift for (a) X ⁻ , and (b) X ⁺ as a function of F^2 for different monolayer TMDCs in vacuum and encapsulated by hBN. Where error bars are not visible they are smaller than the symbols. The solid and dashed lines show the binding energies determined by the polarisabilities in Table 3.2, for monolayer TMDCs in vacuum and encapsulated by hBN.	74
3.4	Diagram illustrating the conduction band filling in both the K (left) and K' (right) valleys in our model of doped MoSe ₂ . Spin-up bands are red, and spin-down bands are blue, the Fermi energy E_F is shown by the orange line with the shading representing the filling of the bands.	79
3.5	Triangular lattice used in the Wigner crystal calculations. In a ferromagnetic Wigner crystal all sites are occupied by spin-up, K valley electrons. In the antiferromagnetic Wigner crystal spin-up, K valley electrons occupy the sites marked with red circles, while spin-down, K' valley electrons occupy the sites marked with blue squares.	80

3.6	Phase diagram of a two-dimensional Keldysh electron gas showing the paramagnetic fluid phase, and the ferromagnetic and antiferromagnetic Wigner crystal phases. The vertical dashed line shows the crystallisation density. The screening length is $r_* = 1.7858 a_0^*$ and corresponds to MoSe ₂ encapsulated by hBN.	82
4.1	DMC total energy per supercell of a MV defect in a 3×3 supercell of graphene against DMC time step $\delta\tau$ at a single, randomly chosen twist \mathbf{k}_s	92
4.2	DMC pure formation energy of a MV defect in a 3×3 supercell of graphene against DMC time step $\delta\tau$ at the twist \mathbf{k}_s used in Fig. 4.1.	93
4.3	Twist-blocked standard error in the twist-averaged DMC pure MV formation energy in a 3×3 supercell against the number of blocks into which the 24 original twists are divided. The standard error in the single-block case is obtained by Gaussian propagation of the Monte Carlo random errors. Standard errors obtained with small numbers of large blocks are relatively unbiased, but suffer significant noise; on the other hand, standard errors obtained with large numbers of small blocks exhibit less noise but are potentially biased.	97
4.4	DFT pure formation energy of a MV defect in graphene against the reciprocal of supercell size N_s . Fine \mathbf{k} -point grids were used in each supercell.	99
4.5	(a) Top-down and (b) in-plane views of the DFT-PBE-relaxed MV structure in a 5×5 supercell. The under-coordinated carbon atom is shown in red.	103
4.6	(a) Top-down and (b) in-plane views of the MV diffusion transition state structure in a 3×3 supercell. The migrating atom is shown in red.	104

4.7 Comparison of different methods for dealing with momentum quantisation errors in both DFT and DMC. Plots are of pure MV defect formation energy against reciprocal of supercell size N_s . Red dashed line show the unweighted least-squares fit to the TB-DMC mean energy data. 105

List of Tables

2.1	De-excitonisation E_D and binding E_b energies for excitonic complexes in the quantum-ring geometry modelling the samples described in Ref. [5].	52
2.2	Temperatures up to which the complexes are stable, in the quantum-ring geometry modelling the samples described in Ref. [5], using the binding energy data from Table 2.1.	54
2.3	Sensitivity of the biexciton binding energy to the electron and hole effective masses m_e^* and m_h^* , the relative permittivity ϵ , the ring volume V , and the mean radius of the ring r_m	56
2.4	DMC binding energies of negative trions (X^-), positive trions (X^+), and biexcitons (XX) in both the SL and MQW geometries. The third row of the table indicates the layers to which the particles in each complex belong; <i>e.g.</i> , XX(e hh e) represents a biexciton with a layer containing one electron, followed by a layer containing two holes, then a layer containing the second electron.	62

2.5	DMC binding energies of a variety of quintons, both negative (XX^-) and positive (XX^+), in both the SL and MQW geometries, where the daughter products are either an exciton and a trion or a biexciton and a single charge carrier. The third row of the table indicates the layers to which the particles in each complex belong; <i>e.g.</i> , $XX^-(e h e h e)$ represents a negative quinton in which each particle occupies a separate layer. Empty entries represent complexes that were unbound in our calculations.	63
2.6	Sensitivity of the SL complex binding energies E_b to the electron and hole effective masses m_e^* and m_h^* , the relative permittivity ϵ , and the electron and hole well widths W_e and W_h	64
3.1	Material parameters for different monolayer TMDCs in vacuum. m_e^* and m_h^* are the electron and hole effective masses, respectively, and r_* is the Keldysh screening parameter for the material in vacuum. r_* should be scaled down by $\epsilon = 4\epsilon_0$ when encapsulated by hBN. . .	72
3.2	Theoretical polarisabilities of X, XX, D^0 , X^- , and X^+ in monolayer TMDCs both in vacuum and encapsulated by hBN.	73
3.3	Calculated binding energy shifts of X, XX, D^0 , X^- , and X^+ using Eq. (3.13) and the polarisabilities in Table 3.2 for monolayer TMDCs in vacuum, and encapsulated by hBN, for $F = 50 \text{ mV nm}^{-1}$. Note, not all complexes remain bound at this electric field strength.	75
4.1	Point groups of relaxed graphene MV structures obtained using non-spin-polarised DFT calculations with different functionals and supercell sizes. Also shown are the energies of the most stable D_{3h} and C_{2v} structures relative to the most stable non-spin-polarised C_s structure. Spin-polarised and non-spin-polarised results are labelled by “(p)” and “(u),” respectively. Empty entries represent structures that were not found to be stable.	102

4.2	Theoretical static-nucleus defect formation energies, together with experimental results, for the MV defect in monolayer graphene. The carbon chemical potential is the energy per atom of graphene. Results without citations were obtained in the present work. Our DMC-corrected DFT result is obtained by taking our DFT-PBE result and applying the mean DMC correction in Table 4.3. Cited DFT works are performed at finite supercell size with no attempt to handle finite-size effects.	106
4.3	DMC corrections to the static-nucleus DFT MV defect formation energy at various supercell sizes. Corrections are evaluated as the average of the difference between energies calculated with TB-DMC and DFT using a fine \mathbf{k} -point grid at different supercell sizes.	107

Introduction

The last century has seen a great many advances in the field of condensed matter physics. The theory of quantum mechanics has been of fundamental importance to many of these advances. There has been much debate over the physical interpretations of the theory but the fact remains: quantum mechanics correctly predicts the properties of a vast range of physical systems. This success has led to the birth of the field of nanotechnology, which in turn has led to the fabrication of many of the electronic devices we see today.

The isolation of graphene from graphite in 2004 [6] only served to further the progress in this area. It led to a boom in the research of low-dimensional materials as, for the first time, such materials were more than just convenient models to study and further develop the theory of condensed matter. Low-dimensional semiconductors became a focal point in the pursuit of ever smaller and more portable electronic devices but also in a host of potential new and exciting applications.

Clearly then, a good understanding of the quantum mechanical properties of low-dimensional systems is crucial in unlocking the full potential of these new materials. The aforementioned development of quantum mechanics has led to a host of new methods with which we can study quantum systems and extract useful information. It is this, in combination with the staggering improvements in the capabilities of the modern computer (which is inextricably linked to our increased understanding of quantum mechanics), that has led to emergence of the quantum Monte Carlo techniques in which this thesis is heavily rooted.

In this work, we harness the power of quantum Monte Carlo methods to study

a range of low-dimensional devices and materials that show interesting electronic properties. To begin with, we study model systems in novel semiconductor nanostructures, as well as in two-dimensional semiconductors, to investigate the optoelectronic properties of such devices and materials. We then move on to investigate the possibility of Wigner crystallisation in a two-dimensional semiconductor, and a study of the properties of defects in graphene itself.

The content of this thesis is organised as follows. In Chapter 1 we give an introduction to the theory of the variational and diffusion quantum Monte Carlo methods. In Chapter 2 we study models of interacting charge-carriers in semiconductor nanostructures embedded in three-dimensional heterostructures, namely quantum rings and superlattices. We start Chapter 3 with a discussion of the same charge-carrier models but this time where the host material is the family of two-dimensional materials known as transition metal dichalcogenides. The remainder of Chapter 3 is devoted to a study of Wigner crystallisation in a two-dimensional electron gas. Finally, in Chapter 4, we perform *ab-initio* calculations to find the defect formation energy of mono-vacancies in graphene. Hartree atomic units ($e = m_e = \hbar = 4\pi\epsilon_0 = 1$) will be used throughout, unless otherwise stated. Here, e is the magnitude of the electron charge, m_e is the free electron mass, \hbar is the Dirac constant, and ϵ_0 is the permittivity of free space.

Chapter 1

Introduction to quantum Monte

Carlo

1.1 The quantum many-body problem

Since Erwin Schrödinger first published his papers on wave mechanics in 1926 [7–10] it has become widely accepted amongst the scientific community that the answers to most of the problems in condensed matter physics (amongst other subjects) can be obtained by solving the many-body Schrödinger equation. It turns out however, that this linear partial differential equation can only be solved analytically for very few, small, and mostly non-interacting systems, such as the hydrogen atom or the quantum harmonic oscillator; the search for (approximate) solutions in more complicated problems is known as the quantum many-body problem. Before we go on to discuss methods of approximately solving the Schrödinger equation, we will first introduce the equation and the many-body Hamiltonian. The non-relativistic time-independent Schrödinger equation can be written as

$$\hat{H}\Psi(\mathbf{R}, \mathbf{N}) = E\Psi(\mathbf{R}, \mathbf{N}), \quad (1.1)$$

where $\Psi(\mathbf{R}, \mathbf{N})$ is a many-body wave function of N_e electrons and N_N nuclei and the coordinate vectors are $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_{N_e})$, $\mathbf{N} = (\mathbf{n}_1, \dots, \mathbf{n}_{N_N})$ with \mathbf{r}_i and \mathbf{n}_I

the coordinates of the $i(I)$ -th electron (nucleus). The many-body Hamiltonian \hat{H} takes the form

$$\begin{aligned} \hat{H} = & -\frac{1}{2} \sum_{i=1}^{N_e} \nabla_{\mathbf{r}_i}^2 - \frac{1}{2} \sum_{I=1}^{N_N} \frac{1}{M_I} \nabla_{\mathbf{n}_I}^2 \\ & + \frac{1}{2} \sum_{i=1}^{N_e} \sum_{\substack{j=1 \\ i \neq j}}^{N_e} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{I=1}^{N_N} \sum_{\substack{J=1 \\ I \neq J}}^{N_N} \frac{Z_I Z_J}{|\mathbf{n}_I - \mathbf{n}_J|} \\ & - \sum_{i=1}^{N_e} \sum_{I=1}^{N_N} \frac{Z_I}{|\mathbf{r}_i - \mathbf{n}_I|}, \end{aligned} \quad (1.2)$$

where M_I is the mass, and Z_I the charge, on the I -th nucleus. The first two terms of Eq. (1.2) describe the kinetic energies of the electrons and the nuclei, the terms on the second line describe the Coulomb electron-electron and nucleus-nucleus interactions, while the final term describes the electron-nucleus Coulomb interactions.

It is common to now make use of the Born-Oppenheimer approximation [11] in which we argue that because the masses of the nuclei are (at least) three orders of magnitude larger than that of the electrons, the timescales in which the nuclei move are much longer than for the electrons. Therefore, the electronic response time to changes in nuclear positions can be considered to be immediate. Making use of this approximation allows us to approximately decouple the nuclear and electronic motion, resulting in a simpler electronic Schrödinger equation

$$\left[-\frac{1}{2} \sum_{i=1}^{N_e} \nabla_{\mathbf{r}_i}^2 + \frac{1}{2} \sum_{i=1}^{N_e} \sum_{\substack{j=1 \\ i \neq j}}^{N_e} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \frac{1}{2} \sum_{i=1}^{N_e} \sum_{I=1}^{N_N} \frac{Z_I}{|\mathbf{r}_i - \mathbf{n}_I|} + C \right] \Psi(\mathbf{R}) = E \Psi(\mathbf{R}), \quad (1.3)$$

where the nuclear kinetic energy is removed (decoupled), the nucleus-nucleus interaction energy is now just a constant offset C , and the positions of the nuclei only appear as parameters. Vibrational effects of the nuclei can then be included as a correction. This Hamiltonian is often simplified further by making use of pseudopotentials to replace the interaction between the electrons and the nuclei

and even some core electrons; see Sec. 1.3.2 for more details. From now on we will just focus on the electronic problem and in doing so relabel N_e as N .

As we turn our attention to the problem of how to solve this equation for large interacting systems, it is first instructive to think about the hydrogen atom. The hydrogen atom (or hydrogen-like atoms), consisting of one atomic nucleus and a single electron, has an analytical solution to the Schrödinger equation $\psi_H(\mathbf{r})$. Consider now the simple approach of including a second electron (in its own orbital) by assuming it does not interact with the first electron; a two-body spatial wave function for this system can be written as $\Psi(\mathbf{r}_1, \mathbf{r}_2) = \psi_H(\mathbf{r}_1)\psi_H(\mathbf{r}_2)$. This form of wave function however has a major flaw that prevents us from just scaling this up to a product of N single-particle orbitals. When we introduce spin, we find that a simple product like this does not satisfy the fermionic anti-symmetry requirements of the Pauli exclusion principle. For spin-independent Hamiltonians (such as the one described above), we can simplify this condition to an anti-symmetry on the spatial wave function when exchanging same-spin electrons, and treat different-spin electrons as distinguishable. The Schrödinger equation is linear, so we can write down a linear combination of these solutions to give a new wave function

$$\begin{aligned}\Psi_S(\mathbf{r}_1, \mathbf{r}_2) &= \frac{1}{\sqrt{2!}} [\psi_{H_1}(\mathbf{r}_1)\psi_{H_2}(\mathbf{r}_2) - \psi_{H_1}(\mathbf{r}_2)\psi_{H_2}(\mathbf{r}_1)] \\ &= \frac{1}{\sqrt{2!}} \begin{vmatrix} \psi_{H_1}(\mathbf{r}_1) & \psi_{H_2}(\mathbf{r}_1) \\ \psi_{H_1}(\mathbf{r}_2) & \psi_{H_2}(\mathbf{r}_2) \end{vmatrix},\end{aligned}\tag{1.4}$$

that does satisfy the anti-symmetry requirements, under exchange of same-spin electrons. This can then be extended to obtain a fully anti-symmetrised product of N single-particle orbitals, known as a Slater determinant. It turns out this Slater determinant description of electrons is equivalent to describing electrons as independent except for the effects of the exchange interaction (due to the anti-symmetry) and an average Coulomb interaction (due to the average position of all the other electrons). This description contains no description of correlation, and so is often called a mean-field theory. By allowing different forms of single-particle

orbital, and inclusion of free parameters, one can make use of the variational principle of quantum mechanics to minimise the energy expectation value of the Hamiltonian for this Slater wave function, resulting in a numerical approximation to the solution of the Schrödinger equation for N electrons. This is known as the Hartree-Fock (HF) method [12–14].

Clearly, we would like to improve upon the HF method and include a better description of correlation. Thanks to the theorems of Hohenberg and Kohn in 1964 [15], we can achieve this by making use of the electron density of the system n , in a method known as density functional theory (DFT). These two theorems tell us that, for a system of interacting particles, the exact ground-state electron density n_0 defines a unique energy functional $E[n]$ of the electron density whose global minimum is the same exact ground-state electron density, $\min_n E[n] = n_0$. In principle then, one just needs to find the form of this unique energy functional; however, finding this energy functional is essentially just a rewriting of the many-problem in Eq. (1.3). The key is the Kohn-Sham ansatz, published a year later in 1965 [16], which assumes that the ground-state electron density is equal to the ground-state electron density of some non-interacting auxiliary system moving in a modified potential. Treating the system as non-interacting, *i.e.* the wave function is a Slater determinant of single-particle orbitals $\{\psi_i(\mathbf{r}_i)\}$ and thus $n = \sum_{i=1}^N |\psi_i(\mathbf{r}_i)|^2$, allows a large part of the energy functional to be computed exactly. The unknown part of the energy functional, dealing with exchange and correlation effects, is then just a small correction to the non-interacting part. Given some approximation for this exchange-correlation functional, the energy can be obtained by varying the single-particle orbitals self-consistently. The simplest approximation to the exchange-correlation functional is the local density approximation, where the functional is determined only by the density local to each point in space. More complicated classes of exchange-correlation functionals include: the generalised-gradient approximation, which also use the gradient of the density at any point, the PBE (Perdew-Burke-Ernzerhof) functional [17] is an example;

and hybrid functionals which include a fraction the exact HF exchange energy, such as B3LYP (Becke, 3-parameter, Lee-Yang-Parr) [18, 19].

While DFT has been hugely successful, there are a few areas in which it struggles to give accurate results, including van der Waals interactions, the description of excitonic complexes, and in calculations of energy differences where cancellation of errors is not guaranteed (*i.e.* where there are large differences in the chemical bonding between two systems) [20]; in this thesis, it is these systems in which we are interested. To study such systems we make use of a powerful class of methods called quantum Monte Carlo (QMC). The biggest advantage of QMC over HF and DFT is the ability to handle large numbers (several hundreds) of explicitly correlated electrons, making them ideal for use in solving the quantum many-body problem. The explicitly correlated nature of QMC means the problem is not separable, as in HF and DFT, hence we rely on stochastic Monte Carlo processes to obtain a solution. Although there are many different variants of QMC methods such as path-integral QMC [21], configuration interaction QMC [22, 23], auxiliary-field QMC [24], and more; in this chapter we focus specifically on just two: variational Monte Carlo (VMC) and diffusion Monte Carlo (DMC) [25, 26]. In VMC, energy expectation values are calculated with respect to some (explicitly correlated) trial wave function, without restriction on the form of the trial wave function. To achieve this, the $3N$ -dimensional integrals in the expectation value are evaluated using Monte Carlo integration, the only method suitable for calculating integrals in a large number of dimensions. The DMC method is a stochastic projector method that simulates an evolution of the trial wave function according to the imaginary-time Schrödinger equation, in order to project out the ground-state component of the trial wave function. Together, that is by using VMC to optimise a trial wave function that will be the starting point in a subsequent DMC calculation, these methods can lead to highly accurate results for large, many-body systems. Indeed, it was QMC studies of the homogeneous electron gas [25] that were used to parametrise the local density approximation exchange-correlation functionals

in DFT. Furthermore, the stochastic nature of QMC methods means they are highly parallelisable and so they are well suited to take advantage of largest scale computers. The CASINO code [27] has been used for all QMC calculations in this thesis.

1.2 Statistical foundations

1.2.1 Monte Carlo integration

Before we discuss any further details of QMC methods, we must first cover some statistics. In a given many-body system consisting of N electrons any integrals that arise, say, in calculating expectation values with respect to a particular wave function are in $3N$ dimensions. In general, no analytic solution exists and so we must turn to numerical methods. For small numbers of electrons ($N < 3$) these integrals can be evaluated using grid-based methods and can be very efficient in doing so. As the electron number rises, however, these can be very inefficient. For example, take a d -dimensional integral; then the error estimate using Simpson's rule scales as $\mathcal{O}(M^{-\frac{4}{d}})$ where M is the number of sampling points, showing extremely poor scaling in high dimensions [28]. In contrast the error in Monte Carlo integration scales as $\mathcal{O}(M^{-\frac{1}{2}})$, *i.e.* independently of d , making it highly suitable for dealing with expectation values of tens or even hundreds of electrons. This remarkable error scaling is achieved by use of random sampling rather than grid-based sampling.

To illustrate this point, take a system of N electrons and let

$$\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N), \quad (1.5)$$

where \mathbf{r}_i is the position of the i -th electron. A particular value of \mathbf{R} , or configuration, is equivalent to a snapshot of the N electrons at a particular time. Let us

now define a function $f(\mathbf{R})$ with mean μ_f and variance σ_f^2 given by

$$\begin{aligned}\mu_f &= \int f(\mathbf{R})\rho(\mathbf{R}) d\mathbf{R}, \\ \sigma_f^2 &= \int [f(\mathbf{R}) - \mu_f]^2 \rho(\mathbf{R}) d\mathbf{R},\end{aligned}\tag{1.6}$$

where $\rho(\mathbf{R})$ denotes the probability density¹ of finding electrons in the configuration \mathbf{R} . Taking a set of M uncorrelated (mutually independent) configurations $\{\mathbf{R}_m : m \in \{1, \dots, M\}\}$ distributed according to the probability density $\rho(\mathbf{R})$ we can define a new quantity

$$\mathbf{Z}_f = \frac{[f(\mathbf{R}_1) + \dots + f(\mathbf{R}_M)]}{M}.\tag{1.7}$$

When M is large enough, the central limit theorem [29] states that \mathbf{Z}_f will be normally distributed with mean μ_f and standard deviation σ_f/\sqrt{M} . Thus, the integral μ_f can be evaluated simply as

$$\mu_f = \lim_{M \rightarrow \infty} \left\{ \frac{1}{M} \sum_{m=1}^M f(\mathbf{R}_m) \right\}.\tag{1.8}$$

Of course, in reality a Monte Carlo estimate is evaluated as

$$\mu_f \approx \frac{1}{M} \sum_{m=1}^M f(\mathbf{R}_m),\tag{1.9}$$

ensuring we use a large enough M . Therefore, we obtain an estimate of the multidimensional integral μ_f with an error that is independent of the dimension of \mathbf{R} and scales as $\mathcal{O}(M^{-\frac{1}{2}})$. One can extend this form of Monte Carlo integration to general integrals by first re-writing the integrand $g(\mathbf{R})$ such that $g(\mathbf{R}) = f(\mathbf{R})\rho(\mathbf{R})$, as we demonstrate later in Chapter 1.3. Here $\rho(\mathbf{R})$ is called an importance function because it has the effect of biasing our sampling to the region of configuration space in which we have the most interest.

¹That is, ρ satisfies $\rho(\mathbf{R}) \geq 0$ and $\int \rho(\mathbf{R}) d\mathbf{R} = 1, \forall \mathbf{R}$.

1.2.2 Metropolis sampling algorithm

The Metropolis sampling algorithm [30–32] is an example of a Markov process [33] and is an efficient way to sample a probability distribution function $\rho(\mathbf{R})$. A Markov process is one in which each step of the process depends only on the current configuration. The algorithm involves taking a set of configurations $\{\mathbf{R}_m\}$ and moving each one according to the following steps:

1. Generate a random position for configuration \mathbf{R}_m .
2. Propose a trial move to some configuration \mathbf{R}'_m . These trial moves are picked from some transition probability density $T(\mathbf{R}_m \rightarrow \mathbf{R}'_m)$ centred on \mathbf{R}_m .
3. Accept the move with probability

$$A(\mathbf{R}_m \rightarrow \mathbf{R}'_m) = \min \left\{ 1, \frac{T(\mathbf{R}'_m \rightarrow \mathbf{R}_m)\rho(\mathbf{R}'_m)}{T(\mathbf{R}_m \rightarrow \mathbf{R}'_m)\rho(\mathbf{R}_m)} \right\}. \quad (1.10)$$

4. Return to step 2.

If the trial move is accepted then \mathbf{R}'_m becomes the next point on the path and becomes the new current configuration, otherwise \mathbf{R}_m is added again. This has the effect of moving configurations away from regions of low probability density to regions where the probability density is higher. After some equilibration period the configuration density will be distributed according to the probability density $\rho(\mathbf{R})$, and further application of the Metropolis algorithm will not alter the configuration density further. In other words, we have reached a stationary point of the Markov process and, provided the system is ergodic², this stationary point is unique and the Markov process is guaranteed to converge to it [34, 35].

To see that $\rho(\mathbf{R})$ is a stationary point, let $n(\mathbf{R})$ be the density of configurations in the volume element $d\mathbf{R}$, then $n(\mathbf{R})d\mathbf{R}$ gives the number of configurations in $d\mathbf{R}$, while $A(\mathbf{R} \rightarrow \mathbf{R}')T(\mathbf{R} \rightarrow \mathbf{R}')d\mathbf{R}'$ gives the probability of configurations moving from $d\mathbf{R}$ to $d\mathbf{R}'$. Therefore, the number of configurations moving from $d\mathbf{R}$ to $d\mathbf{R}'$

²A configuration in one position can reach any other position in a finite number of moves.

is given by

$$A(\mathbf{R} \rightarrow \mathbf{R}')T(\mathbf{R} \rightarrow \mathbf{R}')n(\mathbf{R})d\mathbf{R}'d\mathbf{R}.$$

Now, the detailed balance condition (which is guaranteed to hold in ergodic systems [29]) states that, in equilibrium, the number of configurations moving from $d\mathbf{R}$ to $d\mathbf{R}'$ must equal to the number moving from $d\mathbf{R}'$ to $d\mathbf{R}$, so

$$A(\mathbf{R} \rightarrow \mathbf{R}')T(\mathbf{R} \rightarrow \mathbf{R}')n(\mathbf{R}) = A(\mathbf{R}' \rightarrow \mathbf{R})T(\mathbf{R}' \rightarrow \mathbf{R})n(\mathbf{R}'). \quad (1.11)$$

Using the definition of the Metropolis algorithm to give us the ratio of acceptance probabilities

$$\frac{A(\mathbf{R}' \rightarrow \mathbf{R})}{A(\mathbf{R} \rightarrow \mathbf{R}')} = \frac{T(\mathbf{R} \rightarrow \mathbf{R}') \rho(\mathbf{R})}{T(\mathbf{R}' \rightarrow \mathbf{R}) \rho(\mathbf{R}')}, \quad (1.12)$$

we find that

$$\frac{n(\mathbf{R})}{n(\mathbf{R}')} = \frac{A(\mathbf{R}' \rightarrow \mathbf{R})T(\mathbf{R}' \rightarrow \mathbf{R})}{A(\mathbf{R} \rightarrow \mathbf{R}')T(\mathbf{R} \rightarrow \mathbf{R}')} = \frac{\rho(\mathbf{R})}{\rho(\mathbf{R}')}, \quad (1.13)$$

demonstrating that $\rho(\mathbf{R})$ is indeed a stationary point of this Markov process.

In CASINO the transition probability density used is simply a Gaussian, because it is symmetric and efficient to calculate. The symmetry also allows for the acceptance probability to be simplified to just the ratio of the probability densities. Therefore, when dealing with a wave function $\Psi(\mathbf{R})$ of many quantum particles with probability density $|\Psi(\mathbf{R})|^2 / \int |\Psi(\mathbf{R})|^2 d\mathbf{R}$, the acceptance probability in Eq. (1.10) simplifies to

$$A(\mathbf{R}_m \rightarrow \mathbf{R}'_m) = \min \left\{ 1, \frac{|\Psi(\mathbf{R}'_m)|^2}{|\Psi(\mathbf{R}_m)|^2} \right\}, \quad (1.14)$$

and we see that the normalisation factor of the wave function cancels, another benefit to using this method.

The variance of the Gaussian transition probability density $T(\mathbf{R} \rightarrow \mathbf{R}')$ effectively controls how far a configuration will move on average. When the distance between current and proposed configurations is large, moves are more likely to be rejected and when it is low each step is close to the previous one. In both cases,

recorded configurations are very close to each other and this leads to inefficient exploration of the configuration space. In VMC this variance (also called the time step³ τ) is chosen so that the acceptance probability is approximately 50%, which maximises the efficiency of the algorithm [30, 36].

The first step of the Metropolis algorithm involves generating a position for each configuration. For an *ab-initio* calculation, CASINO biases these random positions to where we expect the electrons to be, *i.e.* around the nuclei. In a model calculation, electrons positions are picked from a uniform distribution in regions where the potential is not infinite, or chosen according to some single-particle orbitals.

Let us now define “some equilibration period”. For an electron on a random, diffusive, walk in d -dimensional space, the root-mean-square distance covered by the electron, over N_{equil} steps of length τ , is given by $\sqrt{dN_{\text{equil}}\tau}$. An equilibration period is deemed sufficient when

$$\sqrt{dN_{\text{equil}}\tau} > L_{\text{large}}, \quad (1.15)$$

where L_{large} is the largest relevant length scale in the problem. For example, the linear size of a molecule, or the size of the periodic supercell in bulk material.

1.2.3 Serial correlation

To increase sampling efficiency, CASINO employs an electron-by-electron scheme for proposing configuration moves rather than configuration-by-configuration [36]. Proposing, and then accepting or rejecting single-electron moves, rather than whole-configuration moves, results in a more efficient algorithm because successive steps are closer together, and so more likely to be accepted. The average distance moved over N single-electron moves will therefore be larger than the distance over one whole-configuration move.

In either case, this presents us with an issue: serial correlation. Choosing

³Note, this is the same τ that appears in Eq. (1.15).

configurations that are close to their predecessor results in serial correlation but our estimate of the error on the Monte Carlo estimate requires uncorrelated samples. One way to deal with this is to use only every k -th sample in the estimate, for example in CASINO, a minimum of only every 3rd sample is used in the Monte Carlo estimate. If the distance between samples that are actually used in the estimate is larger than the decorrelation length (the typical distance between uncorrelated samples) then the data will be uncorrelated.

We can also deal with serial correlation on-the-fly using the reblocking method [37]. Imagine we have some function $f(\mathbf{R})$ sampled at some M sampling points $\{\mathbf{R}_m\}$ generated by the Metropolis algorithm; we then group the data into blocks of successive pairs. Within each of these pairs we average the data; this results in a data set of size $M/2$ but one in which the data points are more independent (less correlated). This process can then be iterated until the serial correlation between the data points is removed. As we increase the block length (*i.e.* reduce serial correlation) we can expect the standard error in the mean of these data points to initially increase, until the block length exceeds the decorrelation length (all the serial correlation is removed from the data). At this point the standard error in the mean remains constant with reblocking length and we have found our true error in the mean. It is necessary to avoid just choosing an overly large block length from the outset in the hope of skipping this analysis as this will result an extremely small data set leading to a poor estimate of the standard error. One can use an automated method to choose the appropriate reblocking length [38].

1.3 Variational Monte Carlo

1.3.1 Expectation values

In the VMC method we use Monte Carlo integration to evaluate many-body expectation values. Taking a trial wave function $\Psi_T(\mathbf{R})$, the expectation value of an

operator \hat{A} can be evaluated as

$$\langle \hat{A} \rangle_{\text{VMC}} = \frac{\int \Psi_{\text{T}}^{\dagger}(\mathbf{R}) \hat{A} \Psi_{\text{T}}(\mathbf{R}) \, \text{d}\mathbf{R}}{\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}} = \frac{\int |\Psi_{\text{T}}(\mathbf{R})|^2 A_{\text{L}}(\mathbf{R}) \, \text{d}\mathbf{R}}{\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}} = \langle A_{\text{L}} \rangle_{\text{VMC}}, \quad (1.16)$$

where we have introduced the local operator $A_{\text{L}}(\mathbf{R}) = \Psi_{\text{T}}^{-1}(\mathbf{R}) \hat{A} \Psi_{\text{T}}(\mathbf{R})$. Hence, we can obtain a Monte Carlo estimate of the expectation value of \hat{A} simply by evaluating the average of the local operator $A_{\text{L}}(\mathbf{R})$ at configurations sampled according to the importance function $|\Psi_{\text{T}}(\mathbf{R})|^2 / \int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}$. The error estimate in this expectation value is given by σ_A / \sqrt{M} where

$$\sigma_A^2 = \frac{\int |\Psi_{\text{T}}(\mathbf{R})|^2 [A_{\text{L}}(\mathbf{R}) - \langle A_{\text{L}} \rangle_{\text{VMC}}]^2 \, \text{d}\mathbf{R}}{\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}}, \quad (1.17)$$

and M is the number of uncorrelated configurations used in the estimate. The only restriction on the form of trial wave function used is that it is a valid wave function, *i.e.* it satisfies the following: that $\Psi_{\text{T}}(\mathbf{R})$ and $\nabla_{\mathbf{R}} \Psi_{\text{T}}(\mathbf{R})$ are continuous everywhere except where the potential is infinite and that the integrals $\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}$ and $\int \Psi_{\text{T}}^{\dagger}(\mathbf{R}) \hat{A} \Psi_{\text{T}}(\mathbf{R}) \, \text{d}\mathbf{R}$ exist. To be able to evaluate the error on the estimate we also require $\int \Psi_{\text{T}}^{\dagger}(\mathbf{R}) \hat{A}^2 \Psi_{\text{T}}(\mathbf{R}) \, \text{d}\mathbf{R}$ to exist. While the form of the trial wave function is not restricted to being real the expectation value of a Hermitian operator⁴ is real in the limit of perfect sampling; therefore we can increase the efficiency of the method by only evaluating the real part of the local operator $\text{Re}(A_{\text{L}}(\mathbf{R}))$.

So far we have avoided mentioning spin. We can include spin without complicating the above discussion, as we will now demonstrate. Take a many-electron wave function $\Psi_{\text{T}}(\mathbf{X})$ where $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$ with $\mathbf{x}_i = (\mathbf{r}_i, \sigma_i)$ for spin $\sigma_i \in \{\uparrow, \downarrow\}$. Then the expectation of a spin-independent operator⁵ \hat{A} is

$$\langle \hat{A} \rangle_{\text{VMC}} = \frac{\sum_{\sigma} \int \Psi_{\text{T}}^{\dagger}(\mathbf{X}) \hat{A}(\mathbf{R}) \Psi_{\text{T}}(\mathbf{X}) \, \text{d}\mathbf{R}}{\sum_{\sigma} \int |\Psi_{\text{T}}(\mathbf{X})|^2 \, \text{d}\mathbf{R}}, \quad (1.18)$$

where we sum over all spin configurations $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$. Now, further

⁴An operator that corresponds to a physical observable.

⁵The only case covered in this thesis.

assume that $\Psi_{\text{T}}(\mathbf{X})$ is an eigenfunction of the total spin operator⁶ $\hat{S}_z = \sum_{i=1}^N \hat{s}_{z_i}$ with eigenvalue $(N_{\uparrow} - N_{\downarrow})/2$ where N_{\uparrow} is the number of spin-up electrons and N_{\downarrow} the number of spin-down electrons (*i.e.* we have $N_{\uparrow} + N_{\downarrow} = N$). For fermionic particles, such as electrons, $\Psi_{\text{T}}(\mathbf{X})$ is anti-symmetric under exchange of electron coordinates so we are free to choose a permutation for \mathbf{X} . Let us choose the permutation \mathbf{X}' where all spin-up electrons occupy the first N_{\uparrow} positions and all spin-down electrons occupy the final N_{\downarrow} positions, then

$$\begin{aligned} \mathbf{X}' &= (\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_N}) \\ &= (\mathbf{r}_{i_1}, \uparrow, \dots, \mathbf{r}_{i_{N_{\uparrow}}}, \uparrow, \mathbf{r}_{i_{N_{\uparrow}+1}}, \downarrow, \dots, \mathbf{r}_{i_N}, \downarrow) \\ &= (\mathbf{r}_1, \uparrow, \dots, \mathbf{r}_{N_{\uparrow}}, \uparrow, \mathbf{r}_{N_{\uparrow}+1}, \downarrow, \dots, \mathbf{r}_N, \downarrow) \end{aligned} \quad (1.19)$$

where in the last line we re-label the dummy variables. The trial wave functions $\Psi_{\text{T}}(\mathbf{X})$ and $\Psi_{\text{T}}(\mathbf{X}')$ now only differ by a factor of $(-1)^n$, where n is the number of electron exchanges we made to obtain our desired permutation \mathbf{X}' , so

$$|\Psi_{\text{T}}(\mathbf{X})|^2 = |(-1)^n \Psi_{\text{T}}(\mathbf{X}')|^2 = |\Psi_{\text{T}}(\mathbf{X}')|^2, \quad (1.20)$$

and we can define the ordered spatial trial wave function

$$\Psi_{\text{T}}(\mathbf{R}) \equiv \Psi_{\text{T}}(\mathbf{X}') = \Psi_{\text{T}}(\mathbf{r}_1, \uparrow, \dots, \mathbf{r}_{N_{\uparrow}}, \uparrow, \mathbf{r}_{N_{\uparrow}+1}, \downarrow, \dots, \mathbf{r}_N, \downarrow), \quad (1.21)$$

which we can use in place of the full space-spin trial wave function without loss. We can now see that the terms in each of the sums of Eq. (1.18) are identical and so they cancel giving the expectation value of the operator \hat{A} as

$$\langle \hat{A} \rangle_{\text{VMC}} = \frac{\int \Psi_{\text{T}}^{\dagger}(\mathbf{R}) \hat{A} \Psi_{\text{T}}(\mathbf{R}) \, \text{d}\mathbf{R}}{\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}}, \quad (1.22)$$

with respect to the spatial wave function $\Psi_{\text{T}}(\mathbf{R})$, which is the same as in Eq. (1.16). Note, the spatial wave function $\Psi_{\text{T}}(\mathbf{R})$ is only anti-symmetric under exchange of

⁶A restriction that is reasonable to impose if $[\hat{A}, \hat{S}_z] = 0$.

same-spin electrons. Spin-up and spin-down electrons are therefore treated as distinguishable.

The operator we are most interested in is, of course, the many-electron Hamiltonian⁷ \hat{H} . With a trial wave function $\Psi_{\text{T}}(\mathbf{R})$, the energy expectation value E can be estimated using VMC and is bounded below by the variational principle

$$\frac{\langle \Psi_{\text{T}} | \hat{H} | \Psi_{\text{T}} \rangle}{\langle \Psi_{\text{T}} | \Psi_{\text{T}} \rangle} = \frac{\int |\Psi_{\text{T}}(\mathbf{R})|^2 E_{\text{L}}(\mathbf{R}) \, \text{d}\mathbf{R}}{\int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}} \approx \frac{1}{M} \sum_{m=1}^M E_{\text{L}}(\mathbf{R}_m) = E_{\text{VMC}} \geq E_0, \quad (1.23)$$

where the local energy is $E_{\text{L}}(\mathbf{R}) = \Psi_{\text{T}}^{-1}(\mathbf{R}) \hat{H} \Psi_{\text{T}}(\mathbf{R})$, the configurations $\{\mathbf{R}_m\}$ are distributed according to $|\Psi_{\text{T}}(\mathbf{R})|^2 / \int |\Psi_{\text{T}}(\mathbf{R})|^2 \, \text{d}\mathbf{R}$, and E_0 is the ground-state energy of the system. Let $\Phi_0(\mathbf{R})$ be the exact ground-state eigenfunction; then equality occurs if, and only if, $\Psi_{\text{T}} = \Phi_0$ and the VMC energy E_{VMC} is equal to E_0 subject to a statistical error that falls away as $\mathcal{O}(M^{-\frac{1}{2}})$, for M uncorrelated configurations $\{\mathbf{R}_m\}$.

1.3.2 Trial wave function

Slater-Jastrow wave functions

The form of trial wave function we can use in VMC is essentially unrestricted, so we now turn our attention to the discussion of what is an appropriate trial wave function. Recall from our earlier discussion (Sec. 1.1) that a Slater determinant effectively imposes the fermionic symmetry requirements on a set of independent indistinguishable particles. Indeed, were we to use this Slater wave function in a VMC calculation we would obtain the HF energy. The most common form of QMC trial wave function (and the form we use in this thesis) makes use of the Slater determinant and multiplies it by a so-called Jastrow factor, which attempts to describe dynamical correlation effects between the electrons, resulting in the Slater-Jastrow wave function

$$\Psi_{\text{SJ}}(\mathbf{R}) = e^{J(\mathbf{R})} \Psi_{\text{S}}(\mathbf{R}) = e^{J(\mathbf{R})} D(\mathbf{R}), \quad (1.24)$$

⁷Which is indeed spin-independent.

where $J(\mathbf{R})$ is called the Jastrow correlation factor, $\Psi_S(\mathbf{R}) = D(\mathbf{R})$ is the Slater part of the wave function.

Slater determinant

As we have already seen, a Slater determinant contains single-particle orbitals of indistinguishable particles. We are treating different spin electrons as distinguishable, so the determinant is often expanded as a product of Slater determinants of each set of distinguishable particles. Hence, for the usual system of spin-up and spin-down electrons

$$D(\mathbf{R}) = D_{\uparrow}(\mathbf{r}_1, \dots, \mathbf{r}_{N_{\uparrow}}) D_{\downarrow}(\mathbf{r}_{N_{\uparrow}+1}, \dots, \mathbf{r}_N), \quad (1.25)$$

where the individual Slater determinants are

$$D_{\uparrow}(\mathbf{r}_1, \dots, \mathbf{r}_{N_{\uparrow}}) = \frac{1}{\sqrt{N_{\uparrow}!}} \begin{vmatrix} \psi_1(\mathbf{r}_1) & \psi_1(\mathbf{r}_2) & \cdots & \psi_1(\mathbf{r}_{N_{\uparrow}}) \\ \psi_2(\mathbf{r}_1) & \psi_2(\mathbf{r}_2) & \cdots & \psi_2(\mathbf{r}_{N_{\uparrow}}) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_{N_{\uparrow}}(\mathbf{r}_1) & \psi_{N_{\uparrow}}(\mathbf{r}_2) & \cdots & \psi_{N_{\uparrow}}(\mathbf{r}_{N_{\uparrow}}) \end{vmatrix}, \quad (1.26)$$

and similarly for $D_{\downarrow}(\mathbf{r}_{N_{\uparrow}+1}, \dots, \mathbf{r}_N)$. Obviously, the normalisation factor cancels out in practice, as we noted in Sec. 1.2.2, and so can be ignored but we include it here as the definition of a Slater determinant. Splitting the Slater determinant in this fashion also acts to increase the efficiency of a QMC calculation as we only need to (repeatedly) evaluate multiple smaller determinants rather than one larger one. The single-particle orbitals can be taken from methods like HF and DFT for *ab-initio* calculations of real systems (as we do in Chapter 4), or specially designed for calculations in model systems (see Chapter 2). The nodal topology defined by these Slater wave functions is actually all that the DMC energy depends on (see Sec. 1.4).

Jastrow factor

Description of dynamical correlation effects is then the job of the Jastrow factor. In order to not interfere with the exchange anti-symmetry provided by the Slater determinant, the Jastrow function $J(\mathbf{R})$ must be symmetric under the exchange of (like) particles. We also require the Jastrow function to be positive definite so as to not disrupt the nodal surface. A typical Jastrow factor for N electrons and N_N nuclei, such as the one developed by Drummond, Towler, and Needs [39], is written as

$$J(\mathbf{R}) = \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ i \neq j}}^N u(r_{ij}) + \sum_{i=1}^N \sum_{I=1}^{N_N} \chi(r_{iI}) + \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ i \neq j}}^N \sum_{I=1}^{N_N} f(r_{iI}, r_{jI}, r_{ij}), \quad (1.27)$$

and is composed of electron-electron (u), electron-nucleus (χ), and electron-electron-nucleus (f) terms, where $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ and $r_{iI} = |\mathbf{r}_i - \mathbf{r}_I|$.

The two-body u term often has the most impact on the quality of the trial wave function. It acts to reduce the magnitude of $\Psi_T(\mathbf{R})$ when electrons approach each other, thereby reducing the Coulomb repulsion energy and increasing the quality of the trial wave function (by the variational principle). In the Drummond-Towler-Needs (DTN) Jastrow factor, the u term has the following form

$$u(r_{ij}) = (r_{ij} - L_u)^C \Theta(L_u - r_{ij}) \sum_{l=0}^{N_u} \alpha_l r_{ij}^l. \quad (1.28)$$

The sum is a power expansion in the inter-particle distances with N_u optimisable free parameters⁸ $\{\alpha_0, \dots, \alpha_{N_u}\}$, which is then smoothly truncated, with $C - 1$ continuous derivatives, at the cut-off distance L_u . Θ is a Heaviside step function. To constitute a valid wave function we must therefore have, the truncation parameter, $C \geq 2$. The reason for the smooth truncation of the term to zero at L_u is twofold. Firstly, by reducing the cut-off we ensure variational freedom is in the right place to facilitate optimisation, *i.e.* at short-range, and simultaneously lowers the cost

⁸One parameter will be determined by ensuring that the wave function satisfies the Kato cusp conditions.

of computing the term, as not every particle pair needs to be calculated. Secondly, in periodic systems inter-particle distances are only evaluated between nearest periodic images, hence we must ensure that L_u is less than the radius L_{WS} of the largest sphere that can be inscribed within the Wigner-Seitz cell of the simulation cell. CASINO also has the ability to apply a separate set of parameters to different sets of particles, for different spins, or for different atomic species, for example.

An unintended consequence of the two-body u term is an altering of the one-electron density, which can be restored (or further altered) by the one-body χ term. The three-body f terms encode additional freedom into the trial wave function. Both the χ and f terms have similar forms to the u term in Eq. (1.28). These are the most common Jastrow factor terms therefore, before discussing any additional terms, we will briefly describe the Kato cusp conditions.

Kato cusp conditions

Consider the situation when two electrons approach other: the Coulomb potential energy between the two diverges. In contrast, for an eigenstate of the state of the system, the local energy remains constant. Therefore, to accurately describe the eigenstates of a system we must enforce on the trial wave function that the kinetic energy has an equal and opposite divergence whenever two charged particles approach other. We can impose this behaviour by satisfying the Kato cusp conditions [40]. For a general pair of charged particles with charges q_i and q_j , respectively, reduced mass μ_{ij} , separated by $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$, the local energy will be finite if we have

$$\left(\frac{1}{\Psi_{\text{T}}} \frac{\partial \Psi_{\text{T}}}{\partial r_{ij}} \right) \Big|_{r_{ij}=0} = \frac{2q_i q_j \mu_{ij}}{d \pm 1} = \Gamma_{ij}, \quad (1.29)$$

where d is the dimensionality of the system, and the plus and minus signs in the denominator are used for indistinguishable and distinguishable particles, respectively. For a Slater-Jastrow wave function this becomes just a condition on the

Jastrow factor

$$\left(\frac{\partial J}{\partial r_{ij}} \right) \Big|_{r_{ij}=0} = \Gamma_{ij}. \quad (1.30)$$

For $d = 3$ the electron-electron cusp conditions become

$$\begin{aligned} \Gamma_{\uparrow\uparrow} = \Gamma_{\downarrow\downarrow} &= \frac{1}{4} \\ \Gamma_{\uparrow\downarrow} = \Gamma_{\downarrow\uparrow} &= \frac{1}{2}, \end{aligned} \quad (1.31)$$

where \uparrow and \downarrow have been used to represent the spins of electrons i and j . In the DTN Jastrow factor, the Kato cusp conditions on electron-electron coalescence are enforced by setting

$$\alpha_1 = \frac{\Gamma_{ij}}{(-L_u)^C} + \frac{C\alpha_0}{L_u}, \quad (1.32)$$

in the electron-electron (u) term from Eq. (1.28).

Equivalently, there exists an electron-nucleus cusp which can be satisfied either by modifying the single-particle orbitals [41], or by setting

$$\left(\frac{\partial J}{\partial r_{iI}} \right) \Big|_{r_{iI}=0} = -Z_I. \quad (1.33)$$

Additional Jastrow terms

While the u , χ , and f Jastrow terms described above are the terms that occur most commonly, it is often necessary to include additional Jastrow terms, such as the H term which is a generalised version of the 3-body f term but without the restriction that one particle is a nucleus. In general, this H term is used in model systems rather than *ab-initio* calculations where the benefit outweighs the additional cost. In periodic systems, the u -term, for example, is truncated at the Wigner-Seitz radius of the simulation cell, meaning we are missing a large number of electron-electron pairs. The p -term is a plane-wave expansion which aims to add variational freedom into the corners of the simulation cell. It takes the following

form

$$p(\mathbf{r}_{ij}) = \sum_s a_s \sum_{\mathbf{G} \in \{\mathbf{G}_s^+\}} \cos(\mathbf{G} \cdot \mathbf{r}_{ij}), \quad (1.34)$$

where $\{a_s\}$ is a set of optimisable parameters, and the set $\{\mathbf{G}_s^+\}$ is the set of simulation-cell reciprocal lattice vectors belonging to the s -th star of vectors equivalent under the symmetry of the lattice⁹ where we exclude $-\mathbf{G}$, if \mathbf{G} is included. Other terms, such as the EXJAS and χ_{cyl} terms are sometimes used and will be introduced later in Chapter 2.

Backflow transformations

The Jastrow factor is positive definite and so has no effect on the nodal surface of the trial wave function, which is defined wholly by the Slater determinants. Backflow transformations of the particle co-ordinates \mathbf{R} are a commonly used method to alter the nodal surface of the Slater determinant. The backflow transformation $\mathbf{R} \rightarrow \mathbf{X}$, where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ is now a vector of quasi-particle co-ordinates and not space-spin vectors, is defined by

$$\mathbf{X}(\mathbf{R}) = \mathbf{R} + \boldsymbol{\xi}(\mathbf{R}), \quad (1.35)$$

where the backflow displacement $\boldsymbol{\xi}$ has a form similar to the DTN Jastrow factor, except that it is a vector function in \mathbf{r}_{ij} instead of a scalar function in r_{ij} . The backflow displacement is, similarly, composed of electron-electron (η), electron-nucleus (μ) terms, and electron-electron-nucleus (Φ , Θ) terms, whose exact form can be found in Ref. [42]. As these are vector functions, the electron-electron-nucleus term requires two functions to span the electron-electron-nucleus plane.

Slater-Jastrow-backflow trial wave functions can be helpful in improving the DMC energy, or evaluating the fixed-node error, but at the expense of adding an extra factor of N to the cost of the calculation [27, 42].

⁹Essentially, a star is set of equivalent lattice vectors that are the same distance away from the origin.

Pseudopotentials

The many-body Hamiltonian of Eq. (1.3) includes sums over all electrons (and nuclei), which works well for light atoms. However, it becomes impractical for heavy atoms; indeed, the cost of an all-electron QMC calculations scales with atomic number Z somewhere in the range $Z^{5.5}$ to $Z^{6.5}$ [43–45]. This arises largely due to two main factors: one, the core electrons move over shorter length scales than the valence electrons and so require a smaller time to ensure we are in the linear time step bias regime (see Sec. 1.4); and second, the (comparatively) large kinetic and potential energies for the core electrons result in large fluctuations in the local energy, reducing the statistical efficiency. In general though, it is only the valence electrons that participate in chemically interesting interactions (such as bonding and excitations), so it would be convenient to remove the core electrons from the problem but retain their effect on the valence electrons. This then is exactly the job of a pseudopotential, removing the core electrons (and nuclei), or ionic cores, from a calculation while reproducing their effect on the valence electrons.

Where ionic cores have been replaced in this thesis we have used Trail-Needs Dirac-Fock pseudopotentials [46, 47]. To ensure accuracy in the pseudopotentials we have used the data in Ref. [48] to choose our plane-wave cut-off energies large enough to ensure minimal basis-set errors, and to choose the angular momentum of the local channel so that we eliminate any potential ghost states.

Two-level sampling

Recall from Sec. 1.2.2 that the Metropolis algorithm acceptance probability is given by

$$A(\mathbf{R} \rightarrow \mathbf{R}') = \min \left\{ 1, \frac{|\Psi(\mathbf{R}')|^2}{|\Psi(\mathbf{R})|^2} \right\}. \quad (1.36)$$

For a Slater-Jastrow trial wave function this becomes

$$A(\mathbf{R} \rightarrow \mathbf{R}') = \min \left\{ 1, \frac{|D_{\uparrow}^2(\mathbf{R}_1^{N_{\uparrow}}) D_{\downarrow}^2(\mathbf{R}_{N_{\uparrow}+1}^N) |e^{2J(\mathbf{R})}|}{|D_{\uparrow}^2(\mathbf{R}'_1^{N_{\uparrow}}) D_{\downarrow}^2(\mathbf{R}'_{N_{\uparrow}+1}^N) |e^{2J(\mathbf{R}')}|} \right\}, \quad (1.37)$$

where $\mathbf{R}_1^{N_\uparrow} = (\mathbf{r}_1, \dots, \mathbf{r}_{N_\uparrow})$ and $\mathbf{R}_{N_\uparrow+1}^N = (\mathbf{r}_{N_\uparrow+1}, \dots, \mathbf{r}_N)$. A more efficient way to evaluate the acceptance probability is to use a two-level sampling algorithm [49]. In CASINO this is done by splitting the acceptance probability into the following two levels

$$A_1(\mathbf{R} \rightarrow \mathbf{R}') = \min \left\{ 1, \frac{|D_\uparrow^2(\mathbf{R}_1^{N_\uparrow}) D_\downarrow^2(\mathbf{R}_{N_\uparrow+1}^N)|}{|D_\uparrow^2(\mathbf{R}'_1^{N_\uparrow}) D_\downarrow^2(\mathbf{R}'_{N_\uparrow+1}^N)|} \right\}, \quad (1.38)$$

and

$$A_2(\mathbf{R} \rightarrow \mathbf{R}') = \min \left\{ 1, \frac{e^{2J(\mathbf{R})}}{e^{2J(\mathbf{R}')}} \right\}. \quad (1.39)$$

Trial moves are then accepted with the probability $A(\mathbf{R} \rightarrow \mathbf{R}') = A_1(\mathbf{R} \rightarrow \mathbf{R}')A_2(\mathbf{R} \rightarrow \mathbf{R}')$. The Slater part of the probability $A_1 = (\mathbf{R} \rightarrow \mathbf{R}')$ is calculated first as it is computationally cheaper and tends to have a lower acceptance probability than for the Jastrow part $A_2(\mathbf{R} \rightarrow \mathbf{R}')$, then only if this first step is accepted do we compute $A_2(\mathbf{R} \rightarrow \mathbf{R}')$.

1.3.3 Optimisation of the trial wave function

The DTN Jastrow factor contains many free parameters and, by making use of the variational principle, we can optimise these parameters¹⁰ to increase the quality of the trial wave function, by which we mean to lower the VMC energy. There are two main classes of methods for optimising trial wave functions. The first involve minimising some measure of the spread of local energies, such as the variance. When $\Psi_T(\mathbf{R})$ is an exact eigenstate of the many-body Hamiltonian, the local energy is constant everywhere and so the variance of the local energy is zero. We can therefore use the variance of the local energy¹¹

$$\sigma_E^2 = \frac{\int |\Psi_T(\mathbf{R})|^2 [E_L(\mathbf{R}) - E_{\text{VMC}}]^2 d\mathbf{R}}{\int |\Psi_T(\mathbf{R})|^2 d\mathbf{R}}, \quad (1.40)$$

¹⁰It is also possible to have free parameters in the single-particle orbitals, determinant coefficients, and backflow terms and optimise these at the same time.

¹¹We need to use the full complex local energy here because it is possible for the real part to be constant and the imaginary part to vary, which is not an eigenstate of the system.

as our cost function for optimisation and attempt to minimise this by using the procedure of Umrigar *et al.* [50]. Almost all of the parameters¹² that appear in the Jastrow terms u , χ , and f appear linearly. For such parameters there exists a far more efficient scheme based on the fact the variance of the local energy is a quartic function of said linear parameters [51]. Given that any eigenstate of the system has zero variance of the local energy it is possible, therefore, to use variance minimisation to obtain excited state wave functions. Instead of the variance one may use other measures of the spread; in particular minimising the mean absolute deviation from the median [52] is sometimes useful in cases where variance minimisation fails, see Chapter 2.2.3 for a discussion of this.

The other class of method is to directly optimise the energy itself, which is bounded below by the true ground-state energy. This is achieved by diagonalising the Hamiltonian in a basis set containing the trial wave function and its derivatives with respect to the free parameters. For a parameter set α , this basis set takes the form

$$\left\{ \Psi_{\text{T}}|_{\alpha}, \left. \frac{\partial \Psi_{\text{T}}}{\partial \alpha_1} \right|_{\alpha}, \left. \frac{\partial \Psi_{\text{T}}}{\partial \alpha_2} \right|_{\alpha}, \dots \right\}.$$

Simply trying to diagonalise the Hamiltonian with respect to this basis set often leads to instability in the algorithm, so energy minimisation needs a robust algorithm, such as the one introduced by Umrigar *et al.* [53], for it to be an effective optimisation technique.

Variance minimisation is generally the most numerically stable approach and it is able to search a large parameter space, while energy minimisation often gets stuck in local minima and so is not very effective when the initial parameter set is far from global minimum in parameter space. However, an energy-minimised trial wave function results in a more efficient DMC algorithm [43]. In this thesis, we optimise trial wave functions first by variance minimisation, to get near the global energy minimum, then by energy minimisation. This process greatly reduces the risk of finding a local energy minimum instead of the global energy minimum,

¹²With the exclusion of the cut-off length and the truncation parameter.

and we obtain a trial wave function that increases the efficiency of the subsequent DMC calculation.

1.4 Diffusion Monte Carlo

1.4.1 Imaginary-time Schrödinger equation

VMC does well at evaluating the energy of a particular trial wave function but this trial wave function completely determines VMC energy. Using DMC we can improve upon this by weakening the dependence on $\Psi_T(\mathbf{R})$, and in some cases even remove this dependence to give the exact ground-state energy. DMC is a stochastic projector method that projects out the ground-state component of the many-body wave function by evolving the wave function through imaginary time, from which we can then calculate the ground-state energy.

We will now introduce the imaginary-time Schrödinger equation. First let us write the many-body time-dependent Schrödinger equation in terms of total energies, and we will include a constant energy offset E_T ,

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) - E_T \right] \Psi(\mathbf{R}, t) = i\frac{\partial}{\partial t}\Psi(\mathbf{R}, t). \quad (1.41)$$

We then arrive at the imaginary-time Schrödinger equation by substituting $\tau = it$, where τ is a real variable that measures progress through imaginary time,

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) - E_T \right] \Psi(\mathbf{R}, \tau) = -\frac{\partial}{\partial \tau}\Psi(\mathbf{R}, \tau). \quad (1.42)$$

To see why this formulation of the Schrödinger equation is useful, take a look at the formal solution. Given a set of eigenstates $\{\Phi_n(\mathbf{R})\}$ with energies $\{E_n\}$, then we can expand the wave function as

$$\Psi(\mathbf{R}, \tau) = \sum_{n=0}^{\infty} c_n(\tau)\Phi_n(\mathbf{R}), \quad (1.43)$$

for some set of (possibly complex) coefficients $\{c_n(\tau)\}$. For time-independent Hamiltonians¹³ the time-dependence is separable, so we can make use of the time-evolution operator and write $\Psi(\mathbf{R}, \tau)$ as

$$\Psi(\mathbf{R}, \tau) = \sum_{n=0}^{\infty} c_n(0) \Phi_n(\mathbf{R}) e^{-(E_n - E_T)\tau}. \quad (1.44)$$

Notice that the exponent is real and so propagation in imaginary time τ causes any eigenstates with $E_n > E_T$ to exponentially decay¹⁴. Choosing E_T such that $E_n > E_T \forall n \geq 1$ results in all excited-state components dying out, leaving $\Psi(\mathbf{R}, \tau \rightarrow \infty) \propto \Phi_0(\mathbf{R})$.

1.4.2 Imaginary-time evolution

To demonstrate how to evolve the wave function through imaginary time let us take a look at the integral form of Eq. (1.42). The wave function $\Psi(\mathbf{R}, \tau)$ at time $\tau + \delta\tau$ is

$$\Psi(\mathbf{R}, \tau + \delta\tau) = \int G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) \Psi(\mathbf{R}', \tau) d\mathbf{R}', \quad (1.45)$$

where the Green's function

$$G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) = \langle \mathbf{R} | e^{-\delta\tau[\hat{H} - E_T]} | \mathbf{R}' \rangle = \langle \mathbf{R} | e^{-\delta\tau[\hat{T} + V(\mathbf{R}) - E_T]} | \mathbf{R}' \rangle \quad (1.46)$$

is a solution to Eq. (1.42), having initial condition $G(\mathbf{R}' \rightarrow \mathbf{R}, 0) = \delta(\mathbf{R}' - \mathbf{R})$, and where we have defined the total kinetic energy operator \hat{T} such that $\hat{T}\Psi(\mathbf{R}, \tau) = (-1/2)\nabla_{\mathbf{R}}^2\Psi(\mathbf{R}, \tau)$. This equation can be interpreted as the probability of a configuration \mathbf{R} occurring at some time $\tau + \delta\tau$ being given by the probability of moving from configuration \mathbf{R}' to configuration \mathbf{R} in the time $\delta\tau$, $G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$, multiplied by the probability of starting at configuration \mathbf{R}' , $\Psi(\mathbf{R}', \tau)$, with the integral accounting for all possible starting configurations.

¹³The only case we encounter in this thesis.

¹⁴A similar solution can be obtained for the real-time Schrödinger equation but the exponent here is imaginary resulting in oscillatory, rather than exponential growth/decay, behaviour.

By the Trotter-Suzuki formula [54] we can approximately factor the Green's function as

$$\begin{aligned}
G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) &= \langle \mathbf{R} | e^{-\delta\tau[\hat{T}+V(\mathbf{R})-E_T]} | \mathbf{R}' \rangle \\
&\approx e^{-\delta\tau[V(\mathbf{R})-E_T]/2} \langle \mathbf{R} | e^{-\delta\tau\hat{T}} | \mathbf{R}' \rangle e^{-\delta\tau[V(\mathbf{R}')-E_T]/2} + \mathcal{O}(\delta\tau^3) \\
&= \langle \mathbf{R} | e^{-\delta\tau\hat{T}} | \mathbf{R}' \rangle e^{-\delta\tau[V(\mathbf{R})+V(\mathbf{R}')-2E_T]/2} + \mathcal{O}(\delta\tau^3). \quad (1.47)
\end{aligned}$$

The part of this Green's function that contains the total kinetic energy operator \hat{T} corresponds to imaginary time Schrödinger equation Eq. (1.42) but where we ignore the potential and reference energy terms on the left hand-side

$$\frac{1}{2} \nabla_{\mathbf{R}}^2 \Psi(\mathbf{R}, \tau) = \frac{\partial}{\partial \tau} \Psi(\mathbf{R}, \tau). \quad (1.48)$$

This, then, is just a diffusion equation and has the following Green's function

$$G_{\text{diff}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) = \frac{1}{(2\pi\delta\tau)^{3N/2}} e^{-\frac{|\mathbf{R}-\mathbf{R}'|^2}{2\delta\tau}}, \quad (1.49)$$

and we can write $G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$ as

$$G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) = \frac{1}{(2\pi\delta\tau)^{3N/2}} e^{-\frac{|\mathbf{R}-\mathbf{R}'|^2}{2\delta\tau}} e^{-\delta\tau[V(\mathbf{R})+V(\mathbf{R}')-2E_T]/2} + \mathcal{O}(\delta\tau^3). \quad (1.50)$$

This Green's function can be used for a DMC calculation; however, there are numerous issues with it, the chief among them being that, for this to be valid, we require that the wave function $\Psi(\mathbf{R}, \tau)$ itself can be interpreted as probability distribution (not just $|\Psi(\mathbf{R}, \tau)|^2$), in particular that it is real and non-negative. However, the need to maintain fermion exchange anti-symmetry for our many-electron wave functions necessitates that $\Psi(\mathbf{R}, \tau)$ take positive and negative values. While naively it may seem the solution would be to introduce, and keep track of, signs for the configurations, it turns out that these methods are particularly inefficient¹⁵ [55]. We will therefore introduce a more efficient method, namely

¹⁵They have an exponentially decaying signal-to-noise ratio.

the importance-sampled method, and then describe how we handle the fermion anti-symmetry problem.

1.4.3 Importance sampling

Importance sampling in DMC makes use of a guiding, or trial, wave function $\Psi_T(\mathbf{R})$ (which we take to be a VMC optimised wave function as described previously) to bias the sampling towards regions of interest. If we, for now, assume the wave function $\Psi(\mathbf{R}, \tau)$ is real (see Sec. 1.4.4 for a discussion of using complex wave functions), we can introduce the mixed distribution

$$f(\mathbf{R}, \tau) = \Psi(\mathbf{R}, \tau)\Psi_T(\mathbf{R}), \quad (1.51)$$

and then the importance-sampled imaginary-time Schrödinger equation is

$$-\frac{1}{2}\nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \nabla_{\mathbf{R}} \cdot [\mathbf{V}(\mathbf{R})f(\mathbf{R}, \tau)] + (E_L(\mathbf{R}) - E_T)f(\mathbf{R}, \tau) = -\frac{\partial}{\partial \tau} f(\mathbf{R}, \tau), \quad (1.52)$$

where $\mathbf{V}(\mathbf{R}) = \Psi_T^{-1}(\mathbf{R})\nabla_{\mathbf{R}}\Psi_T(\mathbf{R})$ is the drift velocity, and $E_L(\mathbf{R})$ is the usual local energy; see Appendix A for a derivation. We can then modify the Green's function of Eq. (1.50) so that

$$G(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) = \underbrace{G_{\text{drift}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)G_{\text{branch}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)}_{G_{\text{DMC}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)} + \mathcal{O}(\delta\tau^2) \quad (1.53)$$

where

$$\begin{aligned} G_{\text{drift}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) &= \frac{1}{(2\pi\delta\tau)^{3N/2}} e^{-\frac{|\mathbf{R}-\mathbf{R}'-\delta\tau\mathbf{V}(\mathbf{R}')|^2}{2\delta\tau}} \\ G_{\text{branch}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau) &= e^{-\delta\tau[E_L(\mathbf{R})+E_L(\mathbf{R}')-2E_T]/2}. \end{aligned} \quad (1.54)$$

The drift-diffusion Green's function $G_{\text{drift}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$ describes a drift-diffusion process and corresponds to removing the third term in Eq. (1.52) (which is then just a Langevin equation). We have assumed that the drift velocity between \mathbf{R}

and \mathbf{R}' is constant, which has increased the error in $G_{\text{DMC}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$ from $\mathcal{O}(\delta\tau^3)$ to $\mathcal{O}(\delta\tau^2)$, this will be discussed further in Sec. 1.4.6. The branching Green's function $G_{\text{branch}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$ describes an exponential growth/decay in the population of configurations and corresponds to removing the first two terms in Eq. (1.52) (which is then just a rate equation). We interpret this branching factor as a survival probability rather than a weight. For example, let $w = G_{\text{branch}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$, then if $w < 1$ the configuration only survives this step with probability w , while if $w \geq 1$ then the configuration survives, and produces a new (independent) configuration with probability $w - 1$. In this way we avoid any one configuration gaining exponentially more weight than any other.

It is clear then, that we can propagate the mixed distribution $f(\mathbf{R}, \tau)$ through imaginary time (according the DMC Green's function) by simulating a population of $M(\tau)$ configurations¹⁶ $\{\mathbf{R}_m\}$ that, over each time step $\delta\tau$, drift by $\delta\tau\mathbf{V}(\mathbf{R}_m)$, randomly diffuse, and then are subject to a survival probability. More rigorously, at any time in a DMC calculation, we represent the mixed distribution $f(\mathbf{R}, \tau)$ by the ensemble average

$$f(\mathbf{R}, \tau) = \left\langle \sum_{i=1}^{M(\tau)} w_m(\tau) \delta(\mathbf{R} - \mathbf{R}_m) \right\rangle, \quad (1.55)$$

where $w_m(\tau)$ is a weight associated with \mathbf{R}_m at time τ . After one time step the mixed distribution $f(\mathbf{R}, \tau)$ becomes

$$f(\mathbf{R}, \tau + \delta\tau) = \left\langle \sum_{i=1}^{M(\tau)} w_m(\tau + \delta\tau) G_{\text{DMC}}(\mathbf{R}_m \rightarrow \mathbf{R}) \right\rangle. \quad (1.56)$$

The weights are updated according to

$$w_m(\tau + \delta\tau) = w_m(\tau) G_{\text{branch}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau), \quad (1.57)$$

and the number of configurations that move from \mathbf{R}_m in the next time step is

¹⁶The number of configurations M depends on τ because the branching/dying interpretation of G_{branch} means the total number of configurations varies with imaginary time τ .

calculated by

$$M_m(\tau + \delta\tau) = \text{INT}\{w(\tau + \delta\tau) + \nu\}, \quad (1.58)$$

where INT extracts the integer part of the real number, and ν is a randomly drawn number in the interval $[0, 1]$. The configurations are then updated¹⁷ by

$$\mathbf{R}_m = \mathbf{R}'_m + \underbrace{\boldsymbol{\chi}}_{\text{diffusion}} + \underbrace{\delta\tau \mathbf{V}(\mathbf{R}'_m)}_{\text{drift}}, \quad (1.59)$$

where $\boldsymbol{\chi}$ is a $3N$ -dimensional vector of random numbers distributed by a Gaussian with zero mean and variance τ . Eq. (1.59) then describes a drift-diffusion process for the N electrons and so, if we assume the drift velocity is constant between \mathbf{R}_m and \mathbf{R}'_m , the transition probability density for the move is

$$T(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau) = G_{\text{drift}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau). \quad (1.60)$$

The importance-sampled Green's function of Eq. (1.53) should satisfy the detailed balance condition;¹⁸ however, the assumption that the drift velocity is constant violates the detailed balance condition. Fortunately, this can be imposed straightforwardly, by inclusion of a Metropolis-style acceptance probability; for a configuration moving from \mathbf{R}'_m to \mathbf{R}_m a move is accepted with probability

$$A_{\text{DMC}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau) = \min \left\{ 1, \frac{G_{\text{drift}}(\mathbf{R}_m \rightarrow \mathbf{R}'_m, \delta\tau) |\Psi(\mathbf{R}_m)|^2}{G_{\text{drift}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau) |\Psi(\mathbf{R}'_m)|^2} \right\}, \quad (1.61)$$

where \mathbf{R}_m becomes the new position of the configuration if the move is accepted, otherwise it remains unchanged. As $\delta\tau$ tends to zero the acceptance probability $A_{\text{DMC}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau)$ tends to 1 and detailed balance is enforced. Therefore, to reduce the time-step error that enters our calculation, we would like to use time steps small enough so that our DMC move acceptance ratio is close to 1, *i.e.*

¹⁷After we introduce the acceptance step in Eq. (1.61) it is actually more efficient to use an electron-by-electron scheme for proposing moves rather than a configuration-by-configuration scheme.

¹⁸That, in equilibrium, the number of configurations moving from \mathbf{R}' to \mathbf{R} matches the number moving from \mathbf{R} to \mathbf{R}' .

$$A_{\text{DMC}}(\mathbf{R}'_m \rightarrow \mathbf{R}_m, \delta\tau) \geq 0.99.$$

1.4.4 Fixed-phase approximation

Finally, we address the issue we hinted at in Sec. 1.4.2, that of fermion exchange anti-symmetry. For both the non-importance-sampled and importance-sampled methods described, it is required that $\Psi(\mathbf{R}, \tau)$ or $f(\mathbf{R}, \tau)$ can be interpreted as a probability distribution, *i.e.* they are real and non-negative everywhere. In the importance-sampled DMC method this is can be easily dealt with.

For a real wave function $\Psi(\mathbf{R}, \tau)$, whenever two same-spin electrons are exchanged, the sign of the wave function must change, which occurs whenever an electron crosses the nodal surface of $\Psi(\mathbf{R}, \tau)$. Therefore, we constrain the nodal surface of $\Psi(\mathbf{R}, \tau)$ to be the same as the nodal surface of the trial wave function $\Psi_{\text{T}}(\mathbf{R})$, this ensures that the mixed distribution $f(\mathbf{R}, \tau) = \Psi(\mathbf{R}, \tau)\Psi_{\text{T}}(\mathbf{R})$ is (real) and non-negative everywhere. Enforcing this condition, the fixed-node approximation [56], in the DMC algorithm outlined above is, in theory, automatic because the drift velocity diverges whenever $\Psi_{\text{T}}(\mathbf{R})$ is zero; however, in practice, due to the use of finite time-steps, some configurations occasionally cross the nodal surface. We can prevent this simply by rejecting any moves that cross the nodal surface.¹⁹

If we allow complex wave functions $\Psi(\mathbf{R}, \tau) = |\Psi(\mathbf{R}, \tau)|e^{i\Omega(\mathbf{R}, \tau)}$, whenever two same-spin electrons are exchanged, the sign of the wave function must again change, and this occurs whenever the phase $\Omega(\mathbf{R}, \tau)$ changes by π . Let us fix the phase of the wave function $\Psi(\mathbf{R}, \tau)$, and let us choose this phase to be the phase of the trial wave function $\Psi_{\text{T}}(\mathbf{R}) = |\Psi_{\text{T}}(\mathbf{R})|e^{i\Omega_{\text{T}}(\mathbf{R})}$ (which has the correct anti-symmetry requirements because of the Slater determinant). We can now only make changes to the modulus of the wave function $|\Psi(\mathbf{R}, \tau)|$, so whatever we do cannot interfere with the fermionic anti-symmetry requirements of the wave function $\Psi(\mathbf{R}, \tau)$. This is the fixed-phase approximation [57], and the fixed-node approximation is the special case where the phase $\Omega_{\text{T}}(\mathbf{R})$ is equal to 0 or π for

¹⁹By checking whether the Slater part of the trial wave function $\Psi_{\text{T}}(\mathbf{R}, \tau)$ has changed sign.

every configuration \mathbf{R} .

The importance-sampled DMC algorithm described above requires only slight modifications to enforce the fixed-phase approximation. To explain these we will first act on the (fixed-phase) complex wave function $\Psi(\mathbf{R}, \tau)$ with the Hamiltonian \hat{H} , finding that

$$\begin{aligned}
\hat{H}\Psi(\mathbf{R}, \tau) &= \left(-\frac{1}{2}\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) \right) e^{i\Omega_{\mathbf{T}}(\mathbf{R})} |\Psi(\mathbf{R}, \tau)| \\
&= e^{i\Omega_{\mathbf{T}}(\mathbf{R})} \left[\frac{1}{2} \left(-\nabla_{\mathbf{R}}^2 + |\nabla_{\mathbf{R}}\Omega_{\mathbf{T}}(\mathbf{R})|^2 + V(\mathbf{R}) \right) \right] |\Psi(\mathbf{R}, \tau)| \\
&+ i e^{i\Omega_{\mathbf{T}}(\mathbf{R})} \left[-\frac{1}{2} \left(2\nabla_{\mathbf{R}}\Omega_{\mathbf{T}}(\mathbf{R}) \cdot \nabla_{\mathbf{R}} + \nabla_{\mathbf{R}}^2\Omega_{\mathbf{T}}(\mathbf{R}) \right) \right] |\Psi(\mathbf{R}, \tau)| \\
&\equiv e^{i\Omega_{\mathbf{T}}(\mathbf{R})} \left[\hat{H}_{\Omega_{\mathbf{T}}} + i\hat{K}_{\Omega_{\mathbf{T}}} \right] |\Psi(\mathbf{R}, \tau)|, \tag{1.62}
\end{aligned}$$

where, in the last line, we define the real and imaginary parts of the complex Hamiltonian $\hat{H}_{\Omega_{\mathbf{T}}}$ and $\hat{K}_{\Omega_{\mathbf{T}}}$, respectively. The expectation of the Hamiltonian becomes

$$\langle \Psi | \hat{H} | \Psi \rangle = \langle |\Psi| | \hat{H}_{\Omega_{\mathbf{T}}} | |\Psi| \rangle + i \langle |\Psi| | \hat{K}_{\Omega_{\mathbf{T}}} | |\Psi| \rangle = \langle |\Psi| | \hat{H}_{\Omega_{\mathbf{T}}} | |\Psi| \rangle, \tag{1.63}$$

because $\langle |\Psi| | \hat{K}_{\Omega_{\mathbf{T}}} | |\Psi| \rangle = 0$. We can now define the fixed-phase Schrödinger equation

$$\hat{H}_{\Omega_{\mathbf{T}}} |\Psi(\mathbf{R})| = E_{\Omega_{\mathbf{T}}} |\Psi(\mathbf{R})|, \tag{1.64}$$

which has a ground-state $|\Phi_{\Omega_{\mathbf{T},0}}(\mathbf{R})|$ with eigenvalue $E_{\Omega_{\mathbf{T},0}}$. Then, by Eq. (1.63), the fixed-phase ground-state energy $E_{\Omega,0}$ is equal to the expectation value of the Hamiltonian \hat{H} when $\Psi(\mathbf{R}) = |\Phi_{\Omega_{\mathbf{T},0}}(\mathbf{R})| e^{i\Omega_{\mathbf{T}}(\mathbf{R})}$. By the variational principle, $E_{\Omega_{\mathbf{T},0}}$ is an upper bound bound on the ground-state energy E_0 of the Hamiltonian \hat{H} , and becomes equal in the limit that $\Omega_{\mathbf{T}}(\mathbf{R})$ is exactly equal to the phase of the ground-state $\Phi_0(\mathbf{R}) = |\Phi_0(\mathbf{R})| e^{i\Omega_0(\mathbf{R})}$. Similarly, we obtain the fixed-phase imaginary-time Schrödinger equation

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 + \frac{1}{2}|\nabla_{\mathbf{R}}\Omega(\mathbf{R})|^2 + V(\mathbf{R}) - E_{\mathbf{T}} \right] |\Psi(\mathbf{R}, \tau)| = -\frac{\partial}{\partial \tau} |\Psi(\mathbf{R}, \tau)|, \tag{1.65}$$

and, by performing an importance-sampling transformation using the mixed distribution

$$f(\mathbf{R}, \tau) = |\Psi(\mathbf{R}, \tau)| |\Psi_T(\mathbf{R})|, \quad (1.66)$$

which is real and non-negative, we obtain a fixed-phase version of Eq. (1.52) (the importance-sampled imaginary-time Schrödinger equation). In the fixed-phase version of Eq. (1.52) only the real parts of the drift velocity $\text{Re}(\mathbf{V}(\mathbf{R})) = |\Psi_T(\mathbf{R})|^{-1} \nabla_{\mathbf{R}} |\Psi_T(\mathbf{R})|$ and the local energy $\text{Re}(E_L(\mathbf{R})) = |\Psi_T(\mathbf{R})|^{-1} \hat{H}_{\Omega_T} |\Psi_T(\mathbf{R})|$ appear, otherwise the two are the same on the surface with the differences in the Hamiltonians hidden away in the local energy. The fixed-phase approximation is automatically enforced in the importance-sampled fixed-phase Schrödinger equation because we only vary the modulus of the wave function and not its phase.

This restriction on the wave function is, in general, the single largest approximation we make in a DMC calculation but it is still the case that the DMC energy only depends on the phase (nodes) of the trial wave function, and not on $\Psi(\mathbf{R}, \tau)$ itself, as in VMC. Fortunately, there exists a variational principle on the ground-state of the fixed-phase DMC energy [55], and for excited states that transform as a one-dimensional irreducible representation of the symmetry group of the Hamiltonian [58]. However, it is important to remember that whenever the phase of the trial wave function is exactly the phase of any eigenstate, the DMC energy will be exact. In model systems where we only have distinguishable particles $\Psi(\mathbf{R}, \tau)$ will be node-less, and hence we obtain the exact ground-state energy.

1.4.5 Expectation values

A typical DMC calculation then proceeds as follows: first, we have an equilibration phase where we allow the excited-state components of $\Psi(\mathbf{R}, \tau)$ to die away, at which point the population of configurations will be distributed as $f_0(\mathbf{R}) = \Phi_0(\mathbf{R}) \Psi_T(\mathbf{R})$; secondly, we have a statistics accumulation phase where we propagate configurations in this steady state and accumulate the value of some operator \hat{A} .

Taking \hat{A} to be the Hamiltonian \hat{H} , the mixed estimate of the DMC energy is then

$$\langle \hat{H} \rangle_{\text{DMC}} = \frac{\langle \Phi_0 | \hat{H} | \Psi_T \rangle}{\langle \Phi_0 | \Psi_T \rangle} = \frac{\int f_0(\mathbf{R}) E_L d\mathbf{R}}{\int f_0(\mathbf{R}) d\mathbf{R}} = E_{\text{DMC}}, \quad (1.67)$$

and we can evaluate this using a Monte Carlo estimate as the average of the local energy

$$E_{\text{DMC}} \approx \sum_{m=1}^M E_L(\mathbf{R}_m), \quad (1.68)$$

over a set of M configurations distributed as $f_0(\mathbf{R}) / \int f_0(\mathbf{R}) d\mathbf{R}$. The error in this estimate decays as $\mathcal{O}(M^{-\frac{1}{2}})$. For operators that commute with the Hamiltonian, unbiased estimates of their expectation values can be evaluated in a similar fashion.

For operators \hat{A} that do not commute with the Hamiltonian we can attempt to remove the bias, which is (in general) linear in the error in the trial wave function, by use of the future-walking method [55]. Alternatively, we can instead evaluate the extrapolated estimate

$$\langle \hat{A} \rangle_{\text{extrap}} = 2\langle \hat{A} \rangle_{\text{DMC}} - \langle \hat{A} \rangle_{\text{VMC}}, \quad (1.69)$$

in which case it can be shown the error is quadratic in the error in the trial wave function [59].

1.4.6 Time-step and population-control biases

The DMC Green's function $G_{\text{DMC}}(\mathbf{R}' \rightarrow \mathbf{R}, \delta\tau)$ of Eq. (1.53) is only valid when we use small time steps, and even then we introduce an error $\mathcal{O}(\delta\tau^2)$ per time step $\delta\tau$, *i.e.* at a rate $\mathcal{O}(\delta\tau)$. The total time-step error Δf is reduced at a rate of $\Delta f / \tau_{\text{decorr}}$, where τ_{decorr} is the typical amount of imaginary time between uncorrelated samples. Once we reach equilibrium these rates cancel to give a total time-step error $\Delta f \propto \tau_{\text{decorr}} \delta\tau$, *i.e.* an error $\mathcal{O}(\delta\tau)$. This time-step bias can be removed post-calculation by linear extrapolation of the DMC energy at two different time steps, provided

small enough time steps have been used.

A good rule for ensuring the time step is small enough is to make sure the root-mean-square distance diffused by each electron, per time step, in a d -dimensional system, $\sqrt{d\delta\tau}$ is smaller than the shortest relevant length scale in the system. With appropriate time steps chosen we can now define the length of the DMC equilibration phase. In this case, we choose the number of equilibration steps N_{equil} such that the root-mean-square distance diffused by an electron, over the whole equilibration period, $\sqrt{dN_{\text{equil}}\delta\tau}$ is greater than the largest length scale in the system. This allows enough time for a configuration to fully explore the space.

The introduction of an acceptance step does not change the scaling of the time-step bias but it massively reduces its magnitude. Nevertheless some time-step bias remains because, as some steps are rejected, the mean distance diffused by electrons at each time step is less than it should be (going as the square root of the time step). This bias can be reduced by using an effective time step to compensate for the rejected moves [60]. A further source of time-step bias arises due to the regions in which the trial wave function goes to zero because then the drift velocity diverges. This can result in very large configuration moves, introducing error in the approximation of the Green's function. To prevent this, one can limit the drift velocity when it becomes too large, using a scheme such as the one by Zen *et al.* [61].

Interpreting the branching factor as a survival probability means that the total number of configurations fluctuates as a function of imaginary time. In importance-sampled DMC these fluctuations (which are reduced anyway because the branching factor contains the, comparatively smoother, local energy rather than the potential energy) are controlled so that the total number of configurations is close to some initial value. This is achieved by changing the reference energy E_T (which also appears in the branching factor) during the calculation, but one needs to be careful to do this smoothly so as not to introduce a large bias [62]. In practice, this population-control bias is often small and can be removed by linear

extrapolation, in tandem with the time-step bias, by varying the target population inversely with time-step. It is possible to remove the population-control bias in a more rigorous fashion [63] but at the expense of introducing extra statistical noise and so is generally disfavoured.

Chapter 2

Excitonic complexes in semiconductor nanostructures

2.1 Models of excitonic complexes in QMC

2.1.1 Effective mass approximation

QMC is an effective way to solve the many-electron Schrödinger equation for many *ab-initio* systems, but it is also very capable of solving the Schrödinger equation for a model system consisting of only a few interacting charge carriers. Such model systems are often a convenient way to extract the key physics of a material or device without the expense of solving the full many-electron Schrödinger equation.

Consider a one-electron (approximate) description in which electrons move independently in effective potentials (which attempt to include the effects of interaction with all other electrons). We can then write down a single-electron Schrödinger equation

$$\hat{H}_1\psi(\mathbf{r}) = \mathcal{E}\psi(\mathbf{r}), \quad (2.1)$$

where \hat{H}_1 is a one-electron Hamiltonian whose potential is an effective one-electron potential, and $\psi(\mathbf{r})$ is a single-particle eigenfunction with eigenvalue \mathcal{E} . In a periodic crystal lattice \hat{H}_1 has the periodicity of the lattice. One can then use

translational invariance to show that

$$\psi(\mathbf{r}) \equiv \psi_{\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} u_{\mathbf{k}}(\mathbf{r}), \quad (2.2)$$

where \mathbf{k} belongs to reciprocal space and can be restricted to points within the first Brillouin zone¹, and $u_{\mathbf{k}}(\mathbf{r})$ has the periodicity of the lattice. These $\psi_{\mathbf{k}}(\mathbf{r})$ (called Bloch orbitals) have eigenvalues $\mathcal{E}(\mathbf{k})$ (called the single-particle energy band²) and form a complete basis. This is Bloch's theorem [64].

If we now introduce a slowly varying (on the scale of the primitive unit cell) external potential $U(\mathbf{r})$ then we obtain the following Hamiltonian

$$\hat{H} = \hat{H}_1 + U(\mathbf{r}). \quad (2.3)$$

The eigenfunctions of this Hamiltonian can be expanded in the basis formed by the Bloch orbitals. If $U(\mathbf{r})$ is constant, then no mixing of the Bloch orbitals occurs and the eigenfunctions of the Hamiltonian in Eq. (2.3) are just the Bloch orbitals; when $U(\mathbf{r})$ is slowly varying, then mixing of the Bloch orbitals happens over a very narrow range of reciprocal space vectors. The eigenfunctions are then linear combinations of Bloch orbitals lying close together in the Brillouin zone, so we can approximate these eigenfunctions in Bloch form, the so-called Bloch wave packets, as

$$\psi_{\text{wp}}(\mathbf{r}) \approx u_{\mathbf{k}}(\mathbf{r}) \sum_{\mathbf{k}' \text{ near } \mathbf{k}} a_{\mathbf{k}'} e^{i\mathbf{k}'\cdot\mathbf{r}} \equiv u_{\mathbf{k}}(\mathbf{r}) \psi_{\text{en}}(\mathbf{r}), \quad (2.4)$$

for a set of coefficients $\{a_{\mathbf{k}'}\}$, and we have defined the envelope function $\psi_{\text{en}}(\mathbf{r})$. The wave packet $\psi_{\text{wp}}(\mathbf{r})$ is described by a short range rapidly varying part $u_{\mathbf{k}}(\mathbf{r})$, which has periodicity of the lattice, modulated by the long range envelope function $\psi_{\text{en}}(\mathbf{r})$. These wave packets consist of a linear combination of Bloch orbitals close in \mathbf{k} so have a (reasonably) well defined crystal momentum and can be thought of as quasi-particles. In an isotropic quadratic band, these quasi-particles can be

¹The set of points in reciprocal space closer to the origin than any other reciprocal lattice vector.

²Here we focus on a single energy band and so have suppressed the band index.

shown to obey the following Schrödinger equation

$$\left[-\frac{1}{2m^*}\nabla_{\mathbf{r}}^2 + U(\mathbf{r}) + \mathcal{E}_{\min} \right] \psi_{\text{en}}(\mathbf{r}) = \mathcal{E}' \psi_{\text{en}}(\mathbf{r}), \quad (2.5)$$

with eigenvalues \mathcal{E}' , where

$$\frac{1}{m^*} = \frac{\partial^2 \mathcal{E}(\mathbf{k})}{\partial k^2}, \quad (2.6)$$

$k = |\mathbf{k}|$, and \mathcal{E}_{\min} is the minimum energy of the band [65–67].

Thus, we have simplified the problem of describing the effects of the electrons in a periodic lattice, to one where we solve the Schrödinger equation for a set of quasi-particles with an effective mass m^* that approximates the effects of the periodic lattice. The same general idea still applies if we instead consider a many-electron Hamiltonian and use the electron-electron interaction in place of an external potential (*i.e.* the complexity of the periodic lattice potential is still buried in the effective mass approximation). Hiding the effects of the periodic potential in the effective mass results in a screening of the interaction between the quasi-particles, and neglects the exchange interaction. At long-range, Eq. (2.5) and its spatially varying envelope functions are still valid, if we assume the potential $U(\mathbf{r})$ to be more-or-less constant over each unit cell.

The effective mass is defined by the curvature of the energy bands at the minimum and maximum for conduction and valence bands, respectively. In a conduction band, the effective mass is positive and we describe negatively-charged particles with effective mass $m_e^* = m^*$, or quasi-electrons. In a valence band, the effective mass is negative; the movement of electrons in the valence band is equivalent to describing the behaviour of positively-charged holes with effective mass $m_h^* = |m^*|$, or quasi-holes. For the remainder of this chapter and in Chapter 3 we refer to these quasi-particles simply as electrons and holes.

Particles of varying mass result in a non-unitary factor of the reciprocal mass in the kinetic energy term of the Schrödinger equation. To accommodate this in the QMC methods described in Chapter 1 is simply a matter of modifying some of the

expressions, such as the drift-diffusion Green's function, to account for the mass (the drift distance now includes a factor of the mass). Most of the modifications are hidden in the algorithm, so let us just focus on the changes to the quantities used to estimate appropriate time steps and equilibration periods. The first, the root-mean-square distance diffused by each particle, per time step $r_{\delta\tau}^{\text{RMS}}$, which allows us to estimate when we are in the linear time-step bias regime, becomes $r_{\delta\tau}^{\text{RMS}} = \sqrt{d\delta\tau/m_{\text{lightest}}}$, for a d -dimensional system. The lightest particle mass is used because lighter particles travel further per time step. The second quantity, which lets us estimate if the equilibration period is long enough, is the root-mean-square distance diffused by a particle in the equilibration period $r_{\text{equil}}^{\text{RMS}}$ and becomes $r_{\text{equil}}^{\text{RMS}} = \sqrt{dN_{\text{equil}}\delta\tau/m_{\text{heaviest}}}$. More massive particles move more slowly, so take longer to explore the system, hence we use the heaviest particle mass.

Positively-charged particles are accommodated by inclusion of the charges (± 1 in Hartree atomic units) in the Coulomb term of the Hamiltonian, and by inclusion of the EXJAS term in the Jastrow factor. The EXJAS term is specific to these type of systems and enforces the general Kato cusp condition given in Chapter 1, see Sec. 2.2.3 for details.

2.1.2 Excitonic complexes

Semiconductors are a class of materials that possess a band gap but where the size of the gap is small enough that electrons can be excited, thermally or otherwise, across the gap. When the conduction band minimum and valence band maximum occur at the same point in the Brillouin zone, the material is a direct-gap semiconductor. In such semiconductors, one can induce a population of electrons and holes by injection of photons (by use of a laser) with energy similar to the band gap, thereby exciting electrons from the valence band to the conduction band, and leaving behind a population of holes in the valence band. These charge carriers can then form bound states called charge-carrier, or excitonic, complexes. One of the most interesting consequences of the existence of excitonic complexes is

the presence of well-defined, atomic-like peaks in the photoluminescence spectra of solid-state materials, that are generally only found in gases of atoms or small molecules. Materials possessing these properties are then promising candidates for devices such as single-photon light sources, photodetectors, and photodiodes.

The simplest such state is the exciton: a two-body complex consisting of a bound electron-hole pair. For a free exciton there is, analogous to the hydrogen atom but using rescaled units, an exact solution to the Schrödinger equation. Three-body complexes called trions (or charged excitons) consist of either two electrons and a single hole (negative trion) or a single hole bound to two electrons (positive trion). The four-body biexcitons are a bound state of a pair of excitons.

While the effective-mass approximation simplifies the discussion of these systems, we are still left with a fully-interacting, inhomogeneous, few-body problem. QMC is then a very attractive method for solving the Schrödinger equation in these model systems. Indeed for such complexes, because similarly-charged particles are of opposite spin we can treat all particles as distinguishable (which results in a node-less trial wave function), the fixed-node DMC method exactly solves the model. We can therefore calculate the ground-state energy of each complex, which allows us to evaluate both the binding energies (for all complexes) and the de-excitonisation energies (for trions and biexcitons). The binding energy determines the temperature at which a complex becomes unstable against dissociation into smaller complexes, while the de-excitonisation energies determine where the trion and biexciton peaks are expected to appear in a photoluminescence experiment. For a free (unconfined) complex the binding and de-excitonisation energies are equivalent, but in systems with type-II band alignment, such as the ones we study here, they can differ.

It is possible to have more than four particles in an excitonic complex. In materials such as monolayer transition-metal dichalcogenides, where the electrons and holes have both spin and valley indices, these larger complexes can still consist of distinguishable particles [68]. Five-particle complexes, or quintons, of this

type are investigated in type-II superlattices in Sec. 2.3; here the distinguishability is spacial separation. Alternatively, one can investigate complexes containing indistinguishable particles, and in this case the trial wave function must be anti-symmetric under exchange of these particles [68].

2.1.3 Excitonic units

It is often sensible to work in units appropriate for the system being studied. For excitonic complexes we can define a set of excitonic units. Given the similarity of an exciton to the hydrogen atom, excitonic units are defined analogously to Hartree atomic units. In Hartree atomic units, masses are in units of the bare electron mass m_e , lengths are in units of the Bohr radius a_0 , and energies are in units of Hartree (Ha). In excitonic units then, masses are in units of the reduced mass of the electron-hole pair

$$\mu = \frac{m_e^* m_h^*}{m_e^* + m_h^*}, \quad (2.7)$$

lengths are in units of the exciton Bohr radius

$$a_0^* = \frac{\epsilon}{\mu} a_0, \quad (2.8)$$

and energies are in units of the exciton Hartree

$$\text{Ha}^* = \frac{\mu}{\epsilon^2} \text{Ha}, \quad (2.9)$$

where ϵ is the permittivity of the material hosting the exciton. The exciton Rydberg R_y^* is $\text{Ha}^*/2$, and is equal to the binding energy of a free exciton in three dimensions (3D) for a given set of material parameters $\{\mu, \epsilon\}$. Therefore, one only needs to know these two parameters of a material to be able to evaluate the binding energy of a free exciton in that material.

We use excitonic units throughout this chapter, and again in Chapter 3, unless

otherwise specified, with the material parameters $\{\mu, \epsilon\}$ specified in each case.

2.2 Type-II quantum rings

2.2.1 Quantum rings

Quantum-dot and quantum-ring heterostructures have long been hailed as “artificial atoms” [69–72] due to their ability to confine charge carriers in all three spatial dimensions. Material combinations exhibiting type-I band alignment produce nanostructures in which both electrons and holes are confined to the same spatial region, and such nanostructures have been studied extensively over the last two decades [73–75]. In type-II nanostructures, on the other hand, only holes but not electrons (or *vice versa*) are confined, presenting a rich variety of new physics [76–78]. For example, GaSb quantum dots or quantum rings in GaAs provide very deep confining potentials for holes [79], while strain in the GaSb raises the conduction-band minimum, expelling the electrons [80]. Excitonic complexes in type-II nanostructures are in fact very much more like artificial atoms than is the case for type-I nanostructures, because the electrons are bound to the holes in the “artificial nuclei” purely by the Coulomb interaction, rather than being confined themselves. Type-II quantum rings are an intriguingly distinct type of artificial atom with no natural analogue due to the radical difference between the ring-shaped “artificial nucleus” and the point-like nucleus of a real atom.

Excitons in type-II quantum dots have been extensively studied both experimentally [81–84] and theoretically [79, 85, 86]; however, while there has been some experimental work on carrier complexes in type-II quantum-ring nanostructures [5, 87–89], there has been little theoretical work to date. The spatial separation of charge carriers allows for a variety of interesting optoelectronic properties [87, 90], including extended recombination times, making type-II quantum rings ideal candidates for applications such as memory devices [91] and solar cells [84]. Binding energies of excitonic complexes reported here are effectively

ionisation energies of these artificial atoms. GaSb quantum rings in GaAs may be produced by molecular beam epitaxy [5, 89, 90, 92] and can form with a variety of different cross-sections ranging from triangular, to semicircular [5], and even trapezoidal [90]. These quantum rings exhibit type-II behaviour, with the holes strongly confined to the rings. Scanning tunnelling microscopy (STM) has been used to investigate the shape and size of the GaSb rings, and their optical properties have been studied in photoluminescence experiments [5, 87–89].

2.2.2 The model

We solve an effective-mass model of excitons (X), positive and negative trions (X⁺ and X⁻), and biexcitons (XX) in type-II quantum-ring heterostructures, focusing on GaSb rings in GaAs. The holes are confined to the ring, which is modelled as an infinite potential well, while the electrons are excluded from the ring but bound to the holes by an isotropically screened Coulomb interaction. The kinetic energy of the tightly confined holes is the dominant contribution to the total energy of each complex; however, the electron-hole attraction is non-negligible, as is the hole-hole repulsion. The ring was chosen to have a rectangular cross-section for computational convenience. The ring is centred on the origin, orientated so that the axis of rotation is the z -axis and the midpoint in the z direction is the x - y plane, see Fig. 2.1. The three parameters defining the ring geometry are the half height of the ring R_z , the inner radius of the ring r_i , and the outer radius r_o . In our model the electron and hole densities do not overlap, so we cannot estimate recombination rates; however our model is reasonable for calculating binding energies.

Within the effective-mass approximation the Hamiltonian for a biexciton is

$$\begin{aligned}
\frac{\hat{H}}{2R_y^*} &= -\frac{(a_0^*)^2\mu}{2m_e^*}(\nabla_{e_1}^2 + \nabla_{e_2}^2) - \frac{(a_0^*)^2\mu}{2m_h^*}(\nabla_{h_1}^2 + \nabla_{h_2}^2) \\
&+ \frac{a_0^*}{r_{e_1e_2}} + \frac{a_0^*}{r_{h_1h_2}} - \frac{a_0^*}{r_{e_1h_1}} - \frac{a_0^*}{r_{e_1h_2}} - \frac{a_0^*}{r_{e_2h_1}} - \frac{a_0^*}{r_{e_2h_2}} \\
&+ \sum_i V_i\left(\frac{\mathbf{r}_i}{a_0^*}\right)
\end{aligned} \tag{2.10}$$

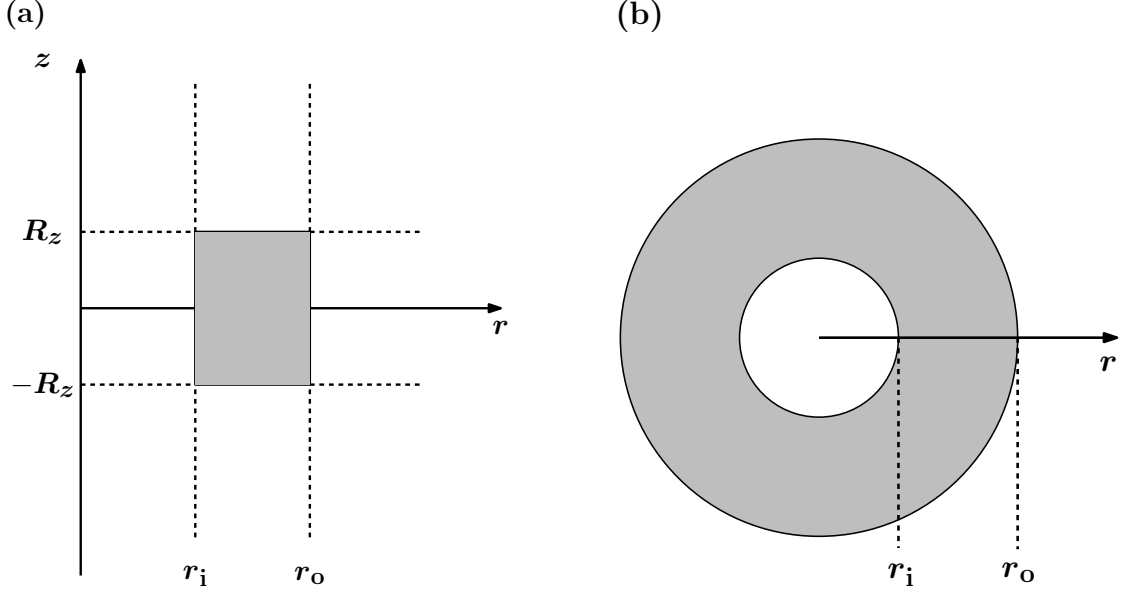


Figure 2.1: Model of the quantum ring shown in (a) cross-section and (b) plan view. The parameters R_z , r_i , and r_o represent the half height, the inner radius, and the outer radius of the ring, respectively.

in excitonic units, where $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ and $V_i(\mathbf{r})$ is the confining potential, which is infinite inside the ring and zero outside for electrons, and *vice versa* for holes. This is an inhomogeneous four-body problem, but the wave function is node-less so the DMC method can be used to calculate the exact ground-state energy for each complex, and hence the de-excitonisation and binding energies. The trion and biexciton de-excitonisation energies $E_D^{X^\pm}$ and E_D^{XX} are

$$\begin{aligned}
 E_D^{X^-} &= E^X - E^{X^-}, \\
 E_D^{X^+} &= E^X + E^{h^+} - E^{X^+}, \\
 E_D^{XX} &= 2E^X - E^{XX},
 \end{aligned} \tag{2.11}$$

where E^i is the ground-state total energy for complex i . These are the energies at which trion and biexciton peaks are expected to appear relative to the exciton peak in the photoluminescence spectrum of a quantum ring [88]. The sign is such that for a free trion or biexciton $E_D > 0$. The binding energies—the energy difference between a complex and its most energetically favourable daughter products,

bearing in mind that the holes are confined to the ring—are

$$\begin{aligned}
E_b^X &= E^{h^+} - E^X, \\
E_b^{X^-} &= E^X - E^{X^-}, \\
E_b^{X^+} &= E^{2h^+} - E^{X^+}, \\
E_b^{XX} &= E^{X^+} - E^{XX},
\end{aligned} \tag{2.12}$$

where E^{2h^+} is the energy of two holes confined to the same ring. The binding energy determines the temperature at which a complex becomes unstable against dissociation into smaller complexes. We also calculate the charge density for each complex, to aid in comparison with experimental data.

2.2.3 Computational details

For excitonic complexes in a type-II quantum ring we take our trial wave functions $\Psi_T(\mathbf{R})$ to be of Slater-Jastrow form. Each particle is distinguishable so the Slater part of the trial wave function $\Psi_T(\mathbf{R})$ is just a product of single particle orbitals. For the biexciton

$$\Psi_T(\mathbf{R}) = e^{J(\mathbf{R})} \psi_e(\mathbf{r}_{e_1}) \psi_e(\mathbf{r}_{e_2}) \psi_h(\mathbf{r}_{h_1}) \psi_h(\mathbf{r}_{h_2}), \tag{2.13}$$

where here $\mathbf{R} = (\mathbf{r}_{e_1}, \mathbf{r}_{e_2}, \mathbf{r}_{h_1}, \mathbf{r}_{h_2})$.

The hole orbital $\psi_h(\mathbf{r})$ was taken to be the exact ground-state solution to the Schrödinger equation for a single hole confined to the ring

$$\psi_h(\mathbf{r}) = \left[\frac{-J_0(\beta r) Y_0(\beta r_i)}{J_0(\beta r_i)} + Y_0(\beta r) \right] \cos\left(\frac{\pi z}{2R_z}\right), \tag{2.14}$$

where J_0 and Y_0 are Bessel functions of the first and second kind, respectively, and we use cylindrical polar coordinates $\mathbf{r} = (r, \theta, z)$. The constant β is determined by imposing the boundary condition $\psi_h(r = r_o) = 0$ numerically for each ring size using the Newton-Raphson method; the other boundary conditions are already

satisfied by Eq. (2.14).

The electronic behaviour is dominated by Coulomb attraction to the positively charged ring together with hard-wall repulsion from the boundary of the ring. At short range the electron orbital $\psi_e(\mathbf{r})$ linearly decreases to zero on the ring boundary, while at long range the electron orbital decays exponentially to keep the electrons localised to the ring; *i.e.*, the behaviour is hydrogenic at long range. The electron orbital was formed piecewise in eight regions about the ring, as shown in Fig. 2.2, with the functions in each region being matched at the boundaries to ensure the value and gradient were smooth everywhere. The orbital was zero inside the ring. The full form of the electron orbital is




Figure 2.2: The nine different regions of space (1–9) used to construct an electron orbital with a continuous first derivative around a quantum ring with a rectangular cross section. The parameters R_z , r_i , and r_o represent the half height, the inner radius, and the outer radius of the ring, respectively.

$$\psi_e(\mathbf{r}) = \begin{cases} \left[a_1(z - R_z) + b_1(z - R_z)^2 \right] e^{-\frac{c_1(z-R_z)^2}{1+d_1(z-R_z)}}, & \text{in region 1} \\ \left[\sqrt{a_1^2(z - R_z)^2 + a_3^2(r - r_o)^2} \right. \\ \quad \left. + b_1(z - R_z)^2 + b_3(r - r_o)^2 \right] e^{-\sqrt{\frac{c_1^2(z-R_z)^4}{[1+d_1(z-R_z)]^2} + c_3^2(r-r_o)^2}}, & \text{in region 2} \\ \left[a_3(r - r_o) + b_3(r - r_o)^2 \right] e^{-c_3(r-r_o)}, & \text{in region 3} \\ \left[\sqrt{a_1^2(z + R_z)^2 + a_3^2(r - r_o)^2} \right. \\ \quad \left. + b_1(z + R_z)^2 + b_3(r - r_o)^2 \right] e^{-\sqrt{\frac{c_1^2(z+R_z)^4}{[1-d_1(z+R_z)]^2} + c_3^2(r-r_o)^2}}, & \text{in region 4} \\ \left[-a_1(z + R_z) + b_1(z + R_z)^2 \right] e^{-\frac{c_1(z+R_z)^2}{1-d_1(z+R_z)}}, & \text{in region 5} \\ \left[\sqrt{a_1^2(z + R_z)^2 + (r_i - r)^2 \left(a_7 + \frac{a_7 r}{r_i} + b_7 r^2 \right)^2} \right. \\ \quad \left. + b_1(z + R_z)^2 \right] e^{-\frac{c_1(z+R_z)^2}{1-d_1(z+R_z)}}, & \text{in region 6} \\ (r_i - r) \left(a_7 + \frac{a_7 r}{r_i} + b_7 r^2 \right), & \text{in region 7} \\ \left[\sqrt{a_1^2(z - R_z)^2 + (r_i - r)^2 \left(a_7 + \frac{a_7 r}{r_i} + b_7 r^2 \right)^2} \right. \\ \quad \left. + b_1(z - R_z)^2 \right] e^{-\frac{c_1(z-R_z)^2}{1+d_1(z-R_z)}}, & \text{in region 8} \\ 0, & \text{in region 9,} \end{cases} \quad (2.15)$$

where the constants $a_1, a_3, a_5, a_7, b_1, b_3, b_7, c_1, c_3,$ and d_1 are variational parameters determined by optimisation, and we again use cylindrical polar coordinates $\mathbf{r} = (r, \theta, z)$. To ensure correct (*i.e.*, hydrogenic) behaviour at long range, it is required that $c_1, c_3, d_1 \geq 0$. The electron orbital $\psi_e(\mathbf{r})$ enforces the correct long- and short-

range behaviour, with the mid-range behaviour determined by the Jastrow factor and variational freedom in the electron orbital.

The Jastrow function $J(\mathbf{R})$ was composed of χ , u and f (for the trions and biexcitons) terms with a cut-off length several times larger than the exciton Bohr radius a_0^* . To include the χ and f terms a mass-less, charge-less, ghost particle is positioned at the origin. The EXJAS term, unique to calculations involving excitonic complexes, was also included. The three-dimensional version of the EXJAS term was developed for this work and is a pairwise sum of terms of the form

$$u_{ij}^{\text{EXJAS}}(r_{ij}) = \frac{\pm\mu_{ij}r_{ij} + b_{ij}r_{ij}^2}{1 + c_{ij}r_{ij}}, \quad (2.16)$$

for each pair of particles i and j . The plus sign is for particles with the same charge and minus for particles with opposite charge. The parameter b_{ij} was set equal to zero here because the long-range behaviour is described by the electron orbitals. c_{ij} is a variational parameter, which was different for each particle-pair type. This term ensured the Kato cusp conditions were satisfied.

Given the cylindrical nature of the ring, it seems feasible that Jastrow terms with cylindrical, rather than spherical symmetry might perform better. To investigate this, the χ_{cyl} term was developed. Like χ , the χ_{cyl} term is a polynomial in the particle–nucleus distance that is smoothly truncated at some cut-off distance, but χ_{cyl} is a product of two such polynomials: one on the radial direction r and one in the axial direction z . Here, we have a ghost particle at the origin so the particle-nucleus distance is just the distance from the origin of each particle. The full form of the χ_{cyl} can be written, using cylindrical polar coordinates $\mathbf{r} = (r, \theta, z)$, as

$$\chi_{\text{cyl}}(\mathbf{r}_i) = (r_i - L_{\chi r})^C (|z_i| - L_{\chi z})^C \Theta(L_{\chi r} - r_i) \Theta(L_{\chi z} - |z_i|) \sum_{l=0}^{N_{\chi r}} \sum_{m=0}^{N_{\chi z}} \omega_{lm} r_i^l |z_i|^m, \quad (2.17)$$

where $L_{\chi r}$ and $L_{\chi z}$ are the cut-off distances in the radial and axial directions, respectively, and C is the truncation parameter. The order of the polynomial ex-

pansion in each direction is given by $N_{\chi r}$ and $N_{\chi z}$, and we require $\omega_{1m} = \omega_{0m}C/L_{\chi r}$ and $\omega_{1l} = \omega_{l0}C/L_{\chi z}$ to ensure the trial wave function has a continuous first derivative. The remaining parameters can be varied freely. In the end, it turns out that the χ term performs better than the χ_{cyl} term, *i.e.* lower VMC energies can be achieved with the χ term than with the χ_{cyl} term. The most likely reason for this is that, while the ring has cylindrical symmetry, the complexes exist on a scale larger than that of the ring, and so favour a spherical symmetry (see the charge density plots in Fig. 2.7).

The scheme suggested in Chapter 1 for optimising trial wave function (variance minimisation, followed by energy minimisation) is not reliable here. In particular, variance minimisation often struggles to optimise the parameters so that we obtain the ground-state of the system. Instead, we can end up with the complex becoming unbound (*i.e.* one or more electrons has moved a very large distance from the ring), which is another valid eigenstate of the system and hence also has zero variance in the local energy. In these cases, one can sometimes find that minimising the mean absolute deviation from the median can be helpful, however, in this case, we find that just using energy minimisation results in well-optimised trial wave functions.

Using energy-minimised trial wave functions DMC calculations are performed for the ground-state energies of excitons, trions, and biexcitons in quantum-ring heterostructures. Pairs of DMC calculations were performed with time steps in a 1 : 4 ratio and target configuration populations in a 4 : 1 ratio and the results were extrapolated linearly to zero time step and infinite population. Charge densities were obtained by binning the radial and axial coordinates of each of the particles sampled during VMC and DMC calculations, cylindrically averaging, and then calculating the extrapolated estimate.

2.2.4 Results and discussion

All energies and charge densities are reported for a ring composed of GaSb surrounded by GaAs. The electron and hole masses are taken to be $m_e^* = 0.063 m_e$

and $m_h^* = 0.4 m_e$, respectively. The former is the effective mass of an electron in bulk GaAs and the latter is the mass of a heavy hole in bulk GaSb [93]. The permittivity is taken to be $\epsilon = 12.9 \epsilon_0$, this is the permittivity of bulk GaAs [93]. Data from Ref. [5] were used to obtain experimentally relevant values for the ring size; these values were $R_z = 2.5 \text{ nm} = 0.199 a_0^*$, $r_i = 6 \text{ nm} = 0.479 a_0^*$, and $r_o = 10 \text{ nm} = 0.799 a_0^*$. This geometry was used as the starting point for our calculations; the aspect ratio $2R_z/(r_o - r_i)$ of the ring was then varied subject to the constraints that the volume of the ring $2\pi R_z(r_o^2 - r_i^2)$ was constant and the centre of the ring in the radial direction $r_m = (r_i + r_o)/2$ was fixed. A ring with aspect ratio much less than 1 is akin to a thin disc with a hole in the centre, while a ring with aspect ratio much greater than 1 resembles a pipe.

The analytically evaluated variation in the hole energy against aspect ratio is shown in Fig. 2.3. The minimum energy occurs when the cross-section is square; away from the minimum, the energy goes roughly as $1/L^2$, where




Figure 2.3: Ground-state total energies per hole of a single hole (h^+), two holes ($2h^+$), an exciton (X), a negative trion (X^-), a positive trion (X^+), and a biexciton (XX) in a quantum ring plotted against the aspect ratio $2R_z/(r_o - r_i)$ of the ring's cross section. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The exciton Rydberg R_y^* is 4.45 meV for the experimentally relevant geometry.

$L = \min\{2R_z, r_o - r_i\}$. Also shown in Fig. 2.3 are DMC ground-state total energies per hole for $2h^+$, X, X^- , X^+ , and XX, all of which are bound. These confirm that the ground-state energies of the single-hole complexes (X and X^-) are very close to the energy of a single hole, while the ground-state energies of the two-hole complexes (X^+ and XX) are comparable with the energy of two confined holes. The ground-state energies for single- and two-hole complexes vary slightly differently as a function of aspect ratio due to the interaction between the holes. The capacitive charging energy $E_{CC} = E^{2h^+} - 2E^{h^+}$ for the experimentally relevant ring geometry [5] is $E_{CC} = 8.8546(8)$ meV; this compares to an experimentally measured value [94] of $E_{CC} = 24(2)$ meV. STM images of quantum rings [5, 88] suggest that the GaSb/GaAs interface is not clean in practice. This disorder could lead to trapping of holes, strongly affecting capacitive charging energies while having relatively little effect on binding energies.

Complex	E_D/R_v^*	E_b/R_v^*	E_D (meV)	E_b (meV)
X	0	0.5004(6)	0	2.226(3)
X^-	+0.0446(4)	0.0446(4)	+0.199(2)	0.199(2)
X^+	-1.111(2)	1.379(2)	-4.944(7)	6.137(7)
XX	-0.911(2)	0.701(2)	-4.052(8)	3.11(1)

Table 2.1: De-excitonisation E_D and binding E_b energies for excitonic complexes in the quantum-ring geometry modelling the samples described in Ref. [5].

The de-excitonisation energies for the trions and biexciton in the geometry modelling the quantum rings described in Ref. [5] can be found in Table 2.1. The de-excitonisation energy is positive for X^- , but negative for X^+ and XX. The negative de-excitonisation energy is a result of the large energy penalty when two holes are confined to the same ring; *e.g.*, two excitons on two separate quantum rings would be the energetically preferred four-particle state rather than a biexciton on a single ring. The expected positions of these peaks in a photoluminescence spectrum are shown in Fig. 2.4. The X^- peak is very close to the X peak, while the peaks for X^+ and XX are separated from the X peak by a few meV. The heights of the peaks indicate the relative stability of the complexes, using binding energy




Figure 2.4: Expected peak positions for the excitonic complexes in a photoluminescence spectrum relative to the exciton peak, for a model of the quantum rings reported in Ref. [5]. The peak heights represent the relative stability of the complexes.

data from Table 2.1. Experimental work has not yet progressed to the point where excitonic complex peak positions have been identified. The only published work showing sharp lines in the photoluminescence spectra of GaSb/GaAs quantum rings is Ref. [88]; however the spectra in this work would likely contain peaks from many, highly positively charged rings, making a direct comparison with theoretical values difficult. The de-excitation energy is plotted against the aspect ratio of the cross-section of the ring for X^- , X^+ , and XX in Fig. 2.5. For each complex it can be seen that there is some slight change in the de-excitation energy as a function of aspect ratio. The de-excitation energies are largely independent of the aspect ratio, and hence exact shape of the ring, somewhat justifying the use of a ring with a rectangular cross-section in our model. Furthermore, the energetic effects of the slight interpenetration of the electron and hole orbitals are likely to be well described by a slight re-normalisation of the cross-section of the ring; however the effects of such small changes in the cross-section appear to be small.

The binding energies for each complex are shown in Table 2.1 for the experimentally relevant geometry [5]. The X binding energy is about half the value for a free X due to the exclusion of the electron from the ring. As expected, X^- is the most weakly bound (against dissociation into a free electron and a neutral exciton), while X^+ is the most stable (against removal of an electron from a ring of charge of $+2e$). From these binding energies we can calculate the temperature




Figure 2.5: De-excitonisation energies against the aspect ratio $2R_z/(r_o - r_i)$ of a quantum ring’s cross-section for different charge-carrier complexes. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The dashed line shows the experimentally relevant aspect ratio [5].

T up to which the complexes are stable by setting

$$E_b = k_B T, \quad (2.18)$$

where k_B is the Boltzmann constant. These temperatures are shown in Table 2.2. As with the de-excitonisation energies, the binding energy depends weakly on the aspect ratio of the ring’s cross-section, but again these differences are much smaller than the differences in binding energy between complexes, see Fig. 2.6. Therefore, the binding energy appears to be largely independent of the exact shape of the

Temperature (K)			
X	X^-	X^+	XX
26	2.3	71	36

Table 2.2: Temperatures up to which the complexes are stable, in the quantum-ring geometry modelling the samples described in Ref. [5], using the binding energy data from Table 2.1.




Figure 2.6: Binding energies against the aspect ratio $2R_z/(r_o - r_i)$ of a quantum ring’s cross-section for different charge-carrier complexes. The mean radius and ring volume are appropriate for the GaSb/GaAs quantum rings reported in Ref. [5]. Error bars are smaller than the size of the symbols. The dashed line shows the experimentally relevant aspect ratio [5].

cross-section of the ring for a given ring volume and mean radius.

Plots of the electronic charge density for each complex in the experimentally relevant geometry are shown in Fig. 2.7. The electrons form a diffuse halo around the ring, with negligible charge density in the ring’s central cavity. The kinetic-energy cost of localising in the ring’s cavity significantly exceeds the gain in electrostatic potential energy. Correlation effects further reduce the probability of finding multiple electrons inside the ring’s cavity. XX and X^+ are the most localised complexes, as reflected in their relatively large binding energies shown in Table 2.1. These two-hole complexes have slightly higher electronic charge densities in the regions directly above and below the centre of the ring compared to the regions to the left and right of the ring. STM images of the electronic density of states in Ref. [5] suggest the electrons are localised to the ring’s cavity, which does not agree with the results presented here. However, in the STM experiments the sample is cleaved in the x - z plane. This is a drastic modification to the system,




Figure 2.7: Electronic charge density ρ for (a) an exciton, (b) a negative trion, (c) a positive trion, and (d) a biexciton in the experimentally relevant quantum-ring geometry [5]. The shaded regions represent the ring and ρ_m is the maximum density across all four plots. The free exciton Bohr radius is $a_0^* = 12.5$ nm.

which is not described by our model. It is plausible that the reduced screening and hence smaller free exciton Bohr radius in the cleaved system allows electrons to localise within, rather than above or below, the quantum ring.

The sensitivity of the XX binding energy to various parameters is presented in Table 2.3. The XX binding energy depends most strongly on the electron effective mass, and is relatively insensitive to the hole effective mass, relative permittivity, ring volume, and mean ring radius. Our conclusions are robust against reasonable uncertainties in model parameters.

$\partial E_b^{XX}/\partial m_e^*$ (meV/ m_e)	$\partial E_b^{XX}/\partial m_h^*$ (meV/ m_e)	$\partial E_b^{XX}/\partial \epsilon$ (meV)	$\partial E_b^{XX}/\partial V$ (meV/nm ³)	$\partial E_b^{XX}/\partial r_m$ (meV/nm)
7.4(3)	0.20(4)	-0.39(1)	-0.0004(2)	-0.07(2)

Table 2.3: Sensitivity of the biexciton binding energy to the electron and hole effective masses m_e^* and m_h^* , the relative permittivity ϵ , the ring volume V , and the mean radius of the ring r_m .

Kehili *et al.* [95] have recently investigated excitons in GaSb rings in GaAs quantum wells using the effective-mass approximation, modelling the ring with a

finite potential, and including strain effects due to lattice-constant mismatch. In their work the interaction between charge carriers is described by a Hartree mean-field approximation, in contrast to the complete treatment of correlation effects used here. Nevertheless, their electronic charge density is qualitatively consistent with our results. Their X binding is slightly larger than our value reported in Table 2.3, however, partly due to their use of slightly different effective masses and mean ring radii. A DMC calculation of the X binding energy using the same ring geometry and effective masses as Kehili *et al.* gives $E_b^X = 2.695(2)$ meV, which is comparable with the binding energy of about 2.6 meV that they report for a GaAs well of width of 40 nm (the largest well width they consider). The X binding energies reported by Kehili *et al.* do not appear to have converged with respect to well width at this point, however, and it looks as if they will be significantly smaller than the DMC exciton binding energy in the limit of large well width. This is consistent with the fact that, by the variational principle, Hartree theory underestimates the magnitude of the X binding energy.

2.2.5 Conclusion

In conclusion, total energies of excitonic complexes in type-II quantum-ring heterostructures are dominated by the confinement energy of the holes in each complex. The de-excitonisation energy is positive for X^- as would be the case for a free trion; however, for X^+ and XX this energy is negative due to the energy penalty associated with confining two holes in the same ring. X^- is the least stable of the complexes studied; it is predicted to be stable only at temperatures below 2.3 K, while the most stable complex, X^+ , endures up to 71 K. De-excitonisation and binding energies were shown to be largely independent of the aspect ratio at fixed ring volume and mean radius, suggesting these energies may also be fairly independent of the precise shape of the cross-section of the ring. The electrons form a halo around the outside of the ring, with a low density in the central cavity. This reflects the fact that the ring size is comparable with the free exciton Bohr

radius. Furthermore, X^+ and XX are the most tightly bound complexes, with a preference for the electrons to position themselves above and below the ring. For X^- , the electronic charge density is much more diffuse, consistent with its very small binding energy. Finally, we note that there has been subsequent research in type-II quantum rings stemming from the publication of this work.

2.3 Type-II superlattices

2.3.1 Superlattices

Type-II superlattices are another type of heterostructure in which the electrons and holes are spatially separated. In essence, these superlattice heterostructures consists of multiple layers of quantum wells with alternating material composition, such that holes are confined in every other layer and electrons are confined in the layers between the hole-confining layers. This spatial separation of the charge carriers gives rise to properties similar to those for quantum dots or quantum rings, such as extended recombination times [96,97]. These heterostructures are therefore suited to a broadly similar set of applications, but there has been particular interest in light-emitting diodes [98,99].

In this section we focus on the superlattice (SL) and multiple-quantum-well (MQW) InAs/InAs_{1-x}Sb_x heterostructures (where $x \approx 0.05$) of Keen *et al.* [98,99]. These materials have a type-II band alignment with electrons confined to the InAs layers, and holes confined to the InAsSb layers. The SL structure studied by Keen *et al.* [99,100] consists of 50 pairs of alternating 14 nm thick InAs and InAsSb layers, while the MQW structure has 10 pairs of 40nm thick InAs layers and 10 nm thick InAsSb layers (see Fig. 2.8). Both the SL and MQW were produced using molecular beam epitaxy, see Ref. [98] for details. Transmission electron microscopy images of the structures show that this procedure produces good but not perfect layers [99].

Among various other experiments, Keen *et al.* [100] performed photolumines-

cence experiments on both the SL and MQW heterostructures and obtained multiple emission peaks. Here, we investigate the energetics of free excitonic complexes in type-II InAs/InAsSb SL and MQW heterostructures using DMC, to aid in identification of the origin of the peaks in the photoluminescence emission spectra. We use DMC to calculate the ground-state energies, and hence binding energies, of excitons and a variety of trions, biexcitons, and quintons and compare them with the experimental results of Keen *et al.* [100].

2.3.2 Computational details

Type-II band alignment ensures that electrons are largely confined to the InAs layers and holes are largely confined to the InAsSb layers. We therefore model the layers in both SLs and MQWs as one-dimensional infinite-square well potentials in the x direction, and allow the charge carriers to move freely in the y - z plane. Charge carriers of the same type (spin and charge) in different layers of these heterostructures are effectively distinguishable. We focus on complexes with at most two charge carriers in each layer; if more than two charge carriers occur in the same layer then the spatial wave function must be antisymmetric under exchange of those particles, destabilising the complex. The consequences of this distinguishability means that not only can we have quintons with 5 distinguishable particles but that we have multiple configurations of the different complexes. We define these by specifying the layers to which each particle belongs. For example: $XX(e|hh|e)$ represents a biexciton with a layer containing a single electron, followed by a layer containing two holes, with a single electron in the next layer; while $XX^-(e|h|e|h|e)$ is a negative quinton where each particle occupies a separate layer, see Fig. 2.8.

The electron and hole masses were taken to be isotropic, with values $m_e^* = 0.023 m_e$ and $m_h^* = 0.41 m_e$, respectively [93]. These are the electron and heavy-hole masses for InAs; the small concentration of Sb in the hole-confining layers has little effect on the hole mass [101]. The relative permittivity was taken to

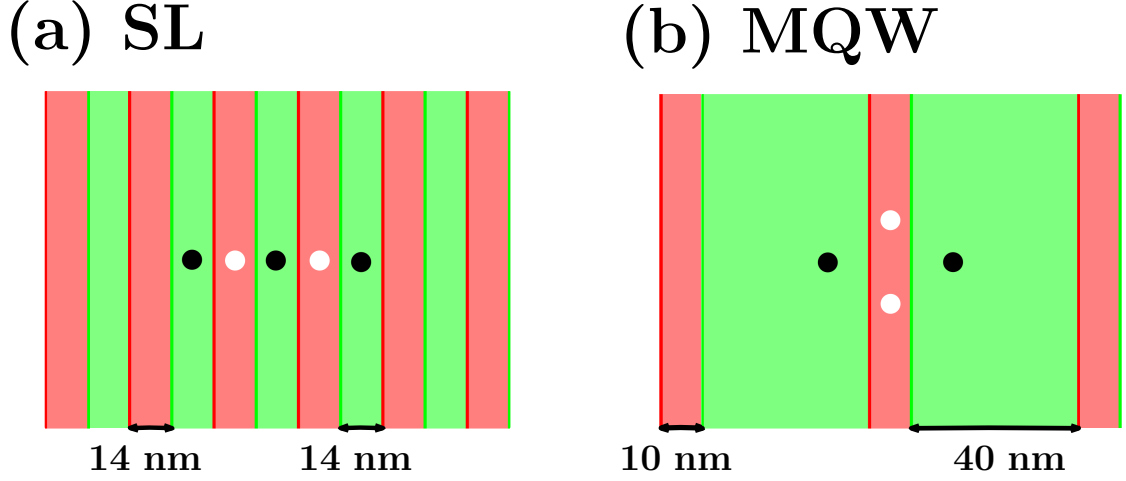


Figure 2.8: Diagrams showing (a) the superlattice (SL) and (b) the multiple quantum well (MQW) heterostructure models. The green coloured layers are the InAs electron-bearing layers and the red coloured layers are the InAsSb hole-bearing layers. The black and white circles represent electrons and holes, respectively. Hence, (a) shows the negative quinton $XX^-(e|h|e|h|e)$ and (b) shows the biexciton $XX(e|hh|e)$.

be isotropic and of value $\epsilon = 15.15 \epsilon_0$, which is the static relative permittivity of InAs [93]. The binding energies of the carrier complexes are given by

$$\begin{aligned}
 E_b^X &= E^{e^-} + E^{h^+} - E^X, \\
 E_b^{X^-} &= E^X + E^{e^-} - E^{X^-}, \\
 E_b^{X^+} &= E^X + E^{h^+} - E^{X^+}, \\
 E_b^{XX} &= 2E^X - E^{XX}, \\
 E_b^{XX^\pm} &= E^A + E^B - E^{XX^\pm},
 \end{aligned} \tag{2.19}$$

where E^i is the ground-state total energy for complex i , and the daughter products E^A and E^B of XX^\pm are defined in Table 2.5. These binding energies determine the thermodynamic stability of the complex with respect to dissociation into daughter products. Furthermore, if an exciton with an electron and a hole in neighbouring layers is one of the daughter products, the binding energy determines the wavelength of a peak in a photoluminescence measurement relative to the exciton peak. Below, our binding energies are quoted to two significant figures. In most cases the statistical error bar from the Monte Carlo calculation is in the third or fourth

significant figure, and so the error bars are omitted from our results. A much larger source of error arises from the approximations made in the model and the uncertainties in parameters such as the effective masses. The sensitivity to the model parameters has been investigated and the results are shown in Table 2.6.

For all complexes, each particle is distinguishable so the Slater part of the trial wave function $\Psi_T(\mathbf{R})$ is just a product of single particle orbitals. For the biexciton,

$$\Psi_T(\mathbf{R}) = e^{J(\mathbf{R})} \psi_e(x_{e_1}) \psi_e(x_{e_2}) \psi_h(x_{h_1}) \psi_h(x_{h_2}), \quad (2.20)$$

with the electron and hole orbitals taking the following form

$$\psi_e(x) = \begin{cases} \cos\left(\frac{\pi}{W_e}(x - x_c(x))\right) & \text{if } x_c(x) - \frac{W_e}{2} \leq x \leq x_c(x) + \frac{W_e}{2}, \\ 0 & \text{otherwise,} \end{cases} \quad (2.21)$$

and

$$\psi_h(x) = \begin{cases} \cos\left(\frac{\pi}{W_h}(x - x_c(x))\right) & \text{if } x_c(x) - \frac{W_h}{2} \leq x \leq x_c(x) + \frac{W_h}{2}, \\ 0 & \text{otherwise,} \end{cases} \quad (2.22)$$

where W_e and W_h are the widths of the electron- and hole-bearing layers, respectively, and $x_c(x)$ just returns the centre (in the x -direction) of the well the particle is occupying. These are just one-dimensional infinite-square-well orbitals confining electrons and holes to separate layers in the x direction and allowing them to move freely in the y and z directions. The Jastrow function $J(\mathbf{R})$ consists of u and EXJAS terms for the excitons, and u , f , and EXJAS terms for all other complexes; the cut-off lengths for these terms are large enough to encompass the entire complex. The EXJAS terms ensure the Kato cusp conditions are satisfied at short range. Energy minimisation is again found to be the best way to optimise the various trial wave functions. Time-step and finite-population-size biases were eliminated by performing pairs of DMC calculations with time steps in a 1 : 4 ratio

and target configuration populations in a 4 : 1 ratio, and extrapolating linearly to zero time step.

2.3.3 Results and discussion

The binding energy of an exciton with the electron and the hole in adjacent layers was found to be $E_b^{X(e|h)} = 1.9$ meV in the SL and $E_b^{X(e|h)} = 1.5$ meV in the MQW. The former of which agrees well with the exciton binding energy measured by Keen *et al.* [100] for SL samples ($E_b^X = 2.1$ meV and $E_b^X = 1.0$ meV, with Sb concentrations of 3.8% and 6.2%, respectively), but less well for MQW samples ($E_b^X = 5.6$ meV with a 4.3% Sb concentration). The binding energies of trions and biexcitons are reported in Table 2.4. In all cases, the complexes in the SL are more strongly bound than in the MQW due to the more 2D-like confinement of the electrons in the SL. Experimental photoluminescence spectra of a SL sample with 3.8% concentration show a peak approximately 6 meV lower in energy than the exciton peak, suggesting the existence of a complex with a 6 meV binding energy (see Fig. 5.25 of Ref. [100]). Keen *et al.* [100] believed this to be a biexciton because the intensity of the peak varied roughly quadratically with laser intensity, rather than the linear relationship between peak and laser intensity of the exciton. However, in our calculations, even the most strongly bound biexciton has a binding energy of only $E_b^{XX(e|hh|e)} = 0.55$ meV, which is considerably less than the energy difference between the exciton peak and the peak in question. It is therefore very

Hetero- structure	Binding energy (meV)							
	X ⁻		X ⁺		XX			
	e h e	ee h	h e h	hh e	e h e h	e hh e	h ee h	ee hh
SL	0.38	0.11	0.49	0.26	0.18	0.55	0.46	0.097(1)
MQW	0.33	0.061	0.42	0.15	0.18	0.36	0.36	0.0030(7)

Table 2.4: DMC binding energies of negative trions (X⁻), positive trions (X⁺), and biexcitons (XX) in both the SL and MQW geometries. The third row of the table indicates the layers to which the particles in each complex belong; *e.g.*, XX(e|hh|e) represents a biexciton with a layer containing one electron, followed by a layer containing two holes, then a layer containing the second electron.

Heterostructure	Products	Binding energy (meV)			
		XX ⁻			
		e h e h e	ee h e h	h ee h e	ee hh e
SL	X ⁻ + X	0.42	0.14	0.34	0.86
	XX + e ⁻	0.62	0.062(1)	0.27	0.41
MQW	X ⁻ + X	0.35	0.12	0.51	0.61
	XX + e ⁻	0.50		0.22	0.67
		XX ⁺			
		h e h e h	hh e h e	e hh e h	hh ee h
SL	X ⁺ + X	0.40	0.054(1)	0.32	0.76
	XX + h ⁺	0.70	0.13	0.25	0.57
MQW	X ⁺ + X	0.34	0.067(1)	0.46	0.62
	XX + h ⁺	0.58	0.036(1)	0.24	0.77

Table 2.5: DMC binding energies of a variety of quintons, both negative (XX⁻) and positive (XX⁺), in both the SL and MQW geometries, where the daughter products are either an exciton and a trion or a biexciton and a single charge carrier. The third row of the table indicates the layers to which the particles in each complex belong; *e.g.*, XX⁻(e|h|e|h|e) represents a negative quinton in which each particle occupies a separate layer. Empty entries represent complexes that were unbound in our calculations.

unlikely that this peak is due to a free biexciton.

Table 2.5 shows the binding energies of a variety of different quintons, or five-particle complexes. Quintons have two possible combinations of daughter products: they can decompose into either a trion plus an exciton or a biexciton plus a free charge carrier. Certain quinton configurations, for example XX⁻(e|h|e|h|e) in the SL, are more strongly bound than biexcitons; however, the magnitude of the binding energy is still not large enough to account for the “biexciton” peak of Keen *et al.* [100]. The data in Table 2.6 show that our binding-energy data are robust against large variations in the well widths, and small variations in the relative permittivity and the effective masses for almost every complex. The only exceptions to this are the single-hole complexes (X and X⁻), which are particularly sensitive to the electron mass. No reasonable adjustments to the model parameters could lead to the biexciton binding energy being consistent with the peak observed by Keen *et al.* [100].

The 3D exciton Bohr radius is 368 Å for InAs, which is significantly larger than

Complex	$\partial E_b/\partial m_e^*$ (meV/ m_e)	$\partial E_b/\partial m_h^*$ (meV/ m_e)	$\partial E_b/\partial \epsilon$ (meV)	$\partial E_b/\partial W_e$ (meV/nm)	$\partial E_b/\partial W_h$ (meV/nm)
X	3600	11	-12	-0.031	-0.034
X ⁻	3500	11	-11	-0.44	0.036
X ⁺	180	0.65	-0.61	0.035	-0.047
XX	-20	0.18	0.023(1)	-0.10	0.029

Table 2.6: Sensitivity of the SL complex binding energies E_b to the electron and hole effective masses m_e^* and m_h^* , the relative permittivity ϵ , and the electron and hole well widths W_e and W_h .

the well widths, so that the motion of the charge carriers in both the SL and MQW heterostructures is strictly 2D to a good approximation. Refs. [102] and [103] provide interpolation formulae for the binding energies of ideal 2D electron-hole-bilayer excitonic complexes as a function of the distance between the layers and the electron-hole mass ratio. As expected, these models show good agreement with the DMC binding energies in Table 2.4. For example, using the same effective masses and relative permittivities as the 3D DMC calculations reported here, and assuming an interlayer distance of 14 nm (appropriate for the SL), the binding energies of ideal 2D bilayer complexes are $E_b^{X(e|h)} = 2.0$ meV, $E_b^{X^-(ee|h)} = 0.14$ meV, $E_b^{X^+(hh|e)} = 0.26$ meV, and $E_b^{XX(ee|hh)} = 0.096$ meV, which are very similar to the SL results in Table 2.4.

While the DMC data rule out a free biexciton causing the “biexciton” photoluminescence peak observed by Keen *et al.* [100], it is in principle possible that a confined biexciton (*e.g.* due to variations in the well widths) might give rise to this peak. We may use the ideal 2D bilayer model to investigate the effects of reducing the layer widths. At an absurdly small interlayer distance of just 1 Å the biexciton binding energy is $E_b^{XX(ee|hh)} = 1.9$ meV, which is still not large enough to explain the “biexciton” peak of Keen *et al.* [100], implying that the peak in question cannot be due to a biexciton that is free to move in two dimensions. On the other hand, studies of InAs quantum dots show biexciton binding energies of up to 9 meV for type-I InAs quantum dots of diameter 25–30 nm in bulk AlAs [104]. While binding energies in a type-II heterostructure will inevitably be lower, it is

possible that an excitonic complex confined in all three dimensions could be the cause of the unexplained photoluminescence peak identified by Keen *et al.* [100].

2.3.4 Conclusion

We have calculated the binding energies of excitons and a variety of trions, biexcitons, and quintons in both SL and MQW heterostructures. We find that our SL exciton binding energy agrees well with experiment, and that SL binding energies of complexes in general are well described by an ideal 2D bilayer model. Furthermore, our binding energies show little dependence on model parameters such as the well widths, and only the single-hole complexes (X and X^-) show a strong dependence on the electron mass. While we were not able to identify origin of the “biexciton” peak observed by Keen *et al.* [100], we were, through use of our binding energy data and the binding energy formulae in Refs. [102] and [103], able to exclude a range of excitonic complexes and, in particular, show that it is very unlikely to be caused by a free biexciton.

Chapter 3

Monolayer transition metal dichalcogenides

3.1 Material properties

Transition metal dichalcogenides (TMDCs) are a class of layered materials that become direct gap semiconductors when in hexagonal monolayer form. The chemical composition of these 2D materials is MX_2 , where M is a transition metal and X a chalcogen; the most common ones (and the ones studied here) are MoS_2 , MoSe_2 , WS_2 and WSe_2 . The presence of a band gap in these materials gives them optoelectronic properties ideal for applications such as photodetectors and light-emitting diodes [105–107]. The conduction-band minimum and valence-band maximum occur at the \mathbf{K} and \mathbf{K}' points of the hexagonal Brillouin zone; charge-carriers in TMDCs therefore have an additional valley degree of freedom, as well the usual spin degree of freedom. Strong spin-orbit coupling in these materials leads to significantly spin-split conduction and valence bands [108]. In MoX_2 materials the conduction-band minimum and the valence-band maximum, in each valley, have the same spin, while in WX_2 materials the minimum and maximum have opposite spins [108], as shown in Fig. 3.1. The degree of spin-splitting in the valence bands is large enough that holes are present only in the upper spin-split




Figure 3.1: Diagram illustrating the conduction band minimum and valence band maximum at the \mathbf{K} (left) and \mathbf{K}' (right) valleys in (a) MoX₂ and (b) WX₂ materials. Spin-up bands are red, and spin-down bands are blue. At room temperature the lower spin-split valence band is not expected to contain any holes.

valence band, but in the conduction band the spin-splitting is small enough that both spin-split bands may contain electrons at room temperature [108]; this allows for 4 species of distinguishable (spin and valley polarised) electrons and 2 species of distinguishable (spin polarised) holes [68].

The presence of these electrons and holes in the bands of TMDCs leads to strong excitonic effects in their photoluminescence spectra [109–111]. The binding and de-excitonisation energies¹ of excitonic complexes in these materials have been well studied using QMC methods [68, 112], and have helped with the identification of lines in the photoluminescence spectra of TMDCs, and even predicted the existence of stable quintons, which was later confirmed experimentally [113, 114].

In this chapter, we investigate the effects of an in-plane electric field on the binding energies of excitonic complexes in TMDCs and calculate their polarisabilities. We then study the properties of a two-dimensional electron gas in doped MoSe₂.

¹For free complexes the two are equivalent.

3.2 The Keldysh interaction

When describing the interactions of quasi-electrons and quasi-holes in semiconductor materials we must account for electrostatic screening effects. In Chapter 2, the bare Coulomb interaction was screened isotropically with magnitude ϵ , and implicitly taken care of by our use of excitonic units. The situation for two-dimensional materials is somewhat different. In this section we show ϵ only for illustrative purposes (it is unity in excitonic units).

The two-dimensional nature of these materials results in a polarisation of the crystal that modifies the interaction between charge carriers. Consider placing a charge density

$$\rho(\mathbf{r}) = \rho(x, y)\delta(z) \quad (3.1)$$

in a two-dimensional semiconductor at $z = 0$. The electric displacement field $\mathbf{D}(\mathbf{r})$ resulting from this charge density is given by

$$\mathbf{D}(\mathbf{r}) = \epsilon\mathbf{E}(\mathbf{r}) + \mathbf{P}(\mathbf{r}) = -\epsilon\nabla_{\mathbf{r}}\phi(\mathbf{r}) + \mathbf{P}_{\perp}(x, y)\delta(z), \quad (3.2)$$

where $\mathbf{E}(\mathbf{r}) = -\nabla_{\mathbf{r}}\phi(\mathbf{r})$ is the electric field with associated electrostatic potential $\phi(\mathbf{r})$, and $\mathbf{P}(\mathbf{r})$ is the polarisation vector. The polarisation vector can be expressed as $\mathbf{P}(\mathbf{r}) = \mathbf{P}_{\perp}(x, y)\delta(z)$ because the charge density $\rho(\mathbf{r})$ has no component in the z direction (*i.e.* the charge lies in plane), and we can write the in-plane polarisation vector as $\mathbf{P}_{\perp}(x, y) = -\kappa\nabla_{\mathbf{r}}\phi(x, y, 0)$, where κ is the in-plane susceptibility of the material. Using Gauss's law

$$\nabla_{\mathbf{r}} \cdot \mathbf{D}(\mathbf{r}) = \rho(\mathbf{r}) = \rho(x, y)\delta(z), \quad (3.3)$$

we can solve for the electrostatic potential $\phi(x, y, 0)$ to obtain the in-plane electrostatic potential due to the charge density $\rho(\mathbf{r})$, see Ref. [68] for more details.

Using this process we can define an effective interaction $v(r)$ for charge carriers interacting in a two-dimensional material. For a pair of particles with charges q_i

and q_j , separated in-plane by a distance $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$, we find

$$v_K(r_{ij}) = \frac{q_i q_j}{\epsilon} \frac{\pi}{2r_*} \left[H_0\left(\frac{r_{ij}}{r_*}\right) - Y_0\left(\frac{r_{ij}}{r_*}\right) \right], \quad (3.4)$$

where, the functions $H_0(r_{ij})$ and $Y_0(r_{ij})$ are, the zeroth-order Struve function and the zeroth-order Bessel function of the second kind, respectively, and we have defined the screening parameter $r_* = \kappa/(2\epsilon)$. This is the (Rytova-)Keldysh interaction [115–117]. At long range, or small susceptibility, ($r_{ij} \gg r_*$) the Keldysh interaction reduces to the (screened) Coulomb interaction

$$v_C(r_{ij}) = \frac{q_i q_j}{\epsilon} \frac{1}{r_{ij}}, \quad (3.5)$$

while at short range, or large susceptibility, ($r_{ij} \ll r_*$) the interaction has a logarithmic form

$$v_{\log}(r_{ij}) = \frac{q_i q_j}{\epsilon} \frac{1}{r_*} \log\left(\frac{2r_*}{e^\gamma r_{ij}}\right), \quad (3.6)$$

where we use the natural logarithm and γ is Euler's constant. Eq. (3.6) is just the solution of a 2D Poisson equation *i.e.* the Coulomb interaction in only two dimensions.

For a system of N electrons and/or holes interacting via the Keldysh interaction, the Schrödinger equation is

$$\left[-\frac{1}{2} \sum_{i=1}^N \frac{1}{m_i} \nabla_{\mathbf{r}_i} + \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ i \neq j}}^N v_K(r_{ij}) \right] \Psi(\mathbf{R}) = E \Psi(\mathbf{R}), \quad (3.7)$$

where $v_K(r_{ij})$ is defined as in Eq. (3.4).

3.3 Electric fields

While experimental techniques are able to observe peaks in the photoluminescence spectra of semi-conducting materials it is somewhat more challenging to identify

the origin of the individual peaks. We briefly discussed this in Chapter 2 when we tried to identify the origin of a particular peak on the photoluminescence spectra of type-II superlattices. Here, we study the effect of an applied electric field on the binding energies of excitonic complexes, in an attempt to understand whether the shift of photoluminescence peaks could be a useful way to experimentally identify the origin of such peaks. We will first discuss the physics, and computational details, of applying an electric field to an excitonic system in a two-dimensional semiconductor.

A bias voltage ΔV applied to a monolayer semiconductor results in an in-plane electric field $F \approx -\Delta V/d$, where d is the distance between the terminals. The excitonic unit of electric field is $F^* = \mu^2 e^5 / [(4\pi\epsilon)^3 \hbar^4]$. The precise form of the electric field depends on the device geometry. Here, we assume a uniform electric field for simplicity. The electric field F will perturb the energies of charge-carrier complexes (in the centre of mass frame for charged complexes). We thus investigate the effects of F on the binding energies of charge-carrier complexes by including an additional term in the Hamiltonian

$$\hat{H}_E = \hat{H} - \sum_{i=1}^N q_i F x_i. \quad (3.8)$$

For neutral-charge complexes, including the effect of the electric field in the VMC and DMC algorithms described in Chapter 1 requires only that we add the contribution to the electrostatic energy due to the field F each time we evaluate the local energy. For charged complexes, we perform the calculation in the centre of mass frame otherwise the total energy of the complex would be undefined; the energy will become arbitrarily low as the complex seeks the potential minimum.

We investigate only complexes containing distinguishable particles but in this case the particles are not confined, therefore our trial wave functions lack a Slater part and is of Jastrow form

$$\Psi_T(\mathbf{R}) = e^{J(\mathbf{R})}. \quad (3.9)$$

In the Jastrow function we use u and EXJAS terms for all the free complexes and also include the H term (a three-body generalisation of the u term) for complexes containing at least three particles. For the donor-atom (D^0) we use the χ term instead of the u term. To account for the difference between the Coulomb and Keldysh interactions at short range, the EXJAS term is slightly modified such that

$$u_{ij}^{\text{EXJAS}}(r_{ij}) = \frac{[a_{ij} + \Gamma_{ij}^{\text{K}} \log(r_{ij}) + b_{ij} r_{ij}^2] r_{ij}^2}{1 + c_{ij} r_{ij}^2}, \quad (3.10)$$

where

$$\Gamma_{ij}^{\text{K}} = -\frac{q_i q_j \mu_{ij}}{2r_*}, \quad (3.11)$$

so that $\Psi_{\text{T}}(\mathbf{R})$ satisfies the Kato cusp conditions [68, 118]. Trial wave functions are first optimised by variance minimisation (ensuring we find a bound state), and then by energy minimisation. These trial wave functions are node-less so DMC exactly solves the model. We perform DMC calculations with time steps in a 1 : 4 ratio and vary target configuration populations inversely to the time step. We then extrapolate the DMC energies to zero time step and infinite populations to obtain ground-state total energies for each complex. Binding energies for each of the complexes are defined by

$$\begin{aligned} E_{\text{b}}^{\text{D}^0} &= -E^{\text{D}^0} \\ E_{\text{b}}^{\text{X}} &= E^{\text{e}^-} + E^{\text{h}^+} - E^{\text{X}} = -E^{\text{X}}, \\ E_{\text{b}}^{\text{X}^-} &= E^{\text{X}} + E^{\text{e}^-} - E^{\text{X}^-} = E^{\text{X}} - E^{\text{X}^-}, \\ E_{\text{b}}^{\text{X}^+} &= E^{\text{X}} + E^{\text{h}^+} - E^{\text{X}^+} = E^{\text{X}} - E^{\text{X}^+}, \\ E_{\text{b}}^{\text{X}^{\text{X}}} &= 2E^{\text{X}} - E^{\text{X}^{\text{X}}}, \end{aligned} \quad (3.12)$$

where E^i represents the ground-state total energy of complex i .

Fig. 3.2 plots the exciton (X) binding energy shift as a function of electric field strengths (Vnm^{-1}) for monolayer MoS_2 , MoSe_2 , WS_2 , WSe_2 , both in vacuum and encapsulated by hexagonal born nitirde (hBN), using the *ab-initio* parameters

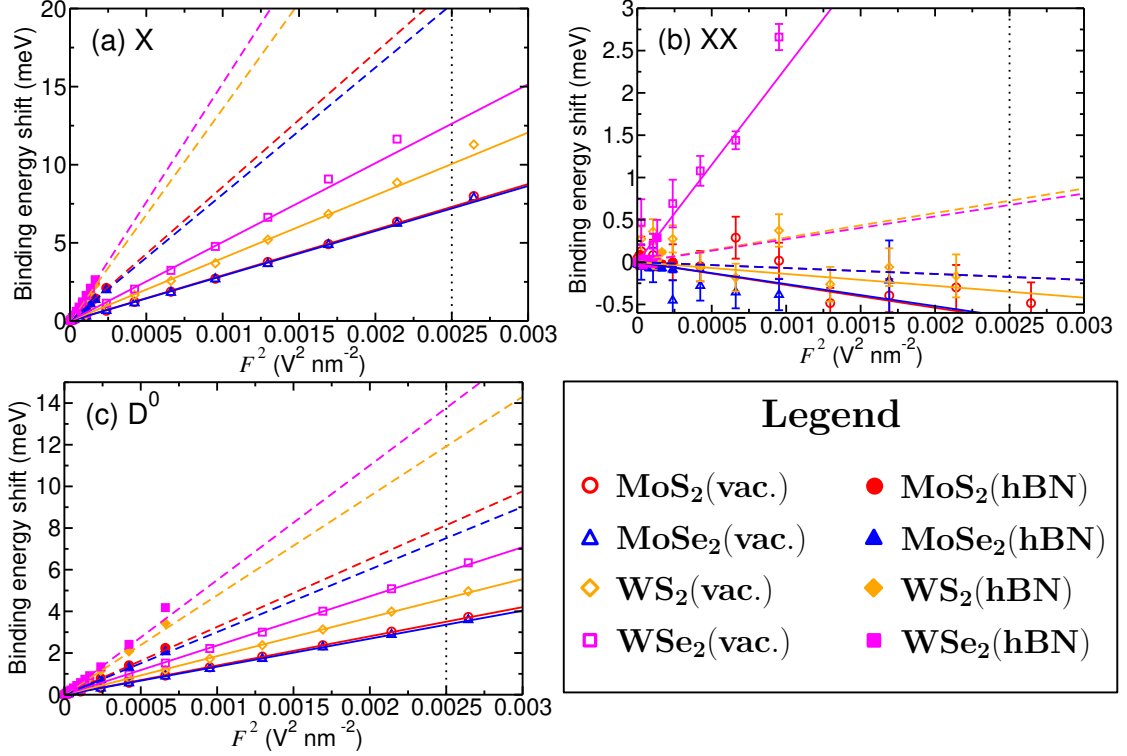


Figure 3.2: DMC binding energy shift for (a) X, (b) XX, and (c) D^0 as a function of F^2 for different monolayer TMDCs in vacuum and encapsulated by hBN. Error bars in (a) and (c) are smaller than the symbols. The solid and dashed lines show the binding energies determined by the polarisabilities in Table 3.2, for monolayer TMDCs in vacuum and encapsulated by hBN. The vertical dotted lines correspond to $F = 50 \text{ mV nm}^{-1}$, beyond which VMC energy minimisation does not result in bound-state wave functions.

TMDC	Material parameter		
	m_e^* (m_e)	m_h^* (m_e)	r_* (a_0)
MoS ₂	0.35 [111]	0.428 [111]	72.98 [111]
MoSe ₂	0.38 [119]	0.44 [119]	75.19 [120]
WS ₂	0.27 [119]	0.32 [119]	71.60 [121]
WSe ₂	0.29 [119]	0.34 [119]	85.25 [121]

Table 3.1: Material parameters for different monolayer TMDCs in vacuum. m_e^* and m_h^* are the electron and hole effective masses, respectively, and r_* is the Keldysh screening parameter for the material in vacuum. r_* should be scaled down by $\epsilon = 4\epsilon_0$ when encapsulated by hBN.

in Table 3.1 and $\epsilon = 4\epsilon_0$ for hBN encapsulation. In each material environment, the X binding energy goes as the square of the magnitude of the in-plane electric field, as expected for a linearly polarisable exciton [122]. The total energy of an isolated neutral complex of polarisability α^i in a uniform external electric field of

magnitude F is

$$E^i = E_{F=0}^i - \alpha^i F^2/2, \quad (3.13)$$

where $E_{F=0}^i$ is the energy of the complex in the absence of external fields [122]. At sufficiently large electric fields ($F \approx 50 \text{ mV nm}^{-1}$) the complexes become unbound in QMC calculations. The variation of energy with electric field strength remains quadratic up to this point. For larger electric field strengths ($F > 50 \text{ mV nm}^{-1}$) we find that optimising wave functions by VMC variance or energy minimisation does not result in bound-state wave functions. If the parameters in the wave function are fixed such that a bound state is forced, the resulting DMC calculations are unstable. It is possible that some, or all, complexes remain bound at these larger electric fields, and for QMC calculations to become unstable due to the choice of trial wave function. Ours is isotropic, so does not allow for the complex to polarise at the VMC level; however, for the complexes studied here, the node-less trial wave functions ensure that when DMC works, it gives the exact ground-state energy of the complex. Use of a polarisable trial wave function may lead to more stable QMC calculations at higher field strengths.

TMDC	Polarisability ($\text{eV nm}^2 \text{V}^{-2}$)				
	X	XX	D ⁰	X ⁻	X ⁺
MoS ₂ (vac.)	5.84(2)	11.14(8)	2.802(9)	66(6)	44(6)
MoSe ₂ (vac.)	5.76(2)	11.0(1)	2.687(9)	80(9)	45(6)
WS ₂ (vac.)	8.04(3)	15.8(1)	3.70(1)	108(10)	72(7)
WSe ₂ (vac.)	10.10(4)	24.8(3)	3.96(1)	130(16)	118(9)
MoS ₂ (hBN)	17.17(4)	34.2(3)	6.51(2)	179(17)	161(22)
MoSe ₂ (hBN)	16.22(4)	32.3(2)	6.89(2)	211(22)	181(23)
WS ₂ (hBN)	27.16(4)	54.9(3)	4.95(1)	316(27)	246(32)
WSe ₂ (hBN)	30.43(4)	61.4(3)	5.29(1)	409(32)	367(32)

Table 3.2: Theoretical polarisabilities of X, XX, D⁰, X⁻, and X⁺ in monolayer TMDCs both in vacuum and encapsulated by hBN.

The XX and D⁰ binding energy vary linearly with F^2 , see Fig. 3.2. However, while the donor-atom binding energies increase with F^2 , the XX binding energies decrease. For a four-particle complex, alignment of charges in the direction of the applied field places like charges closer together, and reduces the binding energy

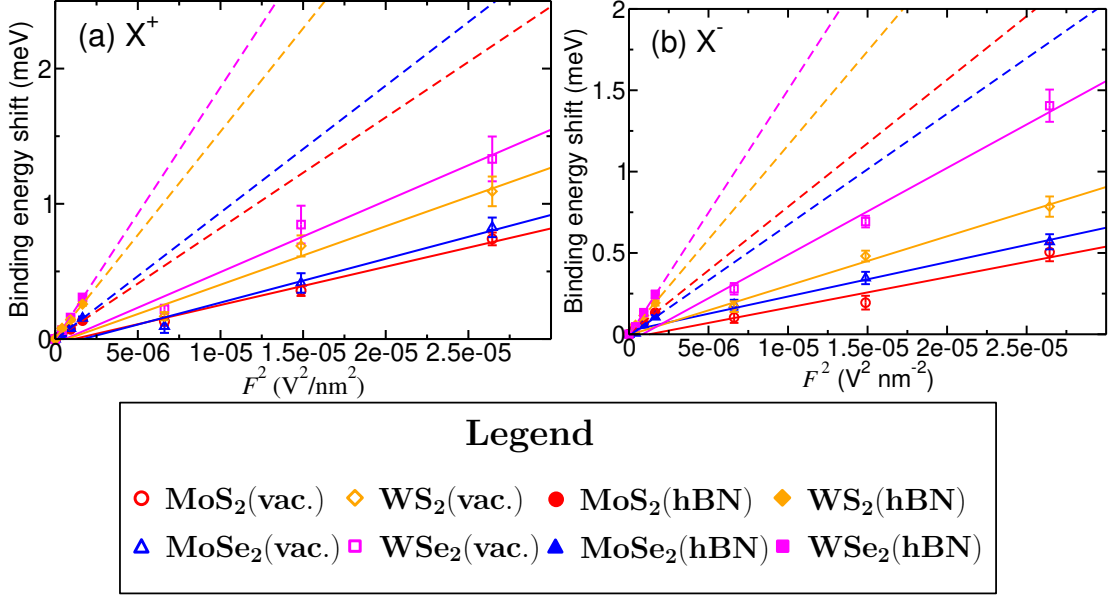


Figure 3.3: DMC binding energy shift for (a) X^- , and (b) X^+ as a function of F^2 for different monolayer TMDCs in vacuum and encapsulated by hBN. Where error bars are not visible they are smaller than the symbols. The solid and dashed lines show the binding energies determined by the polarisabilities in Table 3.2, for monolayer TMDCs in vacuum and encapsulated by hBN.

with respect to dissociation into two-particle complexes. It would be expected then that the polarisability of XX would be slightly less than twice the polarisability of an exciton, as seen by the polarisability values in Table 3.2. Table 3.2 also shows the polarisability of the donor-atom to be less than that of the exciton. Both of the trions binding energies also vary linearly with F^2 , see the plots in Fig. 3.3. However, they vary more strongly and QMC calculations become unstable at much lower F . This is reflected in the correspondingly higher polarisabilities for trions than for the neutral complex in the same material environments. Polarisability values for trions are reported in Table 3.2.

The predicted binding energy shifts of each of the complexes are in Table 3.3 for monolayer TMDCs in vacuum, and encapsulated by hBN, subject to $F = 50 \text{ mV nm}^{-1}$, an achievable order-of-magnitude in-plane electric field strength. The shifts in the peaks of the trions are so large that, at the very least, they should be able to be experimentally identified from the neutral complexes when an electric field is applied. Identification of a positive trion from a negative trion is likely to be possible in some materials/environments but not all. For the neutral

TMDC	Binding energy shift (meV)				
	X	XX	D ⁰	X ⁻	X ⁺
MoS ₂ (vac.)	7.3	< 1	3.5	76	48
MoSe ₂ (vac.)	7.2	< 1	3.4	93	49
WS ₂ (vac.)	10.1	< 1	4.6	125	80
WSe ₂ (vac.)	12.6	5.8	4.9	150	135
MoS ₂ (hBN)	21.5	< 1	8.1	201	180
MoSe ₂ (hBN)	20.3	< 1	8.6	243	206
WS ₂ (hBN)	34.0	< 1	6.2	362	273
WSe ₂ (hBN)	38.0	< 1	6.6	473	408

Table 3.3: Calculated binding energy shifts of X, XX, D⁰, X⁻, and X⁺ using Eq. (3.13) and the polarisabilities in Table 3.2 for monolayer TMDCs in vacuum, and encapsulated by hBN, for $F = 50 \text{ mV nm}^{-1}$. Note, not all complexes remain bound at this electric field strength.

complexes, the differences of only a few meV in the binding energy shifts suggest that complexes are unlikely to be experimentally identified by the shifts of their respective peaks when an electric field is applied.

In conclusion, binding energies of excitonic complexes vary quadratically with the strength of the applied in-plane electric field up to $F = 50 \text{ mV nm}^{-1}$. For neutral complexes, the differences of only a few meV in the binding energy shifts suggest that these complexes are unlikely to be experimentally identified by the shifts of their respective peaks when an electric field is applied. While for trions, the shifts of the binding energies when an in-plane electric field is applied should result in the trion and biexciton peaks shifting apart, significantly so for hBN-encapsulated materials.

3.4 The two-dimensional electron gas

The homogeneous electron gas is fundamental to the study of condensed matter physics. In traditional semiconductor devices there often exists a two-dimensional electron gas at heterostructure interfaces [66, 123, 124], knowledge of the properties of the two-dimensional electron gas is therefore crucial in a proper understanding of such devices.

In low density metallic systems, the electron gas can undergo a phase transi-

tion from a fluid (de-localised) phase to an insulating crystalline (localised) phase, or Wigner crystal [125]. The density at which this transition occurs in two dimensions is much lower than the electron density at which most semiconductor devices operate, therefore most of the interest in the transition density is scientific [126–129]. The study of electron gases in two dimensions has, so far, used a screened Coulomb interaction between the electrons, and while this is suitable for 2D electron gases in 3D semiconductor heterostructures, one wonders whether the use of a Keldysh interaction might be more appropriate for 2D electron gases in 2D systems such as TMDCs. Indeed, the subject of Wigner crystallisation in TMDCs is of considerable current interest [129–131]. In this section, we investigate the Wigner crystallisation density of the 2D electron gas using a periodic Keldysh interaction, rather a periodic Coulomb interaction.

3.4.1 Ewald interaction

When studying systems composed of a thermodynamic quantity of electrons we must, by necessity, restrict our system size to what is computationally possible. We therefore approximate the, essentially infinite, system by applying periodic boundary conditions to a supercell containing a computationally feasible number of electrons. However, we inevitably end up introducing finite-size errors into our results that depend on the size and shape of the supercell. It is desirable, then, to perform calculations in a few different supercells and extrapolate the results to the thermodynamic limit in an attempt to remove these finite-size effects. One type of finite-size effect that arises in periodic supercells is related to the method used to overcome the problematic behaviour of the Coulomb interaction at long range in periodic systems. A periodic Keldysh interaction has similar issues, since the Keldysh interaction reduces to the Coulomb interaction at long range. We will therefore focus our attention on the problem of, and the solution to, using a periodic Coulomb interaction, and then, briefly, how it differs for a periodic Keldysh interaction. Further details on performing QMC calculations in periodic

supercells can be found in Chapter 4, along with a discussion of finite-size effects, with a particular interest in so-called momentum-quantisation errors.

Consider a supercell containing N charged particles interacting via the Coulomb interaction. The potential energy, per supercell, of a periodic lattice of N_s such supercells is given by

$$V(\mathbf{R}) = \frac{1}{2N_s} \sum_{i=1}^N \sum_{j=1}^N \sum'_{\{\mathbf{L}_s\}} v_C(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|), \quad (3.14)$$

where $\{\mathbf{L}_s\}$ is the set of supercell lattice vectors, the prime on the sum over j indicates that we exclude the $i = j$ term when $\mathbf{L}_s = \mathbf{0}$, and $v_C(r)$ is the Coulomb interaction as defined in Eq. (3.5). An equivalent sum can be obtained for the Keldysh interaction simply by replacing $v_C(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|)$ with $v_K(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|)$. There are two things to note here. First, the interactions of a particle with its own periodic images are included; this results in the crystal having a non-zero surface polarisation even for electrically neutral crystals. Second, the lattice sum

$$v(\mathbf{r}_i, \mathbf{r}_j) = \sum_{\{\mathbf{L}_s\}} v_C(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|), \quad (3.15)$$

diverges in a non-neutral supercell, and is only conditionally convergent in a neutral supercell, making it unfeasible for use in a QMC calculation where we need to repeatedly evaluate the potential energy. The most common solution to these problems is to replace the lattice sum in Eq. (3.15) with the Ewald interaction $v_E(\mathbf{r}_i, \mathbf{r}_j)$ [132] plus a Madelung constant.²

The Ewald interaction is the electrostatic potential, and solution of a periodic Poisson equation, for a lattice of charges embedded in a neutralising background. This neutralising background is equivalent to applying “tin-foil” boundary conditions *i.e.* surrounding the crystal with a perfectly metallic medium that cancels out any surface polarisation charges. For computational efficiency, the interaction is

²The Madelung constant is the energy of the interaction between a charge and all its periodic images.

computed in two parts: a short-range real-space sum, and a long-range reciprocal-space sum. The three-dimensional Ewald interaction as implemented in CASINO is

$$v_E(\mathbf{r}_i, \mathbf{r}_j) = \sum_{\{\mathbf{L}_s\}} \frac{\operatorname{erfc}(\zeta|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|)}{|\mathbf{r}_i - \mathbf{r}_j - \mathbf{L}_s|} - \frac{\pi}{\Omega_s \zeta^2} + \frac{4\pi}{\Omega_s} \sum_{\{\mathbf{G}_s\} \setminus \{\mathbf{0}\}} \frac{1}{G_s^2} e^{\left(-\frac{G_s^2}{4\zeta^2} + i\mathbf{G}_s \cdot (\mathbf{r}_i - \mathbf{r}_j)\right)}, \quad (3.16)$$

where $\operatorname{erfc}(r)$ is the complementary error function, Ω_s is the volume of the supercell, and $\{\mathbf{G}_s\}$ denotes the set of supercell reciprocal lattice vectors [62]. The parameter ζ has no effect on the final value of $v_E(\mathbf{r}_i, \mathbf{r}_j)$, and so can be optimised for maximal computational efficiency; in CASINO the default value is $\zeta = (2.8/\Omega_s^{\frac{1}{3}})^2$ [133]. There also exist a two-dimensional version of the Ewald interaction [134].

As the Keldysh interaction reduces to the Coulomb interaction at long range, we require an Ewald-like version of the Keldysh interaction in order to use the Keldysh interaction in periodic systems. A derivation by R. J. Hunt of such a form can be found in Ref. [135], and it is this form we use to study the two-dimensional electron gas. The derivation is based on the idea that the Keldysh interaction can be re-written as a Coulomb interaction (hence we can make use of the Ewald interaction) and a correction term. This correction term, the difference the Keldysh and Coulomb interactions, is absolutely convergent.

3.4.2 Wigner crystallisation

The system we study is that of weakly-to-moderately doped monolayer MoSe₂, both in vacuum and encapsulated by hBN, such that there exists a population of electrons in the lower of the spin-split conduction bands of both the \mathbf{K} and \mathbf{K}' valleys, see Fig. 3.4. As there are no conduction electrons in the upper spin-split conduction band we consider, at most, only two species of electrons. We study three phases of the system, a paramagnetic fluid phase, and ferromagnetic and antiferromagnetic Wigner crystal phases, over a range of densities with the aim of




Figure 3.4: Diagram illustrating the conduction band filling in both the \mathbf{K} (left) and \mathbf{K}' (right) valleys in our model of doped MoSe₂. Spin-up bands are red, and spin-down bands are blue, the Fermi energy E_F is shown by the orange line with the shading representing the filling of the bands.

determining the Wigner crystallisation density. We define our densities using the Wigner-Seitz radius r_s , which is the radius of the circle that (on average) contains one electron. The relationship between r_s in (effective) Hartree atomic units³ to the number density n in (experimentally relevant) units of cm⁻² is

$$r_s = \frac{m_e^* 1.06616 \times 10^8}{\epsilon \sqrt{n}} = \begin{cases} \frac{4.04 \times 10^7 \text{cm}^{-1}}{\sqrt{n}}, & \text{in MoSe}_2 \text{ in vacuum} \\ \frac{1.01 \times 10^7 \text{cm}^{-1}}{\sqrt{n}}, & \text{in hBN-encapsulated MoSe}_2, \end{cases} \quad (3.17)$$

where we use the parameters of MoSe₂ defined in Table. 3.1, and $\epsilon = 4\epsilon_0$ for hBN-encapsulation. The Keldysh screening parameter is $r_* = 28.5728$ and $1.7858 a_0^*$, for MoSe₂ in vacuum and encapsulated by hBN, respectively.

Our trial wave functions are of Slater-Jastrow-backflow form. In both fluid and crystal phases the Jastrow function is composed of u and p terms, while the backflow function consists of just an η term. To satisfy the Kato cusp conditions we add a minimal cusp-satisfying term to the ordinary u term of the following form

$$u_{\mathbf{K}}(r_{ij}) = - \left(1 - \frac{r_{ij}}{L_{u\mathbf{K}}} \right)^C \Theta(L_{u\mathbf{K}} - r_{ij}) \Gamma_{ij}^{\mathbf{K}} r_{ij}^2 \log(r_{ij}), \quad (3.18)$$

where we fix the cut-off distance $L_{u\mathbf{K}} = 1 a_0^*$. For the fluid phase calculations (performed by R. J. Hunt) the single-particle orbitals in the Slater determinant are just plane-wave orbitals. Further details of the fluid phase calculations can be

³Similar to Hartree atomic units but with $e = m_e^* = \hbar = 4\pi\epsilon = 1$.




Figure 3.5: Triangular lattice used in the Wigner crystal calculations. In a ferromagnetic Wigner crystal all sites are occupied by spin-up, \mathbf{K} valley electrons. In the antiferromagnetic Wigner crystal spin-up, \mathbf{K} valley electrons occupy the sites marked with red circles, while spin-down, \mathbf{K}' valley electrons occupy the sites marked with blue squares.

found in Ref. [135]. Here, we focus on the crystal phase calculations, where the single-particle orbitals are Gaussians centred on sites of a triangular lattice⁴

$$\psi_{\mathbf{P}_i}(\mathbf{r}_i) = e^{-g|\mathbf{r}_i - \mathbf{P}_i|^2}, \quad (3.19)$$

where \mathbf{P}_i is the crystal lattice site associated with electron i , and g is an optimisable parameter controlling the width of the Gaussian orbital. For the ferromagnetic crystal, each Wigner crystal lattice site is occupied by a spin-up, \mathbf{K} valley electron, while for antiferromagnetic crystals we have equal numbers of spin-up, \mathbf{K} valley, and spin-down, \mathbf{K}' valley electrons occupying alternate stripes of the Wigner crystal lattice, see Fig. 3.5. In these calculations the Gaussian parameter g is optimised using VMC, however, the VMC-optimised energy minimum does not occur at the same value of g as the DMC-optimised⁵ minimum.

To handle finite-size effects, we perform calculations in two sizes of supercell containing $N = 64$ and $N = 100$ electrons. In each supercell, at each density trial

⁴The sites of this Wigner crystal lattice are not related to the sites of underlying lattice of the TMDC crystal.

⁵By DMC-optimised we mean performing DMC calculations over a range of values of g to find the DMC minimum energy as function of g . Trial wave functions should re-optimised (using VMC variance and energy minimisation) for each new value of g .

wave functions were optimised first by variance minimisation without backflow, and then by energy minimisation with the inclusion of backflow. DMC calculations for the crystal phase calculations were performed at time steps of 2 and 8 (Ha^*)⁻¹, and target configuration populations varied inversely to time step with a minimum size of 512 configurations. DMC energies were then simultaneously extrapolated to zero time step and infinite population. At each density the DMC energy per electron was then extrapolated to the thermodynamic limit using an $\mathcal{O}(N^{-\frac{5}{4}})$ scaling [136].

To extract the crystallisation energy from the fluid and crystal energy data we have performed fits as a function of r_s . In the fluid phase the fit is to the difference between the Keldysh and Coulomb total energies per electron, see Ref. [135] for details. The fit for crystal data is predicated on the basis that the Keldysh interaction reduces to the Coulomb interaction at very low densities (*i.e.* assuming $r_s \gg r_*$). We therefore use the same fit as previous studies of the two-dimensional (Coulomb) Wigner crystal [137, 138]. Our Keldysh total energies per electron E_K/N are fitted to a polynomial in $r_s^{\frac{1}{2}}$

$$\frac{E_K}{N} = \frac{-1.106103}{r_s} + \frac{0.814}{r_s^{3/2}} + \frac{c}{r_s^2} + \frac{d}{r_s^{5/2}} + \frac{e}{r_s^3}, \quad (3.20)$$

where the optimal fitting parameters are found to be $c = 1.81(2)$, $d = 0.2(2)$, and $e = -4.8(8) a_0^*$. The coefficients of the first two terms are fixed by the (Coulomb) Madelung energy of the lattice, and the zero-point energy of the Coulomb 2D Wigner crystal [138], respectively.

For hBN encapsulated MoSe₂, with $r_* = 1.7858 a_0^*$, our energy data and fits, for both fluid and Wigner crystal phases, are shown in the phase diagram in Fig. 3.6. The crystallisation density and its error were obtained using a Monte Carlo bootstrap procedure [139], and found to be $r_s^c = 36.4(2) a_0^*$. The differences in the energies of the ferromagnetic and antiferromagnetic crystal phases are too small for us to obtain an antiferromagnetic to ferromagnetic transition density. For MoSe₂ in vacuum with $r_* = 25.5728 a_0^*$ (these calculations were performed by N.

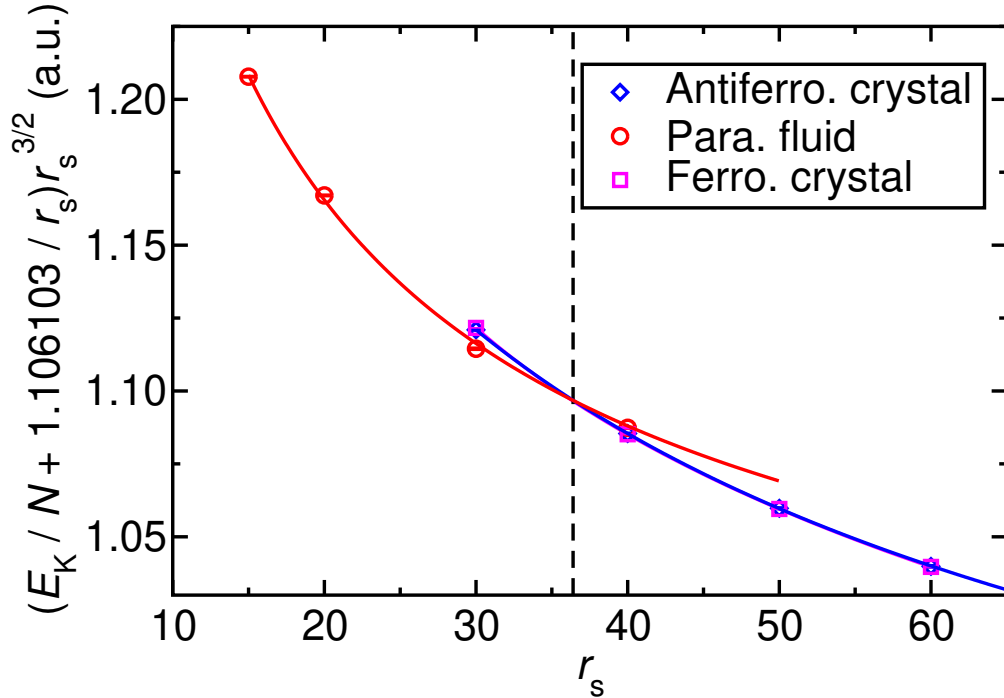


Figure 3.6: Phase diagram of a two-dimensional Keldysh electron gas showing the paramagnetic fluid phase, and the ferromagnetic and antiferromagnetic Wigner crystal phases. The vertical dashed line shows the crystallisation density. The screening length is $r_* = 1.7858 a_0^*$ and corresponds to MoSe₂ encapsulated by hBN.

D. Drummond), the crystallisation density is $r_s^c = 32.5(1) a_0^*$. The crystallisation densities of the Keldysh electron gas for MoSe₂ encapsulated by hBN, and in vacuum, correspond to $n^c = 7.73(8) \times 10^{10} \text{ cm}^{-2}$, and $n^c = 1.55(1) \times 10^{12} \text{ cm}^{-2}$, respectively. Comparing to experiments, Qiu *et al.* [140] achieved an electron gas in a sample of MoS₂ of density $n = 1 \times 10^{11} \text{ cm}^{-2}$, suggesting that experimental detection of Wigner crystals in monolayer TMDCs may be possible.

We can compare our results to the crystallisation density in the 2D Coulomb electron gas, which is essentially the $r_* \rightarrow 0$ limit of the Keldysh interaction, where $r_s^c = 31(1) a_0^*$ [126]. Use of the Keldysh interaction has the effect of lowering the Wigner crystallisation density when compared the Coulomb case. This is expected because the strongly-repulsive short-range Coulomb interaction is screened in the Keldysh electron gas. Consider, however, the behaviour of the crystallisation density r_s^c as a function the screening parameter r_* . One would expect that as the screening is increased, and the strength of the interaction between electrons at

short range decreased, this would raise the value of r_s^c (*i.e.* lower the crystallisation density) monotonically as a function of r_* . This is not what we see in our data. In comparison to do what we do here, the previous study [126], used DMC-optimised parameters g in the Gaussian orbitals for Wigner crystal calculations, while ours are optimised with VMC. A recent investigation by R. J. Hunt has suggested that the difference in DMC energies between optimising the orbital parameter g in VMC and DMC is significant in our calculations. DMC calculations using orbitals with a DMC-optimised orbital parameter g are then desirable, however these calculations are still ongoing.

In summary, we have used a periodic Keldysh interaction to study the Wigner crystal phase transition in monolayer TMDCs. Our results, while still subject to some error, show that Keldysh Wigner crystals form at lower electron densities than their counter parts in Coulomb electron gas. The difference in the Wigner crystallisation densities is however, not enormous, in line with the fact that the Keldysh interaction reduces to the Coulomb interaction at long range.

Chapter 4

Mono-vacancies in graphene

4.1 Point defects in graphene

Graphene, an atomically thin sheet of carbon atoms forming a honeycomb lattice, is one of the most promising materials for future technological applications [6, 141, 142]. However, producing large, defect-free sheets of graphene on insulating substrates remains a significant technological challenge [143]. Point defects may appear naturally during the growth of graphene, or they may be deliberately inserted into pristine graphene by processing [144]. Point defects have a major impact on the electronic and optical properties of graphene [145, 146], so it is necessary to understand their properties to gain a full understanding of the performance of graphene-based devices. High-resolution transmission electron microscopy and related techniques have been employed to obtain clear imaging of defect structures in graphene and to help understand the impacts of defects on the properties of graphene [147, 148], but these methods themselves inevitably introduce defects. There are numerous works in which DFT is used to evaluate defect formation energies and other properties in relation to a range of applications and devices featuring graphene [149–151], graphite [152–154], and other two-dimensional or layered materials [155–157]. The main purpose of this chapter is to provide QMC defect-formation energy data to assess the accuracy of DFT in studies of defects in graphene.

Mono-vacancies (MVs) in graphene have been studied for both their desired and undesired effects on the graphene lattice. Graphite has long been used as a neutron moderator in nuclear reactors, which exposes the graphite to radiation damage [158–160]. It is essential to understand the properties of radiation-induced vacancies and how they may alter or weaken the structure of the graphene layers, and in turn graphite itself [153]. Vacancies in graphene also arise due to damage by electron beams in transmission electron microscopy [161]. Vacancy defects can in fact be useful for some applications and hence may be deliberately introduced into the lattice. Graphite/graphene has commonly been used as an anode material in lithium-ion batteries, with the lithium ions able to intercalate in the lattice [162]. A move towards sodium- or calcium-ion batteries is desirable owing to the greater abundance and lower cost of sodium and calcium. Unfortunately, the larger sizes of calcium and sodium ions compared to lithium prevents intercalation; however, the additional space created by vacancy defects allows larger atoms to intercalate into the anode material [163, 164]. Likewise, sub-nm pores, of which the MV is the smallest possible example, allow ion-selective transport for applications such as desalination of seawater [165, 166]. Many of these studies depend on DFT calculations to explore the behaviour of MVs and their interaction with other defects and chemical species.

The single most important thermodynamic property of a point defect is its formation energy, which is the difference in free energy between the defected material and the pristine material, together with any changes in the energies of reservoirs of the atoms that are added or removed when the defect is formed. For example, with a MV defect, the defect formation energy is the difference between the free energy of a large region of graphene containing a single MV defect and the free energy of the corresponding large region of pristine graphene, plus the free energy per atom of graphene.

MVs have also been observed to migrate across the lattice [167–169]. The MV migration energy barrier E_a is the energy difference between the ground state of

a MV and the saddle point of the energy along the lowest-energy path connecting atomic configurations in which the MV resides on neighbouring atomic sites. The migration energy barrier is an activation energy for MV diffusion; therefore, it can be used to predict the temperature-dependence of the mean time t between successive migrations of the vacancy using the Arrhenius equation

$$\frac{1}{t} = A \exp\left(-\frac{E_a}{k_B T}\right), \quad (4.1)$$

where k_B is the Boltzmann constant and T is the temperature. The pre-factor A in Eq. (4.1) is only weakly temperature-dependent, and is generally of order the optical phonon frequency; it can be estimated using the Vineyard formula [170].

In this chapter, we present QMC calculations of the formation energy and migrations energy barrier of MVs. Our intention is to benchmark the accuracy of the DFT methods that have been widely used in studies of defects in graphene.

4.2 Computational methodology

4.2.1 Defect formation energy

We define the “pure” formation energy E^{pf} of an isolated defect in graphene as the free-energy difference between a large region of graphene containing a single defect and pristine graphene. The defect formation energy E^{f} is the sum of the pure defect formation energy and the changes in the free energies of reservoirs of the atoms that are added or removed. For the MV defect this is

$$E_{\text{MV}}^{\text{f}} = E_{\text{MV}}^{\text{pf}} + \mu_{\text{C}}, \quad (4.2)$$

where the chemical potential μ_{C} is taken to be the Helmholtz free energy per atom of monolayer graphene. The pure defect formation energy is not in general physically meaningful by itself, because it depends on the choice of pseudopotentials. However, it is theoretically useful because it allows us to distinguish

finite-concentration and finite-size effects purely due to defect formation in a periodic supercell from finite-size errors in the chemical potentials. We approximate that the pure defect formation energy is the sum of the difference of static-nucleus electronic ground-state total energies and the temperature-dependent difference of vibrational Helmholtz free energies. Likewise, each chemical potential is taken to be the sum of the static-nucleus electronic ground-state total energy per atom and the temperature-dependent vibrational Helmholtz free energy per atom. Both the pure defect formation energy and the chemical potential are dominated by the temperature-independent electronic total-energy contribution; we therefore evaluate this by DMC calculation. The temperature-dependent vibrational free energy is a relatively small contribution, and therefore we evaluate it within DFT. In both our DFT and QMC calculations we evaluate pure defect formation energies in supercells subject to periodic boundary conditions. The errors that arise due to the use of finite, periodic supercells are discussed in Sec. 4.2.6.

4.2.2 Migration energy barrier

The MV migration energy barrier is the difference between the energy of the transition state, which is the saddle point on the lowest-energy path between configurations in which the MV lies on neighbouring atomic sites, and the energy of the MV ground state. The energy difference should include a difference in vibrational energies between the MV ground state and the transition state, excluding the contribution of the soft mode at the saddle point. We have calculated the transition-state geometry within DFT using the linear synchronous transit method (see Sec. 4.2.4.3).

4.2.3 Twisted periodic boundary conditions

We perform our calculations in supercells subject to periodic boundary conditions. In DFT, we can reduce this to a calculation in a single primitive cell and integrate

over the first Brillouin zone, by using a Monkhorst-Pack \mathbf{k} -point grid¹ [171] corresponding to the supercell we wish to study. We are allowed to do this in DFT because it is a one-electron method; however the explicitly correlated nature of QMC means we do not have the same luxury in QMC. We therefore construct a supercell consisting of a number of primitive unit cells. In the following, we use N to denote the number of electrons in the supercell.

The application of periodic boundary conditions to a supercell means that the Hamiltonian must be invariant whenever we translate an electron through one supercell lattice vector \mathbf{L}_s . This invariance leads to a Bloch-like condition on the many-body wave function, requiring $\Psi(\mathbf{R})$ to be of the form

$$\Psi(\mathbf{R}) = U_{\mathbf{k}_s}(\mathbf{R})e^{(i\mathbf{k}_s \cdot \sum_{i=1}^N \mathbf{r}_i)}, \quad (4.3)$$

where the function $U_{\mathbf{k}_s}(\mathbf{R})$ is invariant under translation of any electron through any supercell lattice vector \mathbf{L}_s [172, 173]. The wave-vector \mathbf{k}_s lies in the first Brillouin zone of the supercell, and is called the offset, or twist, vector; the use of a non-zero twist is known as applying twisted periodic boundary conditions.

A further translational symmetry occurs under simultaneous translation of all electrons through one primitive cell lattice vector \mathbf{L}_p . We then obtain a second many-body Bloch condition on the wave function, which reads

$$\Psi(\mathbf{R}) = W_{\mathbf{k}_p}(\mathbf{R})e^{(i\mathbf{k}_p \cdot \frac{1}{N} \sum_{i=1}^N \mathbf{r}_i)}, \quad (4.4)$$

where $W_{\mathbf{k}_p}(\mathbf{R})$ is invariant under simultaneous translation of all electrons through any primitive cell lattice vector \mathbf{L}_p [172, 173]. The wave-vector \mathbf{k}_p lies in the first Brillouin zone of the primitive cell.

Both of the many-body Bloch conditions in Eqs. (4.3) and (4.4) can be satisfied by: ensuring the single-particle orbitals in the Slater determinant of the trial wave function are of Bloch form; choosing the Bloch wave vector in those single-particle

¹A grid defined by the reciprocal supercell lattice vectors.

orbitals to lie on the grid defined by the supercell reciprocal lattice vectors offset from the origin by \mathbf{k}_s ; and, finally by using Jastrow and backflow functions that are invariant under translation of a single electron through one supercell lattice vector and simultaneous translation of all electrons through one primitive cell lattice vector.

Quantities calculated in periodic supercells are dependent on the grid of Bloch \mathbf{k} vectors used to evaluate it, and vary quasi-randomly with system size (*i.e.* the fine-ness of the grid). Ideally then, we wish to perform calculations in a supercell large enough (\mathbf{k} -point grid fine enough) that these fluctuations are minimal. The expense of QMC calculations, however, often means that supercells of such size are computationally infeasible, resulting in use of sparse \mathbf{k} -point grids and hence sizeable \mathbf{k} -point sampling, or momentum-quantisation, errors. We therefore require some way to reduce the size of the momentum-quantisation errors in our results. One such method is to choose the twist \mathbf{k}_s such that quantities evaluated at \mathbf{k}_s are close to their average value over the whole Brillouin zone (*e.g.* the Baldereschi point for insulators [174]). Another, more general, method is to perform calculations at a few, randomly chosen, twists and take the average of these. This is known as twist-averaging [175], and we will discuss it further in Sec. 4.2.6.2, as we will introduce a new, “twist-blocking”, method.

4.2.4 DFT calculations

4.2.4.1 Total energy, geometry optimisation, and phonon calculations

Our DFT calculations were performed using the PBE generalised gradient approximation exchange-correlation functional [17] and the plane-wave-basis code CASTEP [176]. The total energy, geometry optimisation, and phonon calculations all used ultra-soft pseudopotentials [177] to represent the nuclei and core electrons, and a plane-wave cut-off energy of 556 eV. For pristine graphene, the total energy was calculated using a 51×51 Monkhorst-Pack \mathbf{k} -point grid. The total energy of defective graphene was calculated in supercells of N_s primitive cells of

pristine graphene in a $\sqrt{N_s} \times \sqrt{N_s}$ arrangement containing a single MV defect, using Monkhorst-Pack grids of approximately $(51/\sqrt{N_s}) \times (51/\sqrt{N_s})$ \mathbf{k} -points. The geometry in each of the defective graphene supercells was optimised to a force tolerance of $0.0025 \text{ eV \AA}^{-1}$ with fixed lattice vectors corresponding to a pristine graphene carbon-carbon bond length of 1.42 \AA [178, 179]. All our 2D DFT calculations were performed using an artificial periodicity of $30 a_0$ in the out-of-plane direction. Non-spin-polarised calculations were used unless stated otherwise.

Phonon calculations using the finite displacement method in DFT [180] were used to evaluate the vibrational contributions to the free energy. These calculations were performed using atomic displacements of 0.005 , 0.01 , 0.015 , 0.02 , and $0.025 a_0$, with the final energies obtained by linearly extrapolating to zero phonon displacement. For each supercell, 5×5 Monkhorst-Pack supercell \mathbf{k} -point grids were used. Geometries were first optimised to a force tolerance of $0.0005 \text{ eV \AA}^{-1}$.

4.2.4.2 QMC orbital generation

Our DFT orbital-generation calculations used the PBE functional together with Trail-Needs Dirac-Fock pseudopotentials [46, 47] to represent the nuclei and core electrons, with s being the angular momentum of the local component when the pseudopotentials are re-represented in Kleinman-Bylander form [181]. The geometry was fixed at the DFT-PBE geometry obtained using ultra-soft pseudopotentials. The graphene supercell sizes used for the QMC calculations were 3×3 , 4×4 , and 5×5 , where the plane-wave cut-off energy for the smaller two supercells was 3401 eV , and the plane-wave cut-off energy for the larger supercell was 2231 eV . These cut-off energies are such that the DFT energy per atom is converged to within, respectively, 0.1 mHa and 1.59 mHa (known as chemical accuracy) [48]. An artificial periodicity of $30 a_0$ was used for the graphene calculations. Non-spin-polarised DFT calculations were used in all cases.

4.2.4.3 Transition state calculations

Our transition state calculations used the linear synchronous transit method in DFT [182, 183]. This method generates a set of intermediate images along the pathway between the initial and final states on the basis of the geometry of the system. The maximum energy structure along this path is then refined using a conjugate gradient optimisation [184].

We calculate the MV diffusion transition state in a supercell constructed from a 3×3 arrangement of primitive cells of pristine graphene containing a single MV. With the PBE functional, we use ultra-soft pseudopotentials, to represent the nuclei and core electrons, with a plane-wave cut-off energy of 556 eV, and a 17×17 Monkhorst-Pack \mathbf{k} -point grid. Orbitals for the subsequent QMC calculations were generated in the manner as described above.

In an ideal world, we would perform calculations in multiple supercell sizes and perform an extrapolation to infinite supercell size to remove finite-size effects.

4.2.5 QMC calculations

4.2.5.1 Trial wave functions

The trial wave functions used for the QMC calculations were of Slater-Jastrow (SJ) form, containing a product of determinants of spin-up and spin-down orbitals; see Sec. 4.2.4.2. Separate sets of orbitals were generated for each twist. The plane-wave orbitals were re-represented in a blip (B-spline) basis [185] both for computational efficiency in the QMC calculations and to remove the unwanted periodicity in the out-of-plane direction. The Jastrow function consisted of polynomial u , χ , and f terms, and plane-wave p terms. Trial wave functions were optimised first by minimising the variance of the energy and further by minimising the energy expectation value. For a given supercell, this optimisation was performed at a single, randomly chosen twist, with the resulting Jastrow factor being used at all twists.

The use of twists \mathbf{k}_s not equal to half of a supercell reciprocal-lattice vector means we are required to use a complex wave function, see Eq. (4.3), therefore fermionic antisymmetry is maintained using the fixed-phase approximation.

For some test cases at individual twists, Slater-Jastrow-backflow trial wave functions were used to investigate the fixed-node errors in our SJ-DMC results. These wave functions were obtained by optimising the backflow and Jastrow parameters together using energy minimisation. The backflow functions contained polynomial η and μ terms. Further tests using a long-ranged plane-wave electron-electron backflow function were also carried out: see Sec. 4.2.7.

Trail-Needs Dirac-Fock pseudopotentials [46, 47] were used to represent the ionic cores, with d being the angular momentum of the local component.

4.2.5.2 DMC calculations

To calculate the pure defect formation energy of a MV in graphene, pairs of DMC calculations were carried out at each twist in all of the defective and pristine graphene supercells. Time steps of $\delta\tau = 0.04$ and 0.16 Ha^{-1} were used in these calculations, with the corresponding target configuration populations being varied




Figure 4.1: DMC total energy per supercell of a MV defect in a 3×3 supercell of graphene against DMC time step $\delta\tau$ at a single, randomly chosen twist \mathbf{k}_s .

in inverse proportion to the time step. In all cases the target population was at least 256 configurations. The energies were then extrapolated linearly to zero time step. For the total energies of defective and pristine graphene we would not expect these time steps to be small enough to be in the linear bias regime (and indeed Fig. 4.1 shows this is the case); however, as shown in Fig. 4.2, the non-linear parts of the time-step bias cancel out of the pure defect formation energy for total energies calculated at the same twist.




Figure 4.2: DMC pure formation energy of a MV defect in a 3×3 supercell of graphene against DMC time step $\delta\tau$ at the twist \mathbf{k}_s used in Fig. 4.1.

To calculate the chemical potential (energy per atom of graphene) we used smaller time steps of $\tau = 0.01$ and 0.04 Ha^{-1} , allowing time-step bias in the total energy per atom to be largely removed by linear extrapolation. Again, we varied the target configuration population inversely with time step.

4.2.6 Finite-concentration and finite-size effects

4.2.6.1 Periodic supercells

Our QMC calculations of defect formation energies were performed in finite supercells subject to periodic boundary conditions, with a single point defect in the

supercell. This leads to a number of physical differences from the dilute limit of isolated point defects in which we are interested. Firstly, there are finite-concentration effects due to the fact that we are simulating a periodic array of point defects rather than an isolated defect. Leading-order systematic finite-concentration effects are due to screened electrostatic interactions between periodic images of defects and elastic interactions between defects [186]. There are also non-systematic finite-size effects due to interactions between charge-density oscillations around defects, *etc.* We remove the systematic effects and average out the non-systematic effects by extrapolation to infinite cell size using an appropriate fitting function. In principle, there could also be finite-concentration effects arising from the unwanted dispersion of defect states; however, none of the defects we study here supports a bound state. Secondly, there are finite-size effects arising from the simulation of periodic supercells rather than infinite crystals. These include quasi-random, oscillatory effects due to momentum quantisation, which we address by averaging over twisted boundary conditions on the supercell [175]. There are further quasi-random effects due to Ruderman-Kittel oscillations around defects being forced to be commensurate with the supercell, which are approximately averaged out when we extrapolate to the limit of large cell size.

To calculate the carbon chemical potential we must find the ground-state energy per atom of graphene. In a finite supercell this suffers from quasi-random momentum-quantisation effects, as well as long-range effects due to the evaluation of the interaction between each electron and the surrounding exchange-correlation hole using the Ewald interaction rather than $1/r$ [187] and the neglect of long-range two-body correlations [136, 188]. Long-range finite-size effects largely cancel out of the pure defect formation energies: the expressions for the leading-order corrections are the same for pristine and defective cells.

4.2.6.2 Twist averaging

Unlike DFT, only a single \mathbf{k} point can be used in each QMC calculation. We use twist averaging in the canonical ensemble (*i.e.* keeping the number of electrons fixed) [175] to reduce momentum-quantisation errors in our results. All our graphene DMC calculations were carried out at 24 random twists. Since momentum quantisation is a single-particle effect, it is well described by DFT, so that the QMC and DFT energies are correlated as a function of twist. DFT energies can therefore be used as a control variate when evaluating the twist-averaged DMC energy. The twist-averaged (TA) energy $E_{\text{DMC}}^{\text{TA}}$ is found by fitting

$$E_{\text{DMC}}(\mathbf{k}_s) = E_{\text{DMC}}^{\text{TA}} + b [E_{\text{DFT}}(\mathbf{k}_s) - E_{\text{DFT}}^{\text{fine}}] \quad (4.5)$$

to the DMC energy $E_{\text{DMC}}(\mathbf{k}_s)$ at twist \mathbf{k}_s , where b is a fitting parameter and $E_{\text{DFT}}(\mathbf{k}_s)$ is the corresponding DFT energy, and $E_{\text{DFT}}^{\text{fine}}$ is the DFT energy evaluated using a fine \mathbf{k} -point grid (using the same pseudopotentials as the DMC calculations). Eq. (4.5) simultaneously removes most of the quasi-random noise due to momentum quantisation and corrects for residual errors in the twist-averaged energy, by virtue of the fact that the correlator is the DFT energy relative to the DFT energy with a fine \mathbf{k} -point grid rather than the twist-averaged DFT energy. When calculating the pure defect formation energy, we use DMC and DFT pure defect formation energies in Eq. (4.5) rather than total energies. This is important because the pure and defective graphene calculations are performed at identical twists, so the twist-sampling error in the difference is much smaller than the twist-sampling errors in the total energies.

There are two very different sources of (quasi-)random error in the twist-averaged DMC energy for a given supercell: the statistical error from the Monte Carlo simulation, and the residual momentum quantisation error that is not fully removed by fitting Eq. (4.5). The statistical error can easily be accounted for in the twist-averaged energy by Gaussian propagation of errors; however, the resid-

ual momentum quantisation error is unknown at the outset. By choosing the same twists for the pristine and defective supercells it is reasonable to assume the momentum-quantisation errors will largely cancel out of the defect formation energy. To quantify the remaining error, a twist-blocking (TB) procedure has been used. For example, 24 twists can be grouped into 6 blocks of 4 twists, and within each block the twist-averaged energy can be calculated by fitting Eq. (4.5). An estimate of the true twist-averaged energy is then given by the mean of the 6 independent values of $E_{\text{DMC}}^{\text{TA}}$, while the standard error in the mean quantifies both the quasi-random momentum quantisation error and the random Monte Carlo errors. The mean energy obtained by this procedure is a biased estimate of the twist-averaged energy due to the small number of twists used in each fit; however, we can check for bias in both the mean and the standard error in the mean by increasing the block size. In fact we minimise the bias in the mean by using Eq. (4.5) with all the twists to obtain the twist-averaged energy, and only use the twist-blocking method to estimate the error bar (including twist-sampling errors) in the twist-averaged energy.

Figure 4.3 shows the twist-blocked standard error in the mean pure MV formation energy in a 3×3 supercell against the number blocks into which the 24 original twists are divided. Figure 4.3 does provide evidence that there are significant \mathbf{k} -point sampling errors after fitting Eq. (4.5) to all 24 twists. Furthermore, there is no evidence to suggest that the random error obtained by Gaussian propagation of the Monte Carlo errors in the fit to all 24 twists is unreliable. We also find that the behaviour of the twist-blocked standard error is similar in the cases of the other two supercell sizes here. An investigation using twist-blocking in the homogeneous electron gas would likely provide a clearer picture as to when twist-blocking provides an advantage over twist-averaging.

To our knowledge this is the first work to use twist-averaging to evaluate a defect formation energy. The approach is valid, since twist-averaged and non-twist-averaged finite-size effects agree in the infinite-system-size limit. Twist-averaging

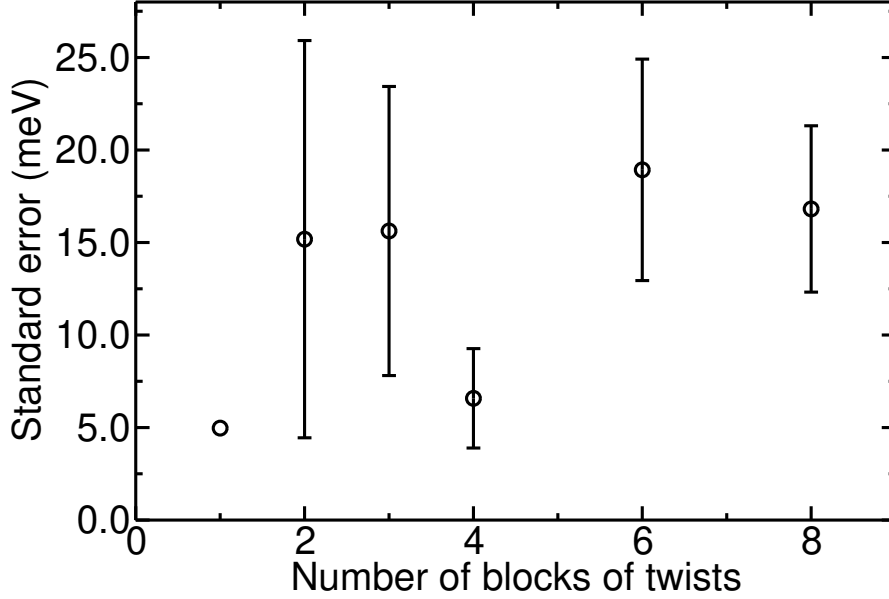


Figure 4.3: Twist-blocked standard error in the twist-averaged DMC pure MV formation energy in a 3×3 supercell against the number of blocks into which the 24 original twists are divided. The standard error in the single-block case is obtained by Gaussian propagation of the Monte Carlo random errors. Standard errors obtained with small numbers of large blocks are relatively unbiased, but suffer significant noise; on the other hand, standard errors obtained with large numbers of small blocks exhibit less noise but are potentially biased.

has the considerable advantage of greatly reducing a non-systematic finite-size effect by turning a sum over supercell reciprocal lattice vectors into an integral over \mathbf{k} , aiding extrapolation to the thermodynamic limit. For example, the standard deviation of the DMC pure MV defect formation energies as a function of twist is 0.3 eV, 0.2 eV, and 0.1 eV in 3×3 , 4×4 , and 5×5 supercells, respectively, indicating the size of the unquantified quasi-random error in the pure defect formation energy for each supercell size that would arise from using a non-twist-averaged calculation.

On the other hand, it is possible to find examples in which twist-averaging introduces finite-size errors that would not otherwise be present. Consider adding a single, non-interacting, impurity particle to a supercell containing a homogeneous fluid of host particles, which are distinguishable from the impurity particle. The ground state of the impurity particle is simply its zero-kinetic-energy plane-wave state. The zero-temperature immersion energy is therefore zero, both in the infinite-system limit and in a finite supercell subject to periodic boundary con-

ditions. Now suppose that twist-averaging is used to calculate the energy of the host fluid and the energy of the combined system in a finite supercell. The energy of the host fluid exactly cancels out of the immersion energy if identical twists are used, but twist averaging gives the impurity particle a positive kinetic energy, leading to a non-zero immersion energy. This problem can be solved in this case by also twist-averaging the energy of the “isolated” impurity particle in an otherwise empty periodic cell, using the same set of twists. In our defect-formation calculations we have calculated the carbon chemical potential as the twist-averaged energy per atom of pristine graphene. The systematic finite-size effects in the pure defect formation energies and the chemical potential scale differently as a function of system size; however, given the size of the quasi-random error in the non-twist-averaged defect formation energy at a given supercell size, it is clearly essential to use twist averaging in the calculation of the chemical potential. Any finite-size errors introduced by twist averaging are largely removed by the fit of Eq. (4.5) and extrapolation to infinite system size.

4.2.6.3 Long-range effects

To deal with long-range finite-size effects, defect-formation energies are calculated at various supercell sizes N_s , where N_s is the number of pristine primitive cells contained in the supercell. These results can then be extrapolated to infinite system size using an appropriate scaling law.

The MV defect is neutral and does not involve charge transfer between atoms, so has no dipole moment. In principle there exists a quadrupole moment associated with the defect, giving rise to weak electrostatic interactions between periodic images falling off rapidly as $\mathcal{O}(N_s^{-2} - N_s^{-\frac{5}{2}})$.

In addition to electrostatic interactions between periodic images, there are elastic finite-size effects due to the stress arising from the change in the size and shape of the supercell around the MV defect [186]. Assuming the defect results in nearly isotropic stress, the elastic finite-size effects in the energy go as $\mathcal{O}(N_s^{-1})$.

In summary the scaling of the elastic finite-size error suggest that twist-blocked pure MV defect formation energy $\langle E_{\text{MV}}^{\text{pf}}(N_s) \rangle_{\text{TB}}$ should be extrapolated to the thermodynamic limit by fitting

$$\langle E_{\text{MV}}^{\text{pf}}(N_s) \rangle_{\text{TB}} = E_{\text{MV}}^{\text{pf}}(N_s \rightarrow \infty) + cN_s^{-1}, \quad (4.6)$$

where $E_{\text{MV}}^{\text{pf}}(N_s \rightarrow \infty)$ and c are a fitting parameters. Using DFT calculations, we show in Fig. 4.4 that the $\mathcal{O}(N_s^{-1})$ finite-size error is dominant in the MV defect in the graphene sheet. We therefore extrapolate our twist-blocked DMC pure formation energies to the thermodynamic limit using an order $\mathcal{O}(N_s^{-1})$ scaling.




Figure 4.4: DFT pure formation energy of a MV defect in graphene against the reciprocal of supercell size N_s . Fine \mathbf{k} -point grids were used in each supercell.

The pristine graphene energy per atom (for the chemical potential) was extrapolated to infinite system size by fitting the twist-blocked supercell energies per atom $\langle e_{\text{P}}(N_s) \rangle_{\text{TB}}$ to

$$\langle e_{\text{P}}(N_s) \rangle_{\text{TB}} = e_{\text{P}}(N_s \rightarrow \infty) + c'N_s^{-\gamma}, \quad (4.7)$$

where $e_{\text{P}}(N_s \rightarrow \infty)$ and c' are fitting parameters. For pristine graphene, $\gamma = 5/4$ [136].

After separately dealing with the finite-size effects in the pure defect formation energy and chemical potential, the MV defect formation energy was calculated using Eq. (4.2).

4.2.7 Backflow

DMC calculations in a 3×3 supercell for the MV at a single twist show that the inclusion of backflow correlations with polynomial η and μ terms lowers the pure defect formation energy by 41(30) meV, while the addition of the plane-wave electron-electron term further lowers the pure defect formation energy by 16(29) meV, giving a total lowering of 58(31) meV. These differences are statistically insignificant, and are an order of magnitude smaller than the error bars on the SJ-DMC twist-averaged pure defect formation energies reported in Table 4.2; fixed-node errors are therefore well controlled. A backflow function with η and μ terms lowers the chemical potential of carbon by 46(9) meV, and the inclusion of the plane-wave electron-electron term does not have a statistically significant effect. The effects of backflow are insignificant on the 0.1 eV scale of the error bars on our SJ-DMC defect-formation energies.

4.3 Results and discussion

4.3.1 Atomic structures

4.3.1.1 Pristine graphene

All our pristine graphene calculations have used a carbon-carbon bond length of 1.42 Å [178, 179], and we have used exactly the same supercell sizes in our pristine and defective graphene calculations.

4.3.1.2 MV

It has previously been shown that a graphene MV undergoes a Jahn-Teller distortion, with two neighbours of the missing atom moving together to form a weak, reconstructed bond, and lowering the symmetry from D_{3h} point group [153]. The resulting structure is of either C_{2v} point group (a planar structure with a single horizontal mirror plane and a single vertical mirror plane) or C_s point group (a non-planar structure with just a single vertical mirror plane). Some DFT works on MVs have found and used the C_{2v} structure [189]; other DFT works have found that when two neighbours of the missing atom form a reconstructed bond, the third neighbour moves out of plane [153, 190, 191], resulting in a structure of C_s point group.

As shown in Table 4.1, non-spin-polarised DFT-LDA and DFT-PBE calculations with and without many-body dispersion (MBD*) corrections [192, 193] (performed by N. D. Drummond) show the C_s MV structure to be favoured (with the exception of DFT-LDA in a 3×3 supercell). The energy differences between the different non-spin-polarised MV structures are less than 0.3 eV. This is just about large enough to be non-negligible on the scale of our DMC error bars (see Sec. 4.3.2).

The difference between the C_s DFT-PBE and DFT-PBE-MBD* structures is small. For example, in a 3×3 supercell, the DFT-PBE energy is only increased by 1.5 meV when the DFT-PBE-MBD* structure is used instead of the DFT-PBE structure. We have used the non-spin-polarised C_s -symmetry structures obtained by relaxing within DFT-PBE in our QMC calculations; this is shown Fig. 4.5.

Previous DFT calculations have found that the MV has a magnetic moment of around 1.04–1.2 μ_B , where μ_B is the Bohr magneton [194, 195]. We examine the effect of performing spin-polarised DFT calculations in Table 4.1. Within DFT-LDA, the MV is non-magnetic. Within DFT-PBE and DFT-PBE-MBD*, spin-polarised structures of C_{2v} and C_s symmetry are found to be stable. There are evidently multiple minima in the spin-polarised energy, because Table 4.1 includes

Functional	Supercell	Pt. gp.	Energy relative to non-spin-polarised C_s defect (meV)				
			$E_{D_{3h}(u)} - E_{C_s(u)}$	$E_{C_{2v}(u)} - E_{C_s(u)}$	$E_{C_{2v}(p)} - E_{C_s(u)}$	$E_{C_s(p)} - E_{C_s(u)}$	$E_{C_s(u)}$
LDA	3×3	D_{3h}	0	0			0
LDA	5×5	C_s	94.6	92.4			0
LDA	7×7	C_s	120.2	115.4			0
PBE	3×3	C_s	21.5	21.5	-36.4		0
PBE	5×5	C_s	214.8	211.6	-96.8		-73.4
PBE	7×7	C_s	274.2	268.6	85.6		-113.8
PBE-MBD*	3×3	C_s	17.5	17.6	13.4		17.8
PBE-MBD*	5×5	C_s	195.6	193.7	22.2		0
PBE-MBD*	7×7	C_s	245.9	250.8	72.8		-122.3

Table 4.1: Point groups of relaxed graphene MV structures obtained using non-spin-polarised DFT calculations with different functionals and supercell sizes. Also shown are the energies of the most stable D_{3h} and C_{2v} structures relative to the most stable non-spin-polarised C_s structure. Spin-polarised and non-spin-polarised results are labelled by “(p)” and “(u),” respectively. Empty entries represent structures that were not found to be stable.




Figure 4.5: (a) Top-down and (b) in-plane views of the DFT-PBE-relaxed MV structure in a 5×5 supercell. The under-coordinated carbon atom is shown in red.

a case in which the spin-polarised energy is higher than the non-spin-polarised energy for the same point symmetry (in the final column), and cases in which higher-symmetry spin-polarised structures have lower energy than lower-symmetry structures (PBE-MBD* in the 3×3 supercell, comparing the final two columns). Our DFT-PBE calculations in a 7×7 supercell confirm that the spin-polarised C_s MV structure is more stable than the non-spin-polarised C_s structure by about 114 meV, in agreement with previous DFT calculations [194], with magnetic moment $1.4 \mu_B$. The ~ 0.1 eV DFT-PBE energy difference between magnetic and non-magnetic MV structures is less than the error bars on our DMC defect-formation energies reported in Sec. 4.3.2. For consistency, we have used non-magnetic MV structures and non-spin-polarised orbitals in our QMC calculations.

4.3.1.3 MV diffusion transition state

For MV diffusion there are multiple reported transition states, some of which are planar [153, 196], and others which are non-planar [189, 197]. Using DFT-PBE without MBD* corrections we find our transition state to be non-planar, in agreement with Wadey *et al.* [189]. A thorough investigation of the transition state structure using MBD* corrections, or similar, is desirable but the required DFT calculations proved to be difficult, expensive, and challenging. In the non-planar structure the under-coordinated carbon atom migrates to the centre of the space




Figure 4.6: (a) Top-down and (b) in-plane views of the MV diffusion transition state structure in a 3×3 supercell. The migrating atom is shown in red.

usually occupied by the vacant atom and the migrating atom itself, and forms bonds with the four neighbouring carbon atoms, as shown in Fig. 4.6. These four neighbouring atoms are displaced out-of-plane, with opposite pairs displaced in the same direction.

4.3.2 Defect formation energy

Figure 4.7 shows the DMC pure formation energy using twist-averaging (TA-DMC) and twist-blocking (TB-DMC), grouping the data into 6 blocks of 4 twists with, and without, using a DFT control variate (CV). It also shows the DFT pure formation energy obtained by twist-averaging (TA-DFT) and using a fine \mathbf{k} -point grid (DFT fine). TA-DFT results are obtained in the same way as TA-DMC results without a CV.

We see here that TA has the biggest effect on improving the error estimate on the DMC energy data, and that twist-blocking has little effect in comparison to twist-averaging (as we have shown in Fig. 4.3). However, there is significant quasi-random noise in the pure formation energies at different supercell sizes that is not quantified by the twist-blocked errors. The apparent quantification of this noise in the TA-DMC data by the twist-averaged error bars when we do not use a CV is then just coincidence. This remaining finite-size noise is, however, also present in the DFT results² in Fig. 4.4, and we see that our DMC energies deviate

²The difference between the DFT results in Fig. 4.4 and the DFT fine results in Fig. 4.7 is




Figure 4.7: Comparison of different methods for dealing with momentum quantisation errors in both DFT and DMC. Plots are of pure MV defect formation energy against reciprocal of supercell size N_s . Red dashed line show the unweighted least-squares fit to the TB-DMC mean energy data.

from a linear fit by a similar degree to the DFT results. The obvious way to reduce this would be to average both DMC and DFT energies over results obtained in a larger range of supercell sizes and possibly shapes.

In theory, the most accurate way to obtain the TA energy would be to use the TA mean energy with the TB error bars but here the difference between the TA and TB mean energies is negligible, and so we just use the TB-DMC mean energies to perform finite-size extrapolation. Error bars on the pure defect formation energy are obtained by performing a unweighted least-squares fit on the TB-DMC mean energy data. The DMC MV formation energy extrapolated to the dilute, infinite-system-size limit is shown in Table 4.2, along with DFT results from the literature. The large error on our DMC result is due to the finite-size noise in our pure formation energies at different supercell sizes.

While the errors on our DMC energies are too large to give an accurate DMC defect formation energy, the differences between the TB-DMC and DFT fine pure defect formation energies at each supercell size are significant. Per supercell, DFT

the choice of pseudopotentials.

Method	MV defect formation energy (eV)
DFT-PBE	7.64 [189], 7.65 [198], 7.97
DFT-LDA	8.02 [189], 7.40 [197]
DFT-B3LYP-D*	8.05 [190]
DMC	9.1(5)
DMC-corrected DFT	9.1(2)
Experiment	7.0(5) [167]

Table 4.2: Theoretical static-nucleus defect formation energies, together with experimental results, for the MV defect in monolayer graphene. The carbon chemical potential is the energy per atom of graphene. Results without citations were obtained in the present work. Our DMC-corrected DFT result is obtained by taking our DFT-PBE result and applying the mean DMC correction in Table 4.3. Cited DFT works are performed at finite supercell size with no attempt to handle finite-size effects.

consistently, and significantly, underestimates the pure formation energy for the MV defect. Therefore, we propose the best scheme for evaluating defect formation energies is to use DFT calculations in as large a supercell as possible and use our DMC data to apply a correction to the final DFT result, using the data provided in Table 4.2. Indeed, this is the same conclusion reached by Ma *et al.* [199] in their study of the Stone-Wales defect. Our DMC-corrected DFT defect formation energy is reported in Table 4.3, and is around 1 eV higher than the DFT MV defect formation energy. On the other hand, the DFT and DMC-corrected DFT defect formation energies for silicon substitutions and Stone-Wales defects differ by only around 0.3 eV [4]. The sizeable correction for the MV is consistent with our expectation that DFT performs poorly when evaluating differences in energies between structures where the chemical bonding is different.

The values for the DFT-PBE zero-point vibrational energy and Helmholtz free energy at 300 K in the pure MV defect formation energy are -0.74 eV and -0.68 eV, respectively. The Helmholtz free energy should be added to the static-nucleus DMC-corrected DFT defect formation energy in Table 4.3. This gives a MV defect formation energy of $8.4(2)$ eV.

Supercell	DMC correction (eV)
3×3	1.09(2)
4×4	0.81(2)
5×5	1.33(2)
Mean	1.1(2)

Table 4.3: DMC corrections to the static-nucleus DFT MV defect formation energy at various supercell sizes. Corrections are evaluated as the average of the difference between energies calculated with TB-DMC and DFT using a fine \mathbf{k} -point grid at different supercell sizes.

4.3.3 Migration energy barrier

We calculate the MV migration energy barrier as 1.18 eV using DFT-PBE. This is somewhat higher than the 0.87 eV DFT-PBE result of Wadey *et al.* [189] for a non-planar transition state structure; however it compares much better with the 1.29 eV result of Zobelli *et al.* [197] obtained for a non-planar transition state structure using the density functional tight binding method [200].

We have only studied the MV diffusion transition state in one (small) supercell, so we can not give an accurate estimate for the DMC MV migration energy. Indeed, the TB-DMC migration energy barrier calculated using only a 3×3 supercell is 0.937(9) eV, which has an absurdly small error bar due to the neglect of quasi-random finite-size effects for the migration energy barrier. We expect the finite-size error to be dominated by the elastic $\mathcal{O}(N_s^{-1})$ contribution and that the pre-factor is likely to be larger than the MV itself given the increased distortion in the lattice.

The aim of this work is to benchmark the accuracy of DFT calculations of properties of defects, so a comparison to the DFT MV migration energy barrier is still useful. For defect formation energies we found the best approach is to provide a DMC correction to the DFT energy. We can therefore evaluate the DMC correction to the DFT MV migration energy barrier in the same manner as for our defect formation energies. Doing so we find the DMC correction to the MV migration barrier is $-0.066(9)$ eV, and our DMC-corrected DFT result is 1.114(9) eV, which is smaller than the correction we obtained for the MV defect formation energy.

4.4 Conclusion

A new twist-blocking method is introduced to investigate how much of the momentum-quantisation error is recovered by simple twist-averaging. For defect formation energies at least, we find no evidence to suggest that there remains significant momentum-quantisation error after twist-averaging while making use of a DFT control variate.

We have investigated the lowest energy structure of the MV, where there is some disagreement in the literature. The out-of-plane C_s structure is found to be the most stable, rather than the often-cited C_{2v} structure. However, the differences in energy between the various structures are small enough that they are insignificant on the scale of our QMC error bars, therefore it is unlikely that these energy differences can be resolved in a QMC calculation.

The defect formation energy for MVs in graphene has been calculated using DMC. We find that, while great care was taken to account for finite-size effects, there still exists significant quasi-random finite-size effects in our DMC formation energy. We propose then, that the best way to evaluate defect formation energies in graphene is to use DFT in large supercells and then apply the difference between the DMC and DFT defect formation energies, averaged over twists and supercell sizes, as a correction to the DFT defect formation energy. Our DMC-corrected DFT results provide our best estimate for the MV defect formation energy. Our results suggest that DFT underestimates the MV defect formation energy by as much as 1 eV.

Also calculated is the MV diffusion transition state and the associated migration energy barrier. A non-planar structure is found for the transition state, consistent with the non-planar MV structure. A proper treatment of finite-size effects is lacking, however a comparison to DFT results in a similar calculation is still valid. The DMC-correction for the DFT MV migration energy barrier is less than 0.1 eV.

Since overall defect concentrations, and migration rates, depend exponentially

on the defect formation energies, and migration energy barriers, respectively, the errors on our DMC-corrected DFT results are too large to give accurate estimates for the absolute concentration, or migration rate, of MVs in graphene. More realistically, we can aim to give (semi-)quantitative estimates for the relative concentrations of different types of defect. The typical eV energy scale of point defects in graphene suggests that we need to reduce our error bars on defect formation energies to a (sub-)0.1 eV scale. To do so, it is clear from our results that we need to obtain a better understanding of the finite-size effects that occur in defect formation energies. A good place to start would be to perform calculations in a larger range of supercell sizes and shapes.

Conclusion

From the humble beginnings of the many-body Schrödinger equation, Chapter 1 saw us present, in detail, the methods of variational and diffusion QMC; methods which have the potential to further our understanding of low-dimensional quantum systems, and indeed condensed matter physics in general. We have made use of these to extend our knowledge of the properties of a variety of low-dimensional devices and materials.

Firstly, in Chapter 2, we used DMC to exactly solve models of excitonic complexes in novel semiconductor nanostructures, extracting information about the optoelectronic properties of such systems. Type-II quantum rings give rise to an interesting distinction between the binding and de-excitonisation energies of excitonic complexes, due to the large potential well confining holes inside the ring. Calculation of de-excitonisation energies allows us to predict the relative peak position in photoluminescence spectra, and these energies are shown to only weakly depend on the exact shape of the ring. The positive trion is the most strongly bound of the complexes, but even this is expected to dissociate well below room temperature. Perhaps the most intriguing find is the electrons' tendency to exist in a halo surrounding the ring, rather localising in the ring's central cavity. In superlattices, binding energies of a variety of complexes were found to be well described by an ideal two-dimensional bilayer model. The focus of the study into superlattices was to aid in the identification of an unknown peak in the experimental photoluminescence spectra of particular superlattice and multiple-quantum-well heterostructures. Through use of our binding energy data for a range of excitons

up to quintons, and ideal bilayer fits we were not able to identify the origin of the experimental peak, however we managed to prevent an incorrect identification of the peak's origin being due to a free biexciton.

The first part of Chapter 3 continued the theme of exactly solving models of excitonic complexes but this time in the TMDC family of materials. We demonstrated that the application of an in-plane electric field acts to massively shift the binding energies (and their respective photoluminescence peaks) of trions, but only to slightly perturb the binding energies of neutral complexes. These shifts in binding energy represent another tool that can be used for experimental identification of charged complexes from neutral ones, but the differences in shifts for neutral complexes are too small to (yet) be experimentally resolvable. In the second part of the chapter we investigated the effect on the Wigner crystallisation density, in a two-dimensional electron gas, of using a periodic Keldysh interaction in our calculations, rather than a periodic Coulomb interaction. This is a more realistic model of the electron gas in two-dimensional semiconductors compared to using a Coulomb interaction. Use of a Keldysh interaction was shown to slightly lower the Wigner crystallisation density compared with the Coulomb counter part, and reinforced the position that such systems require careful handling when optimising trial wave functions.

Finally, in Chapter 4, we considered the properties of MVs in graphene. Our main focus was to benchmark the accuracy of the DFT method, which has been widely used to study defects in graphene, rather than to provide experimentally relevant data. Great care was taken to control finite-size effects in our calculations; indeed, we introduced a new twist-blocking method to evaluate the performance of the previously-used twist-averaging method, but found no evidence to suggest that twist-averaging is unreliable. We were able to show that DFT underestimates the MV defect formation energy by up to 1 eV, with our (finite-size-treatment lacking) MV migration energy barrier showing much closer agreement with DFT. We proposed the best way to calculate defect formation energies is to use the

difference between DMC and DFT energies as a correction to the DFT defect formation energy.

In conclusion, we have demonstrated the applicability of QMC methods to a variety of low-dimensional systems, and thus pushed out the boundary of our scientific knowledge, with the hope that our findings will be of advantage in future developments in the field.

Bibliography

- [1] D. M. Thomas, R. J. Hunt, N. D. Drummond, and M. Hayne, “Binding energies of excitonic complexes in type-II quantum rings from diffusion quantum Monte Carlo calculations,” *Phys. Rev. B*, vol. 99, p. 115306, 2019. [iii](#)
- [2] E. Mostaani, R. J. Hunt, D. M. Thomas, M. Szyniszewski, A. R. P. Montblanch, M. Barbone, M. Atatüre, N. D. Drummond, and A. C. Ferrari, “Charge-carrier complexes in monolayer semiconductors,” in preparation. [iii](#)
- [3] R. J. Hunt, D. M. Thomas, and N. D. Drummond, “Quantum Monte Carlo study of Wigner crystallization in weakly doped monolayer MoSe₂,” in preparation. [iii](#)
- [4] D. M. Thomas, Y. Asiri, and N. D. Drummond, “Defect formation energies in graphene from diffusion quantum Monte Carlo and density functional theory,” in preparation. [iii](#), [106](#)
- [5] R. J. Young, E. P. Smakman, A. M. Sanchez, P. Hodgson, P. M. Koentraad, and M. Hayne, “Optical observation of single-carrier charging in type-II quantum ring ensembles,” *Appl. Phys. Lett.*, vol. 100, p. 082104, 2012. [x](#), [xi](#), [xv](#), [43](#), [44](#), [51](#), [52](#), [53](#), [54](#), [55](#), [56](#)
- [6] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov, “Electric field effect in atomically thin carbon films,” *Science*, vol. 306, p. 666, 2004. [1](#), [84](#)

- [7] E. Schrödinger, “Quantisierung als eigenwertproblem,” *Annalen der Physik*, vol. 384, p. 361, 1926. [3](#)
- [8] E. Schrödinger, “Quantisierung als eigenwertproblem,” *Annalen der Physik*, vol. 384, p. 489, 1926. [3](#)
- [9] E. Schrödinger, “Quantisierung als eigenwertproblem,” *Annalen der Physik*, vol. 385, p. 437, 1926. [3](#)
- [10] E. Schrödinger, “Quantisierung als eigenwertproblem,” *Annalen der Physik*, vol. 386, p. 109, 1926. [3](#)
- [11] M. Born and R. Oppenheimer, “Zur quantentheorie der molekeln,” *Annalen der Physik*, vol. 389, p. 457, 1927. [4](#)
- [12] V. Fock, “Näherungsmethode zur lösung des quantenmechanischen mehrkörperproblems,” *Z. Physik*, vol. 61, p. 126, 1930. [6](#)
- [13] J. C. Slater, “Note on Hartree’s method,” *Phys. Rev.*, vol. 35, p. 210, 1930. [6](#)
- [14] D. R. Hartree and W. Hartree, “Self-consistent field, with exchange, for beryllium,” *P. Roy. Soc. Lond. A Mat.*, vol. 150, p. 9, 1935. [6](#)
- [15] P. Hohenberg and W. Kohn, “Inhomogeneous electron gas,” *Phys. Rev.*, vol. 136, p. B864, November 1964. [6](#)
- [16] W. Kohn and L. J. Sham, “Self-consistent equations including exchange and correlation effects,” *Phys. Rev.*, vol. 140, p. A1133, 1965. [6](#)
- [17] J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized gradient approximation made simple,” *Phys. Rev. Lett.*, vol. 77, p. 3865, 1996. [6](#), [89](#)
- [18] A. D. Becke, “Density-functional exchange-energy approximation with correct asymptotic behavior,” *Phys. Rev. A*, vol. 38, p. 3098, 1988. [7](#)

- [19] C. Lee, W. Yang, and R. G. Parr, “Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density,” *Phys. Rev. B*, vol. 37, p. 785, 1988. [7](#)
- [20] K. Burke, “Perspective on density functional theory,” *J. Chem. Phys.*, vol. 136, p. 150901, 2012. [7](#)
- [21] D. M. Ceperley, “Path integrals in the theory of condensed helium,” *Rev. Mod. Phys.*, vol. 67, p. 279, 1995. [7](#)
- [22] G. H. Booth, A. J. W. Thom, and A. Alavi, “Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space,” *J. Chem. Phys.*, vol. 131, p. 054106, 2009. [7](#)
- [23] D. Cleland, G. H. Booth, and A. Alavi, “Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum Monte Carlo,” *J. Chem. Phys.*, vol. 132, p. 041103, 2010. [7](#)
- [24] S. Sorella, S. Baroni, R. Car, and M. Parrinello, “A novel technique for the simulation of interacting fermion systems,” *EPL (Europhysics Letters)*, vol. 8, p. 663, 1989. [7](#)
- [25] D. M. Ceperley and B. J. Alder, “Ground state of the electron gas by a stochastic method,” *Phys. Rev. Lett.*, vol. 45, p. 566, 1980. [7](#)
- [26] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Rev. Mod. Phys.*, vol. 73, p. 33, 2001. [7](#)
- [27] R. J. Needs, M. D. Towler, N. D. Drummond, P. López Ríos, and J. R. Trail, “Variational and diffusion quantum Monte Carlo calculations with the CASINO code,” *J. Chem. Phys.*, vol. 152, p. 154106, 2020. [8](#), [21](#)
- [28] S. Weinzierl, “Introduction to Monte Carlo methods,” *arXiv:hep-ph/0006269*, 2000. [8](#)

- [29] W. Feller, *An introduction to probability theory and its applications Vol. I*. John Wiley & Sons Inc., 1968. 9, 11
- [30] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of state calculations by fast computing machines,” *J. Chem. Phys.*, vol. 21, p. 1087, 1953. 10, 12
- [31] W. K. Hastings, “Monte Carlo sampling methods using Markov chains and their applications,” *BioMetrika*, vol. 57, p. 97, 1970. 10
- [32] D. Ceperley, G. V. Chester, and M. H. Kalos, “Monte Carlo simulation of a many-fermion study,” *Phys. Rev. B*, vol. 16, p. 3081, 1977. 10
- [33] P. A. Gagniuc, *Markov Chains: From Theory to Implementation and Experimentation*. John Wiley and Sons, Ltd, 2017. 10
- [34] O. Perron, “Zur theorie der matrices,” *Mathematische Annalen*, vol. 64, p. 248, 1907. 10
- [35] G. Frobenius, *Über Matrizen aus nicht negativen Elementen*. Ber. Akad. Wiss. Berlin, 1912. 10
- [36] R. M. Lee, G. J. Conduit, N. Nemec, P. López Ríos, and N. D. Drummond, “Strategies for improving the efficiency of quantum Monte Carlo calculations,” *Phys. Rev. E*, vol. 83, p. 066706, 2011. 12
- [37] H. Flyvbjerg and H. G. Petersen, “Error estimates on averages of correlated data,” *J. Chem. Phys.*, vol. 91, p. 461, 1989. 13
- [38] U. Wolff, “Monte Carlo errors with less errors,” *Comput. Phys. Commun.*, vol. 156, p. 143, 2004. 13
- [39] N. D. Drummond, M. D. Towler, and R. J. Needs, “Jastrow correlation factor for atoms, molecules, and solids,” *Phys. Rev. B*, vol. 70, p. 235119, 2004. 18

- [40] T. Kato, “On the eigenfunctions of many-particle systems in quantum mechanics,” *Communications on Pure and Applied Mathematics*, vol. 10, p. 151, 1957. [19](#)
- [41] A. Ma, M. D. Towler, N. D. Drummond, and R. J. Needs, “Scheme for adding electron-nucleus cusps to gaussian orbitals,” *J. Chem. Phys.*, vol. 122, p. 224322, 2005. [20](#)
- [42] P. López Ríos, A. Ma, N. D. Drummond, M. D. Towler, and R. J. Needs, “Inhomogeneous backflow transformations in quantum Monte Carlo calculations,” *Phys. Rev. E*, vol. 74, p. 066701, 2006. [21](#)
- [43] D. Ceperley, “The statistical error of green’s function Monte Carlo,” *J. Stat. Phys.*, vol. 43, p. 815, 1986. [22](#), [24](#)
- [44] B. L. Hammond, P. J. Reynolds, and W. A. Lester, “Valence quantum Monte Carlo with *ab initio* effective core potentials,” *J. Chem. Phys.*, vol. 87, p. 1130, 1987. [22](#)
- [45] A. Ma, N. D. Drummond, M. D. Towler, and R. J. Needs, “All-electron quantum Monte Carlo calculations for the noble gas atoms he to xe,” *Phys. Rev. E*, vol. 71, p. 066704, 2005. [22](#)
- [46] J. R. Trail and R. J. Needs, “Norm-conserving Hartree-Fock pseudopotentials and their asymptotic behavior,” *J. Chem. Phys.*, vol. 122, p. 014112, 2005. [22](#), [90](#), [92](#)
- [47] J. R. Trail and R. J. Needs, “Smooth relativistic Hartree-Fock pseudopotentials for H to Ba and Lu to Hg,” *J. Chem. Phys.*, vol. 122, p. 174109, 2005. [22](#), [90](#), [92](#)
- [48] N. D. Drummond, J. R. Trail, and R. J. Needs, “Trail-Needs pseudopotentials in quantum Monte Carlo calculations with plane-wave/blip basis sets,” *Phys. Rev. B*, vol. 94, p. 165170, 2016. [22](#), [90](#)

- [49] M. Dewing, “Improved efficiency with variational Monte Carlo using two level sampling,” *J. Chem. Phys.*, vol. 113, p. 5123, 2000. [23](#)
- [50] C. J. Umrigar, K. G. Wilson, and J. W. Wilkins, “Optimized trial wave functions for quantum Monte Carlo calculations,” *Phys. Rev. Lett.*, vol. 60, p. 1719, 1988. [24](#)
- [51] N. D. Drummond and R. J. Needs, “Variance-minimization scheme for optimizing Jastrow factors,” *Phys. Rev. B*, vol. 72, p. 085124, 2005. [24](#)
- [52] D. Bressanini, G. Morosi, and M. Mella, “Robust wave function optimization procedures in quantum Monte Carlo methods,” *J. Chem. Phys.*, vol. 116, p. 5345, 2002. [24](#)
- [53] C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Hennig, “Alleviation of the fermion-sign problem by optimization of many-body wave functions,” *Phys. Rev. Lett.*, vol. 98, p. 110201, 2007. [24](#)
- [54] M. Suzuki, “Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems,” *Commun. Math. Phys.*, vol. 51, p. 183, 1976. [27](#)
- [55] B. L. Hammond, W. A. Lester, and P. J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry*. WORLD SCIENTIFIC, 1994. [27](#), [33](#), [34](#)
- [56] J. B. Anderson, “Quantum chemistry by random walk,” *J. Chem. Phys.*, vol. 65, p. 4121, 1976. [31](#)
- [57] G. Ortiz, D. M. Ceperley, and R. M. Martin, “New stochastic method for systems with broken time-reversal symmetry: 2d fermions in a magnetic field,” *Phys. Rev. Lett.*, vol. 71, p. 2777, 1993. [31](#)
- [58] W. M. C. Foulkes, R. Q. Hood, and R. J. Needs, “Symmetry constraints and variational principles in diffusion quantum Monte Carlo calculations of excited-state energies,” *Phys. Rev. B*, vol. 60, p. 4558, 1999. [33](#)

- [59] M. H. Ceperley D. M., Kalos, “Quantum many-body problems,” in *Monte Carlo Methods in Statistical Physics* (K. Binder, ed.), vol. 7, p. 145, Springer-Verlag, Heidelberg, 2nd ed., 1979. [34](#)
- [60] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, “Fixed-node quantum Monte Carlo for molecules,” *J. Chem. Phys.*, vol. 77, p. 5593, 1982. [35](#)
- [61] A. Zen, S. Sorella, M. J. Gillan, A. Michaelides, and D. Alfè, “Boosting the accuracy and speed of quantum Monte Carlo: Size consistency and time step,” *Phys. Rev. B*, vol. 93, p. 241118, 2016. [35](#)
- [62] R. Needs, M. Towler, N. Drummond, and P. López Ríos, *User’s Guide Version 2.13 (2019)*, 2021. [35](#), [78](#)
- [63] C. J. Umrigar, M. P. Nightingale, and K. J. Runge, “A diffusion Monte Carlo algorithm with very small time-step errors,” *J. Chem. Phys.*, vol. 99, p. 2865, 1993. [36](#)
- [64] F. Bloch, “Über die quantenmechanik der elektronen in kristallgittern,” *Zeitschrift für Physik*, vol. 52, p. 555, 1929. [38](#)
- [65] S. Datta, *Electronic transport in mesoscopic systems*. Cambridge studies in semiconductor physics and microelectronic engineering, Cambridge University Press, 1995. [39](#)
- [66] T. Heinzl, *Mesoscopic electronics in solid state nanostructures*. Wiley-VCH, 3rd completely revised and enlarged ed., 2010. [39](#), [75](#)
- [67] N. W. Ashcroft, *Solid state physics*. Holt, Rinehart and Winston, Brooks/Cole, Thomson Learning, 1976. [39](#)
- [68] E. Mostaani, M. Szyniszewski, C. H. Price, R. Maezono, M. Danovich, R. J. Hunt, N. D. Drummond, and V. I. Fal’ko, “Diffusion quantum Monte Carlo

- study of excitonic complexes in two-dimensional transition-metal dichalcogenides,” *Phys. Rev. B*, vol. 96, p. 075431, 2017. [41](#), [42](#), [67](#), [68](#), [71](#)
- [69] R. C. Ashoori, “Electrons in artificial atoms,” *Nature*, vol. 379, p. 413, 1996. [43](#)
- [70] D. Gammon, “Electrons in artificial atoms,” *Nature*, vol. 405, p. 899, 2000. [43](#)
- [71] R. J. Warburton, C. Schäfflein, D. Haft, F. Bickel, A. Lorke, K. Karrai, J. M. Garcia, W. Schoenfeld, and P. M. Petroff, “Optical emission from a charge-tunable quantum ring,” *Nature*, vol. 405, p. 926, 2000. [43](#)
- [72] M. Bayer, O. Stern, P. Hawrylak, S. Fafard, and A. Forchel, “Hidden symmetries in the energy levels of excitonic ‘artificial atoms’,” *Nature*, vol. 405, p. 923, 2000. [43](#)
- [73] U. Banin, Y. W. Cao, D. Katz, and O. Millo, “Identification of atomic-like electronic states in indium arsenide nanocrystal quantum dots,” *Nature*, vol. 400, p. 542, 1999. [43](#)
- [74] A. Fuhrer, S. Lüscher, T. Ihn, T. Heinzel, K. Ensslin, W. Wegscheider, and M. Bichler, “Energy spectra of quantum rings,” *Nature*, vol. 413, p. 822, 2001. [43](#)
- [75] F. Nagasawa, D. Frustaglia, H. Saarikoski, K. Richter, and J. Nitta, “Control of the spin geometric phase in semiconductor quantum rings,” *Nat. Commun.*, vol. 4, p. 2526, 2013. [43](#)
- [76] B. Bansal, M. Hayne, M. Geller, D. Bimberg, and V. V. Moshchalkov, “Excitonic Mott transition in type-II quantum dots,” *Phys. Rev. B*, vol. 77, p. 241304, 2008. [43](#)

- [77] B. Bansal, S. Godefroo, M. Hayne, G. Medeiros-Ribeiro, and V. V. Moshchalkov, “Extended excitons and compact heliumlike biexcitons in type-II quantum dots,” *Phys. Rev. B*, vol. 80, p. 205317, 2009. [43](#)
- [78] M. Kroutvar, Y. Ducommun, D. Heiss, M. Bichler, D. Schuh, G. Abstreiter, and J. J. Finley, “Optically programmable electron spin memory using semiconductor quantum dots,” *Nature*, vol. 432, p. 81, 2004. [43](#)
- [79] T. Nowozin, A. Marent, L. Bonato, A. Schliwa, D. Bimberg, E. P. Smakman, J. K. Garleff, P. M. Koenraad, R. J. Young, and M. Hayne, “Linking structural and electronic properties of high-purity self-assembled GaSb/GaAs quantum dots,” *Phys. Rev. B*, vol. 86, p. 035305, 2012. [43](#)
- [80] M. Hayne, J. Maes, S. Bersier, V. V. Moshchalkov, A. Schliwa, L. Müller-Kirsch, C. Kapteyn, R. Heitz, and D. Bimberg, “Electron localization by self-assembled GaSb/GaAs quantum dots,” *Appl. Phys. Lett.*, vol. 82, p. 4355, 2003. [43](#)
- [81] K. Suzuki, R. A. Hogg, and Y. Arakawa, “Structural and optical properties of type II GaSb/GaAs self-assembled quantum dots grown by molecular beam epitaxy,” *J. Appl. Phys.*, vol. 85, p. 8349, 1999. [43](#)
- [82] F. Hatami, M. Grundmann, N. N. Ledentsov, F. Heinrichsdorff, R. Heitz, J. Böhrer, D. Bimberg, S. S. Ruvimov, P. Werner, V. M. Ustinov, P. S. Kopev, and Z. I. Alferov, “Carrier dynamics in type-II GaSb/GaAs quantum dots,” *Phys. Rev. B*, vol. 57, p. 4635, 1998. [43](#)
- [83] M. Geller, C. Kapteyn, L. Müller-Kirsch, R. Heitz, and D. Bimberg, “450 meV hole localization in GaSb/GaAs quantum dots,” *Appl. Phys. Lett.*, vol. 82, p. 2706, 2003. [43](#)
- [84] R. B. Laghumavarapu, A. Moscho, A. Khoshakhlagh, M. El-Emawy, L. F. Lester, and D. L. Huffaker, “GaSbGaAs type II quantum dot solar cells for

- enhanced infrared spectral response,” *Appl. Phys. Lett.*, vol. 90, p. 173125, 2007. [43](#)
- [85] A. J. Williamson, L. W. Wang, and A. Zunger, “Theoretical interpretation of the experimental electronic structure of lens-shaped self-assembled InAs/GaAs quantum dots,” *Phys. Rev. B*, vol. 62, p. 12963, 2000. [43](#)
- [86] A. D. Andreev and E. P. O’Reilly, “Theory of the electronic structure of GaN/AlN hexagonal quantum dots,” *Phys. Rev. B*, vol. 62, p. 15851, 2000. [43](#)
- [87] P. J. Carrington, R. J. Young, P. D. Hodgson, A. M. Sanchez, M. Hayne, and A. Krier, “Long-wavelength photoluminescence from stacked layers of high-quality type-II GaSb/GaAs quantum rings,” *Cryst. Growth and Des.*, vol. 13, p. 1226, 2013. [43](#), [44](#)
- [88] M. P. Young, C. S. Woodhead, J. Roberts, Y. J. Noori, M. T. Noble, A. Krier, E. P. Smakman, P. M. Koenraad, M. Hayne, and R. J. Young, “Photoluminescence studies of individual and few GaSb/GaAs quantum rings,” *AIP Adv.*, vol. 4, p. 117127, 2014. [43](#), [44](#), [45](#), [52](#), [53](#)
- [89] P. D. Hodgson, M. Hayne, A. J. Robson, Q. D. Zhuang, and L. Danos, “GaSb quantum rings in GaAs/Al_xGa_{1-x}As quantum wells,” *J. Appl. Phys.*, vol. 119, p. 044305, 2016. [43](#), [44](#)
- [90] R. Timm, H. Eisele, A. Lenz, L. Ivanova, V. Vossebürger, T. Warming, D. Bimberg, I. Farrer, D. A. Ritchie, and M. Dähne, “Confined states of individual type-II GaSb/GaAs quantum rings studied by cross-sectional scanning tunneling spectroscopy,” *Nano Lett.*, vol. 10, p. 3972, 2010. [43](#), [44](#)
- [91] M. Hayne, R. J. Young, E. P. Smakman, T. Nowozin, P. Hodgson, J. K. Garleff, P. Rambabu, P. M. Koenraad, A. Marent, L. Bonato, A. Schliwa, and D. Bimberg, “The structural, electronic and optical properties of

- GaSb/GaAs nanostructures for charge-based memory,” *J. Phys. D: Appl. Phys.*, vol. 46, p. 264001, 2013. [43](#)
- [92] S. Kobayashi, C. Jiang, T. Kawazu, and H. Sakaki, “Self-assembled growth of GaSb type II quantum ring structures,” *Jpn. J. Appl. Phys.*, vol. 43, p. L662, 2004. [44](#)
- [93] M. E. Levinshthen, S. L. Rumyantsev, and M. Shur, *Handbook Series on Semiconductor Parameters, Volume 1: Si, Ge, C (Diamond), GaAs, GaP, GaSb, InAs, InP, InSb*, vol. 1. World Scientific, London, 1996. [51](#), [59](#), [60](#)
- [94] P. D. Hodgson, R. J. Young, M. Ahmad Kamarudin, P. J. Carrington, A. Krier, Q. D. Zhuang, E. P. Smakman, P. M. Koenraad, and M. Hayne, “Blueshifts of the emission energy in type-II quantum dot and quantum ring nanostructures,” *J. Appl. Phys.*, vol. 114, p. 073519, 2013. [52](#)
- [95] M. S. Kehili, A. Ben Mansour, R. Sellami, and A. Melliti, “Electronic and optical properties of charge carriers in a GaSb quantum ring in a GaAs/Al_{0.6}Ga_{0.4}As quantum well,” *Semicond. Sci. Technol.*, vol. 33, p. 115019, 2018. [56](#)
- [96] G. Kiršanskė, P. Tighineanu, R. S. Daveau, J. Miguel-Sánchez, P. Lodahl, and S. Stobbe, “Observation of the exciton mott transition in the photoluminescence of coupled quantum wells,” *Phys. Rev. B*, vol. 94, p. 155438, 2016. [58](#)
- [97] B. Deveaud, L. Kappei, J. Berney, F. Morier-Genoud, M. Portella-Oberli, J. Szczytko, and C. Piermarocchi, “Excitonic effects in the luminescence of quantum wells,” *Chem. Phys.*, vol. 318, p. 104, 2005. [58](#)
- [98] J. A. Keen, D. Lane, M. Kesaria, A. R. J. Marshall, and A. Krier, “InAs/InAsSb type-II strained-layer superlattices for mid-infrared LEDs,” *Journal of Physics D: Applied Physics*, vol. 51, p. 075103, 2018. [58](#)

- [99] J. Keen, E. Repiso, Q. Lu, M. Kesaria, A. Marshall, and A. Krier, “Electroluminescence and photoluminescence of type-II InAs/InAsSb strained-layer superlattices in the mid-infrared,” *Infrared Physics & Technology*, vol. 93, p. 375, 2018. 58
- [100] J. Keen, *Characterisation and simulation of InAs/InAsSb structures for mid-infrared LEDs*. PhD thesis, Lancaster University, 2018. 58, 59, 62, 63, 64, 65
- [101] M. E. Levinshen, S. L. Rumyantsev, and M. Shur, *Handbook Series on Semiconductor Parameters, Volume 2: Ternary And Quaternary III-V Compounds*, vol. 2. World Scientific, London, 1999. 59
- [102] M. Y. J. Tan, N. D. Drummond, and R. J. Needs, “Exciton and biexciton energies in bilayer systems,” *Phys. Rev. B*, vol. 71, p. 033303, 2005. 64, 65
- [103] O. Witham, R. J. Hunt, and N. D. Drummond, “Stability of trions in coupled quantum wells modeled by two-dimensional bilayers,” *Phys. Rev. B*, vol. 97, p. 075424, 2018. 64, 65
- [104] D. Sarkar, H. P. van der Meulen, J. M. Calleja, J. M. Becker, R. J. Haug, and K. Pierz, “Exciton fine structure and biexciton binding energy in single self-assembled InAs/AlAs quantum dots,” *Journal of Applied Physics*, vol. 100, p. 023109, 2006. 64
- [105] C. Palacios-Berraquero, M. Barbone, D. M. Kara, X. Chen, I. Goykhman, D. Yoon, A. K. Ott, J. Beitner, K. Watanabe, T. Taniguchi, A. C. Ferrari, and M. Atatüre, “Atomically thin quantum light-emitting diodes,” *Nature communications*, vol. 7, p. 12978, 2016. 66
- [106] J. Dang, S. Sun, X. Xie, Y. Yu, K. Peng, C. Qian, S. Wu, F. Song, J. Yang, S. Xiao, L. Yang, Y. Wang, M. A. Rafiq, C. Wang, and X. Xu, “Identifying defect-related quantum emitters in monolayer WSe₂,” *npj 2D Mater. and Applications*, vol. 4, p. 2, 2020. 66

- [107] H. Tian, M. L. Chin, S. Najmaei, Q. Guo, F. Xia, H. Wang, and M. Dubey, “Optoelectronic devices based on two-dimensional transition metal dichalcogenides,” *Nano Research*, vol. 9, p. 1543, 2016. [66](#)
- [108] A. Kormányos, G. Burkard, M. Gmitra, J. Fabian, V. Zólyomi, N. D. Drummond, and V. Fal’ko, “ $\mathbf{k}\cdot\mathbf{p}$ theory for two-dimensional transition metal dichalcogenide semiconductors,” *2D Mater.*, vol. 2, p. 022001, 2015. [66](#), [67](#)
- [109] A. Splendiani, L. Sun, Y. Zhang, T. Li, J. Kim, C.-Y. Chim, G. Galli, and F. Wang, “Emerging photoluminescence in monolayer MoS₂,” *Nano Letters*, vol. 10, p. 1271, 2010. [67](#)
- [110] A. Ramasubramaniam, “Large excitonic effects in monolayers of molybdenum and tungsten dichalcogenides,” *Phys. Rev. B*, vol. 86, p. 115409, 2012. [67](#)
- [111] T. Cheiwchanchamnangij and W. R. L. Lambrecht, “Quasiparticle band structure calculation of monolayer, bilayer, and bulk MoS₂,” *Phys. Rev. B*, vol. 85, p. 205302, 2012. [67](#), [72](#)
- [112] M. Szyniszewski, E. Mostaani, N. D. Drummond, and V. I. Fal’ko, “Binding energies of trions and biexcitons in two-dimensional semiconductors from diffusion quantum Monte Carlo calculations,” *Phys. Rev. B*, vol. 95, p. 081301, 2017. [67](#)
- [113] M. Barbone, A. R. Montblanch, D. M. Kara, C. Palacios-Berraquero, A. R. Cadore, D. De Fazio, B. Pingault, E. Mostaani, H. Li, B. Chen, K. Watanabe, T. Taniguchi, S. Tongay, G. Wang, A. C. Ferrari, and M. Atatüre, “Charge-tuneable biexciton complexes in monolayer WSe₂,” *Nat. Commun.*, vol. 9, p. 3721, 2018. [67](#)
- [114] Z. Ye, L. Waldecker, E. Y. Ma, D. Rhodes, A. Antony, B. Kim, X. Zhang, M. Deng, Y. Jiang, Z. Lu, D. Smirnov, K. Watanabe, T. Taniguchi, J. Hone,

- and T. F. Heinz, “Efficient generation of neutral and charged biexcitons in encapsulated WSe₂ monolayers,” *Nat. Commun.*, vol. 9, p. 3718, 2018. [67](#)
- [115] N. S. Rytova, “Coulomb interaction of electrons in a thin film,” *Dokl. Akad. Nauk. SSSR*, vol. 163, p. 1118, 1965. [69](#)
- [116] N. S. Rytova, “Screened potential of a point charge in the thin film,” *Vestn. Mosk. Univ. Fiz. Astron.*, vol. 3, p. 30, 1967. [69](#)
- [117] L. V. Keldysh, “Coulomb interaction in thin semiconductor and semimetal films,” *JETP*, vol. 29, p. 658, 1979. [69](#)
- [118] B. Ganchev, N. Drummond, I. Aleiner, and V. Fal’ko, “Three-particle complexes in two-dimensional semiconductors,” *Phys. Rev. Lett.*, vol. 114, p. 107401, 2015. [71](#)
- [119] H. Shi, H. Pan, Y.-W. Zhang, and B. I. Yakobson, “Quasiparticle band structures and optical properties of strained monolayer MoS₂ and WS₂,” *Phys. Rev. B*, vol. 87, p. 155304, 2013. [72](#)
- [120] A. Kumar and P. Ahluwalia, “Tunable dielectric response of transition metals dichalcogenides MX₂ (M=Mo, W; X=S, Se, Te): Effect of quantum confinement,” *Physica B: Condensed Matter*, vol. 407, p. 4627, 2012. [72](#)
- [121] T. C. Berkelbach, M. S. Hybertsen, and D. R. Reichman, “Theory of neutral and charged excitons in monolayer transition metal dichalcogenides,” *Phys. Rev. B*, vol. 88, p. 045318, 2013. [72](#)
- [122] J. R. Reitz, *Foundations of electromagnetic theory*. Addison-Wesley series in physics, Addison-Wesley Pub. Co, 1960. [72](#), [73](#)
- [123] H. Tampo, H. Shibata, K. Matsubara, A. Yamada, P. Fons, S. Niki, M. Yamagata, and H. Kanie, “Two-dimensional electron gas in Zn polar Zn-MgO/ZnO heterostructures grown by radical source molecular beam epitaxy,” *Appl. Phys. Lett.*, vol. 89, p. 132113, 2006. [75](#)

- [124] J. D. Ye, S. Pannirselvam, S. T. Lim, J. F. Bi, X. W. Sun, G. Q. Lo, and K. L. Teo, “Two-dimensional electron gas in Zn-polar ZnMgO/ZnO heterostructure grown by metal-organic vapor phase epitaxy,” *Appl. Phys. Lett.*, vol. 97, p. 111908, 2010. [75](#)
- [125] E. Wigner, “On the interaction of electrons in metals,” *Phys. Rev.*, vol. 46, p. 1002, 1934. [76](#)
- [126] N. D. Drummond and R. J. Needs, “Phase diagram of the low-density two-dimensional homogeneous electron gas,” *Phys. Rev. Lett.*, vol. 102, p. 126402, 2009. [76](#), [82](#), [83](#)
- [127] B. Tanatar and D. M. Ceperley, “Ground state of the two-dimensional electron gas,” *Phys. Rev. B*, vol. 39, p. 5005, 1989. [76](#)
- [128] F. Rapisarda and G. Senatore, “Diffusion Monte Carlo study of electrons in two-dimensional layers,” *Aust. J. Phys.*, vol. 49, p. 161, 1996. [76](#)
- [129] M. Zarenia, D. Neilson, B. Partoens, and F. M. Peeters, “Wigner crystallization in transition metal dichalcogenides: A new approach to correlation energy,” *Phys. Rev. B*, vol. 95, p. 115438, 2017. [76](#)
- [130] Y. Zhou, J. Sung, E. Brutschea, I. Esterlis, Y. Wang, G. Scuri, R. J. Gelly, H. Heo, T. Taniguchi, K. Watanabe, G. Zaránd, M. D. Lukin, E. Kim, P. Demler, and H. Park, “Bilayer Wigner crystals in a transition metal dichalcogenide heterostructure,” *Nature*, vol. 595, p. 48, 2021. [76](#)
- [131] J. Knörzer, M. J. A. Schuetz, G. Giedke, D. S. Wild, K. De Greve, R. Schmidt, M. D. Lukin, and J. I. Cirac, “Wigner crystals in two-dimensional transition-metal dichalcogenides: Spin physics and readout,” *Phys. Rev. B*, vol. 101, p. 125101, 2020. [76](#)
- [132] P. P. Ewald, “Die berechnung optischer und elektrostatischer gitterpotentiale,” *Annalen der Physik*, vol. 369, p. 253, 1921. [77](#)

- [133] V. Saunders, C. Freyria-Fava, R. Dovesi, L. Salasco, and C. Roetti, “On the electrostatic potential in crystalline systems where the charge density is expanded in gaussian functions,” *Mol. Phys.*, vol. 77, p. 629, 1992. [78](#)
- [134] D. Parry, “The electrostatic potential in the surface region of an ionic crystal,” *Surface Science*, vol. 49, p. 433, 1975. [78](#)
- [135] R. Hunt, *Ab initio modelling of two-dimensional semiconductors*. PhD thesis, Lancaster University, 2019. [78](#), [80](#), [81](#)
- [136] N. D. Drummond, R. J. Needs, A. Sorouri, and W. M. C. Foulkes, “Finite-size errors in continuum quantum Monte Carlo calculations,” *Phys. Rev. B*, vol. 78, p. 125106, 2008. [81](#), [94](#), [99](#)
- [137] D. Ceperley, “Ground state of the fermion one-component plasma: A Monte Carlo study in two and three dimensions,” *Phys. Rev. B*, vol. 18, p. 3126, 1978. [81](#)
- [138] L. Bonsall and A. A. Maradudin, “Some static and dynamical properties of a two-dimensional Wigner crystal,” *Phys. Rev. B*, vol. 15, p. 1959, 1977. [81](#)
- [139] B. Efron and R. Tibshirani, *An introduction to the bootstrap*. Chapman & Hall, 1993. [81](#)
- [140] H. Qiu, T. Xu, Z. Wang, W. Ren, H. Nan, Z. Ni, Q. Chen, S. Yuan, F. Miao, F. Song, G. Long, Y. Shi, L. Sun, J. Wang, and X. Wang, “Hopping transport through defect-induced localized states in molybdenum disulphide,” *Nat. Commun.*, vol. 4, p. 2642, 2013. [82](#)
- [141] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim, “The electronic properties of graphene,” *Rev. Mod. Phys.*, vol. 81, p. 109, 2009. [84](#)
- [142] S. Das Sarma, S. Adam, E. H. Hwang, and E. Rossi, “Electronic transport in two-dimensional graphene,” *Rev. Mod. Phys.*, vol. 83, p. 407, 2011. [84](#)

- [143] R. Muñoz and C. Gómez-Aleixandre, “Review of CVD synthesis of graphene,” *Chem. Vap. Depos.*, vol. 19, p. 297, 2013. 84
- [144] L. Vicarelli, S. J. Heerema, C. Dekker, and H. W. Zandbergen, “Controlling defects in graphene for optimizing the electrical properties of graphene nanodevices,” *ACS Nano*, vol. 9, p. 3428, 2015. 84
- [145] P. Y. Huang, J. C. Meyer, and D. A. Muller, “From atoms to grains: Transmission electron microscopy of graphene,” *MRS Bulletin*, vol. 37, p. 1214, 2012. 84
- [146] S. J. Zhang, S. S. Lin, X. Q. Li, X. Y. Liu, H. A. Wu, W. L. Xu, P. Wang, Z. Q. Wu, H. K. Zhong, and Z. J. Xu, “Opening the band gap of graphene through silicon doping for the improved performance of graphene/GaAs heterojunction solar cells,” *Nanoscale*, vol. 8, p. 226, 2016. 84
- [147] A. Hashimoto, K. Suenaga, A. Gloter, K. Urita, and S. Iijima, “Direct evidence for atomic defects in graphene layers,” *Nature*, vol. 430, p. 870, 2004. 84
- [148] A. W. Robertson and J. H. Warner, “Atomic resolution imaging of graphene by transmission electron microscopy,” *Nanoscale*, vol. 5, p. 4079, 2013. 84
- [149] M. T. Lusk and L. D. Carr, “Nanoengineering defect structures on graphene,” *Phys. Rev. Lett.*, vol. 100, p. 175503, 2008. 84
- [150] G. M. Yang, H. Z. Zhang, X. F. Fan, and W. T. Zheng, “Density functional theory calculations for the quantum capacitance performance of graphene-based electrode material,” *J. Phys. Chem. C*, vol. 119, p. 6464, 2015. 84
- [151] Y. Okamoto, “Density functional theory calculations of lithium adsorption and insertion to defect-free and defective graphene,” *J. Phys. Chem. C*, vol. 120, p. 14009, 2016. 84

- [152] E. Kaxiras and K. C. Pandey, “Energetics of defects and diffusion mechanisms in graphite,” *Phys. Rev. Lett.*, vol. 61, p. 2693, 1988. [84](#)
- [153] A. A. El-Barbary, R. H. Telling, C. P. Ewels, M. I. Heggie, and P. R. Briddon, “Structure and energetics of the vacancy in graphite,” *Phys. Rev. B*, vol. 68, p. 144107, 2003. [84](#), [85](#), [101](#), [103](#)
- [154] L. Li, S. Reich, and J. Robertson, “Defect energies of graphite: Density-functional calculations,” *Phys. Rev. B*, vol. 72, p. 184109, 2005. [84](#)
- [155] S. Azevedo, J. R. Kaschny, C. M. C. de Castilho, and F. de Brito Mota, “A theoretical investigation of defects in a boron nitride monolayer,” *Nanotechnology*, vol. 18, p. 495707, 2007. [84](#)
- [156] J. R. Reimers, A. Sajid, R. Kobayashi, and M. J. Ford, “Understanding and calibrating density-functional-theory calculations describing the energy and spectroscopy of defect sites in hexagonal boron nitride,” *J. Chem. Theory Comput.*, vol. 14, p. 1602, 2018. [84](#)
- [157] W. H. Blades, N. J. Frady, P. M. Litwin, S. J. McDonnell, and P. Reinke, “Thermally induced defects on WSe₂,” *J. Phys. Chem. C*, vol. 124, p. 15337, 2020. [84](#)
- [158] F. Banhart, “Irradiation effects in carbon nanostructures,” *Rep. Prog. Phys.*, vol. 62, p. 1181, 1999. [85](#)
- [159] J. R. Hahn and H. Kang, “Vacancy and interstitial defects at graphite surfaces: Scanning tunneling microscopic study of the structure, electronic property, and yield for ion-induced defect creation,” *Phys. Rev. B*, vol. 60, p. 6007, 1999. [85](#)
- [160] R. H. Telling and M. I. Heggie, “Radiation defects in graphite,” *Philos. Mag.*, vol. 87, p. 4797, 2007. [85](#)

- [161] A. Bachmatiuk, J. Zhao, S. M. Gorantla, I. G. G. Martinez, J. Wiedermann, C. Lee, J. Eckert, and M. H. Rummeli, “Low voltage transmission electron microscopy of graphene,” *Small*, vol. 11, p. 515, 2015. 85
- [162] S. Goriparti, E. Miele, F. De Angelis, E. Di Fabrizio, R. Proietti Zaccaria, and C. Capiglia, “Review on recent progress of nanostructured anode materials for Li-ion batteries,” *J. Power Sources*, vol. 257, p. 421, 2014. 85
- [163] A. H. Farokh Niaei, T. Hussain, M. Hankel, and D. J. Searles, “Hydrogenated defective graphene as an anode material for sodium and calcium ion batteries: A density functional theory study,” *Carbon*, vol. 136, p. 73, 2018. 85
- [164] V. Palomares, P. Serras, I. Villaluenga, K. B. Hueso, J. Carretero-González, and T. Rojo, “Na-ion batteries, recent advances and present challenges to become low cost energy storage systems,” *Energy Environ. Sci.*, vol. 5, p. 5884, 2012. 85
- [165] J. Lee, Z. Yang, W. Zhou, S. J. Pennycook, S. T. Pantelides, and M. F. Chisholm, “Stabilization of graphene nanopore,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 111, p. 7522, 2014. 85
- [166] H. Qi, Z. Li, Y. Tao, W. Zhao, K. Lin, Z. Ni, C. Jin, Y. Zhang, K. Bi, and Y. Chen, “Fabrication of sub-nanometer pores on graphene membrane for ion selective transport,” *Nanoscale*, vol. 10, p. 5350, 2018. 85
- [167] P. A. Thrower and R. M. Mayer, “Point defects and self-diffusion in graphite,” *Phys. Status Solidi A*, vol. 47, p. 11, 1978. 85, 106
- [168] J. I. Paredes, P. Solís-Fernández, A. Martínez-Alonso, and J. M. D. Tascón, “Atomic vacancy engineering of graphitic surfaces: Controlling the generation and harnessing the migration of the single vacancy,” *J. Phys. Chem. C*, vol. 113, p. 10249, 2009. 85

- [169] E. Asari, M. Kitajima, K. G. Nakamura, and T. Kawabe, “Thermal relaxation of ion-irradiation damage in graphite,” *Phys. Rev. B*, vol. 47, p. 11143, 1993. [85](#)
- [170] G. H. Vineyard, “Frequency factors and isotope effects in solid state rate processes,” *J. Phys. Chem. Solids*, vol. 3, p. 121, 1957. [86](#)
- [171] H. J. Monkhorst and J. D. Pack, “Special points for Brillouin-zone integrations,” *Phys. Rev. B*, vol. 13, p. 5188, 1976. [88](#)
- [172] G. Rajagopal, R. J. Needs, S. Kenny, W. M. C. Foulkes, and A. James, “Quantum Monte Carlo calculations for solids using special k points methods,” *Phys. Rev. Lett.*, vol. 73, p. 1959, 1994. [88](#)
- [173] G. Rajagopal, R. J. Needs, A. James, S. D. Kenny, and W. M. C. Foulkes, “Variational and diffusion quantum Monte Carlo calculations at nonzero wave vectors: Theory and application to diamond-structure germanium,” *Phys. Rev. B*, vol. 51, p. 10591, 1995. [88](#)
- [174] A. Baldereschi, “Mean-value point in the Brillouin zone,” *Phys. Rev. B*, vol. 7, p. 5212, 1973. [89](#)
- [175] C. Lin, F. H. Zong, and D. M. Ceperley, “Twist-averaged boundary conditions in continuum quantum Monte Carlo algorithms,” *Phys. Rev. E*, vol. 64, p. 016702, 2001. [89](#), [94](#), [95](#)
- [176] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. J. Probert, K. Refson, and M. C. Payne, “First principles methods using CASTEP,” *Z. Kristallogr.*, vol. 220, p. 567, 2005. [89](#)
- [177] D. Vanderbilt, “Soft self-consistent pseudopotentials in a generalized eigenvalue formalism,” *Phys. Rev. B*, vol. 41, p. 7892, 1990. [89](#)
- [178] B. T. Kelly, *Physics of graphite*. Applied Science, 1981. [90](#), [100](#)

- [179] M. S. Dresselhaus, G. Dresselhaus, K. Sugihara, I. L. Spain, and H. A. Goldberg, *Graphite Fibers and Filaments*. Springer-Verlag, 1988. 90, 100
- [180] K. Kunc and R. M. Martin, “Ab initio force constants of gaas: A new approach to calculation of phonons and dielectric properties,” *Phys. Rev. Lett.*, vol. 48, p. 406, 1982. 90
- [181] L. Kleinman and D. M. Bylander, “Efficacious form for model pseudopotentials,” *Phys. Rev. Lett.*, vol. 48, p. 1425, 1982. 90
- [182] T. A. Halgren and W. N. Lipscomb, “The synchronous-transit method for determining reaction pathways and locating molecular transition states,” *Chem. Phys. Lett.*, vol. 49, p. 225, 1977. 91
- [183] N. Govind, M. Petersen, G. Fitzgerald, D. King-Smith, and J. Andzelm, “A generalized synchronous transit method for transition state location,” *Comput. Mater. Sci.*, vol. 28, p. 250, 2003. 91
- [184] J. E. Sinclair and R. Fletcher, “A new method of saddle-point location for the calculation of defect migration energies,” *Journal of Physics C: Solid State Physics*, vol. 7, p. 864, 1974. 91
- [185] D. Alfè and M. J. Gillan, “Efficient localized basis set for quantum Monte Carlo calculations on condensed matter,” *Phys. Rev. B*, vol. 70, p. 161101, 2004. 91
- [186] C. W. M. Castleton and S. Mirbt, “Finite-size scaling as a cure for supercell approximation errors in calculations of neutral native defects in InP,” *Phys. Rev. B*, vol. 70, p. 195202, 2004. 94, 98
- [187] L. M. Fraser, W. M. C. Foulkes, G. Rajagopal, R. J. Needs, S. D. Kenny, and A. J. Williamson, “Finite-size effects and Coulomb interactions in quantum Monte Carlo calculations for homogeneous systems with periodic boundary conditions,” *Phys. Rev. B*, vol. 53, p. 1814, 1996. 94

- [188] S. Chiesa, D. M. Ceperley, R. M. Martin, and M. Holzmann, “Finite-size error in many-body simulations with long-range interactions,” *Phys. Rev. Lett.*, vol. 97, p. 076404, 2006. [94](#)
- [189] J. D. Wadey, A. Markevich, A. Robertson, J. Warner, A. Kirkland, and E. Besley, “Mechanisms of monovacancy diffusion in graphene,” *Chem. Phys. Lett.*, vol. 648, p. 161, 2016. [101](#), [103](#), [106](#), [107](#)
- [190] C. Ronchi, M. Datteo, D. Perilli, L. Ferrighi, G. Fazio, D. Selli, and C. Di Valentin, “ π magnetism of carbon monovacancy in graphene by hybrid density functional calculations,” *J. Phys. Chem. C*, vol. 121, p. 8653, 2017. [101](#), [106](#)
- [191] S. T. Skowron, I. V. Lebedeva, A. M. Popov, and E. Bichoutskaia, “Energetics of atomic scale structure changes in graphene,” *Chem. Soc. Rev.*, vol. 44, p. 3143, 2015. [101](#)
- [192] A. Tkatchenko, R. A. DiStasio, R. Car, and M. Scheffler, “Accurate and efficient method for many-body van der Waals interactions,” *Phys. Rev. Lett.*, vol. 108, p. 236402, 2012. [101](#)
- [193] A. Ambrosetti, A. M. Reilly, R. A. DiStasio, and A. Tkatchenko, “Long-range correlation energy calculated from coupled atomic response functions,” *J. Chem. Phys.*, vol. 140, p. 18A508, 2014. [101](#)
- [194] Y. Ma, P. O. Lehtinen, A. S. Foster, and R. M. Nieminen, “Magnetic properties of vacancies in graphene and single-walled carbon nanotubes,” *New J. Phys.*, vol. 6, p. 68, 2004. [101](#), [103](#)
- [195] W. S. Paz, W. L. Scopel, and J. C. Freitas, “On the connection between structural distortion and magnetism in graphene with a single vacancy,” *Solid State Commun.*, vol. 175, p. 71, 2013. [101](#)
- [196] C. D. Latham, M. I. Heggie, M. Alatalo, S. Öberg, and P. R. Briddon, “The contribution made by lattice vacancies to the Wigner effect in radiation-

- damaged graphite,” *Journal of physics. Condensed matter*, vol. 25, p. 135403, 2013. 103
- [197] A. Zobelli, V. Ivanovskaya, P. Wagner, I. Suarez-Martinez, A. Yaya, and C. P. Ewels, “A comparative study of density functional and density functional tight binding calculations of defects in graphene,” *Phys. Status Solidi B*, vol. 249, p. 276, 2012. 103, 106, 107
- [198] M. M. Ervasti, Z. Fan, A. Uppstu, A. V. Krasheninnikov, and A. Harju, “Silicon and silicon-nitrogen impurities in graphene: Structure, energetics, and effects on electronic transport,” *Phys. Rev. B*, vol. 92, p. 235412, 2015. 106
- [199] J. Ma, D. Alfè, A. Michaelides, and E. Wang, “Stone-Wales defects in graphene and other planar sp^2 -bonded materials,” *Phys. Rev. B*, vol. 80, p. 033407, 2009. 106
- [200] D. Porezag, T. Frauenheim, T. Köhler, G. Seifert, and R. Kaschner, “Construction of tight-binding-like potentials on the basis of density-functional theory: Application to carbon,” *Phys. Rev. B*, vol. 51, p. 12947, 1995. 107

Appendix A

Importance sampling transformation of the imaginary-time Schrödinger equation

For a wave function $\Psi(\mathbf{R}, \tau)$ the imaginary-time Schrödinger equation, with some constant energy offset E_T , is

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) - E_T \right] \Psi(\mathbf{R}, \tau) = -\frac{\partial}{\partial \tau} \Psi(\mathbf{R}, \tau). \quad (\text{A.1})$$

Then, using the mixed distribution

$$f(\mathbf{R}, \tau) = \Psi(\mathbf{R}, \tau)\Psi_T(\mathbf{R}), \quad (\text{A.2})$$

for some trial wave function $\Psi_T(\mathbf{R})$, we can transform Eq. (A.1) into its importance-sampled version by substituting

$$\Psi(\mathbf{R}, \tau) = \frac{f(\mathbf{R}, \tau)}{\Psi_T(\mathbf{R})}, \quad (\text{A.3})$$

to obtain

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} + V(\mathbf{R}) \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} - E_{\mathbf{T}} \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} \right] = -\frac{\partial}{\partial \tau} \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})}. \quad (\text{A.4})$$

The trial wave function $\Psi_{\mathbf{T}}(\mathbf{R})$ is independent of τ , so we can pull this out of the time derivative on the right

$$\left[-\frac{1}{2}\nabla_{\mathbf{R}}^2 \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} + V(\mathbf{R}) \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} - E_{\mathbf{T}} \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} \right] = -\frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau), \quad (\text{A.5})$$

and multiply both sides by $\Psi_{\mathbf{T}}(\mathbf{R})$ to give

$$\left[-\frac{1}{2}\Psi_{\mathbf{T}}(\mathbf{R})\nabla_{\mathbf{R}}^2 \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} + V(\mathbf{R})f(\mathbf{R}, \tau) - E_{\mathbf{T}}f(\mathbf{R}, \tau) \right] = -\frac{\partial}{\partial \tau} f(\mathbf{R}, \tau). \quad (\text{A.6})$$

Now, focussing on the Laplacian and using the product rule we find

$$\begin{aligned} \nabla_{\mathbf{R}}^2 \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} &= \nabla_{\mathbf{R}} \cdot \nabla_{\mathbf{R}} \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} \\ &= \nabla_{\mathbf{R}} \left[\frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) - \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right] \\ &= \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) - \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \\ &+ \frac{2}{\Psi_{\mathbf{T}}^3(\mathbf{R})} f(\mathbf{R}, \tau) \left(\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right)^2 - \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \\ &- \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}). \end{aligned} \quad (\text{A.7})$$

Collecting like terms and multiplying by $-\frac{1}{2}\Psi_{\mathbf{T}}(\mathbf{R})$ we obtain

$$\begin{aligned} -\frac{1}{2}\Psi_{\mathbf{T}}(\mathbf{R})\nabla_{\mathbf{R}}^2 \frac{f(\mathbf{R}, \tau)}{\Psi_{\mathbf{T}}(\mathbf{R})} &= -\frac{1}{2}\nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \\ -\frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} f(\mathbf{R}, \tau) \left(\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right)^2 &+ \frac{1}{2\Psi_{\mathbf{T}}(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}). \end{aligned} \quad (\text{A.8})$$

Substituting Eq. (A.8) into Eq. (A.6) we now have

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) - \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} f(\mathbf{R}, \tau) \left(\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right)^2 \\
& + \underbrace{\frac{1}{2\Psi_{\mathbf{T}}(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) + V(\mathbf{R}) f(\mathbf{R}, \tau) - E_{\mathbf{T}} f(\mathbf{R}, \tau)}_K \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau). \tag{A.9}
\end{aligned}$$

Then, adding $\frac{1}{2}K - \frac{1}{2}K$ gives

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \\
& + \left[- \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} f(\mathbf{R}, \tau) \left(\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right)^2 + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) \right] \\
& - \frac{1}{2\Psi_{\mathbf{T}}(\mathbf{R})} f(\mathbf{R}, \tau) \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) + V(\mathbf{R}) f(\mathbf{R}, \tau) - E_{\mathbf{T}} f(\mathbf{R}, \tau) \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau), \tag{A.10}
\end{aligned}$$

and pulling $f(\mathbf{R}, \tau)$ out as a common factor (and swapping the order of the terms in square brackets), we arrive at

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \\
& + f(\mathbf{R}, \tau) \left[\frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) - \frac{1}{\Psi_{\mathbf{T}}^2(\mathbf{R})} \left(\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) \right)^2 \right] \\
& + \left(- \frac{1}{2\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) + V(\mathbf{R}) - E_{\mathbf{T}} \right) f(\mathbf{R}, \tau) \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau). \tag{A.11}
\end{aligned}$$

The terms in square brackets are now just a product rule expansion of $\nabla_{\mathbf{R}}(\Psi_{\mathbf{T}}^{-1}(\mathbf{R}) \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}))$ and we multiply the potential term $V(\mathbf{R})$ by $\Psi_{\mathbf{T}}(\mathbf{R})/\Psi_{\mathbf{T}}(\mathbf{R})$

to find

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) \cdot \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R}) + f(\mathbf{R}, \tau) \nabla_{\mathbf{R}} \left(\frac{\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R})}{\Psi_{\mathbf{T}}(\mathbf{R})} \right) \\
& + \left(- \frac{1}{2 \Psi_{\mathbf{T}}(\mathbf{R})} \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) + V(\mathbf{R}) \frac{\Psi_{\mathbf{T}}(\mathbf{R})}{\Psi_{\mathbf{T}}(\mathbf{R})} - E_{\mathbf{T}} \right) f(\mathbf{R}, \tau) \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau).
\end{aligned} \tag{A.12}$$

Now, the second and third terms on the first line are just another product rule expansion, and we can pull a factor of $\Psi_{\mathbf{T}}^{-1}(\mathbf{R})$ from the first two terms on the second line to give

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \nabla_{\mathbf{R}} \cdot \left[\frac{\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R})}{\Psi_{\mathbf{T}}(\mathbf{R})} f(\mathbf{R}, \tau) \right] \\
& + \left(\frac{1}{\Psi_{\mathbf{T}}(\mathbf{R})} \left[\underbrace{- \frac{1}{2} \nabla_{\mathbf{R}}^2 \Psi_{\mathbf{T}}(\mathbf{R}) + V(\mathbf{R}) \Psi_{\mathbf{T}}(\mathbf{R})}_{\hat{H} \Psi_{\mathbf{T}}(\mathbf{R})} \right] - E_{\mathbf{T}} \right) f(\mathbf{R}, \tau) \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau),
\end{aligned} \tag{A.13}$$

and we recognise that the square brackets on the second line are now just the Hamiltonian \hat{H} acting on $\Psi_{\mathbf{T}}(\mathbf{R})$, so

$$\begin{aligned}
& - \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \nabla_{\mathbf{R}} \cdot \left[\underbrace{\frac{\nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R})}{\Psi_{\mathbf{T}}(\mathbf{R})}}_{\mathbf{V}(\mathbf{R})} f(\mathbf{R}, \tau) \right] + \left(\underbrace{\frac{\hat{H} \Psi_{\mathbf{T}}(\mathbf{R})}{\Psi_{\mathbf{T}}(\mathbf{R})}}_{E_{\mathbf{L}}(\mathbf{R})} - E_{\mathbf{T}} \right) f(\mathbf{R}, \tau) \\
& = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau).
\end{aligned} \tag{A.14}$$

Finally, by substituting the drift velocity $\mathbf{V}(\mathbf{R}) = \Psi_{\mathbf{T}}^{-1}(\mathbf{R}) \nabla_{\mathbf{R}} \Psi_{\mathbf{T}}(\mathbf{R})$ and the local energy $E_{\mathbf{L}}(\mathbf{R}) = \Psi_{\mathbf{T}}^{-1} \hat{H} \Psi_{\mathbf{T}}(\mathbf{R})$, we arrive at the final result

$$- \frac{1}{2} \nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + \nabla_{\mathbf{R}} \cdot [\mathbf{V}(\mathbf{R}) f(\mathbf{R}, \tau)] + (E_{\mathbf{L}}(\mathbf{R}) - E_{\mathbf{T}}) f(\mathbf{R}, \tau) = - \frac{\partial}{\partial \tau} f(\mathbf{R}, \tau), \tag{A.15}$$

which is the importance-sampled imaginary-time Schrödinger equation.