

Ab initio modelling of two-dimensional semiconductors



Ryan James Hunt
Department of Physics
Lancaster University

This thesis is submitted for the degree of
Doctor of Philosophy

September 2019

This thesis is dedicated to my mother, and to the rest of my family. I am here
because of you all.

Abstract

We study excited-state phenomena in a variety of semiconductor systems, with use of the variational and diffusion quantum Monte Carlo (QMC) methods. Firstly, we consider the formation of charge-carrier complexes in the Mott-Wannier model, for systems of restricted geometrical freedom (the coupled quantum well bilayer, and the quantum ring). We find in such systems that geometrical constraints lead to the characteristic formation of certain charge-carrier complexes, and highlight how such effects are of relevance to the interpretation of recent experiments.

Secondly, we illuminate a key difference between two-dimensional systems formed from geometrical restriction, and those which are truly two-dimensional in extent, by introduction of the Keldysh interaction. We then study the formation of charge-carrier complexes in two-dimensional semiconductors and their heterostructures in the so-called Mott-Wannier-Keldysh model, deriving appropriate extensions of the Keldysh interaction as necessary.

Thirdly, we undertake a comprehensive survey of the use of continuum QMC methods to evaluate excited-state properties in a truly *ab initio* fashion, establishing best-practices, and presenting energy gap calculations for several real materials. This includes the first published QMC calculation of the electronic energy gaps of a two-dimensional semiconductor, phosphorene.

Finally, we propose an extension of the Keldysh interaction which permits the study of continuum phases, the so-called “periodic Keldysh interaction”, and use it to probe the possible Wigner crystallisation of electrons in a weakly-doped two-dimensional semiconductor.

Declaration

The content of this thesis defines work undertaken as a member of the Graphene NOWNANO Centre for Doctoral Training from April 2016 to September 2019 in the Department of Physics at Lancaster University, and under the supervision of Dr. N. D. Drummond. The following articles, given in order of appearance of their contents, are products, or products in part, of this work:

Chapter 2.2 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, Feb 2018.

Chapter 2.3 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, Mar 2019.

Chapter 3.1 Mostaani, E. and Szyniszewski, M. and Price, C. H. and Maezono, R. and Danovich, M. and Hunt, R. J. and Drummond, N. D. and Fal’ko, V. I., “Diffusion quantum Monte Carlo study of excitonic complexes in two-dimensional transition-metal dichalcogenides,” *Phys. Rev. B*, vol. 96, p. 075431, Aug 2017, and E. Mostaani, R. J. Hunt, D. M. Thomas, M. Szyniszewski, and N. D. Drummond, “Charge-carrier complexes in monolayer transition metal dichalcogenides: quintons, external fields, and the accuracy of the Keldysh interaction potential.” In preparation.

Chapter 3.2 M. Danovich, D. A. Ruiz-Tijerina, R. J. Hunt, M. Szyniszewski, N. D. Drummond, and V. I. Fal’ko, “Localized interlayer complexes in hetero-

bilayer transition metal dichalcogenides,” *Phys. Rev. B*, vol. 97, p. 195452, May 2018, and

F. Vialla, M. Danovich, D. A. Ruiz-Tijerina, M. Massicotte, P. Schmidt, T. Taniguchi, K. Watanabe, R. J. Hunt, M. Szyniszewski, N. D. Drummond, T. G. Pedersen, V. I. Fal’ko, and F. H. L. Koppens, “Tuning of impurity-bound interlayer complexes in a van der Waals heterobilayer,” *2D Materials*, vol. 6, p. 035032, May 2019, and

R. J. Hunt, M. Szyniszewski, D. M. Thomas, E. Mostaani, N. D. Drummond, and V. I. Fal’ko, “Diffusion Monte Carlo study of charge-carrier complexes in multilayer van der Waals heterostructures.” In preparation.

Chapter 4.1 R. J. Hunt, M. Szyniszewski, G. I. Prayogo, R. Maezono, and N. D. Drummond, “Quantum Monte Carlo calculations of energy gaps from first principles,” *Phys. Rev. B*, vol. 98, p. 075122, Aug 2018, and

R. J. Hunt, N. D. Drummond, B. Monserrat, V. Zólyomi, and V. I. Fal’ko, “Diffusion quantum Monte Carlo and *GW* study of the electronic properties of monolayer and bulk hexagonal boron nitride.” In preparation.

I declare that this dissertation 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 recommended maximum word count of 80,000 words.

Ryan James Hunt

Lancaster, September 2019

Acknowledgements

Though this thesis bears the name of one author, it would not exist were it not for the guidance, camaraderie, and support offered by a series of other individuals.

The privileged position which I have occupied within post-graduate research has been created primarily by my supervisor, Neil Drummond. Few human beings are simultaneously capable of great patience, kindness, and cognition. I owe him greatly. I also owe thanks to my second supervisor, Volodya Fal'ko, whose physical intuition is something to behold, and whose input has been critical in various facets of the work presented in this thesis. I know that I am incredibly lucky to have had the chance to stand on these particular pairs of shoulders.

I am indebted to my collaborators: Marcin Szyniszewski, Elaheh Mostaani, Cameron Price, Mark Danovich, David Ruiz-Tijerina, Bartomeu Monserrat, Viktor Zólyomi, Ryo Maezono, Genki Prayogo, David Thomas, and Manus Hayne for their valuable inputs to our shared efforts. I would like to thank them all for their countless hours of discussion and (scientific, of course) argument. Further, I would like to thank collectively all of the members of the Condensed Matter theory group in the Physics Department at Lancaster University. Our meetings and discussions have often been a highlight of my week. Special thanks goes to Marcin Szyniszewski, Simon Malzard, and Marjan Famili, without whom I would have had spent lonely days, and smiled much less often.

Finally, whilst knowing that a lowly sentence does her disproportionately vast contribution to my life no justice at all, I thank Emily Alderson for all of her love, support, and patience in these last few years.

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) facility, the N8 HPC (EPSRC grant number EP/K000225/1), the ARCHER UK National Supercomputing Service, and by the facilities of the Centre for Information Science at the Japan Advanced Institute for Science and Technology (JAIST).

Table of Contents

List of Abbreviations	13
List of Figures	14
List of Tables	20
Introduction	24
1 Theoretical underpinnings of quantum Monte Carlo	26
1.1 Two-dimensional semiconductors	26
1.2 The quantum many-body problem	28
1.3 Statistics	29
1.3.1 Monte Carlo integration	29
1.3.2 Markov chains	30
1.3.3 The Metropolis-Hastings algorithm	33
1.3.4 Metropolis-Hastings caveats	36
1.4 Variational Monte Carlo (VMC)	38
1.4.1 The many-particle Schrödinger equation	39
1.4.2 VMC expectation values	41
1.4.3 Explicitly correlated trial wave functions	43
1.4.4 The Kato cusp conditions	47
1.4.5 Non-standard Jastrow terms	49
1.4.6 Backflow transformations	49

1.4.7	Optimisation of many-electron wave functions	50
1.5	Diffusion Monte Carlo (DMC)	52
1.5.1	Imaginary time Schrödinger equation (ITSE)	53
1.5.2	Importance sampling transformation and the ISITSE	54
1.5.3	Expectation values in DMC	58
1.6	Pseudopotentials in QMC	59
2	Charge-carrier complexes in models of III-V semiconductors	61
2.1	The effective mass approximation	61
2.1.1	Quasiparticle dynamics	61
2.1.2	Excitonic units	64
2.1.3	Charge-carrier complexes	64
2.2	Coupled quantum wells in III-V semiconductors	66
2.2.1	Trion binding energies	68
2.2.2	Biexciton binding energies	70
2.3	Quantum rings in III-V semiconductors	71
3	Charge-carrier complexes in models of two-dimensional semi-	
	conductors	77
3.1	Charge complexes in two-dimensional TMDs	77
3.1.1	The Keldysh interaction	78
3.1.2	The Mott-Wannier-Keldysh model	82
3.1.3	Quintons and related heavy complexes in TMDs	85
3.1.4	Beyond the Keldysh interaction: <i>ab initio</i> dielectric functions	88
3.2	Charge complexes in TMD heterobilayers	91
3.2.1	Bilayer Keldysh interaction	91
3.2.2	Free and donor-bound complexes in hBN-encapsulated MoSe ₂ /	
	WSe ₂ heterobilayers	93
3.2.2.1	Sensitivity to model parameters	98

3.2.2.2	Intralayer binding energies: comparison with experiment	100
3.3	Complexes in multilayers of two-dimensional semiconductors	101
3.3.1	Obtaining real-space interaction potentials	103
3.3.2	Excitons and Biexcitons in few-layer TMDs	103
3.3.2.1	Excitons	105
3.3.2.2	Biexcitons	107
3.3.2.3	Bulk crossover	109
3.3.2.4	General trends	111
4	Excited-states by QMC	114
4.1	Excited-state QMC	114
4.1.1	Quasiparticle and excitonic gaps	114
4.1.2	Singlet and triplet excitations	117
4.1.3	Wave-function nodes and variational principles	118
4.1.4	Nodal topology	121
4.1.5	Finite-size effects	122
4.1.6	QMC band structures	129
4.1.7	Excitations in metallic systems	130
4.1.8	Computational expense	131
4.1.9	Nuclear relaxation and vibrational effects	132
4.2	Computational details	133
4.2.1	DFT orbital generation	133
4.2.2	QMC calculations	134
4.2.2.1	Slater-Jastrow(-backflow) wave functions	134
4.2.2.2	Multideterminant trial wave functions	134
4.3	Results and discussion	136
4.3.1	Atoms	136
4.3.1.1	H atom: a model of excited-state fixed-node errors	136
4.3.1.2	Ne atom: VMC and backflow?	140

4.3.2	Molecules	143
4.3.2.1	H ₂ dimer	143
4.3.2.2	O ₂ dimer	145
4.3.2.3	Non-dimer molecules	147
4.3.3	Three-dimensional solids	149
4.3.3.1	Diamond Si	149
4.3.3.2	Hexagonal boron nitride	155
4.3.3.3	Cubic boron nitride	160
4.3.3.4	α -quartz: SiO ₂	161
4.3.4	Two-dimensional phosphorene	162
4.3.5	Summary	166
5	Electron gases in doped two-dimensional semiconductors	169
5.1	Periodic Keldysh interaction	170
5.1.1	Lattice sum	170
5.1.2	Computational implementation	172
5.1.2.1	Cusp conditions: $\log-u$ Jastrow term	172
5.1.2.2	Precomputation and runtime evaluation	173
5.1.2.3	Blip representation of pre-computed sum	174
5.1.2.4	Integral correction	176
5.1.3	Finite-size errors in first-principles calculations on two-dimensional materials	178
5.2	Crystallisation of the electron gas in doped two-dimensional semiconductors	179
5.2.1	Trial wave functions	181
5.2.1.1	Crystal phase	181
5.2.1.2	Fluid phase	183
5.2.2	DMC calculations	184
5.2.3	Finite-size effects	186
5.2.4	Phase diagram	186

5.2.5	Some comments on experimental detection	189
6	Future Work	191
6.1	Superfluidity in electron-hole bilayers: revisited	191
6.2	Holes immersed in a doped two-dimensional semiconductor	192
7	Conclusions	194
A	Monte Carlo bootstrapping	232
B	Approximate bilayer interactions	233
C	Biexcitons in TMD homostructures	235
D	Hartree-Fock theory of Keldysh-screened electron gases	240

List of Abbreviations

QMC	Quantum Monte Carlo
VMC	Variational Monte Carlo
DMC	Diffusion Monte Carlo
DFT	Density functional theory
TMD	Transition metal dichalcogenide
BO	Born-Oppenheimer
BZ	Brillouin zone
FN	Fixed-node
GS	Ground-state
PP	Pseudopotential
SC	Supercell

List of Figures

2.1	Schematic depiction of an indirect negative trion (X^-) in a coupled quantum well system. The yellow(blue) layer is the electron(hole)-bearing layer.	67
2.2	Negative-trion binding energy as a function of interlayer spacing d and electron-hole mass ratio σ . The inset shows the edge of the region of stability for the negative trion in greater detail.	69
2.3	Schematic representations of a quantum ring. Grey areas represent hole-confining regions in the type-II ring (Sb-rich regions in the case of GaSb/GaAs grown rings).	72
3.1	Schematic band structure of MoX_2 -type TMDs ($X=\text{S, Se}$). In WX_2 -type TMDs, the VBM and CBM states host electrons with <i>different</i> spin quantum numbers. Inequivalent valleys are designated $\xi = \pm 1$, and the quantities Δ_{Ex} , Δ_{QP} , and $\Delta_{\sigma\text{C}}$ are the excitonic and quasiparticle energy gaps (see Chapter 4), and the conduction band spin splitting respectively.	78
3.2	Equipotential surfaces from a charge at the origin in the Coulomb and Keldysh interactions.	81
3.3	Possible daughter configurations of a biexciton decay.	84
3.4	VMC and DMC results for the binding energy of a electron in a MoS_2 layer bound to a defect at $z = 2\text{\AA}$ with an effective charge Z , compared to the results of Ref. [10]. VMC and DMC error bars are shown, but are $\mathcal{O}(10^{-8}$ a.u.).	89

3.5	Comparison of the <i>ab initio</i> AML quasiparticle-quasiparticle interaction with the Keldysh interaction. (Left panel) Inside a layer ($z = 0$). (Right panel) Outside a layer ($z = 15$ a.u.).	90
3.6	Schematic diagram depicting type-II (staggered) band alignment for a heterobilayer system. Orange (blue) shaded pockets highlight bands which host holes (electrons). In the MoSe ₂ /WSe ₂ case, “Layer 1” is identified as MoSe ₂ (where holes dominantly reside, in band v'), and “Layer 2” as WSe ₂ (where electrons dominantly reside, in band c). We have also labelled the minority bands c' , v , where electrons and holes <i>may</i> reside, and do in some of our calculations.	93
3.7	Schematic diagram for an $n = 3$ heterostructure of two-dimensional semiconductors. An interlayer exciton is also depicted.	101
3.8	Encapsulation models for an $n = 8$ homostructure. (left) dielectric screening model in cases (1) and (2), grey space corresponds to dielectric medium. (right) dielectric screening model in case (3), with extra (purple) hBN layers for visibility.	104
3.9	Binding energies of excitons in few-layer MoSe ₂ homostructures. The upper branch of black squares represent intralayer excitons, with “c”, “e”, and “ce” labelling excitons in central, edge, and centre-edge layers, respectively. $X_{e(c)}(\infty)$ is defined in the text. Red circles, green diamonds, light-red upward-point triangles, and a single light-green downward-pointing triangle label interlayer complexes separated by 1, 2, and 3 empty MoSe ₂ layers. The light-green downward-pointing triangle, for example, is the exciton formed in a five-layer system with electron and hole occupying opposite edge layers.	106

- 4.1 Scaled difference of exciton binding energy E_X^b and the exciton Rydberg against the lattice parameter L in an effective-mass model of a three-dimensional exciton confined in a periodic FCC cell. $R_y^* = m_e^* m_h^* / [2\epsilon^2 (m_e^* + m_h^*)]$ and $a_0^* = \epsilon (m_e^* + m_h^*) / (m_e^* m_h^*)$ are the exciton Rydberg and the exciton Bohr radius, respectively, where m_e^* and m_h^* are the electron and hole masses and ϵ is the permittivity. The exciton Rydberg is the binding energy of a free exciton (or, one in a cell of infinite extent). The gradient on this log-log plot gives the scaling exponent of the FS error in the exciton binding energy. 128
- 4.2 Approximations to the first-excited-state energy of an H atom using the ψ^γ excited-state trial wave function of Eqn. (4.11) as a function of γ , obtained by various means. DMC errors are smaller than the thickness of the lines. The pocket eigenvalues outside and inside the nodal surface, $E_{\text{outside}}^{\text{pocket}}$ and $E_{\text{inside}}^{\text{pocket}}$, were determined by numerical integration of the Schrödinger equation with fixed-node boundary conditions, and $\langle \hat{H} \rangle = \langle \psi^\gamma | \hat{H} | \psi^\gamma \rangle$, where \hat{H} is the Hamiltonian. . . 138
- 4.3 DMC first-excited-state energies of an H atom with the trial wave function $\psi_L^{\alpha,\beta}$ [see Eqn. (4.12)] for $L = 2$ and 4 at various amplitudes α of wrinkling of the nodal surface [see Eqn. (4.12)]. DMC error bars are of order the size of the symbols. 139
- 4.4 Nodal surface of the SJB trial wave function $\psi_{L=4}^{\alpha=0.1,\beta}$ [see Eqn. (4.12)] for the first excited state of an H atom. The wave function is optimised by VMC energy minimisation. SJB- n labels the nodal surface of the SJB wave function after the n^{th} cycle of energy minimisation. The $n = 3$ and 4 cases are indistinguishable from each other, and correspond to the termination of the optimisation process. 139

4.5	Non-dimer molecules whose energy gaps we have calculated. From left to right: anthracene ($C_{14}H_{10}$), tetracyanoethylene (C_6N_4), benzothiazole (C_7H_5NS), and boron trifluoride (BF_3).	147
4.6	Finite-size errors in uncorrected SJ-DMC quasiparticle and excitonic gaps Δ_{QP} and Δ_{Ex} of Si as a function of the number of primitive cells N_P in the supercell. The dashed lines show the infinite-system gaps estimated by subtracting the supercell Madelung constant from the gaps obtained in finite cells and averaging over the different cells.	150
4.7	Uncorrected SJ-DMC quasiparticle and excitonic energy gaps of Si at Γ , plotted against DFT analogue “quasiparticle” (AQP) gaps, obtained using the defect formation energies for positive and negative charged defects. The results were obtained in different sizes and shapes of periodic cell. The straight lines are linear fits of SJ-DMC gap against DFT AQP gaps.	151
4.8	Time-step bias in (a) SJ-DMC energy gaps and (b) SJ-DMC total energies for ground (Γ^0), excitonic ($\Gamma \rightarrow \Gamma$), cationic (Γ^-), and anionic (Γ^+) states of Si. All calculations have been performed in a $2 \times 2 \times 2$ supercell with a target population of 256 walkers. The Madelung correction is not included (and would only offset the gaps by a constant).	152
4.9	SJ-DMC quasiparticle gaps Δ_{QP} and excitonic gaps Δ_{Ex} of bulk hBN against $1/N_P$, where N_P is the number of primitive cells in the supercell. The quasiparticle gaps include the Madelung correction as described in the text. The statistical error bars show the random error in the SJ-DMC gap in a particular supercell; the noise due to quasirandom FS effects clearly exceeds the noise due to the Monte Carlo calculation.	157

4.10	(a) DFT-PBE charge density of the state at Γ_c as a function of lattice parameter c for bulk hBN. $c = 12.5878$ a.u. is the experimental lattice parameter [11]. The charge density is plotted along a straight line in the z direction, passing through a boron atom at $z/c = 0.25$ and a nitrogen atom at $z/c = 0.75$. At large c the CBM at Γ_c is an arbitrary linear combination of the degenerate monolayer CBMs.	
	(b) DFT-PBE bulk hBN band structure at three large values of the lattice parameter c . The inset to (b) displays a close-up of the two near-degenerate states at Γ_c	159
4.11	DMC quasiparticle gaps Δ_{QP} , excitonic gaps Δ_{Ex} , and exciton binding energies E_X^b at Γ against the inverse of the number N_P of primitive cells in the supercell for a free-standing phosphorene monolayer. The Keldysh Madelung constant correction has been applied to the quasiparticle gaps; no FS correction has been applied to the excitonic gaps. The non-diagonal supercell results (filled symbols) have been slightly shifted relative to the 3×2 supercell result for readability.	163
5.1	Blip grid points (purple crosses) within the simulation cell (shaded grey), for a two-dimensional rectangular lattice with simulation cell lattice vectors \mathbf{a}_1 and \mathbf{a}_2	176
5.2	A schematic representation of the integral correction scheme for lattice sums appearing in the Ewald-Keldysh interaction. The green region Ω defines the integration domain.	177
5.3	Lattice points for the striped configuration of electrons in our crystal phase trial function. Black squares (blue circles) denote up(down)-spin, K(')-valley electrons (<i>i.e.</i> the $\{P_\uparrow\}$ ($\{P_\downarrow\}$) set). The grey shaded region denotes one standard deviation of an individual electron orbital, as related to the orbital parameter g	182

5.4 DMC energy as a function of density parameter r_s for paramagnetic
fluid and Wigner crystal phases of our model. 188

List of Tables

2.1	Critical layer separations for negative trions ($d_{\text{X}^{\pm}}^{\text{crit}}$) and biexcitons ($d_{\text{XX}}^{\text{crit}}$) over a range of electron-hole mass ratios σ . For the final three mass ratios, the critical distance lies outside of our chosen range of calculation parameters, and we are only able to quote a bound. The biexciton critical layer separations were obtained from fits given in Ref. [12].	69
2.2	De-excitonisation (E^{D}) and binding (E^{b}) energies, in meV, for excitonic complexes in the quantum ring geometry of Ref. [13]. . . .	75
3.1	Material parameters for various monolayer TMDs. m_e , m_h , r_{\star} , and d are the electron effective mass (for the conduction band at K), the hole effective mass (for the valence band at K), the Keldysh screening parameter (see Section 3.1.1), and the interlayer spacing of the bulk parent crystal, respectively.	79
3.2	Binding energies of larger charge-carrier complexes in different TMDs. The binding energies were evaluated using the effective mass and in-plane permittivity parameters reported in bold in Table 3.1. The fitting functions of Ref. [3] were used to evaluate negative donor ion and donor-bound trion total energies.	87
3.3	Binding energies of excitonic charge complexes in MoS ₂ as modelled by different charge-carrier interaction potentials. ‘‘R-AML’’ is the rounded AML interaction, with the rounding distance given in brackets. Keldysh values are obtained from fits given in Ref. [3]. . .	91

3.4	DMC total energies of various charge-carrier complexes in the hBN/MoSe ₂ /WSe ₂ /hBN heterostructure calculated using the monolayer Keldysh approximation to the bilayer potential (Eqn. (B.3) and Eqn. (B.4)) and using the full bilayer interaction (Eqn. (3.16) and Eqn. (3.17)). Interlayer complexes in which all the electrons are in the MoSe ₂ layer and all the holes are in the WSe ₂ layer are listed in the upper section of the table; complexes in which some of the electrons are in the WSe ₂ layer are listed in the lower section of the table.	96
3.5	Decays of complexes and their associated binding energies in hBN/MoSe ₂ /WSe ₂ /hBN.	97
3.6	Changes in binding energies per fractional parameter changes for DMC binding energies $E_{D_{c'}^0 X_{vc'}}^b$ of the interlayer donor-bound trion under variations δP of different model parameters P . With the exception of the dielectric constant ϵ , the parameters are varied by $\pm 10\%$ about the values listed in Table 3.1 and the central difference approximation is used to estimate the derivative with respect to the parameter value. The binding energy when all the parameters take the values listed in Table 3.1 is $E_{D_{c'}^0 X_{vc'}}^b = 22.516(4)$ meV. Note that $r_{*,1}$ and $r_{*,2}$ are here the screening lengths for monolayers isolated in vacuum. The derivative of the binding energy with respect to the dielectric constant ϵ was evaluated by the chain rule, as described in the text.	99
3.7	Binding energies in meV of intralayer biexcitons in few-layer TMD homostructures. Indices of c, e and ce denote biexcitons confined to central, edge, and centre-edge (that between the centre and the edge) layers.	107

- 4.1 n^{th} ionisation potential of an all-electron Ne atom at various levels of QMC theory, together with corrected non-relativistic experimental values [14]. The mean absolute errors (MAEs) have been calculated over all ionisation potentials obtained within a given level of theory. 141
- 4.2 DMC ionisation potentials of the H_2 and O_2 dimers. All-electron (AE) and pseudopotential (PP) calculations have been performed with Gaussian (G) and plane-wave (PW) bases. Calculations employing relaxed excited-state geometries are denoted “ER.” The “J-DMC ($p^+p^+e^-e^-$)” calculations used a Jastrow wave function to describe the ground state of two distinguishable quantum protons and two distinguishable electrons for parahydrogen H_2 , and the ground state of two distinguishable protons and one electron for the parahydrogen cation H_2^+ . Self-consistent quasiparticle GW results are denoted “ $QSGW$,” coupled cluster results with single, double, and (triple) excitations “CCSD(T)” (“EPT” means electron propagator theory), second-order Møller-Plesset perturbation theory results “MP2,” quadratic configuration interaction “QCI” (with levels of excitations as with coupled cluster), and results obtained by means of the generalised James-Coolidge expansion “JCE.” 144
- 4.3 SJ-DMC ionisation potentials and electron affinities of various non-dimer molecules. Calculations employing relaxed excited-state geometries are designated with “(ER).” Adiabatic gaps are given the subscript “A” and vertical gaps the subscript “V.” 147
- 4.4 DMC ionisation potentials (IPs) and electron affinities (EAs) of C_6N_4 at various levels of QMC theory, compared to experiment and other methods. Calculations employing relaxed excited-state geometries are designated “(ER),” and those employing reoptimised backflow functions “(R).” Adiabatic gaps are given the subscript “A,” vertical gaps the subscript “V.” 148

4.5	Uncorrected quasiparticle and excitonic energy gaps Δ_{QP} and Δ_{Ex} of Si evaluated in SJ-DMC for different simulation supercells.	149
4.6	Finite-size-corrected SJ-DMC vertical quasiparticle gaps Δ_{QP} and SJ-DMC vertical excitonic gaps Δ_{Ex} at the Γ point in Si for various non-cubic supercells. After correction, the Δ_{QP} and Δ_{Ex} data sets have lower variances by a factor of more than two.	151
4.7	Static-nucleus quasiparticle gaps for bulk hBN, determined by different methods, compared with experimental results. Recall that DFT vibrational-renormalisation calculations indicate that the static-nucleus gaps should be renormalised by -0.40 eV at 300 K. An asterisk (*) denotes the SJ-DMC energy gap from $K_v \rightarrow M_c$	157
4.8	Uncorrected quasiparticle and excitonic energy gaps Δ_{QP} and Δ_{Ex} of cubic BN evaluated in SJ-DMC for different simulation supercells.	161
4.9	Comparison of the SJ-DMC energy gaps and exciton binding of monolayer phosphorene with results available in the literature for a free-standing monolayer and a monolayer on an SiO_2 substrate. EMA labels results of an effective mass approximation analysis.	164
C.1	Charge-carrier complex binding energies (in meV) in a MoSe_2 monolayer, in various dielectric environments.	236
C.2	Charge-carrier complex binding energies (in meV) in an $n = 2$ MoSe_2 homostructure, in various dielectric environments.	236
C.3	Charge-carrier complex binding energies (in meV) in an $n = 3$ MoSe_2 homostructure, in various dielectric environments.	236
C.4	Charge-carrier complex binding energies (in meV) in an $n = 4$ MoSe_2 homostructure, in various dielectric environments.	237
C.5	Charge-carrier complex binding energies (in meV) in an $n = 5$ MoSe_2 homostructure, in various dielectric environments.	239

Introduction

From the vantage point of condensed matter physics, quantum mechanics is arguably a “theory of everything”. Dirac was perhaps the first to take such a stance, saying in 1929 [15]

“The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems without too much computation.”

Since Dirac’s day, however, some things have changed. Firstly, the potency of modern statistical and mathematical methods mean that “solving” these complicated equations, and the ability to obtain useful insight from them are two separate things. Secondly, the last hundred years has seen a huge technological advance in the advent of the computer (and indeed, the supercomputer). What one considers “too much computation” is fundamentally changed. Modern quantum Monte Carlo (QMC) methods, and their application to condensed matter systems, are a consequence of both of these advances.

“Condensed matter” can be defined somewhat generically as any system comprised of an extremely large (often macroscopic) number of strongly interacting constituents. Solids and liquids, but not weakly interacting gases, are examples of

condensed phases, held together by strong Coulombic forces between their atomic, ionic, or molecular constituents. “Condensed matter physics” is then the study of such phases, and their description in terms of underlying physical laws.

Lurches in our understanding of low dimensional semiconductors in particular were critical to the realisation of the modern computer, and helped grow “nanotechnology” into the immense research field and umbrella term which it now is. Presently, research into low dimensional semiconductors is almost entirely focussed on the newly discovered variety of atomically thin semiconductors, the interest in which boomed following the first isolation of graphene from graphite in 2004 [16].

In this work, we have used QMC methods to study a variety of two-dimensional semiconductor systems, with a focus on their excited state properties. Such excited state properties effectively determine the response of a semiconducting system to light, and are integral to the operating principles of a variety of known (and future) optoelectronic components: photodiodes, solar cells, light emitting diodes, *etc.*

The content of this thesis is as follows. In Chapter 1, we will introduce the theoretical background of the variational and diffusion quantum Monte Carlo methods. In Chapter 2 we will apply these methods in studies of charge-carrier complexes in some model systems, aimed at describing realistic semiconductor heterostructures formed from III-V materials (the coupled quantum well, and the quantum ring). Chapter 3 introduces a series of modelling strategies for the study of truly two-dimensional semiconductors – those formed from two-dimensional crystals, and their heterostructures. In Chapter 4 we present a comprehensive review of the use of continuum QMC methods in the evaluation of energy gaps from first-principles, by a full treatment of the many-particle Schrödinger equation, and finally in Chapter 5 we present an extension of the work discussed in Chapter 3, and study the Wigner crystallisation transition of the electron gas in an n -doped two-dimensional semiconductor. Hartree atomic units ($4\pi\epsilon_0 = \hbar = e = m_e = 1$) will be assumed throughout, unless otherwise stated.

Chapter 1

Theoretical underpinnings of quantum Monte Carlo

1.1 Two-dimensional semiconductors

Throughout this thesis, our focus will be on two-dimensional semiconductors. It is instructive, given the size and scope of the literature in this general field, to narrow our scope appropriately.

Firstly, let us define what is meant by “two-dimensional”. We live in a universe of three (or possibly more) spatial dimensions, and as such it is impossible to realise a truly two-dimensional object. The closest such system we can hope to realise would be a single layer of atoms. Such systems do exist, notably graphene [16], but nevertheless each atomic nucleus, and each electron, remains capable of motion in more than two directions. The system as a whole cannot be described in terms of two-dimensional position vectors. However, the degree to which the motion of electrons, for instance, move in one of the three spatial directions in graphene is severely limited. There is a strong anisotropy to the properties of graphene: conduction in the plane of confinement is possible, but conduction perpendicular to that plane is not. Graphene is a two-dimensional electronic system.

A more general, succinct, statement would be as follows: a two-dimensional

electronic system is one in which electrons are free to move in two directions, but tightly confined in the third direction.

Confinement in this way acts to create a kind of one-dimensional particle in a box problem, leading to a discrete spectrum of energy levels which, throughout this thesis, we will assume to be much higher in energy than those accessed by the two-dimensional dynamics of the system. *I.e.* our two-dimensional systems will be tightly confined three-dimensional systems, in the lowest confinement subband.

Secondly, let us give an overview of semiconductors. A semiconductor is an electrical insulator with a small energy gap – such that carriers can be thermally excited across the gap at reasonable temperatures (hundreds of Kelvin, say).

Whilst thermal excitation is often invoked in the definition of a semiconductor, a more common means of excitation is optical. When a photon of light impinges on a semiconductor, if it possesses enough energy, it may *excite* an electron from the lower energy “valence” band to the higher energy (normally unoccupied) conduction band. This process of photoexcitation occurs *resonantly* if the incoming light has enough energy to excite an electron directly from the valence to the conduction band, but non-resonant so-called “excitonic” absorption is also possible. Because of interaction effects, the energy of the system after excitation is slightly lower than the energy gap, and the formation of “excitons” leads to a weak, characteristic, photoabsorption signature in semiconductors which is *below* resonant photoabsorption in terms of incoming photon energy. An “exciton” is the name given to the bound state formed from an excited electron and the remnant (net positively charged) manifold of valence states it leaves behind when it is excited. Excitons relax by the process of “recombination”, where the energy required for their creation is once again released. If momentum and energy can be conserved by doing so, the recombination will be “radiative”, and the energy will be released directly as a photon. If momentum and energy cannot be conserved, then the exciton must relax by a so-called “non-radiative” process. Such processes often involve the excitation of lattice vibrations (phonons) or other scattering events

(with lattice defects, for example).

The formation of excitons and related bound states are the excited state phenomena which we will be interested in throughout much of this thesis.

1.2 The quantum many-body problem

The “quantum many-body problem” is the generic name for any physical problem which concerns the microscopic description of a large number of interacting particles. “Microscopic” here means that we are interested in a complete description of each and every particle making up the system, which in turn implies a quantum mechanical treatment of the many-body system at hand is necessary. For the vast majority of physical problems encountered in daily life, the system will be a collection of electrons and nuclei (formed from protons and neutrons, which are themselves formed from first-generation quarks), described by a many-particle version of the Schrödinger equation, laid out in detail in Section 1.4.

The repetitive, complicating interactions experienced by particles in a quantum many-body system, make analytical solution of the quantum many-body problem wholly unlikely. Dirac himself noted this in 1929 [15] with a now famous quote on the “underlying physical laws” (which we gave in full in the Introduction) and their solutions as pertaining to physics and chemistry. Later, Van Vleck [17] and Kohn [18] added their own quantitative analyses. These are now known as the “van Vleck catastrophe”; the exponential rise in the number of parameters necessary to achieve a given level of accuracy if the many-body wave function is represented on a grid.

In this chapter, and throughout the remainder of this thesis, we will describe and use methods for the solution of the quantum many-body problem which are *not* reliant on any “illegitimacy” of the many-body wave function. Quantum Monte Carlo methods sidestep direct solution of the many-body Schrödinger equation, instead relying on stochastic sampling, and can (particularly in diffusion Monte Carlo) elegantly avoid the van Vleck catastrophe.

1.3 Statistics

1.3.1 Monte Carlo integration

Consider the mean (\bar{f}) and variance (σ_f^2) of some possibly multi-dimensional function $f(x)$, whose domain is the real numbers, weighted by a probability distribution function $P(x)$

$$\bar{f} = \int dx f(x)P(x), \quad (1.1)$$

$$\sigma_f^2 = \int dx [f(x) - \bar{f}]^2 P(x). \quad (1.2)$$

The mean and variance may be approximated by drawing N independent random variables X_i from the distribution $P(x)$, and forming

$$\bar{f}_N = \frac{1}{N} \sum_{i=1}^N f(X_i), \quad (1.3)$$

$$\sigma_{f,N}^2 = \frac{1}{N-1} \sum_{i=1}^N [f(X_i) - \bar{f}_N]^2, \quad (1.4)$$

where we have used Bessel's correction to the variance estimate $\sigma_{f,N}^2$. By the central limit theorem, \bar{f}_N is distributed as a Gaussian with mean \bar{f} , and a variance σ_f^2/N at large N . Hence in the limit that $N \rightarrow \infty$, $\bar{f}_N = \bar{f}$. At finite N , we can take the standard error on the mean (or, the standard deviation of the underlying Gaussian distribution of \bar{f}_N), $\bar{\sigma}_f = \sigma_f/\sqrt{N} \simeq \sigma_{f,N}/\sqrt{N}$ to quantify the error in \bar{f}_N , and write

$$\bar{f} = \int dx f(x)P(x) = \frac{1}{N} \sum_{i=1}^N f(X_i) + \mathcal{O}(\bar{\sigma}_f). \quad (1.5)$$

Monte Carlo integration is the use of Eqn. (1.5) to evaluate integrals written in the form of Eqn. (1.1). Monte Carlo integration differs from conventional deterministic and grid-based methods of integration – the error estimate in Monte Carlo integration is $\mathcal{O}(N^{-\frac{1}{2}})$, irrespective of the dimension of x . In d -dimensional grid-based methods, for example, the error estimate is often $\mathcal{O}(n^{-p})$, with n the number of

grid points in each dimension, and p a rule-specific positive integer ($p = 2$ for the trapezium rule, $p = 4$ for Simpson’s rule). Such a method requires $N = n^d$ function evaluations so, in terms of N , the error scales like $\mathcal{O}(N^{-p/d})$. For any given multidimensional grid-based integration rule, then, there is some d beyond which Monte Carlo integration scales better with the total number of function evaluations N . Taking the example of an integral over the position vectors of a series of particles in three-dimensional space, Monte Carlo integration is more efficient than Simpson’s (grid-based) rule when $3 \times (\text{the number of particles}) > 8$ (*i.e.* for three particles or more).

1.3.2 Markov chains

A discrete stochastic process can be represented by a sequence of states $\{X_0, X_1, \dots, X_N\}$, with $N > 0$, which are realised by an entity referred to as a “walker”. The random variable X_t is a member of the *state space* of the process at a *time* t . The probability that a walker realises any given series of states $\{x_0, x_1, \dots, x_N\}$ is given by

$$\mathcal{P}\{X_0 = x_0, X_1 = x_1, \dots, X_N = x_N\}, \quad (1.6)$$

where the $\{x_i\}$ are realisations of the (*continuous*) state space. If the state space is the space of (N) particle coordinates, then the $\{x_i\}$ are vectors in \mathbb{R}^{3N} . If the state space is the principal quantum number of a hydrogen-like atom, then the $\{x_i\}$ are the natural numbers, \mathbb{N} . Such a stochastic process is a *Markov chain* if

$$\mathcal{P}\{X_{N+1} = x | X_0 = x_0, X_1 = x_1, \dots, X_N = x_N\} = \mathcal{P}\{X_{N+1} = x | X_N = x_N\}, \quad (1.7)$$

which says that the conditional probability of the walker realising the state $X_{N+1} = x$ in the $(N + 1)^{\text{th}}$ iteration of the process, given a history, is independent of all past states $\{x_0, x_1, \dots, x_{N-1}\}$ defining that history, depending only on the current state of the system, x_N . Markov chains are *memoryless*.

One can define *transition probabilities* for the Markov chain as

$$p_{xx'} = \mathcal{P}\{X_{N+1} = x' | X_N = x\} = \mathcal{P}(x' \leftarrow x), \quad (1.8)$$

and clearly these probabilities must satisfy

$$\begin{aligned} \sum_{x'} p_{x,x'} &= 1 \quad \forall x, \\ p_{xx'} &\geq 0 \quad \forall (x, x'). \end{aligned} \quad (1.9)$$

Let the probability of a Markov process being in state j at time n be denoted \mathcal{P}_j^n . After k further steps of the Markov chain, where intermediate states of the system are labelled by α , and by definition of the transition probabilities, the probability that the system is in state l is given by

$$\mathcal{P}_l^{n+k} = \sum_{j, \alpha_2, \dots, \alpha_{k-1}} p_{l, \alpha_{k-1}} p_{\alpha_{k-1}, \alpha_{k-2}} \cdots p_{\alpha_3, \alpha_2} p_{\alpha_2, x_j} \mathcal{P}_j^n, \quad (1.10)$$

or, in vector notation, where $\mathcal{P}^n = (\mathcal{P}_1, \dots, \mathcal{P}_{N_s})$ is a vector containing the state probabilities after n steps,¹ and p is a matrix of transition probabilities,

$$\mathcal{P}^{n+k} = p^k \mathcal{P}^n. \quad (1.11)$$

We define the stationary probability distribution, $\tilde{\mathcal{P}}$, of the Markov process as follows

$$\tilde{\mathcal{P}}^{n+1} = p \tilde{\mathcal{P}}^n = \tilde{\mathcal{P}}^n = \tilde{\mathcal{P}}^n(x). \quad (1.12)$$

Once a Markov process reaches a stationary distribution, it must therefore stay there. Successive realisations of the state of a stationary Markov system, x , are then identically distributed according to that stationary distribution. The stationary distribution is unique if the corresponding Markov process is *ergodic*, as per

¹Note that N_s need not be finite.

the Perron-Frobenius theorem [19, 20].² Ergodicity is guaranteed by the following conditions

- All elements of $p_{xx'}^n$ are finite, for some $n < \infty$. This ensures that all states of the system are accessible in a finite number of steps, and that there are no “trapping regions” for the process.³
- The average return time to any state is finite. Always true in cases of a finite state space, but not necessarily true in an infinite state space.
- The matrix $p_{xx'}$ is non-periodic: $p_{xx} > 0 \forall x$.⁴

Detailed balance

Transition probabilities often satisfy, or are designed to satisfy, the detailed balance property. In words, detailed balance means that the probability flux into any given state x' (from any other state) equals the probability flux out of x' (and into any other state). In symbols,

$$\tilde{\mathcal{P}}(x)\mathcal{P}(x' \leftarrow x) = \tilde{\mathcal{P}}(x')\mathcal{P}(x \leftarrow x') \forall (x, x'). \quad (1.13)$$

If the stationary distribution and the transition probabilities satisfy detailed balance, and the Markov process is itself ergodic, then the stationary distribution is unique. Detailed balance is a stronger condition on uniqueness of the stationary distribution than ergodicity; it is possible for a random walk to be ergodic, and to

²In the context of stochastic matrices, the Perron-Frobenius eigenvalue is that associated with the conservation of probability (it is unity, and its associated eigenvector is a row of ones). The implication of the Perron-Frobenius theorem that all other eigenvalues are ≤ 1 constitutes the uniqueness of the stationary distribution.

³As a concrete example, consider a Markov process on a two state system, with a transition matrix $\begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$. The process either endlessly samples the first state (if it starts there), or makes one transition to the first state from the second (if it starts there), and then endlessly samples the first state. There is no probability of transition out of the first state - it is a “trapping region”.

⁴This is similar to trapping. One can imagine scenarios where a Markov process endlessly loops through two or more states, so-called “ n -cycles”, which are effectively trapping regions for the subset of states involved in the n -cycle. As an example, consider a three-state Markov process where the only non-zero transition probabilities are from states $1 \rightarrow 2$, $2 \rightarrow 3$, and $3 \rightarrow 1$. This system, regardless of the initial distribution, three-cycles indefinitely.

have a unique stationary distribution, all *without* satisfying detailed balance. The requirement of detailed balance is often a simple way of ensuring uniqueness of the stationary distribution in practical developments of computer algorithms. A simple way of guaranteeing detailed balance is to work with transition probabilities which are symmetric: $\mathcal{P}(x \leftarrow x') = \mathcal{P}(x' \leftarrow x)$.

1.3.3 The Metropolis-Hastings algorithm

The Markov chain is an incredibly useful mathematical concept. Perhaps the most useful of its applications is that concerned with sampling probability distributions. The need for such sampling has already been made clear in the context of Monte Carlo integration. For *invertible* probability distributions $P_{\text{inv}}(x)$, with x a realisation of the random variable X , we may generate realisations of X which are distributed according to $P_{\text{inv}}(x)$ by feeding the inverse of P_{inv} uniformly distributed random variates u

$$x = P_{\text{inv}}^{-1}(u), \quad (1.14)$$

however, very few interesting probability distribution functions are invertible. Alternative means of efficiently sampling non-invertible probability distribution functions are therefore desirable.

One method for sampling arbitrary non-negative probability distributions is the so-called Metropolis-Hastings algorithm [21, 22].⁵ Here, one need only know *a priori* the value of a function which is *proportional* to the target probability distribution in order to generate samples of said distribution. Such a feature does not immediately sound tremendously useful, but in high-dimensional state spaces, the computation of a normalisation factor for a probability distribution is in itself a non-trivial task. The avoidance of such a difficulty makes the usefulness of the proposed algorithm plain to see.

Suppose that we wish to generate realisations x of a random variable X which

⁵We acknowledge the contributions of *all* of the authors of the referenced article to the development and popularisation of this algorithm, but have chosen to keep it named as such owing to its ubiquity.

are distributed according to a general (invertible or non-invertible, in arbitrary dimension) probability distribution function $P(x)$. Let $f(x)$ be a function such that $f(x) \propto P(x)$. We aim to construct, and generate samples of, a Markov process whose equilibrium distribution is the target distribution P . If we require that this Markov process satisfies detailed balance, then we have

$$P(x)\mathcal{P}(x' \leftarrow x) = P(x')\mathcal{P}(x \leftarrow x'), \quad (1.15)$$

or equivalently

$$\frac{P(x')}{P(x)} = \frac{\mathcal{P}(x' \leftarrow x)}{\mathcal{P}(x \leftarrow x')} = \frac{f(x')}{f(x)}. \quad (1.16)$$

Now, breaking the transition probability \mathcal{P} down into a proposal distribution \mathcal{T} times an acceptance distribution \mathcal{A} , we have

$$\frac{f(x')}{f(x)} = \frac{\mathcal{T}(x' \leftarrow x)\mathcal{A}(x' \leftarrow x)}{\mathcal{T}(x \leftarrow x')\mathcal{A}(x \leftarrow x')}. \quad (1.17)$$

The proposal distribution $\mathcal{T}(x' \leftarrow x)$ gives the probability that a move to x' will be proposed, from a current state of x . The acceptance distribution $\mathcal{A}(x' \leftarrow x)$ gives the probability that such a move will be accepted. If a move is not accepted, then the state of the process does not change (equivalently, x is re-sampled). We now require an acceptance distribution that fulfils Eqn. (1.17). The choice of Metropolis *et al.* does, and is given by

$$\mathcal{A}(x' \leftarrow x) = \min \left(1, \underbrace{\frac{f(x')\mathcal{T}(x \leftarrow x')}{f(x)\mathcal{T}(x' \leftarrow x)}}_{\text{AR}} \right). \quad (1.18)$$

Moves are accepted automatically if the so-called *acceptance ratio* (AR) is greater than 1, but accepted with a finite probability if this is not true. If the AR is close to one, then a move is likely. If it is close to zero, a re-sampling of the previous state of the system is likely. A simple proof that Eqn. (1.18) satisfies detailed

balance proceeds as follows. Write

$$\begin{aligned}
\mathcal{P}(x' \leftarrow x) &= \mathcal{T}(x' \leftarrow x)\mathcal{A}(x' \leftarrow x) \\
&= \mathcal{T}(x' \leftarrow x)\min\left(1, \frac{P(x')\mathcal{T}(x \leftarrow x')}{P(x)\mathcal{T}(x' \leftarrow x)}\right) \\
&= \frac{1}{P(x)}\min(P(x)\mathcal{T}(x' \leftarrow x), P(x')\mathcal{T}(x \leftarrow x')). \quad (1.19)
\end{aligned}$$

Then

$$P(x)\mathcal{P}(x' \leftarrow x) = \min(P(x)\mathcal{T}(x' \leftarrow x), P(x')\mathcal{T}(x \leftarrow x')), \quad (1.20)$$

but the right hand side is now invariant under the replacement $x \leftrightarrow x'$, and hence so must the left hand side be invariant. Therefore

$$P(x)\mathcal{P}(x' \leftarrow x) = P(x')\mathcal{P}(x \leftarrow x'), \quad (1.21)$$

which is a recurrence of the detailed balance condition (Eqn. (1.13)), and the proof is therefore complete.

The Metropolis-Hastings algorithm exploits the fact that the accept-reject step defined by Eqn. (1.18) forces the detailed balance condition to obtain a sequence of random numbers distributed according to the target distribution $P(x)$. The algorithm proceeds as follows:

0. Seed an initial point (randomly, or deterministically), x_0 , chosen to be the first sample. Select a proposal distribution, $\mathcal{T}(x' \leftarrow x)$. This is normally a Gaussian distribution of the desired dimension, centred on x , with some given width (see later section on autocorrelation).

1. Until a termination criterion:

1. Generate a candidate move, x' , by sampling the proposal distribution, $\mathcal{T}(x' \leftarrow x)$.⁶

⁶Although not quite invertible, the Gaussian distribution may be sampled efficiently by an alternate means: the Box-Muller transformation.

2. Calculate the acceptance probability $\mathcal{A}(x' \leftarrow x)$ (from Eqn. (1.18)).
3. Generate a uniform random variate v on the interval $[0, 1]$, if:
 - $v \leq \mathcal{A}(x' \leftarrow x)$: Accept the move, next sample is x' .
 - $v > \mathcal{A}(x' \leftarrow x)$: Reject the move, next sample is (again) x .
4. Check termination criterion. If not terminated, increment the time step.

1.3.4 Metropolis-Hastings caveats

The Metropolis-Hastings algorithm has a pair of particularly important features.

1. Serial correlation: Sequential samples are *correlated* with one another. Because the proposal distribution $\mathcal{T}(x' \leftarrow x)$ typically has a characteristic width,⁷ samples taken one after another are likely to be closer together than would be suggested from an unbiased sampling from the target distribution. This leads to an eventual underestimate of the variance of computed quantities obtained from Metropolis-Hastings sampling.
2. Equilibration: The equilibrium distribution of the Markov chain defining the Metropolis-Hastings algorithm is the target distribution, but only after a finite number of “equilibration” steps.

Both of these issues require attention.

The former issue, serial correlation, is rather benign. There are two means by which it can be overcome. The first is the implementation of so-called “decorrelation loops”. In such a loop, new moves are proposed and accepted, but samples are not taken. A decorrelation loop of finite length therefore acts to dilute the impact of serial correlation. The second is the re-blocking procedure [23]. Assuming that decorrelation loops are not used, or that they have been ineffective in completely nullifying the effects of serial correlation, from a set of S serially correlated data samples, only S/t_{corr} are independent, with t_{corr} the mean correlation time. The

⁷In all quantum Monte Carlo calculations presented in this thesis, the proposal distribution *will* be a Gaussian with a characteristic width.

mean correlation time is the average number of steps, or amount of time taken, until two walkers are independent of each other. Unbiased estimates of the mean, variance, and standard error of a quantity calculated with S serially-correlated Metropolis-Hastings samples may be obtained from the re-blocking procedure. Here we will consider a VMC calculation.⁸ In this case data are grouped into successive blocks and averaged. In the first step of the procedure, single data points are collected into pairs and averaged, with new estimates of the mean, variance and standard error formed. In the n^{th} step, sets of 2^n data points are collected and averaged. When the block lengths are large enough so that they contain at least a mean correlation time, the block averages are uncorrelated, and by definition further repetitions of the re-blocking procedure should return the same mean energy, and the same error in the mean. In practice, the formation of a “plateau” in the standard error of the mean, where successive re-blocking transformations do not change the standard error in the mean, is used to signal the termination of the re-blocking procedure [24]. Repeated re-blocking transformations eventually lead to significant problems in their own right. Ultimately one ends up dealing with effectively tiny data sets formed from averages over large amounts of serially correlated data.

Serial correlation can be present over multiple scales in any given sampling run. The effects of serial correlation should be removed for the longest correlation period present in a given problem. This is usually determined by the smallest physical length scale. Additionally, whilst we will not give details of this approach, one may calculate the mean autocorrelation time for a data set, and perform re-blocking based solely on this period. For a large enough data set, the autocorrelation time method and the reblocking method tend to yield uncertainty estimates which agree within error bars. Both are calculated in the `CASINO` code, and all results in this thesis will be quoted for cases where reblocking has converged successfully and where the correlation time and reblocking uncertainty estimates are in reasonable

⁸The extension to DMC is simple, but that case is only slightly complicated by the fact that different steps are weighted differently in the DMC algorithm.

agreement.

The latter issue, equilibration, is a necessary part of any QMC calculation. One can estimate the necessary length of the equilibration stage of a calculation by considering the simulation to constitute a random walk. In this random walk, any single electron should have enough time to adequately explore the phase space associated with the calculation. Assuming a single electron diffuses randomly throughout d -dimensional space over a period of N (imaginary) time steps each of length τ (see Section 1.5), the root-mean-square diffusion distance is given by

$$\langle x \rangle = \sqrt{\alpha d N \tau}, \quad (1.22)$$

with α the move acceptance ratio (kept close to $\frac{1}{2}$ in VMC; 1 in DMC). If N is large enough that $\langle x \rangle \gtrsim L$, with L the longest physical length-scale in a problem (the spatial extent of a simulation supercell in a calculation on a periodic solid, for example), then N is our estimate for the number of necessary simulation equilibration steps for simulations on that system.

1.4 Variational Monte Carlo (VMC)

Quantum mechanics is the physical theory describing the smallest constituents of matter. Ultimately, all physical laws or rules find their roots in quantum mechanics, and condensed matter physics is no different. In contrast to classical physics, quantum mechanics is a probabilistic theory: outcomes are not generally guaranteed, and quantum mechanical predictions or trajectories are only ever possible *with some probability*.

The dynamics of quantum mechanical systems are encoded in the Schrödinger equation, whose solution is the so-called “wave function” of the system at hand. Probabilistic outcomes of quantum mechanical experiments are encoded in this wave function, and follow (with numerous possible interpretations) from the Born rule [25].

Variational Monte Carlo is a means by which one may form approximate solutions to the Schrödinger equation, and subsequently extract certain properties of the (possibly large, possibly strongly interacting) system at hand.

1.4.1 The many-particle Schrödinger equation

A general non-relativistic many-particle quantum system may be described in terms of the position coordinates of all of the (N_e) electrons and (N_n , with the I^{th} nucleus having charge Z_I) nuclei that comprise it. Let $\mathbf{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_{N_e}\}$ denote a “configuration” of electron positions, and $\mathbf{N} = \{\mathbf{n}_1, \dots, \mathbf{n}_I, \dots, \mathbf{n}_{N_n}\}$ denote a “configuration” of nuclear positions. The many-particle Schrödinger equation describing such a system reads

$$\hat{\mathcal{H}}_{\text{full}}\Phi_n(\mathbf{R}, \mathbf{N}, t) = i\frac{\partial}{\partial t}\Phi_n(\mathbf{R}, \mathbf{N}, t) = E_n\Phi_n(\mathbf{R}, \mathbf{N}), \quad (1.23)$$

where Φ_n and E_n are the eigenfunctions and energy eigenvalues of the (here time-independent) operator $\hat{\mathcal{H}}_{\text{full}}$, which is given by

$$\begin{aligned} \hat{\mathcal{H}}_{\text{full}} = & -\frac{1}{2}\sum_{i=1}^{N_e}\nabla_{\mathbf{r}_i}^2 - \frac{1}{2}\sum_{I=1}^{N_n}\nabla_{\mathbf{n}_I}^2 \\ & + \sum_{\langle i,j \rangle} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_{\langle I,J \rangle} \frac{Z_I Z_J}{|\mathbf{n}_I - \mathbf{n}_J|} \\ & - \sum_{i=1}^{N_e} \sum_{I=1}^{N_n} \frac{Z_I}{|\mathbf{r}_i - \mathbf{n}_I|}, \end{aligned} \quad (1.24)$$

The first line of Eqn. (1.24) contains electron and nuclear kinetic energies, the second line electron-electron and nucleus-nucleus interactions ($\langle i, j \rangle$ denotes a single-counted sum over pairs of particles i and j), and the final line contains the electron-nucleus interactions. The eigenvectors $\Phi_n(\mathbf{R}, \mathbf{N})$ and eigenvalues E_n correspond to a proper quantum mechanical treatment of electrons, nuclei, and all of their associated interactions. In practice, however, it is common to make use of the Born-Oppenheimer (BO) approximation [26, 27]. In the BO approximation, the

fact that nuclear masses far exceed electron masses is used to argue that, in the adiabatic approximation, nuclei move on much smaller time scales than electrons do (*i.e.* at lower velocities).⁹ Under these assumptions, the relevant independent variables in Eqn. (1.23) are simply the electron coordinates, collected together in \mathbf{R} . The problem decouples to one in which the nuclear coordinates, \mathbf{N} , play only a parametric role

$$\hat{\mathcal{H}}\Phi_n(\mathbf{R}) = E_n(\mathbf{N})\Phi_n(\mathbf{R}), \quad (1.25)$$

with Hamiltonian

$$\begin{aligned} \hat{\mathcal{H}} = & -\frac{1}{2} \sum_{i=1}^{N_e} \nabla_{\mathbf{r}_i}^2 + c_1 \\ & + \sum_{\langle i,j \rangle} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + c_2 \\ & - \sum_{i=1}^{N_e} \sum_{I=1}^{N_n} \frac{Z_I}{|\mathbf{r}_i - \mathbf{n}_I|}, \end{aligned} \quad (1.26)$$

which is identical to Eqn. (1.24), save for the fact that the nuclear kinetic energy and nucleus-nucleus interaction terms are now constants, c_1 ($=0$) and c_2 , respectively. We note that it is common for modern electronic structure codes to check compatibility by comparing values of c_2 , the nucleus-nucleus potential energy.¹⁰ If c_2 is explicitly included in $\hat{\mathcal{H}}$, then that Hamiltonian gives the BO potential for the nuclear configuration \mathbf{N} . The remainder of this thesis will deal explicitly with the eigenstates and eigenvalues of equations of the form in Eqn. (1.26), and will also assume the following shorthands for terms in Eqn. (1.26)

$$-\frac{1}{2} \sum_{i=1}^{N_e} \nabla_{\mathbf{r}_i}^2 \rightarrow -\frac{1}{2} \nabla_{\mathbf{R}}^2 \text{ or } \hat{T}, \quad (1.27)$$

$$\sum_{\langle i,j \rangle} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{N_e} \sum_{I=1}^{N_n} \frac{Z_I}{|\mathbf{r}_i - \mathbf{n}_I|} + c_2 \rightarrow U(\mathbf{R}) \text{ or } \hat{U}. \quad (1.28)$$

⁹The mass argument is clearly always true, but the subsequent statement regarding time scales does not always follow.

¹⁰The CASINO code, for example, often performs comparative checks of the electron-nucleus potential energy in calculations where the Slater part of a trial wave function has originated from another electronic structure code.

Additionally, we will henceforth let $N_e \rightarrow N$.

1.4.2 VMC expectation values

In VMC, samples of so-called *trial* many-electron wave functions are generated by the Metropolis-Hastings algorithm in order to evaluate their corresponding expectation values. Suppose we wish to consider taking the expectation value of an operator (\hat{A}) with respect to a trial many-electron wave function $\Psi_T(\{\mathbf{r}_i, \sigma_i\})$, where the $\{\sigma_i\}$ denote the set of spin projections of the electrons ($i = 1, \dots, N$). If, as will always be the case in this thesis, \hat{A} is spin-independent, we can write

$$A = \langle \hat{A} \rangle = \frac{\langle \Psi_T | \hat{A} | \Psi_T \rangle}{\langle \Psi_T | \Psi_T \rangle} = \frac{\sum_{\{\sigma_i\}} \int d\mathbf{R} \Psi_T^*(\{\mathbf{r}_i, \sigma_i\}) \hat{A}(\mathbf{R}) \Psi_T(\{\mathbf{r}_i, \sigma_i\})}{\sum_{\{\sigma_i\}} \int d\mathbf{R} |\Psi_T(\{\mathbf{r}_i, \sigma_i\})|^2}. \quad (1.29)$$

If we further restrict ourselves to the case where $\Psi_T(\{\mathbf{r}_i, \sigma_i\})$ is an eigenstate of $\hat{S}_z = \sum_{i=1}^N \hat{s}_{z_i}$,¹¹ we are free to make a choice of the ordering of the spin variables. *I.e.* we can re-label the $\{\mathbf{r}_i\}$ such that

$$\begin{aligned} \Psi_T(\{\mathbf{r}_i, \sigma_i\}) &= \Psi_T((\mathbf{r}_1, \uparrow), \dots, (\mathbf{r}_{N_\uparrow}, \uparrow), (\mathbf{r}_{N_\uparrow+1}, \downarrow), \dots, (\mathbf{r}_N, \downarrow)), \\ &\equiv \Psi_T(\mathbf{r}_1, \dots, \mathbf{r}_N), \\ &= \Psi_T(\mathbf{R}), \end{aligned} \quad (1.30)$$

without changing $\langle \hat{A} \rangle$. Thus we no longer need to keep track of individual spin quantum numbers. This ordering convention will be assumed in all further discussions of many-electron trial functions, and we will only ever discuss collinear spin states in this thesis. For spin-independent operators, the spin sums in Eqn. (1.29) are then over N identical terms, and this factor cancels. The expectation value

¹¹Having eigenvalue $\frac{1}{2}(N_\uparrow - N_\downarrow)$, with N_\uparrow the number of up-spin electrons, and $N = N_\uparrow + N_\downarrow$.

then collapses to

$$A = \frac{\int d\mathbf{R} |\Psi_T(\mathbf{R})|^2 A_L(\mathbf{R})}{\int d\mathbf{R} |\Psi_T(\mathbf{R})|^2}, \quad (1.31)$$

where the *local* quantity A_L has been defined as

$$A_L(\mathbf{R}) = \frac{\langle \mathbf{R} | \hat{A} | \Psi_T \rangle}{\langle \mathbf{R} | \Psi_T \rangle}. \quad (1.32)$$

It is then clear that Eqn. (1.31) and Eqn. (1.32) together establish VMC expectation values of quantum mechanical operators \hat{A} as averages over the associated *local* functions $A_L(\mathbf{R})$, taken over configurations \mathbf{R} which are distributed according to the probability distribution function ($|\Psi_T(\mathbf{R})|^2 / \int d\mathbf{R} |\Psi_T(\mathbf{R})|^2$). The preferential sampling of configurations from this distribution, rather than from the uniform distribution, is known as *importance sampling*, as sampling is biased to regions of higher importance (physically, in this example these are regions where the electron probability density is large).

An important set of expectation values in QMC are those associated with the Hamiltonian operator, $\hat{\mathcal{H}}$. If $|\Phi_0\rangle$ is the exact ground state, and E_0 is its associated eigenvalue, then the Rayleigh-Ritz variational principle reads

$$E_0 = \frac{\langle \Phi_0 | \hat{\mathcal{H}} | \Phi_0 \rangle}{\langle \Phi_0 | \Phi_0 \rangle} \leq \frac{\langle \Psi_T | \hat{\mathcal{H}} | \Psi_T \rangle}{\langle \Psi_T | \Psi_T \rangle} = E, \quad (1.33)$$

with equality if and only if $\Psi_T = \Phi_0$. The expectation value of the energy is the average of the local energy, $E_L(\mathbf{R})$. A second important expectation value is the variance of the local energy

$$\sigma^2 = \frac{\int d\mathbf{R} |\Psi_T(\mathbf{R})|^2 [E_L(\mathbf{R}) - E]^2}{\int d\mathbf{R} |\Psi_T(\mathbf{R})|^2}. \quad (1.34)$$

Note that *all* eigenstates of $\hat{\mathcal{H}}$ have $\sigma^2 = 0$, but that only the ground state (Φ_0) simultaneously has the minimal energy (E_0). Further note that the variance esti-

mate presented here is the “unweighted” variance.

At the start of a VMC calculation, a series of configurations are generated by (initially) randomly locating electrons, but with weight centred on atomic positions, in proportion to atomic number. From this starting configuration, a series of Metropolis steps are taken, with moves chosen from a symmetric multi-dimensional Gaussian distribution ($\mathcal{T}(\mathbf{R} \leftarrow \mathbf{R}') = \mathcal{T}(\mathbf{R}' \leftarrow \mathbf{R})$) centred on the current configuration, and of width determined by a parameter τ_{VMC} (the VMC time step). Moves are therefore accepted with probability

$$P(\mathbf{R}' \leftarrow \mathbf{R}) = \min \left(1, \frac{|\Psi_{\text{T}}(\mathbf{R}')|^2}{|\Psi_{\text{T}}(\mathbf{R})|^2} \right), \quad (1.35)$$

with $\Psi_{\text{T}}(\mathbf{R})$ the trial wave function. Throughout the course of a VMC calculation, τ_{VMC} is adjusted to enforce that approximately half of all proposed moves are accepted (which lessens the effects of serial correlation) [21]. After an equilibration criterion is met (see Eqn. (1.22)), the sequence of sampled configurations $\{\mathbf{R}\}$ defines a Markov chain at equilibrium, and from this point on we can begin to accumulate local energies for averaging. Local energy accumulation need not happen at every configuration move (decorrelation loops may be present).

The parameter τ_{VMC} acts to amplify or dampen the effects of serial correlation between moves. If it is too small, the move distribution is too narrow, and as a consequence subsequent configuration samples are too close (more correlated). If it is too large, the move distribution is too broad, and as a consequence most moves are rejected, and existing samples are routinely re-sampled. This, again, leads to a series of correlated (indeed, often duplicated, or maximally correlated) configurations.

1.4.3 Explicitly correlated trial wave functions

Both VMC and (see 1.5) DMC take as their basic inputs a so-called trial wave function. In the case of VMC, energies and other expectation values are entirely

determined by this function.

The trial wave function is a quantity which we need to be able to evaluate repeatedly, and so we would like it to be compact and easy to represent. We would also like to avoid large multideterminant expansions (which are common to post-Hartree-Fock methods in quantum chemistry, for example) as a means of describing correlation effects. Such determinant expansions, whilst offering significant further variational freedom at a given system size, are a non-compact means of retrieving correlation energy. The number of important multideterminantal configurations needed to achieve a certain level of accuracy grows very rapidly with system size, making the use of multideterminant expansions beyond the realm of small molecules very computationally demanding. For these reasons, by far the most common choice of trial wave function in modern continuum quantum Monte Carlo calculations is the so-called Slater-Jastrow form

$$\Psi_{\text{SJ}}(\mathbf{R}) = \exp [J(\mathbf{R})] \times \mathcal{D}(\mathbf{R}), \quad (1.36)$$

where $J(\mathbf{R})$ is the Jastrow exponent ($\exp [J(\mathbf{R})]$ the ‘‘Jastrow factor’’), and $\mathcal{D}(\mathbf{R})$ is an antisymmetric product of spin-up and spin-down single-particle orbitals usually taken from a Hartree-Fock (HF) or density functional theory (DFT) calculation. In cases without a spin dependent Hamiltonian, we can choose¹² to populate spin-up and spin-down determinants such that the first N_{\uparrow} electrons occupy the first (spin-up) determinant and the remaining N_{\downarrow} (from $N_{\uparrow} + 1$ to N) populate the second (spin-down) determinant *i.e.*

$$\mathcal{D}(\mathbf{R}) = \mathcal{D}_{\uparrow}(\mathbf{r}_1, \dots, \mathbf{r}_{N_{\uparrow}}) \times \mathcal{D}_{\downarrow}(\mathbf{r}_{N_{\uparrow}+1}, \dots, \mathbf{r}_N), \quad (1.37)$$

¹²As antisymmetry need only be maintained between pairs of like-spin electrons – unlike spins are distinguishable.

with

$$\mathcal{D}_\uparrow(\mathbf{r}_1, \dots, \mathbf{r}_{N_\uparrow}) = \begin{vmatrix} \phi_1(\mathbf{r}_1) & \phi_1(\mathbf{r}_2) & \phi_1(\mathbf{r}_3) & \cdots & \phi_1(\mathbf{r}_{N_\uparrow}) \\ \phi_2(\mathbf{r}_1) & \phi_2(\mathbf{r}_2) & \phi_2(\mathbf{r}_3) & \cdots & \phi_2(\mathbf{r}_{N_\uparrow}) \\ \phi_3(\mathbf{r}_1) & \phi_3(\mathbf{r}_2) & \phi_3(\mathbf{r}_3) & \cdots & \phi_3(\mathbf{r}_{N_\uparrow}) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \phi_{N_\uparrow}(\mathbf{r}_1) & \phi_{N_\uparrow}(\mathbf{r}_2) & \phi_{N_\uparrow}(\mathbf{r}_3) & \cdots & \phi_{N_\uparrow}(\mathbf{r}_{N_\uparrow}) \end{vmatrix} \quad (1.38)$$

and similarly for \mathcal{D}_\downarrow .

Determinants allow for the automatic satisfaction of the antisymmetry constraint (which follows from the Pauli exclusion principle). If the overall wave function is to inherit this exchange property of the determinant, however, we must also ensure that the Jastrow exponent J is a symmetric function under particle exchange, so that the Jastrow factor does not interfere with the exchange antisymmetry of the determinant part of the wave function.

In this thesis, we will usually only be concerned with closed-shell systems. Likewise, we will almost always populate our determinant with single-particle orbitals $\{\phi_i\}$ which have been taken from DFT calculations. For spin-degenerate cases we will often neglect to reference explicitly the spins of electrons involved in calculations. A case where we consider explicit spin-polarisation in the orbitals, and the formation of a multideterminant trial state, is presented in Section 4.3.2.2).

The HF energy is (by definition) the *minimal* total energy one can obtain for a state described by a single Slater determinant. A VMC evaluation of the energy of a state of the form in Eqn. (1.36), with single-particle orbitals taken from an underlying HF calculation, and with $J = 0$, would return the HF energy (up to an error bar). The utility of the explicitly correlated trial wave function in QMC is to be able to systematically improve on the HF energy by direct inclusion of particle-particle correlation factors in the trial wave function. By inclusion of the Jastrow factor, we can form a better description of the many-electron system, and our associated VMC or DMC energy will be *lower* than the HF energy. We can

define the “correlation energy” of an electronic system to be the difference between the HF energy and the true ground state energy (which is lower because of many-particle correlations). Although we can never know the true correlation energy, it is common in the field of QMC to abuse nomenclature, and call the difference between the HF energy and the best available DMC energy the “correlation energy”. If a trial wave function captures a large fraction of the correlation effects present in an electronic system, then the difference between the VMC energy obtained with that wave function and the best DMC energy will be small.

A particularly well-tested scheme for generating compact Jastrow factors which capture significant amounts of ground-state correlation energy is due to Drummond *et al.* [28]. The so-called Drummond-Towler-Needs Jastrow factor will be used exclusively in this thesis, though we note the existence of the more general scheme of López Ríos *et al.* [29], and of the recently developed “ ν ” Jastrow factor of Whitehead *et al.* [30]. In the Drummond-Towler-Needs prescription, the Jastrow exponent for a system of electrons and nuclei reads

$$J(\mathbf{R}) = \sum_{\langle i,j \rangle} u(r_{ij}) + \sum_I \sum_i^{N_n} \chi(r_{iI}) + \sum_{\langle i,j \rangle} \sum_I^{N_n} f(r_{ij}, r_{iI}, r_{jI}) + \sum_{\langle i,j \rangle} p(\mathbf{r}_{ij}), \quad (1.39)$$

which is a sum of electron-electron (u), electron-nucleus (χ), electron-electron-nucleus (f) and (periodic) electron-electron (p) correlation factors, whose expressions may be found in the original paper [28]. The vector $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ has magnitude r_{ij} and is an electron-electron separation. The terms r_I and r_{iI} are nucleus and electron-nucleus distances, respectively. Similar expansions may be used to study electron-hole systems, or electron-positron systems, or more general ensembles of fermions having varying types (*e.g.* in Chapter 3 we will consider electrons having additional “valley” degrees of freedom). The p -term is a plane-wave expansion in the electron-electron separation vectors \mathbf{r}_{ij} , and is only valid for periodic systems. In cases with periodicity, the remaining (u, χ, f) terms are appropriately truncated at the edge of the periodic cell. It is precisely this truncation which

necessitates the p -term. The u -term, for example, is truncated at a radius which is less than or equal to the radius of the largest sphere which may be inscribed within the Wigner-Seitz cell of the periodic system in question, and is rendered ineffective in the corners of the simulation cell. The p -term has variational freedom in these corners, and remedies this particular shortcoming of the standard u -term. This shortfall is precisely the issue which the aforementioned “ ν ” Jastrow term [30] also aims to address.

1.4.4 The Kato cusp conditions

An important property of the exact wave function which is hard to satisfy with a single Slater determinant (or by a sum of many determinants, for that matter) is its behaviour at particle-particle coalescence. For systems of charged particles interacting via the bare Coulomb potential, the potential energy of two particles diverges as the distance between them tends towards zero. The local energy of an exact eigenstate, however, is a constant. It is therefore the case that if we are to properly describe particle-particle coalescence, we must require that the kinetic energy features a cancelling singularity at coalescence. As this behaviour is difficult to describe with Slater determinants, it is often the case that the divergence in the kinetic energy is ensured by an appropriate choice of two-body term in the Jastrow exponent.

Consider, without loss of generality,¹³ a single two-electron system where the electrons interact via the Coulomb interaction. Suppose initially that the electrons have opposite spins, and that their separation is \mathbf{r} . A general two-particle wave function for these spins can be written as a sum over spherical harmonics $\mathcal{Y}_{\ell,m_\ell}(\theta, \phi)$

$$\psi(\mathbf{r}) = \sum_{\ell=0}^{\infty} \sum_{m_\ell=-\ell}^{\ell} \psi_{\ell,m_\ell}(r) \times \mathcal{Y}_{\ell,m_\ell}(\theta, \phi) \times r^\ell, \quad (1.40)$$

where the particular form of the functions ψ_{ℓ,m_ℓ} is to be determined. The contri-

¹³The same argument will apply generally to *every* such pairing of particles interacting via a divergent interaction potential in an N-electron system.

bution to the local energy E_L from such a pair is given by

$$E_L(\mathbf{r}) = \frac{-\nabla^2\psi}{\psi} + \frac{1}{r} = \frac{-\nabla^2\psi_{00}(r)}{\psi_{00}(r)} + \frac{1}{r} + \mathcal{O}(r). \quad (1.41)$$

At small values of r , $\psi_{00}(r) \approx \psi_{00}(0) + r \times \left[\frac{\partial\psi_{00}(r)}{\partial r} \right]_{r=0}$, and substitution into Eqn. (1.41) gives

$$E_L(\mathbf{r}) = \frac{-2}{r\psi_{00}(0)} \left[\frac{\partial\psi_{00}(r)}{\partial r} \right]_{r=0} + \frac{1}{r} + \mathcal{O}(r). \quad (1.42)$$

If, in the limit of coalescence ($r \rightarrow 0$), the divergence in the local energy from the Coulomb interaction is to be cancelled by that arising from the kinetic energy, then we require

$$\left[\frac{\partial\psi_{00}(r)}{\partial r} \right]_{r=0} = \frac{1}{2}\psi_{00}(0), \quad (1.43)$$

which is the (antiparallel) Kato cusp condition [31]. The parallel condition, for like-spin particles, is obtained by a similar derivation, but where the two-particle wave function is a sum over *odd* ℓ only, and where each of the $\ell = 1$ wave function components is accounted for separately. The result is similar, with the fraction $\frac{1}{4}$ replacing the $\frac{1}{2}$ in Eqn. (1.43).

Again, a similar analysis also applies to the electron-nucleus interaction (when bare nuclei are present in a calculation). In summary, the Kato cusp conditions for two-particle coalescences are each given by conditions on the Jastrow factor (by choice), and these conditions read

$$\left[\frac{\partial J}{\partial r_{ij}} \right]_{r_{ij}=0} = \frac{1}{2}, \text{ Antiparallel spin electrons } (i, j) \quad (1.44)$$

$$\left[\frac{\partial J}{\partial r_{ij}} \right]_{r_{ij}=0} = \frac{1}{4}, \text{ Parallel spin electrons } (i, j) \quad (1.45)$$

$$\left[\frac{\partial J}{\partial r_{iI}} \right]_{r_{iI}=0} = -Z_I, \text{ Nucleus } (I) \text{ and electron } (i). \quad (1.46)$$

and must be held true for all electron-electron pairs $\langle i, j \rangle$ and all electron-nucleus pairs $\langle i, I \rangle$. The first two conditions are typically enforced by parameter choices in the Drummond-Towler-Needs u -term. The third may be enforced by the orbitals

themselves, or by parameter choices in the χ term.

1.4.5 Non-standard Jastrow terms

Throughout this thesis, Jastrow terms will be required, for one reason or another, that are *not* present in the Drummond-Towler-Needs Jastrow factor. These will be presented when necessary, and used only for the calculations that they are described in the context of. Examples include the “ex2D” Jastrow term (necessitated by the difference of the Keldysh and Coulomb interactions at short range; utilised in Chapter 3, and in a periodic form in Chapter 5), and the Jastrow “H” term (necessitated by the attractive fermion-fermion interactions in large charge-complexes in two-dimensional semiconductors; and used in Chapters 2, 3 and 5).

1.4.6 Backflow transformations

A further means of explicitly describing correlation effects in the trial wave function is the so-called backflow transformation. Initially the brainchild of Feynman and Cohen in their study of the excitation properties of ^4He [32], backflow correlations are, in a somewhat rigorous sense, the “next tier up” from the Slater-Jastrow form in the hierarchy of natural trial wave function forms for continuum QMC [33]. In this thesis, we have exclusively used the parameterisation of the backflow function given by López Ríos *et al.* [29]. The backflow transformation amounts to the replacement $\mathbf{R} \rightarrow \mathbf{X}$ in the Slater part of a Slater-Jastrow trial wave function form, where $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ is a set of so-called “quasiparticle” coordinates, related to the true electron coordinates by the backflow function, $\boldsymbol{\xi}(\mathbf{R}) = \{\boldsymbol{\xi}_1(\mathbf{R}), \dots, \boldsymbol{\xi}_N(\mathbf{R})\}$. In equations,

$$\mathbf{R} \rightarrow \mathbf{X}(\mathbf{R}) = \mathbf{R} + \boldsymbol{\xi}(\mathbf{R}), \quad (1.47)$$

and, in the parameterisation of López Ríos *et al.*,

$$\boldsymbol{\xi}(\mathbf{R}) = \sum_{\langle i,j \rangle} \eta(r_{ij}) \mathbf{r}_{ij} + \sum_{I=1}^{N_n} \sum_{i=1}^N \mu(r_{iI}) \mathbf{r}_{iI} + \sum_{\langle i,j \rangle} \sum_I^{N_n} (\Phi(r_{iI}, r_{jI}, r_{ij}) \mathbf{r}_{ij} + \Theta(r_{iI}, r_{jI}, r_{ij}) \mathbf{r}_{iI}). \quad (1.48)$$

Here, analogously with the Drummond-Towler-Needs Jastrow exponent, the backflow function is a sum of electron-electron (η), electron-nucleus (μ), and electron-electron-nucleus (Φ, Θ) terms, whose expressions may be found in the original paper [29]. Such backflow terms are truncated in periodic calculations in the same way that the corresponding Jastrow terms of the same order (electron-electron, for example) are. Two-body backflow terms are normally smoothly truncated to zero in the presence of bare nuclei, so as not to interfere with the two-body electron-nucleus cusp condition. In the earlier referenced work of Holzmann *et al.* [33], the backflow functions generated were gradient fields ($\boldsymbol{\xi}_i(\mathbf{R}) = \nabla_{\mathbf{r}_i} Y(\mathbf{R})$, for some “backflow potential” $Y(\mathbf{R})$). There is no general need for this constraint, and it is a matter of empirical wisdom, and objective fact [29], that backflow functions that do not satisfy it often lead to more flexible variational trial wave functions. For this reason, all of the results generated with backflow transformations in this thesis have been obtained without applying such a constraint.¹⁴

1.4.7 Optimisation of many-electron wave functions

Trial wave functions Ψ_T often contain optimisable parameters. For example, the Drummond-Towler-Needs two-body u term reads

$$u(r_{ij}) = (r_{ij} - L_u)^C \Theta(L_u - r_{ij}) \times \left(\alpha_0 + \left[\frac{\Gamma_{ij}}{(-L_u)^C} + \frac{C\alpha_0}{L_u} \right] r_{ij} + \sum_{l=2}^{N_u} \alpha_l r_{ij}^l \right), \quad (1.49)$$

where C is the truncation order (controlling the smoothness of the derivatives of u), and L_u is a cut-off distance, which itself may be varied (and/or optimised).

Γ_{ij} is a constant chosen to satisfy the Kato cusp conditions for both parallel and

¹⁴Incidentally, this constraint is often satisfied in practice by forcing $\nabla_{\mathbf{r}_i} \times \boldsymbol{\xi}_i(\mathbf{R}) = \mathbf{0}$.

antiparallel spins (dependent on the particular i and j in question). The particular values of $\{L_u, \alpha_0, \dots, \alpha_{N_u}\}$ largely determine the expectation values of any observable taken with respect to Ψ_T (*i.e.* they completely determine observables in VMC). We may subvert this fact, and use obtained VMC observables themselves to determine the best set of optimisable parameters for our trial wave functions. Since we are normally interested in ground state properties, or those of low-lying excited states, a reasonable target state would be the ground state itself. “Optimising” a many-electron trial wave function is then the act of minimising some cost or objective function. Our objective functions should therefore be those which are minimal for the ground state. The exact ground state, by definition, is that of lowest energy. It is also an eigenstate of the many-electron Schrödinger equation, and hence has zero variance. Other measures of spread are also zero for the eigenstates, *e.g.* the mean absolute deviation (from the median; or the “MAD”) [34] (see diary at <https://vallico.net/casinoqmc/diary/> for discussion of the MAD and similar measures of spread as important in QMC calculations), or the recently proposed Ω -functional [35]. The Ω -functional is similar in practice to the variance, but it has recently been shown that, whilst state-specific (minimal for eigenstates), minimisation of the Ω -functional does not lead to a size-consistent description of electronic correlation in ground and excited states. We will neglect to discuss the Ω -functional in the remainder of this thesis. Instead, we will employ four methods for the optimisation of trial wave functions.

The first three methods are measure-of-spread minimisation methods. The computationally efficient scheme of Drummond *et al.* (which is only capable of optimising linear parameters in the Jastrow parameter, those which simply multiply other functions), relies on analytic properties of the local energy to find optimal parameter sets [36]. The less-efficient, but more general, scheme of Umrigar *et al.* instead relies on correlated sampling of different parameter sets to find a variance minimum [37, 38]. This scheme may be used to optimise Jastrow parameters which are not linear, as well as other parameters (multideterminant expansion

coefficients, cutoff lengths, backflow parameters). The same scheme for variance minimisation may be used to minimise other measures of spread, for example the MAD. The MAD is another measure of spread of local energies. We only consider minimisation of the MAD for the preliminary optimisation cycles in model calculations on charge-carrier complexes (Chapters 2 and 3), as we have empirically observed this to be useful for cases of near-stability.¹⁵ All production calculations are performed with wave functions that minimise the VMC energy.

The second method is energy minimisation [39]. Energy minimisation relies on a variant of an algorithm of Nightingale *et al.*, which is exact (in the limit of perfect sampling) for linear parameters (multideterminant expansion coefficients, for example) and works very well in practice for non-linear parameters [40].

In periodic solids, throughout this thesis, the values of any cutoff lengths are fixed to physically motivated values. For example, L_u should be such that the two-body Jastrow factor covers as much of the supercell as possible. If r_{SC} is the radius of the largest sphere which may be inscribed within the periodic supercell, then we would have $L_u < r_{SC}$ ($L_u = 0.99 \times r_{SC}$, say).

1.5 Diffusion Monte Carlo (DMC)

Whilst VMC is in itself a framework which is capable of retrieving adequate approximate solutions to the many-particle Schrödinger equation, it has one main limitation: the quality of one's results are determined solely by the quality of one's trial wave function. As long as the caveats of finite sampling are appropriately managed (equilibration, serial correlation), and that a large enough number of VMC samples are taken, the value of an observable evaluated in VMC is, to a

¹⁵Variance minimisation techniques are often used to “seed” an optimisation process, as the algorithm used to perform energy minimisation is effectively perturbative - relying on parameter values which are not too different from their optimal values. In cases of near-stability, which are more frequently encountered for charge-carrier complexes (some of which are unbound in our models), the variance may be effectively minimised if the trial wave function is able to describe the “fragmented” system formed by two or more products of dissociation. This is undesirable, and it has been found in practice that in cases where such dissociation is observed that MAD minimisation can be a helpful tool to ensure description of the bound system.

given uncertainty, determined only by Ψ_T . Diffusion Monte Carlo (DMC) is a means by which one can remove such an explicit dependence on the form of Ψ_T .

1.5.1 Imaginary time Schrödinger equation (ITSE)

The imaginary time Schrödinger equation reads

$$\left(\hat{\mathcal{H}} - E_r\right) \Psi(\mathbf{R}, \tau) = -\frac{\partial}{\partial \tau} \Psi(\mathbf{R}, \tau), \quad (1.50)$$

which differs from that in real time (in Eqn. (1.23)) by the transformation $\tau = it$, and the inclusion of an arbitrary shift in the total energy of E_r (the so-called “reference energy”).¹⁶ Or, written in integral form,

$$\Psi(\mathbf{R}, \tau + \delta\tau) = \int d\mathbf{R}' G(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau) \Psi(\mathbf{R}', \tau), \quad (1.51)$$

with $G(\mathbf{R} \leftarrow \mathbf{R}')$ the Green’s function $\langle \mathbf{R} | \exp \left[-\delta\tau \left(\hat{\mathcal{H}} - E_r \right) \right] | \mathbf{R}' \rangle$ (a real-space representation of the imaginary time evolution operator) and $\Psi(\mathbf{R}, \tau)$ an initial wave function. The Green’s function satisfies the boundary condition $G(\mathbf{R} \leftarrow \mathbf{R}', 0) = \delta(\mathbf{R} - \mathbf{R}')$. Suppose we expand $\Psi(\mathbf{R}', \tau)$ in eigenstates Φ_n (having, as earlier, energies E_n) of the many-particle Schrödinger equation, such that $\Psi(\mathbf{R}', \tau) = \sum_{n=0}^{\infty} c_n(\tau) \Phi_n(\mathbf{R}')$, with $\{c_n(\tau)\}$ a set of complex (for arguments sake) coefficients. Then Eqn. (1.51) reduces to

$$\Psi(\mathbf{R}, \tau + \delta\tau) = \sum_{n=0}^{\infty} c_n(\tau) \Phi_n(\mathbf{R}) \exp[-\delta\tau(E_n - E_r)]. \quad (1.52)$$

Imaginary time evolution, enacted by a Hamiltonian operator with a shifted energy, then leads to the exponential proliferation of components of the initial wave function which have energies $E_n < E_r$, and exponential decay of components of the initial wave function which have energies $E_n > E_r$. This is the basis of the DMC

¹⁶We will still work in the BO approximation, and have hence suppressed the parametric dependence on nuclear coordinates \mathbf{N} . We omit for brevity the index j – though its presence will be implied when excited states are discussed in Chapter 4

method. If the reference energy is chosen low enough, then imaginary time evolution ensures that all but the ground state component of the initial wave function exponentially decays (assuming that $\langle \Psi(\tau) | \Phi_0 \rangle \neq 0$).

1.5.2 Importance sampling transformation and the ISITSE

The DMC method outlined above may be turned into a computer algorithm which allows for imaginary time propagation of an initial state $\Psi(\mathbf{R}, \tau)$. However, this so-called “non-importance-sampled” approach by itself is almost never implemented, as far more efficient approaches are known. For this reason, we will neglect to describe the less efficient algorithm, and here discuss the importance-sampled ITSE (ISITSE) and its algorithmic solution. Similarly with the evaluation of VMC expectation values discussed earlier (see Section 1.4.2), one consequence of the importance sampling transformation considered here will be to improve the efficiency of configuration space sampling.

Consider the multiplication of an initial state $\Psi(\mathbf{R}, \tau)$ by a “trial” or “guiding” function $\Psi_T(\mathbf{R})$ to form the so-called “mixed distribution” $\mathcal{M}(\mathbf{R}, \tau) = \Psi_T(\mathbf{R})\Psi(\mathbf{R}, \tau)$. By writing $\Psi(\mathbf{R}, \tau) = \mathcal{M}(\mathbf{R}, \tau)/\Psi_T(\mathbf{R})$, we can show that the mixed distribution obeys an (importance-sampled) imaginary-time Schrödinger equation

$$-\frac{1}{2}\nabla_{\mathbf{R}}^2\mathcal{M}(\mathbf{R}, \tau) + \nabla_{\mathbf{R}} \cdot (\mathbf{V}(\mathbf{R})\mathcal{M}(\mathbf{R}, \tau)) + (E_L(\mathbf{R}) - E_T)\mathcal{M}(\mathbf{R}, \tau) = -\frac{\partial}{\partial\tau}\mathcal{M}(\mathbf{R}, \tau), \quad (1.53)$$

where $\mathbf{V}(\mathbf{R}) = \Psi_T(\mathbf{R})^{-1}\nabla_{\mathbf{R}}\Psi_T(\mathbf{R})$ is the “drift velocity” or “quantum force”. Eqn. (1.53) has a formal solution

$$\mathcal{M}(\mathbf{R}, \tau + \delta\tau) = \int d\mathbf{R}' G_{\text{IS}}(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau)\mathcal{M}(\mathbf{R}', \tau), \quad (1.54)$$

with an “importance-sampled” Green’s function defined as

$$G_{\text{IS}}(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau) = \langle \mathbf{R} | \exp \left[-\delta\tau \left(\hat{F} + E_L(\hat{\mathbf{R}}) - E_T \right) \right] | \mathbf{R}' \rangle, \quad (1.55)$$

and with \hat{F} the so-called Fokker-Planck operator,

$$\hat{F} = \hat{T} + \nabla_{\mathbf{R}} \cdot \mathbf{V}(\hat{\mathbf{R}}). \quad (1.56)$$

By insertion of $\Psi(\mathbf{R}, \tau) = \mathcal{M}(\mathbf{R}, \tau)/\Psi_{\text{T}}(\mathbf{R})$ into Eqn. (1.51), one can also show that $G_{\text{IS}}(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau) = \Psi_{\text{T}}(\mathbf{R})G(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau)\Psi_{\text{T}}(\mathbf{R}')^{-1}$.

By use of the Trotter-Suzuki formula [41], one may show that the exponent featured in the definition of G_{IS} factorises approximately as

$$\begin{aligned} \langle \mathbf{R} | \exp \left[-\delta\tau \left(\hat{F} + E_{\text{L}}(\hat{\mathbf{R}}) - E_{\text{r}} \right) \right] | \mathbf{R}' \rangle &= \exp \left[-\delta\tau (E_{\text{L}}(\mathbf{R}) - E_{\text{r}}) / 2 \right] \langle \mathbf{R} | \exp \left[\delta\tau \hat{F} \right] | \mathbf{R}' \rangle \times \\ &\exp \left[-\delta\tau (E_{\text{L}}(\mathbf{R}') - E_{\text{r}}) / 2 \right] + \mathcal{O}(\delta\tau^3). \end{aligned} \quad (1.57)$$

Henceforth, we will assume $\delta\tau$ to be finite, and small (the “short-time approximation”). If we further assume that the drift velocity between \mathbf{R} and \mathbf{R}' is constant, then we can approximate $G_{\text{IS}} = G_{\text{DMC}}$ ¹⁷ with

$$\begin{aligned} G_{\text{DMC}}(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau) &= \underbrace{\frac{1}{(2\pi\delta\tau)^{3N/2}} \exp \left[-\frac{|\mathbf{R} - \mathbf{R}' - \delta\tau\mathbf{V}(\mathbf{R}')|^2}{2\delta\tau} \right]}_{G_{\text{drift}}} \times \\ &\underbrace{\exp \left[-\delta\tau (E_{\text{L}}(\mathbf{R}) + E_{\text{L}}(\mathbf{R}') - E_{\text{r}}) / 2 \right]}_{G_{\text{branch}}} \end{aligned} \quad (1.58)$$

The first “drift” term is named as such because it is the Green’s function of the ISITSE *without* the position-dependent source term (which is a Langevin equation). Notice that G_{drift} is *not* symmetric in \mathbf{R} and \mathbf{R}' . This equation alone would describe the evolution of the mixed distribution subject to random diffusion with an underlying drift velocity. The second “branch” term is named as such because it is the solution of the ISITSE *without* the spatial derivatives (which is a rate equation). This equation alone would describe the growth and decay of the mixed

¹⁷The dominant error here is now $\mathcal{O}(\delta\tau^2)$, owing to the drift velocity assumption, which can also be seen as equivalent to taking the normal-ordered version of the exponential of the Fokker-Planck operator (times $\delta\tau$) in place of the exponential of the operator itself.

distribution over periods of imaginary time (dependent on the local and reference energies).

In any given step, the use of the DMC Green’s function to propagate the mixed distribution (as per Eqn. (1.54)) introduces an error ($\Delta\mathcal{M}$, say) of $\mathcal{O}(\delta\tau^2)$ – or at a *rate* of $\mathcal{O}(\delta\tau)$. Error is removed from a simulation over a significant number of individual steps by statistical means, at an approximate rate of $\Delta\mathcal{M}/t_{\text{corr}}$. After equilibration, these two rates balance, and $\Delta\mathcal{M} = \mathcal{O}(\delta\tau)$. This is “time step bias”, and it may be removed from DMC estimates of expectation values evaluated with samples from the mixed distribution (see Section 1.5.3) by linear extrapolation [42].

The DMC Green’s function is then descriptive of the evolution of the mixed distribution subject to random diffusion of configurations, with an overall drift, and with configuration growth and decay mediated by the relative magnitudes of the local and reference energies.

More concretely, throughout a DMC calculation, the mixed distribution is represented as the ensemble average over a population

$$\mathcal{M}(\mathbf{R}, \tau) = \left\langle \sum_i w_i(\tau) \delta(\mathbf{R} - \mathbf{R}_i) \right\rangle, \quad (1.59)$$

where $w_i(\tau)$ is the “weight” of a configuration \mathbf{R}_i at imaginary time τ . In the first step of a DMC calculation, it is normal to take a population of configurations from a preceding VMC calculation performed for the trial function $\Psi_T(\mathbf{R})$ (*i.e.* a finite number of $\{\mathbf{R}\}$ which are distributed as $|\Psi_T(\mathbf{R})|^2$). An upper limit on the sum is not imposed, for reasons which will shortly become clear. Substitution into Eqn. (1.54) gives

$$\mathcal{M}(\mathbf{R}, \tau + \delta\tau) = \left\langle \sum_i w_i(\tau) G_{\text{branch}}(\mathbf{R} \leftarrow \mathbf{R}_i, \delta\tau) G_{\text{drift}}(\mathbf{R} \leftarrow \mathbf{R}_i, \delta\tau) \right\rangle. \quad (1.60)$$

We have chosen to write this separately as $G_{\text{branch}}G_{\text{drift}}$ and not G_{DMC} to make what follows more obvious. The ensemble average can be maintained in the form of Eqn. (1.59) if we now interpret $G_{\text{drift}}(\mathbf{R} \leftarrow \mathbf{R}_i)$ as a transition-probability density

for a configuration \mathbf{R}_i , and define

$$w_i(\tau + \delta\tau) = w_i(\tau)G_{\text{branch}}(\mathbf{R} \leftarrow \mathbf{R}', \delta\tau), \quad (1.61)$$

such that one step of the DMC algorithm involves (for each configuration) an update of the weight w_i as above, and the proposal of a configuration move from \mathbf{R}_i to \mathbf{R}'_i with

$$\mathbf{R}'_i = \mathbf{R}_i + \boldsymbol{\xi} + \delta\tau\mathbf{V}(\mathbf{R}_i), \quad (1.62)$$

with $\boldsymbol{\xi}$ a $3N$ -dimensional vector of Gaussian-distributed random variates of zero mean and variance $\delta\tau$. The update to the weights could, in principle, lead to the excessive proliferation or destruction of a population of initial configurations. To remedy this, the reference energy is often updated throughout a simulation to ensure that the total simulation weight is (on average) equal to some target weight.¹⁸ This introduces a small, but systematic population control bias which may be removed simultaneously with time step bias [42]. A configuration move need not be accepted, for two possible reasons.

Firstly, as noted earlier, the DMC Green's function (actually, G_{drift}) is not symmetric in \mathbf{R} and \mathbf{R}' . The interpretation of G_{drift} as a transition-probability density is at odds with this fact, as it leads to a violation of detailed-balance for proposed configuration moves. Detailed balance can be imposed on the potential move by means of an accept/reject step. Moves of a configuration from $\mathbf{R} \rightarrow \mathbf{R}'$ are accepted with a probability

$$\mathcal{A}(\mathbf{R} \leftarrow \mathbf{R}') = \min \left(1, \frac{G_{\text{drift}}(\mathbf{R}' \leftarrow \mathbf{R}) |\Psi_{\text{T}}(\mathbf{R})|^2}{G_{\text{drift}}(\mathbf{R} \leftarrow \mathbf{R}') |\Psi_{\text{T}}(\mathbf{R}')|^2} \right). \quad (1.63)$$

Secondly, suppose a move is proposed and subsequently accepted, but that such a move causes a configuration \mathbf{R}_i to move across a *node* of the mixed distribution

¹⁸The means for doing this varies from code to code. In CASINO [34], the reference energy is often updated such that the total weight would hit some target after a fixed amount of imaginary time, but this behaviour is configurable in various ways (see the output of `$ casinohelp search EREF`, for example).

(such that the sign of \mathcal{M} changes). Our representation of \mathcal{M} , and the means by which moves are proposed and accepted, relies on the interpretation of \mathcal{M} as a probability density (which is non-negative). Fermion antisymmetry requires the exact many-body wave function to have nodes across which its sign changes. If we allow that the nodes of Ψ (the “DMC wave function”) may differ from those of the trial wave function Ψ_T , then \mathcal{M} must too have nodes. The existence of nodes of the mixed distribution ruins our algorithm as proposed, and further, any extensions to the case of “signed” configurations are invariably plagued by exponentially growing signal-to-noise problems (see *e.g.* [43], and references therein). Therefore, we force the scenario that \mathcal{M} is positive-definite, by assuming that Ψ and Ψ_T have the same nodal surface. This is the fixed-node approximation (FNA), and its only algorithmic consequence is that moves which would result in crossing a node are discarded or blocked. The FNA has far more serious consequences for the DMC method in general – it is the *only* uncontrolled approximation which we will henceforth employ.

1.5.3 Expectation values in DMC

The evaluation of expectation values in DMC differs from that in VMC in two ways. Firstly, when importance sampling is used, samples of the *mixed* distribution, \mathcal{M} , are generated, and so we are only (without extra care) able to extract so-called mixed estimators, of the form

$$A_{\text{mixed}} = \frac{\langle \Psi | \hat{A} | \Psi_T \rangle}{\langle \Psi | \Psi_T \rangle}, \quad (1.64)$$

the estimation of “pure” estimators in DMC is possible, but significantly more computationally demanding. That is discussed elsewhere, and can be done with use of the forward/future walking methods [43], or reptation Monte Carlo [44], or by so-called “Hellmann-Feynman operator sampling” [45].

The second complication concerns how populations are controlled. In each step

of a DMC calculation the population is adjusted by an amount which is exponential in the reference energy. A control mechanism for this feature (which avoids excessive growth or decay of the configuration population) involves the adjustment of the reference energy (see Section 13.4 of Ref. [46] for details). The reference energy is adjusted throughout a DMC calculation in order that the population size is approximately equal to some specified target population size, introducing a bias which is linear, for large enough population sizes, in the fixed population size and is often removed simultaneously with time step bias [42].

There are other methods available to negate this bias in mixed DMC estimators, *e.g.* the bias negation method of Umrigar, Nightingale and Runge [47] (the so-called “II-weights” scheme). Many other aspects of the DMC calculations presented here also originate in Ref. [47], but again, more recent strategies aimed at improving computational efficiency are also routinely applied [42].

1.6 Pseudopotentials in QMC

An important practical simplification of Eqn. (1.24) involves the electron-nucleus potential. As written, the electron-nucleus potential energy constitutes a problem for electronic structure calculations. Its inclusion implies that we always account for *all* of the electrons in any given system of interest (whereas chemical intuition tells us that often the relevant electrons are in valence states). The computational expense of an all-electron QMC calculation for an atom scales with atomic number Z as $\propto Z^{5.5}$ [48–50], making calculations on heavy atoms significantly more expensive than those on light atoms. This has the further consequence that we must attempt to describe (in a finite basis of states; plane waves for example) the core states themselves. Close to the atomic core, such states are wildly oscillatory (a requirement of orthogonality with other atomic eigenstates), and it is computationally expensive to try to represent them in a finite basis.

The idea behind the pseudopotential approximation is to replace the troublesome bare Coulomb potential with a potential which is identical beyond some

cutoff radius, but augmented within that radius. The details of this augmentation vary from scheme to scheme, but the salient point is that the pseudopotential is a means of (a) removing core electrons from a calculation, but including their effect on valence electrons in an effective core potential and (b) lessening the computational expense required to represent core and even valence states within the cutoff radius. The latter point remains important when one considers the formation of pseudopotentials where core states are not neglected (*e.g.* in forming pseudopotentials for the hydrogen atoms or in “pseudising” other quantities [51]).

In this thesis, when pseudopotentials are used in VMC and DMC calculations, they will be non-local, and of the Trail-Needs prescription [52, 53]. In DMC calculations, the so-called “T-move” scheme of Casula [54] will be used throughout. The “T-move” scheme is an important (necessary) means by which FN-DMC variational principles continue to exist for calculations involving non-local pseudopotentials.

We have ensured adequate accuracy in representing trial functions generated using Trail-Needs pseudopotentials by choosing DFT plane-wave cutoff energies which achieve a basis set convergence tolerance of 0.1 mHa per atom, as determined in Ref. [55]. Additionally, we have chosen the local pseudopotential channel in our DFT and QMC calculations such that no pseudopotential ghost states exist.¹⁹

¹⁹Fully non-local (angular momentum dependent) pseudopotentials are often necessary to accurately reproduce the scattering properties of atoms, a desirable property of any pseudopotential. Such pseudopotential operators may be written in the form of a local (angular momentum independent) operator plus a series of projectors onto individual angular momentum components. An arbitrary function can be subtracted from *all* angular momentum projectors, and added to the local operator, without changing the action of the pseudopotential for physical single-electron states. Choosing to do this so as to completely remove a given angular momentum projector amount to selecting that angular momentum component as the “local channel”. This is important if a given channel features *unphysical* or *ghost* states.

Chapter 2

Charge-carrier complexes in models of III-V semiconductors

2.1 The effective mass approximation

The effective mass approximation is a means of describing the dynamics of electrons and holes in solid semiconductors as those of free particles, whilst *partially* accounting for the existence of the crystal lattice, and other particles in the solid. Here, we will give an account of such dynamics from the point of view of a single-particle (independent electron band theory) picture.

The effective mass approximation then serves as a middle ground between fully quantum and fully classical dynamics, where we solve the Schrödinger equation (not Newton's equations), for a collection of quasiparticles whose properties are determined on some level by the nature of the host material to which they belong, and not those of free particles in vacuum.

2.1.1 Quasiparticle dynamics

The single-particle eigenfunctions for electrons in a periodic solid are Bloch waves [56], indexed by wave vector \mathbf{k} , of the form $\psi_{\mathbf{k}}(\mathbf{r}) = \exp[i\mathbf{k} \cdot \mathbf{r}]u_j(\mathbf{r})$. Their associated eigenvalues, $E_j(\mathbf{k})$, define the band structure of the solid (here, j labels the “band”

to which the electron belongs) [57–59]. Note that the concept of a band structure has parallels in many-body theory. This idea is discussed in Section 4.1.6.

Consider now, for simplicity, that we are interested in the dynamics of electrons in one band in particular (eigenvalue $E(\mathbf{k})$), for a given, small, range of wave vectors $\delta\mathbf{k}$. We could represent the electron as a packet of Bloch waves having, via the Heisenberg uncertainty relation, a spatial extent $\delta\mathbf{r}$. If we are to focus specifically on a given region of the first Brillouin zone (BZ), we require that $|\delta\mathbf{k}| \ll 1/a$,¹ with a the lattice constant of our chosen host material² and hence that $\delta\mathbf{r} \gg a$. *I.e.* our wave packet is delocalised over several unit cells of the crystal.

Such a wave packet moves through the crystal at a group velocity $\mathbf{v}(\mathbf{k}) = \nabla_{\mathbf{k}}E(\mathbf{k})$. In the absence of external fields, \mathbf{k} is constant, and the group velocity for a given electron can never change (this is an immediate issue for the semiclassical theory of electron transport; an injected current is dissipationless at this point). In an applied (constant) electric field $\boldsymbol{\mathcal{E}}$, however, the potential energy of an electron ($-\boldsymbol{\mathcal{E}} \cdot \mathbf{r}$) changes at a rate given by $-\boldsymbol{\mathcal{E}} \cdot \mathbf{v}$ (assuming no transitions between the single-particle bands occur). Equating this to the time derivative of the band-energy of the electron, we find

$$\begin{aligned} \frac{dE(\mathbf{k})}{dt} &= \frac{d\mathbf{k}}{dt} \cdot \nabla_{\mathbf{k}}E(\mathbf{k}) = \frac{d\mathbf{k}}{dt} \cdot \mathbf{v}(\mathbf{k}) \\ \implies \frac{d\mathbf{k}}{dt} &= -\boldsymbol{\mathcal{E}} \end{aligned} \tag{2.1}$$

which is reminiscent of Newton’s second law, but this time the role of “momentum” is played by the wave vector or crystal momentum \mathbf{k} . Taking this analogy further,

¹The full extent of the first BZ is defined by the set of points closer to the origin than any other reciprocal lattice vector - whose length scale is defined by $\sim 1/a$ (for a cubic material, the first BZ is a cube of side exactly $2\pi/a$), the resolution of a small fraction of the first BZ implies our stated constraint.

²Or the shortest length scale associated with the unit cell, if our material is not simple cubic.

we can consider the acceleration of the Bloch wave packet,

$$\begin{aligned}
\mathbf{a} &= \frac{d\mathbf{v}}{dt}, \\
&= \left(\frac{d\mathbf{k}}{dt} \cdot \nabla_{\mathbf{k}} \right) \mathbf{v}(\mathbf{k}), \\
&= (-\boldsymbol{\mathcal{E}} \cdot \nabla_{\mathbf{k}}) \nabla_{\mathbf{k}} E(\mathbf{k}),
\end{aligned} \tag{2.2}$$

or, in index notation

$$a_i = - \sum_j \left(\frac{\partial^2 E(\mathbf{k})}{\partial k_i \partial k_j} \right) \mathcal{E}_j, \tag{2.3}$$

and identify the (inertial) effective mass tensor for Bloch wave packets as

$$m_{ij}^{-1} = \left(\frac{\partial^2 E(\mathbf{k})}{\partial k_i \partial k_j} \right). \tag{2.4}$$

For the case of electrons in a perfectly isotropic quadratic band, which will be assumed in all examples within this thesis, we have a simpler form

$$m_{ij}^{-1} = \delta_{ij} \frac{\partial^2 E(\mathbf{k})}{\partial k^2}. \tag{2.5}$$

For an inverted band, m can be negative, and then the response of an associated wave packet would be to move in an opposite direction to that which one might expect from the arguments above for an electron. This is often the case when one considers a nearly-full band (*i.e.* a band with a single vacant state), and it is this notion that justifies the treatment of such a configuration as a “hole” with a *positive charge* and an effective mass $|m|$.

The dynamics of excitons and related charge-carrier complexes are then simply an extension of the quasiparticle dynamics presented above, with the inclusion of (screened) charge-charge interaction terms in an associated model Hamiltonian [60]. For the simple-case of a Coulomb interaction potential, exact solutions are possible for excitons, and such excitonic states are dubbed “Mott-Wannier”

excitons [61]. Other classes of exciton do exist, but not necessarily within the confines of the effective mass approximation. One particular issue related with the breakdown of the effective mass approximation is discussed in Section 3.1.4.

2.1.2 Excitonic units

In calculations employing the effective mass approximation, it is often convenient to work in a scaled set of units which are physically sensible for the system at hand. For the hydrogen atom, we have atomic units (Ha for energies, Bohr for lengths). For the Mott-Wannier exciton (and related problems), we have excitonic atomic units (Ha^{*} and Bohr^{*}). Such units simplify greatly the form of the model Hamiltonian, and can be defined as

$$\begin{aligned} \text{Ha}^* &= \frac{\mu}{\epsilon^2} \text{Ha}, \\ \text{Bohr}^* &= \frac{\epsilon}{\mu} \text{Bohr} \end{aligned} \tag{2.6}$$

for energies and lengths respectively, where the parameters μ , ϵ are the electron-hole reduced mass, and the dielectric constant of the material in question, respectively. In such units, and in centre of mass coordinates, the two-body Schrödinger equation for an interacting electron and hole is transformed to a one-body (Wannier) equation in relative coordinates. The Wannier equation for an exciton is identical to the Schrödinger equation for the hydrogen atom, but in excitonic units.

In the remainder of this chapter, and again in Chapters 3 and 5, excitonic atomic units will be assumed, with material parameters specified in each case. We will, in Chapters 3 and 5, refer to quasielectrons (quasiholes) simply as electrons (holes).

2.1.3 Charge-carrier complexes

Charge-carrier complexes are bound-states formed by the attractive (screened) Coulomb interaction between electrons and holes in semiconductors. The simplest

example is the exciton: a paired bound state of a single electron and a single hole whose effective mass description mirrors that of the exact solution for the hydrogen atom (in centre of mass coordinates).

Higher complexes are possible, and have been studied for several decades. The trion (or charged exciton) is a charged three-fermion complex, formed from either two electrons and a hole (negative trion), or two holes and a electron (positive trion). The biexciton is a neutral four-quasifermion complex. The formation of such states is always linked with some kind of excitation of the ground state in the underlying host material.

Normally, this is photoexcitation. In the single-particle picture, light from an external source (with some characteristic wavelength, and energy) illuminates a sample, and promotes electrons from the valence band to the conduction band. This may happen *en masse*, and as such generally leads to the formation of a population (whose extent is dependent on the intensity of the source) of excited electrons and valence band holes. The distribution of holes and electrons may be *biased* by the inclusion, intentional or otherwise, of dopants in a sample. Suppose a given sample contains some donor atom dopants (P or As in Si act as donor atoms, for example). Then, the population of electrons and holes in a sample will be biased towards the formation of electrons. Scenarios like this favour the formation of negatively charged charge-carrier complexes (such as the negative trion).

The trade off between attractive interaction of unlike quasiparticles, and the repulsive interaction of like quasiparticles, determines whether or not a given charge-carrier complex will form. Complexes that can form are characterised by their binding energies – the energy needed to break them apart in the most energetically favourable way (we will use the analogy with nuclear physics, and call such a process a “decay”). Binding energies are positive for bound states (which do not spontaneously decay), and negative if no bound state is possible (*i.e.* if the state spontaneously decays).

In the remainder of this chapter, simple charge-carrier complexes whose possible decays are trivial will be studied. Non-trivial decays become important in Chapter 3, where binding energies are discussed more thoroughly.

2.2 Coupled quantum wells in III-V semiconductors

Coupled quantum wells (CQWs) are extended heterostructures formed from a pair of close-by (in the z -direction, say) 2DEGs. Such heterostructures are of great experimental interest, and experimentalists and theorists alike have eagerly studied the possibility of proving the existence of a Bose-Einstein condensate (BEC) of the spatially indirect excitons which may be formed within them [62–68]. The detection of such a state is not as unambiguously obvious as one might expect, as it is expected to be *dark*, not coupling directly to light [62].

Specifically, upon photoexcitation, carriers in electrically biased CQWs are preferentially created in one of the two wells (electrons in one, holes in the other), and the subsequent mutual interaction of electrons and holes can lead to the formation of a bound interlayer exciton. Such interlayer excitons have a much larger lifetime than their intralayer counterparts, and it is their prolonged existence which leads to the possible formation of a BEC, as a non-equilibrium population of interlayer excitons can be created and (crucially) thermalise before all of the constituent members have started to decay [69, 70]. This gas of weakly-interacting bosonic particles may subsequently condense.

In this study, we have parameterised the binding energy of the interlayer trion (depicted in Fig. 2.1) for a range of material parameters, and identified the region of stability for the trion in terms of a critical layer separation (beyond which the trion is unstable) for a series of electron-hole mass ratios. Additionally, we have parameterised the biexciton binding energy over a wide range of material parameters.

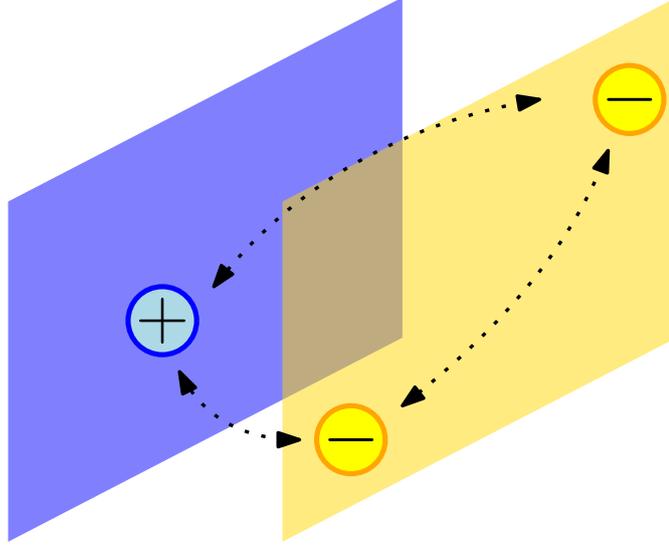


Figure 2.1: Schematic depiction of an indirect negative trion (X^-) in a coupled quantum well system. The yellow(blue) layer is the electron(hole)-bearing layer.

Here, a negative (positive) trion is formed by considering the addition of an electron (hole) to the electron (hole)-bearing layer of an existing indirect exciton. In such a case, the exact wave function is nodeless. The like-charges are of opposite spin quantum number. The trion binding energy is defined as $(E_X^T - E_{X^-}^T)$, where $E_{X^-}^T$ and E_X^T are total energies of the indirect trion and indirect exciton state states respectively. The exciton wave function is of the same form as used in previous work [12] and in later parts of this thesis (Drummond-Towler-Needs form, with two-body correlation factors), whilst the trion wave function in the present work was developed by O. Witham (and is described fully in Ref. [1]). The key property that this trial wave function form captures is the $\sim \frac{1}{\sqrt{r}}$ dependence of the wave function for a single electron far from a dipole source.³ This is the physics which binds a spatially indirect trion. Such physics is also of relevance to the family of anionic systems where an excess electron is bound to a molecule of fixed dipole moment [71–73]. In the case of an electron bound in the dipole field of a molecule, however, rigorous bounds on stability were proven to exist by Fermi and Teller in 1947 [74], though this result was famously overlooked and rediscovered by

³This may be proven in various ways. Series methods with a Frobenius *Ansatz* yield the correct leading order behaviour easily, as does exact solution in the limit of large- r (here the Schrödinger equation reduces to Bessel’s equation, having solutions which are asymptotically $r^{-1/2} \times \exp[-kr]$, k a constant).

various others (an interesting account of this is given in Ref. [75]). Such systems are different, however, as there the excess charge is not subject to a geometrical constraint, as in the present case.

2.2.1 Trion binding energies

The negative trion binding energy was calculated for values of the electron-hole mass ratio $\sigma = \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1, \frac{4}{3}, 2$, and 4. The corresponding total and binding energies of the positive trion may be generated by the transformation $\sigma \rightarrow \sigma^{-1}$ (charge conjugation symmetry applies, with a mass switch; and the total energy of the exciton is preserved under this operation). We have fitted these binding energies to a function of Padé form

$$E_{X^-}^b(d)/2 = \frac{E_{X^-}^b(0) + \sum_{i=1}^3 a_i d^i}{1 + \sum_{j=1}^4 b_j d^j}, \quad (2.7)$$

where $\{a_i\}$ and $\{b_j\}$ are fitting parameters, and subsequently used a Monte Carlo bootstrapping procedure to extract critical layer separations at a given mass ratio (with an associated error bar). Monte Carlo bootstrapping is outlined generally in Appendix A. Note the repeated factor of 1/2 feature in our quoted binding energies, as our fits are performed to data in excitonic Rydberg (1 Ry = 2 Ha) [1]. As previously stated, all of the binding energies given here, and those featuring in equations, are in excitonic units.

Aside from our critical distance determinations, we have also fitted our trion binding energies to a two-dimensional Padé function of the form

$$E_{X^-}^b(\sigma, d)/2 = \frac{\sum_{i=0}^3 \sum_{j=0}^3 f_{ij} (1 + \sigma)^{-i/2} d^j}{1 + \sum_{k=1}^4 g_k d^k}, \quad (2.8)$$

where f_{ij} and g_k are fitting parameters, and the σ dependence is motivated by the harmonic approximation within the BO approximation, in the limiting case that $\sigma \rightarrow \infty$. [3, 76, 77]

Unlike Eqn. (2.7), Eqn. (2.8) defines a fitted surface on all data points simul-

taneously, and hence may be used to evaluate accurately the trion binding energy away from the calculation data set. This fit is given mainly so that experimentalists who may have arbitrary experimental parameter sets within the range we have studied can use our data to form conclusions about trion binding in their samples. We note that we have tested the inclusion of higher cross-terms in Eqn. (2.8), but find that the currently included terms sufficiently balance accuracy, whilst maintaining a relatively low total number of parameters. In any case, we find that additional terms in fits of the form in Eqn. (2.8) yield corresponding fit parameters that are of low statistical significance.

Our critical distances obtained from Eqn. (2.7) are given in Table 2.1, whilst a corresponding plot of binding energies is displayed in Fig. 2.2.

σ	1/4	1/2	3/4	1	4/3	2	4
$d_{X^-}^{\text{crit}}$	4.52(4)	5.26(6)	6.63(9)	7.69(7)	> 8	> 8	> 8
d_{XX}^{crit}	0.48	0.43	0.42	0.42	0.42	0.43	0.48

Table 2.1: Critical layer separations for negative trions ($d_{X^-}^{\text{crit}}$) and biexcitons (d_{XX}^{crit}) over a range of electron-hole mass ratios σ . For the final three mass ratios, the critical distance lies outside of our chosen range of calculation parameters, and we are only able to quote a bound. The biexciton critical layer separations were obtained from fits given in Ref. [12].

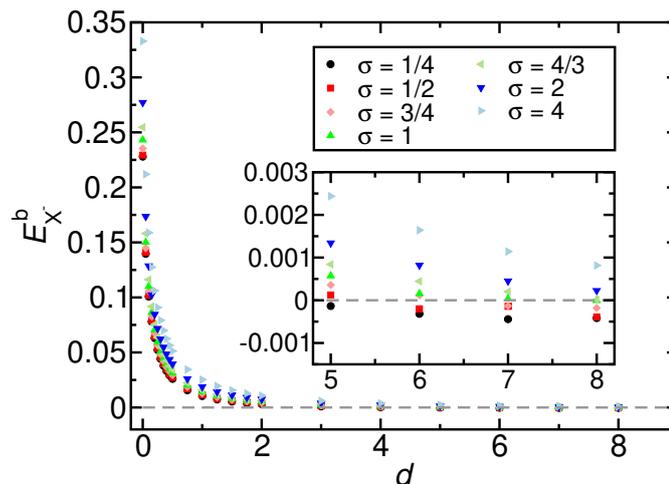


Figure 2.2: Negative-trion binding energy as a function of interlayer spacing d and electron-hole mass ratio σ . The inset shows the edge of the region of stability for the negative trion in greater detail.

The best-fit parameters⁴ in our trion binding energy fit are

$$f = \begin{pmatrix} 1.408 & 21.53 & 25.25 & 1.676 \\ -2.340 & -40.43 & -36.22 & -11.51 \\ 1.617 & 30.47 & 5.803 & 17.36 \\ -0.2129 & -0.5492 & 7.423 & -8.694 \end{pmatrix}, \text{ and } \mathbf{g} = \begin{pmatrix} 26.16 \\ 147.7 \\ 186.4 \\ 29.45 \end{pmatrix}, \quad (2.9)$$

where our fit is over 151 data points, having a maximal error of 0.25 mHa*, and with 90% of the data falling within 0.1 mHa* of the fit surface. The fit applies for $\sigma \in [1/4, 4]$ and $d \in [0, \min(8, d_{\text{X}^-}^{\text{crit}})]$.

2.2.2 Biexciton binding energies

Similarly with the trion binding energies, we have calculated the binding energy of biexcitons for the same range of material parameters. The biexciton critical distances are already known [12]. Again, we have fit to a two-dimensional function of Padé form, with the slight difference that the fit function is now appropriately symmetric under charge conjugation ($\sigma \rightarrow \sigma^{-1}$), reflecting the fact that in the case of the biexciton, the charge conjugated system is physically identical to the original.

$$E_{\text{XX}}^{\text{b}}(\sigma, d)/2 = \frac{\sum_{i=0}^2 \sum_{j=0}^2 F_{ij} (\sigma + \sigma^{-1})^{i/2} d^j}{\sum_{k=0}^3 G_k d^k}. \quad (2.10)$$

We find that the best-fit parameters in this case are

$$F = \begin{pmatrix} 0.03495 & -0.9822 & 2.437 \\ 0.07670 & 0.3303 & -1.786 \\ -0.005277 & -0.02931 & 0.2942 \end{pmatrix}, \text{ and } \mathbf{G} = \begin{pmatrix} 0.1726 \\ 3.256 \\ 1.567 \\ 29.95 \end{pmatrix}, \quad (2.11)$$

where in this case our fit is over 188 data points, having a maximal error of 10 mHa*, and with over 90% of the data falling within 2.5 mHa* of the fit surface.

⁴Note that in the matrix representation for f (and similar for g , and F, G later), the top-left-most entry corresponds to f_{00} , not f_{11} , and so on.

The fit applies for $\sigma \in [1/4, 4]$ and for $d \in [0, \min(8, d^{\text{crit}_{\text{xx}}})]$ (critical distances given in Ref. [12]).

In conclusion, we have established the trion region of stability in a model for charges in CQW heterostructures. Such trions are generally far more stable than biexcitons in the same systems, remaining bound for critical layer separations around an order of magnitude greater than biexcitons for the same model parameters. For example, taking model parameters appropriate for a CQW device studied in experiments by Butov *et al.* [63, 78] (see Section IIIA of Ref. [1] for a comprehensive list of model parameters), we find that the critical distance for negative trion formation is $821(9)\text{\AA}$, whereas for the biexciton, Lee *et al.* find that the critical distance is $67(8)\text{\AA}$.

We have presented fits to the trion and biexciton binding energies that well-represent our data in the space of model parameters that we have studied, and where the complexes are bound. Our biexciton data are in agreement with the results of earlier works [12], and (although we have not studied the stability boundary directly) do not contradict the conclusions reached therein regarding stability of the biexciton.

2.3 Quantum rings in III-V semiconductors

Quantum rings are localised heterostructures, akin to artificial atoms, which may spontaneously form under molecular beam epitaxy growth conditions of certain III-V semiconductors. GaSb rings tend to form in GaAs, when layers are grown on top of each other, for example, as a consequence of lattice strain and Sb segregation effects [79]. Specifically, Type-II quantum rings are those in which holes are confined to an annular region (comprised dominantly of GaSb, for example), and electrons are free to exist in the regions around the annulus (in the region comprised of GaAs, for example). Such features have been of experimental interest in recent years [13, 80–82] for their unique properties as artificial atoms, which may find uses in (opto)electronic devices such as quantum dot-based memory [83], for

example.

From the point of view of the effective mass approximation, the description of a confined hole or a free electron in a quantum ring heterostructure is rather simple. The only complicating factor is geometry. In reality, of course, various other effects are present, including those of lattice mismatch and subsequent strain fields, finite well-width effects (which are necessary for a realistic study of carrier recombination), and inhomogeneous dielectric background to name a few.

Here, we give a derivation of the form of hole orbital that can be (and was, in Ref. [2]) used in subsequent QMC studies of charge-carrier complexes in and around a model for type-II quantum rings. Such an orbital describes a single non-interacting hole bound inside an annular-shaped region of space (see Fig. 2.3). This orbital can be combined with other hole orbitals, or an accompanying electron orbital, defined in the region surrounding the ring, to form QMC trial wave functions for charge-carrier complexes comprised of distinguishable (opposite spin) or indistinguishable (same-spin) fermions. In the present case, only carrier complexes formed of distinguishable particles were ever considered, and DMC hence provides a numerically exact solution to the corresponding few-body Schrödinger equations.

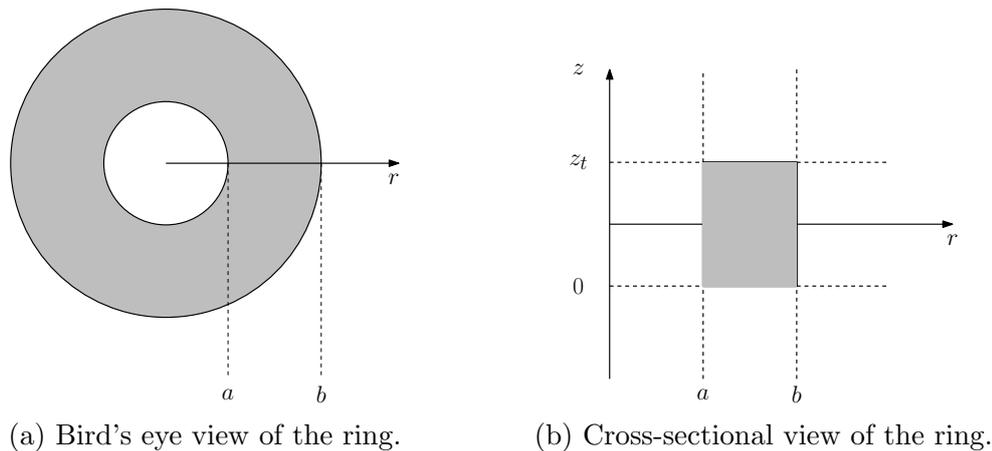


Figure 2.3: Schematic representations of a quantum ring. Grey areas represent hole-confining regions in the type-II ring (Sb-rich regions in the case of GaSb/GaAs grown rings).

The single hole orbital ψ_h satisfies the Schrödinger equation (in a.u.; with

masses in units of the hole effective mass)

$$-\frac{\nabla^2}{2}\psi_h(\mathbf{r}) + V(\mathbf{r})\psi_h(\mathbf{r}) = E\psi_h(\mathbf{r}), \quad (2.12)$$

with energy eigenvalue E , and with a confining potential

$$V(\mathbf{r}) = \begin{cases} 0, & a < \rho < b \text{ and } 0 < z < z_t \\ \infty, & \text{elsewhere} \end{cases}. \quad (2.13)$$

Assuming a separable form of $\psi_h(\mathbf{r}) = R(r)\Theta(\theta)Z(z)$, we can determine R , Θ and Z . Substituting the Laplacian in cylindrical coordinates, and shuffling terms, we find that Eqn. (2.12) becomes

$$\begin{aligned} -\frac{1}{Z}\frac{\partial^2 Z}{\partial z^2} &= \frac{1}{rR}\left[\frac{\partial^2 R}{\partial r^2} + r\frac{\partial R}{\partial r}\right] + \frac{1}{r^2\Theta}\frac{\partial^2 \Theta}{\partial \theta^2} + 2E, \\ &\equiv k_z^2, \end{aligned} \quad (2.14)$$

in the ring region, and where the constant k_z^2 is to be determined. Solving for Z we find

$$Z(z) = c_z \sin(k_z z), \quad k_z = \frac{n\pi}{z_t}, \quad n \in \mathbb{Z}, \quad (2.15)$$

where we have invoked the boundary conditions $\psi_h(z = 0) = \psi_h(z = z_t) = 0$. Similar separation is possible for the RHS in Eqn. (2.14),

$$\begin{aligned} \frac{r}{R}\left[r\frac{\partial^2 R}{\partial r^2} + \frac{\partial R}{\partial r}\right] + 2r^2(E - k_z^2) &= -\frac{1}{\Theta}\frac{\partial^2 \Theta}{\partial \theta^2}, \\ &\equiv -\ell^2, \end{aligned} \quad (2.16)$$

with ℓ a constant, to be determined. Again, invoking boundary conditions ($\Theta(\theta) = \Theta(\theta + 2\pi)$), we can solve for Θ

$$\Theta(\theta) = c_\theta e^{i\ell\theta}, \quad \ell \in \mathbb{Z}. \quad (2.17)$$

Finally, the radial equation reduces to

$$r^2 \frac{\partial^2 R}{\partial r^2} + r \frac{\partial R}{\partial r} + [r^2[2(E - k_z^2)] - \ell^2] R = 0, \quad (2.18)$$

which is Bessel's equation in a scaled coordinate. Letting $x = r\sqrt{2(E - k_z^2)}$, we recover Bessel's equation exactly, and as such the general solution to our radial equation is given by

$$R(r) = c_r J_\ell(r\sqrt{2(E - k_z^2)}) + \tilde{c}_r Y_\ell(r\sqrt{2(E - k_z^2)}). \quad (2.19)$$

In this case, the boundary conditions lead to a linear problem in the coefficients.

We require

$$\begin{pmatrix} J_\ell(a\sqrt{2(E - k_z^2)}) & Y_\ell(a\sqrt{2(E - k_z^2)}) \\ J_\ell(b\sqrt{2(E - k_z^2)}) & Y_\ell(b\sqrt{2(E - k_z^2)}) \end{pmatrix} \begin{pmatrix} c_r \\ \tilde{c}_r \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad (2.20)$$

which has unique solutions when the determinant of the matrix on the LHS has zero determinant. *I.e.* the allowed values of $\sqrt{2(E - k_z^2)}$ are given by the m^{th} roots $\xi_{m,\ell}$ of

$$f(\xi) = J_\ell(a\xi)Y_\ell(b\xi) - J_\ell(b\xi)Y_\ell(a\xi). \quad (2.21)$$

The constants c_r and \tilde{c}_r are related by Eqn. (2.20) and finally, the eigenstates of the ring problem are given by

$$\psi_{h;n,\ell,m}(\mathbf{r}) = c \times e^{i\ell\theta} \sin\left(\frac{n\pi}{z_t}\right) [J_\ell(r\xi_{m,\ell})Y_\ell(a\xi_{m,\ell}) - Y_\ell(r\xi_{m,\ell})J_\ell(a\xi_{m,\ell})], \quad (2.22)$$

where c is an (irrelevant) overall normalisation factor. The associated energy eigenvalues are

$$E = \frac{\xi_{m,\ell}^2}{2} + \frac{n^2\pi^2}{z_t^2}. \quad (2.23)$$

The ground state orbital relevant for QMC simulations is that of the lowest-lying radial eigenstates, $\psi_{h;0,0,1}$. Secondary holes are assumed (in forming trial wave

Complex	E^D	E^b
X	0	-2.226(3)
X ⁻	0.199(2)	-0.199(2)
X ⁺	-4.944(7)	-6.137(7)
XX	-4.052(8)	-3.11(1)

Table 2.2: De-excitonisation (E^D) and binding (E^b) energies, in meV, for excitonic complexes in the quantum ring geometry of Ref. [13].

functions) to be non-interacting with the first, and of opposing spin quantum number, such that the trial wave function is nodeless. Numerous further details of QMC calculations (performed by D. M. Thomas) employing the hole orbital derived here are available in Ref. [2]. Here, we will give a short summary of the important results, which are predicated on the hole orbital as derived here.

Firstly, we note that binding energies do not necessarily determine the position of a spectral peak in a photoluminescence experiment (*i.e.* when a charge-carrier complex decays, the peak position on an absolute scale is not directly related with the binding energy). It is true, however, that spectral lines are often measured *relative to the exciton line*. To this end, we also define the de-excitonisation energy (E^D) of a complex as the energy needed to remove an exciton from the complex. This is the energy at which (measured relative to the exciton line) a spectral peak attributed to might be observed in an experiment. With this in mind, complex binding energies and de-excitonisation energies for a quantum ring whose geometry mimics that of the GaSb/GaAs quantum rings studied in Ref. [13] are given in Table 2.2.

Table 2.2 therefore predicts that the lowest-energy photoluminescence peak for emission from a quantum ring corresponds to that from the negative trion, which is redshifted from the exciton line by a fraction of an meV. Next is the exciton line itself, likely to be the most intense peak, followed by the blueshifted biexciton peak around 4 meV from the exciton line, and then at the highest energy the positive trion peak, which is blueshifted by around 5 meV (see Fig. 2 of Ref. [2] for a pictorial representation of the same information). The binding energies tell

us a little more about what experiments are likely to actually see, however. Even for the most stable complex (the positive trion), the binding energy is only around 6 meV, and we would expect thermal smearing to destroy any well-defined signal originating from the decay of such a complex at a temperature of around 70K (where $k_B T \sim 6$ meV). These conclusions are maintained under variation of the model parameters, where we find only a tiny dependence of the binding and de-excitonisation energies on ring geometry. The derivatives of the biexciton binding energy w.r.t. the ring volume and the mean radius of the ring are 0.0004(4) meV/nm³ and 0.07(2) meV/nm, respectively, whereas the ring volume and mean ring radius corresponding to the systems in Ref. [13] are of order hundreds of nm³, and tens of nm, respectively. Hence, there is only scope for small geometry effects to change binding energies by a fraction of an meV.

An obvious extension to the work described here would be to explicitly consider the formation of pair orbitals for holes bound in doubly-occupied rings. Such an approach would likely not yield a closed-form solution, and owing to the intended use in QMC calculations (which are able to describe interaction effects exactly, anyway), this has not been attempted. Furthermore, experimental evidence suggests that such interaction effects are not very important in realistic ring geometries [13], where charging energies appear approximately constant.

Chapter 3

Charge-carrier complexes in models of two-dimensional semiconductors

In this chapter, various models are derived and solved (with QMC techniques) which pertain to the formation of charge-carrier complexes in two-dimensional semiconductors and their heterostructures. The formation of charge-carrier complexes dominates the optical response of these materials at low enough temperatures (such that binding energies are $\gg k_B T$).

3.1 Charge complexes in two-dimensional TMDs

Two-dimensional transition metal dichalcogenides (specifically WS_2 , MoS_2 , WSe_2 , and MoSe_2) are a family of two-dimensional semiconductors. They are, as monolayers, direct band gap semiconductors, with not-insignificant spin-orbit effects leading to a sizeable conduction band spin splitting (see Fig. 3.1), [84, 85]. They have been identified as ideal candidates for use in next-generation optoelectronic devices [86–88], owing to their unique, but also tunable, optical properties.

In our various calculations on TMDs, we have always used the same base mate-

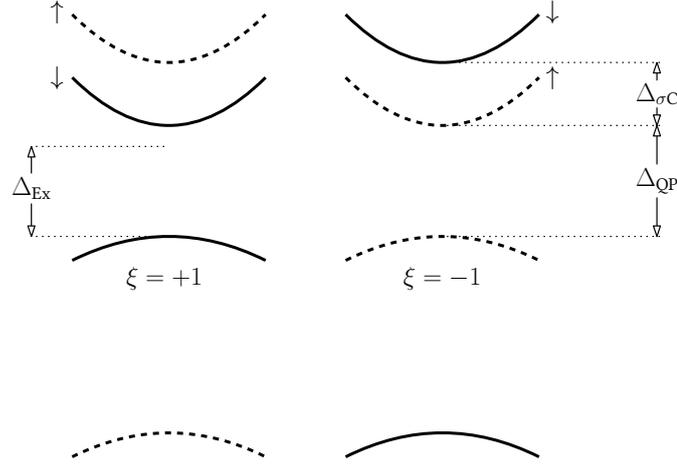


Figure 3.1: Schematic band structure of MoX_2 -type TMDs ($X=\text{S}, \text{Se}$). In WX_2 -type TMDs, the VBM and CBM states host electrons with *different* spin quantum numbers. Inequivalent valleys are designated $\xi = \pm 1$, and the quantities Δ_{Ex} , Δ_{QP} , and $\Delta_{\sigma\text{C}}$ are the excitonic and quasiparticle energy gaps (see Chapter 4), and the conduction band spin splitting respectively.

rial parameters,¹ sometimes with modification depending on the dielectric environment, or when layering to describe a heterostructure. These are summarised in Table 3.1. The electronic properties of monolayers (masses and screening parameters) are taken from recent theoretical works. The interlayer separation is important for the later treatment of multilayer systems, and here we have sought experimentally derived values. Theoretical calculations based on density functional theory fail to describe the interlayer bonding in such layered systems, as this is dominated by (non-local) dispersion and van der Waals interactions, which are poorly accounted for in many DFT exchange-correlation functionals. See *e.g.* Ref. [89] for a recent benchmark study comparing various flavours of DFT to first-principles QMC calculations for the related case of bilayer graphene.

3.1.1 The Keldysh interaction

An important aspect of the model calculations whose results are reported in this thesis chapter regards the nature of the electrostatic screening in two-dimensional semiconductors. In the models for charge-carriers in III-V heterostructures which

¹We have explored the sensitivity of our results to our choice of model parameters in Section 3.2.

TMD	Vacuum Parameter Value			
	$m_e (m_0)$	$m_h (m_0)$	r_* (a.u.)	d (a.u.)
MoS ₂	0.35 [90]	0.428 [90]	72.98 [90]	11.617 [91]
MoSe ₂	0.38 [84]	0.44 [84]	75.19 [92]	12.20 [93]
WS ₂	0.27 [84]	0.32 [84]	71.60 [84]	11.644 [94]
WSe ₂	0.29 [84]	0.34 [84]	85.25 [84]	12.25 [94]

Table 3.1: Material parameters for various monolayer TMDs. m_e , m_h , r_* , and d are the electron effective mass (for the conduction band at K), the hole effective mass (for the valence band at K), the Keldysh screening parameter (see Section 3.1.1), and the interlayer spacing of the bulk parent crystal, respectively.

were presented in Chapter 2, the Coulomb interaction, screened isotropically by the bulk dielectric constant ϵ of the host material, provided a reasonable description of the interactions between particles. In two-dimensional materials, however, conduction charges are ultimately confined within a layer, and the response of the rest of the two-dimensional material is important. When a conduction charge is present in a layer, the valence charge may be polarised within the plane of the two-dimensional material, altering the electrostatic potential produced by said conduction charge. The self-consistent treatment of this polarisation effect culminates in the so-called Keldysh interaction, which we will now derive.

Consider a two-dimensional material having two-dimensional polarisability κ , suspended at $z = 0$ and which is “sandwiched” between two identical dielectric media, having dielectric tensor $\tilde{\epsilon} = \text{diag}(\epsilon_{\parallel}, \epsilon_{\parallel}, \epsilon_{\perp})$. Further, consider placing a test charge density ρ within the plane of the two-dimensional material (modelling a conduction charge, for example), such that

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

where $\mathbf{r} = (x, y, z)$ and $\mathbf{s} = (x, y)$. The electric displacement field $\mathbf{D}(\mathbf{r})$ due to this charge, and due to the polarisation of the surrounding charge within the two-dimensional layer $\mathbf{P}_{\perp}(\mathbf{r})$, is given by

$$\mathbf{D}(\mathbf{r}) = -\tilde{\epsilon} \cdot \nabla\phi(\mathbf{s}, z) - \kappa[\nabla_{\parallel}\phi(\mathbf{s}, z = 0)]\delta(z), \quad (3.2)$$

where ϕ is the electrostatic potential, and the operator $\nabla_{\parallel} \equiv (\partial_x, \partial_y, 0)$. Gauss' law ($\nabla \cdot \mathbf{D} = \rho$) gives

$$\rho(\mathbf{s})\delta(z) = \nabla \cdot \mathbf{D} = -\epsilon_{\parallel}\nabla_{\parallel}^2\phi(\mathbf{r}, z) - \epsilon_{\perp}\partial_{zz}\phi(\mathbf{r}, z) - \kappa[\nabla_{\parallel}^2\phi(\mathbf{s}, z=0)]\delta(z). \quad (3.3)$$

Taking Fourier transforms of either side,

$$f(\mathbf{s}, z) = \int \frac{d^2\mathbf{q}}{(2\pi)^2} e^{i\mathbf{s}\cdot\mathbf{q}} \int \frac{dk}{2\pi} e^{ikz} f(\mathbf{q}, k), \quad (3.4)$$

we then find,

$$\rho(\mathbf{q}) = (\epsilon_{\parallel}q^2 + \epsilon_{\perp}k^2)\phi(\mathbf{q}, k) + \kappa q^2\phi(\mathbf{q}, z=0), \quad (3.5)$$

or,

$$\phi(\mathbf{q}, k) = \frac{\rho(\mathbf{q}) - \kappa q^2\phi(\mathbf{q}, z=0)}{(\epsilon_{\parallel}q^2 + \epsilon_{\perp}k^2)}. \quad (3.6)$$

We then perform a partial inverse Fourier transform in k , obtaining

$$\begin{aligned} \int_0^{\infty} \frac{dk}{2\pi} \phi(\mathbf{q}, k) e^{ikz} &= \frac{[\rho(\mathbf{q}) - \kappa q^2\phi(\mathbf{q}, z=0)]}{2\pi} \int_0^{\infty} \frac{dk e^{ikz}}{(\epsilon_{\parallel}q^2 + \epsilon_{\perp}k^2)}, \\ &= \frac{[\rho(\mathbf{q}) - \kappa q^2\phi(\mathbf{q}, z=0)]}{2q\sqrt{\epsilon_{\parallel}\epsilon_{\perp}}} e^{-\alpha q|z|}, \\ &= \phi(\mathbf{q}, z), \end{aligned} \quad (3.7)$$

with $\alpha = \sqrt{\epsilon_{\parallel}/\epsilon_{\perp}}$, and which can be solved for $\phi(\mathbf{q}, z=0)$,

$$\begin{aligned} \phi(\mathbf{q}, z=0) &= \frac{\rho(\mathbf{q})}{q(2\sqrt{\epsilon_{\parallel}\epsilon_{\perp}} + \kappa q)}, \\ &= \frac{Q}{2\epsilon q(1 + r_{\star}q)}. \end{aligned} \quad (3.8)$$

where in the last equality we have specialised to the case of a point charge at the origin ($\rho(\mathbf{q}) = Q$), and defined $\epsilon = \sqrt{\epsilon_{\parallel}\epsilon_{\perp}}$ and $r_{\star} = \frac{\kappa}{2\epsilon}$.

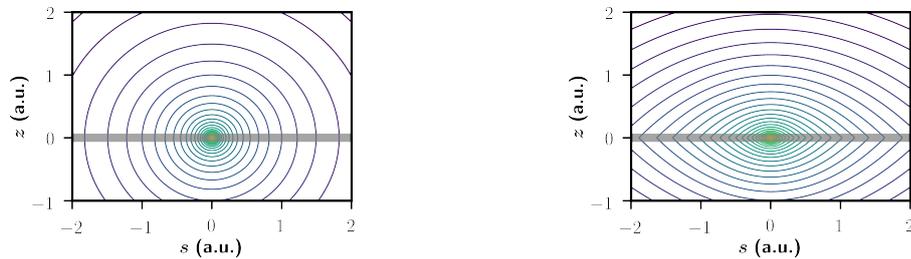
Finally, taking the inverse Fourier transform in \mathbf{q} , we find that the electrostatic

potential around a point charge in a polarisable two-dimensional medium is

$$\begin{aligned}
\phi(\mathbf{s}, z = 0) &= \int \frac{d^2\mathbf{q}}{(2\pi)^2} e^{i\mathbf{s}\cdot\mathbf{q}} \frac{Q}{2\epsilon q(1+r_*q)}, \\
&= \int \int \frac{q \, dq \, d\theta}{(2\pi)^2} \frac{Q}{2\epsilon q(1+r_*q)}, \\
&= \int_0^\infty \frac{q \, dq}{(2\pi)^2} \frac{Q}{2\epsilon q(1+r_*q)} \int_0^{2\pi} d\theta e^{iqs \cos(\theta)}, \\
&= \frac{Q}{4\pi\epsilon} \int_0^\infty dq \frac{J_0(qs)}{1+r_*q}, \\
&= \frac{Q}{4\pi\epsilon} \cdot \frac{\pi}{2r_*} \left[H_0\left(\frac{s}{r_*}\right) - Y_0\left(\frac{s}{r_*}\right) \right]. \tag{3.9}
\end{aligned}$$

H_0 is a zeroth order Struve function, and Y_0 a zeroth order Bessel function of the second kind. This result was first derived by Rytova [95, 96], but has since become known as the ‘‘Keldysh’’ interaction owing to the more widely known, later, work of Keldysh [97]. In the limit that $s \gg r_*$, screening becomes irrelevant and this potential reduces to the expected (Coulomb) $1/r$ form. In the limit that $s \ll r_*$, the potential is logarithmic (as would be expected from true two-dimensional electromagnetism). The logarithmic form of the potential at short range ($s \rightarrow 0$) has consequences for the Kato cusp condition which applies between particles interacting via the Keldysh interaction (see *e.g.* Ref. [98]).

Note that we could evaluate the Keldysh potential for $z \neq 0$, if we wished to. A qualitative² comparison of the resulting equipotential surfaces is given in Fig. 3.2.



(a) The Coulomb interaction. (b) The Keldysh interaction ($r_* = 50$ a.u.).

Figure 3.2: Equipotential surfaces from a charge at the origin in the Coulomb and Keldysh interactions.

²The equipotential lines in the two subfigures do not necessarily correspond to the same values of electrostatic potential.

3.1.2 The Mott-Wannier-Keldysh model

The use of the Keldysh interaction combined with an effective mass treatment of carrier dynamics constitutes what is known as the Mott-Wannier-Keldysh model for charge-carrier complex total energies. If a charge-carrier complex (denoted C) is formed of n distinguishable quantum particles, the few-body Schrödinger equation describing the system reads

$$\left[-\sum_{i=1}^n \frac{\nabla_i^2}{2m_i} + \sum_{\langle i,j \rangle} q_i q_j V(\mathbf{r}_i, \mathbf{r}_j) \right] \Psi(\mathbf{R}) = E_C^T \Psi(\mathbf{R}), \quad (3.10)$$

where $V(\mathbf{r}_i, \mathbf{r}_j)$ is the interaction potential between charges i and j (having masses m_i and m_j), $\sum_{\langle i,j \rangle}$ denotes a single-counted sum over particle pairs, $\mathbf{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$ is the $3n$ -dimensional vector of particle positions, and $\Psi(\mathbf{R})$ is the few-body wave function. E_C^T is the total energy of the complex C . For the Mott-Wannier-Keldysh model, the charge-charge interaction potential is now of Keldysh form (or, in later sections, bilayer or multilayer Keldysh form). The same equation describes complexes formed from indistinguishable fermions, but we explicitly exclude such cases from the present thesis. Such a situation requires the few-body wave function to have a nodal surface. Biexcitons formed from indistinguishable fermions were studied by another author in Ref. [3].

Throughout the remainder of this section, when solving the Mott-Wannier-Keldysh model, we have used trial wave functions of Jastrow form, where the Jastrow exponents contained smoothly truncated polynomial particle-particle terms, ion-particle terms, ion-particle-particle, and particle-particle-particle terms [28]. Additional terms satisfying the analogues of the Kato cusp conditions [31, 98, 99] were applied to the trial wave function between pairs of particles wherever there was a logarithmic divergence in the interaction between them, including the unphysical divergences in the approximate bilayer Keldysh interaction (see Appendix B). Free parameters were optimised using VMC with variance [36, 37] and, subsequently, energy minimisation [100] as implemented in the CASINO code [34].

In DMC calculations, we invariably use two DMC time steps in the ratio 1:4 and corresponding target populations in the ratio 4:1, allowing a simultaneous extrapolation to zero time step and infinite population. Since we will only consider distinguishable charge carriers, there is no fixed-node error and hence DMC provides exact ground-state solutions to the Mott-Wannier-Keldysh model. However, practical issues do exist. Weakly bound systems can experience ergodicity problems in random sampling, where a subset of charges separate appreciably from the rest, and sampling of some regions of configuration space (particularly the regions which correspond to all charges being nearby) becomes very rare, allowing the possibility for the formation of unreasonable features in trial wave functions during VMC optimisation, for example (*e.g.* the formation of “pseudonodes”, regions of extremely small wave function magnitude which act, effectively, as nodes).

For cases where charge-carrier complexes are well bound, we expect that the major sources of error in this chapter are those relating to the validity of the model itself.

The binding energy of a complex C is defined as the minimum energy necessary to separate it into a series of smaller constituents (daughter configurations). If E_d^T is the total energy of a possible daughter configuration d , then the binding energy can be written as

$$E_C^b = \min_{\{d\}} \left[\sum_{\alpha} E_{d_{\alpha}}^T - E_C^T \right]. \quad (3.11)$$

Following this definition, binding energies are *positive* quantities. Unbound complexes C are characterised by *non-positive* values of E_C^b .³ For a lone exciton in this model the total energy is equal (but of opposite sign) to the binding energy. The two constituent particles have, separately, zero total energy. A schematic diagram giving some possible daughter configurations for a biexciton is given in Fig. 3.3.

Binding energies do not determine the positions of spectral lines in, for example, photoluminescence experiments. In a photoluminescence experiment, an

³Binding energies can become negative if one uses a trial wave function which artificially binds a complex, but in that case the results themselves are highly sensitive to the particular form of trial wave function, and the negative value is, in itself, meaningless.

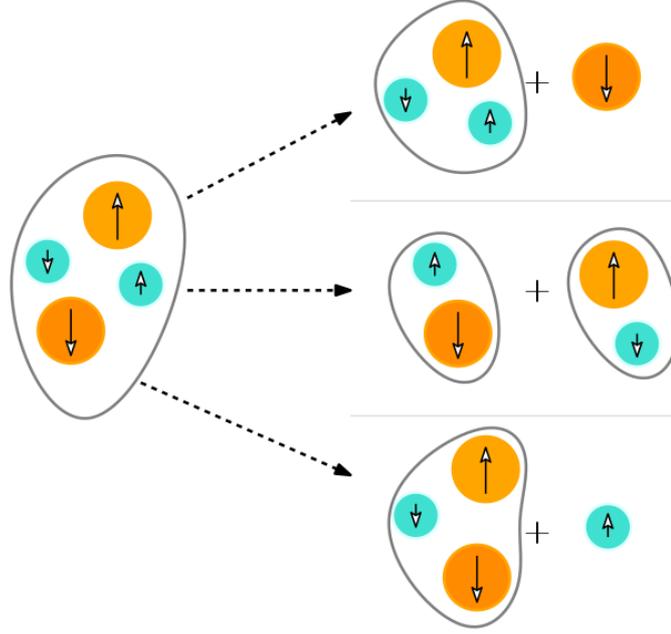


Figure 3.3: Possible daughter configurations of a biexciton decay.

exciton line (attributed to an exciton, X , of binding energy E_X^b) would appear at an (incoming laser) excitation energy of $\Delta_{QP} - E_X^b$. Binding energies only determine whether or not a charge-complex can form upon photoexcitation, and a positive binding energy is only a necessary condition for the corresponding observance of a spectral line.

The position of spectral lines are, as mentioned in Chapter 2, still determined by the de-excitation energy. Even then, the binding energy condition is a necessary (but not sufficient) condition, owing to the angular momentum selection rules which must be obeyed in order for photon emission to occur. If angular momentum rules forbid excitonic recombination in a charge-carrier complex, then the complex is termed “dark”.⁴

In TMDs, there are two further pieces of physics to consider when it comes to the optical activity of excitonic complexes. The first is the nature of the optical gap: do the electron and hole exist with a common crystal momentum, or does the exciton necessarily have a net momentum? This has ramifications for the

⁴The complex will form upon photoexcitation, but direct (fast) recombination through electron-hole recombination is forbidden by selection rules. Other (slow; typically defect- or phonon-assisted) recombination mechanisms exist that do allow for the recombination of dark complexes.

allowed recombination mechanism. Excitons with non-zero momentum must (in order to conserve that momentum) recombine via an indirect recombination process, *i.e.* with the assistance of an external energy/momentum carrier (a phonon, perhaps). Monolayer TMDs have direct optical gaps, but this doesn't hold in few-layers. The second regards an often important aspect of hexagonal two-dimensional materials: inequivalent valleys (ξ/ξ'), and the scattering between them. We will not consider this possibility, but such scattering events do, in reality, occur. An obvious consequence of such effects is to determine the polarisation of emitted light as measured in a photoluminescence experiment performed with a circularly polarised source. Excitons in such systems are created in one corner of the BZ (ξ or ξ'), and intervalley scattering effects lead to the production of an output signal is consistent with excitons in *both* valleys, which are populated unequally, and with populations that evolve characteristically in time with the intervalley scattering rate [101–103].

3.1.3 Quintons and related heavy complexes in TMDs

Recently, experimentalists have observed narrow resonances in the photoluminescence emission spectra of TMD monolayers. The dominant resonance (or “peak”) is, normally, that which corresponds to the recombination of an optically active exciton (X).⁵ However, the formation of higher complexes is possible. Trions (X^\pm ; three-quasi-fermion complexes) and biexcitons (XX four-quasi-fermion complexes) are obvious examples of such higher complexes, but there are also a plethora of observed defect-bound complexes known to form in defected monolayers under optical excitation.

The formation of further purely quasi-fermion complexes seems inhibited by considerations of the band structure: further electron states are separated by a large spin-orbit splitting, and further hole states are significantly lower-lying in energy. However, if we allow for the consideration of the valley degree of freedom,

⁵Note that excitons and excitonic complexes which do not recombine by emission of a single photon, owing to angular momentum selection rules, are termed “dark”.

we can in principle consider the formation of complexes comprised of (ultimately) twice as many quasi-fermions. Here for example, a quinton (XX^- a negatively charged biexciton) would be composed of two distinguishable holes having opposite valley and spin degrees of freedom, and three distinguishable electrons: two having opposite spin, but identical valley quantum numbers, and the other having a different valley quantum number and a spin which is either up or down. Ignoring inter-valley scattering effects, we can then treat the quinton as we would an intravalley exciton, trion, or biexciton.

Here, we have studied the energetics (in the Mott-Wannier-Keldysh) model of such possible large complexes. Specifically, we have studied the aforementioned quinton (XX^-), the donor-bound double-negative exciton (D^-X), the donor-bound quinton (D^0XX), and the (largest) donor-bound double-negative biexciton (D^-XX).⁶ Our notation in monolayers is as follows: X denotes an exciton (an electron and hole), D denotes a donor atom (an electron and a donor atom D^+). Additional electrons are inferred from any charge superscripts. In all cases, valley and spin indices are such that the complex under study is comprised entirely of distinguishable particles. Hence, the trial and exact wave functions are nodeless, and the DMC method yields a numerically exact solution to the Mott-Wannier-Keldysh model.

Taking the quinton as an example, we will give a determination of a binding energy. Given that the boundedness of the trions and biexciton are established (both empirically, and within the Mott-Wannier-Keldysh model), the only decay possibilities for the quinton are

$$\begin{aligned}
XX^- &\xrightarrow{1} X^- + X, \\
XX^- &\xrightarrow{2} X^+ + e + e, \\
XX^- &\xrightarrow{3} XX + e,
\end{aligned}
\tag{3.12}$$

⁶Results pertaining to the donor-bound double-negative biexciton are not included in Ref. [3]. There, it is claimed that such a complex was found to be unstable. After discussion, the authors of that work decided collectively not to issue an erratum correcting this minor point.

where the overset numbers label decay paths. The associated binding energies are

$$\begin{aligned}
E_{\text{XX}^-}^{\text{b}}(1) &= E_{\text{X}^-}^{\text{T}} + E_{\text{X}}^{\text{T}} - E_{\text{XX}^-}^{\text{T}}, \\
E_{\text{XX}^-}^{\text{b}}(2) &= E_{\text{X}^+}^{\text{T}} + 2E_{\text{e}}^{\text{T}} - E_{\text{XX}^-}^{\text{T}}, \\
E_{\text{XX}^-}^{\text{b}}(3) &= E_{\text{XX}}^{\text{T}} + E_{\text{e}}^{\text{T}} - E_{\text{XX}^-}^{\text{T}}.
\end{aligned}
\tag{3.13}$$

Substituting expressions for the trion and biexciton binding energies here, we then find that the favoured decay is determined by $\max(E_{\text{X}^-}^{\text{b}}, (E_{\text{X}^+}^{\text{b}} - E_{\text{X}}^{\text{b}}), E_{\text{XX}}^{\text{b}})$ (the first argument pertains to decay 1, and so on; the maximum here defines the *minimal* quinton binding energy). For the TMDs studied in this project (and for a large portion of the entire model parameter space; see Fig. 4 of Ref. [98]), it is found that decay 1, to a trion and an exciton, is most favourable. The positive trion decay is displaced by a huge margin (E_{X}^{b}), and is unlikely to ever be competitive with the negative trion and biexciton decays. The biexciton decay (3) could, in principle, become the dominant decay mechanism in families of two-dimensional materials which lie in a vastly different pocket of model parameter space (see earlier referenced figure; no two-dimensional semiconductors that we are aware of lie in this region of the parameter space).

Similar algebraic considerations are used to define the binding energies of the other, larger, complexes here studied. The results of our calculations are displayed in Table 3.2.

TMD	Binding energy (meV)			
	XX ⁻	D ⁻ X	D ⁰ XX	D ⁻ XX
MoS ₂	58.6(6)	84.4(4)	61.6(6)	58.3(5)
MoSe ₂	57.0(4)	57.9(2)	56.9(9)	78.6(3)
WS ₂	57.4(3)	59.2(4)	58.2(6)	80.4(5)
WSe ₂	52.5(7)	51.3(4)	51(1)	70.5(7)

Table 3.2: Binding energies of larger charge-carrier complexes in different TMDs. The binding energies were evaluated using the effective mass and in-plane permittivity parameters reported in bold in Table 3.1. The fitting functions of Ref. [3] were used to evaluate negative donor ion and donor-bound trion total energies.

For the case of a quinton itself, the decay mechanism which we have identified as the dominant one (to X and X⁻) means that the position of a quinton peak in PL is somewhat unambiguous: since an exciton is a decay product, the quinton peak ought to be observed at an energy which is $E_{XX^-}^b$ *lower* than the exciton line. In the cases where an exciton is not a decay product, the energy at which the observation of a PL peak would be observed is *not* simply displaced from the exciton line by the binding energy. In these cases, the peak position relative to the exciton line is determined by the de-excitonisation energy. In the case of the large complexes presented here, each of the observed decays are of the former class (with X as a daughter product), and the binding energy equals the de-excitonisation energy.

The importance of the exciton line in determining peak position in this work is only because of our choice to evaluate line positions *relative* to the exciton line (which is often most visible in experiments). The absolute energy scale for PL phenomena is in reality set by the quasiparticle energy gap.

Interestingly, after the publication of this work, an experimental group [104] claimed the observation of a five-particle complex of electrons and holes in an ultra-clean single-crystal of MoSe₂ encapsulated between graphite and hBN. However, many peaks are observed, and, despite the additional evidence that the authors were able to provide, unambiguous and unequivocal assignment of a given PL peak to a given charge-carrier complex is far from a reality.

3.1.4 Beyond the Keldysh interaction: *ab initio* dielectric functions

The Keldysh interaction is a model in which the two-dimensional susceptibility of a material is approximated by a constant. In fact, it is known that this is not the case. More realistic dielectric functions $\epsilon(\mathbf{q})$ for two-dimensional materials can in fact be calculated by various methods [105–107].⁷ Here, we consider the replace-

⁷Implicitly, any *GW* method has its own associated $\epsilon(\mathbf{q})$, defining the screened Coulomb interaction, though we won't list the plethora of *GW*-based studies on two-dimensional semiconductors here.

ment of the Keldysh interaction by an *ab initio* dielectric function of the form presented by Aghajanian, Mostofi, and Lischner (henceforth collectively referred to as “AML”), which was originally used to study the formation of bound states about charged adatoms on TMD surfaces [10]. The interaction presented by AML is parameterised (in momentum space) from first-principles, using the RPA expression with Kohn-Sham DFT orbitals and eigenvalues (as is done in forming the W in a GW calculation), from DFT calculations on MoS_2 . The AML interaction is compared to the Keldysh interaction in Fig. 3.5. The interaction potential varies significantly from the Keldysh interaction at short-range, with the strongest modulation occurring at distances of order of a single MoS_2 bond length (~ 5 a.u.), but ultimately it has the same tail at large r . This is the Coulomb limit, where charges are so far apart that polarisation of the host layer is, physically speaking, irrelevant. At short-range, the AML interaction recovers the physical divergence ($1/r$) for charges in a real (inhomogeneous) material.

Before performing any new QMC calculations using the RPA-parameterised AML interaction, we have performed QMC test calculations for a electron in a layer bound to a donor charge at a fixed distance above a layer, as a function of the effective defect charge. Our findings at the VMC and DMC levels are presented in Fig. 3.4, alongside the results of Aghajanian *et al.* (Ref. [10]).

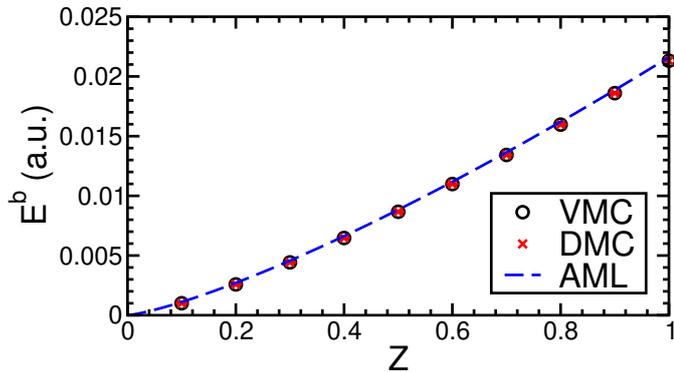


Figure 3.4: VMC and DMC results for the binding energy of a electron in a MoS_2 layer bound to a defect at $z = 2\text{\AA}$ with an effective charge Z , compared to the results of Ref. [10]. VMC and DMC error bars are shown, but are $\mathcal{O}(10^{-8}$ a.u.).

However, as mentioned in Section 2.1.1, the effective mass approximation is

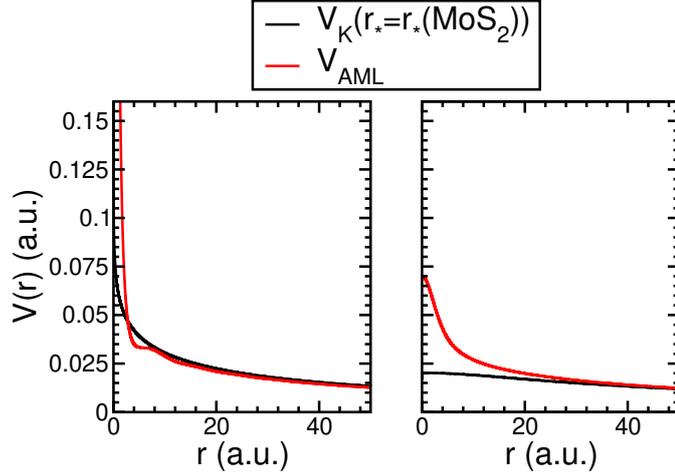


Figure 3.5: Comparison of the *ab initio* AML quasiparticle-quasiparticle interaction with the Keldysh interaction. (Left panel) Inside a layer ($z = 0$). (Right panel) Outside a layer ($z = 15$ a.u.).

only valid where charge carriers have a spatial extent well exceeding the lattice constant of their host material. That said, we should be wary of the bare singularity in the AML interaction, as if this deep potential energy feature were to cause strong over-binding of charge complexes, the validity of the effective mass approximation in the present work would be in question. In reality, the physical ($1/r$) divergence of the potential between quasiparticles in an inhomogeneous material would be compensated somewhat by short-range exchange effects. The inclusion of such effects is in itself a challenge, as has been discussed elsewhere [3, 98, 108, 109], and will not be attempted here.

As is evident in Table 3.3, this singularity *does* cause significant overbinding of charge-carrier complexes. As a means of remedying this issue, we have considered “rounding” the AML potential at a distance which is approximately a single bond length ($d_{\text{Mo-S}}$) – such that no bare singularity is present in our calculations. This procedure takes some inspiration from considerations of the success of the Keldysh model. There, the singularity is logarithmic, and has a weakened impact on binding. Rounding takes this to the extreme, completely (and somewhat arbitrarily) removing the singularity. As a test of the unreasonableness of this approach, we have varied the rounding distance to half, and twice, a bond length. Whilst there remains some variation in the AML binding energies, they are very small (some-

times not significant, as in the case of the negative trion), and they are far smaller than those witnessed in the replacement from the Keldysh interaction to the full AML interaction with its bare Coulomb singularity.

$v(r)$	Binding energy (meV)		
	X	X ⁻	XX
Bare AML	765.45(2)	77.7(7)	184.5(7)
Keldysh	546.5	35.0	23.5
R-AML ($0.5 \times d_{\text{Mo-S}}$)	492.47(3)	30.9(2)	20.4(4)
R-AML ($d_{\text{Mo-S}}$)	483.8(4)	30.7(4)	19.4(8)
R-AML ($2 \times d_{\text{Mo-S}}$)	454.5(1)	30.0(3)	16.0(4)

Table 3.3: Binding energies of excitonic charge complexes in MoS₂ as modelled by different charge-carrier interaction potentials. “R-AML” is the rounded AML interaction, with the rounding distance given in brackets. Keldysh values are obtained from fits given in Ref. [3].

3.2 Charge complexes in TMD heterobilayers

3.2.1 Bilayer Keldysh interaction

In a multilayer of two-dimensional semiconductors, the effective interaction between charges is once again altered. The (displaced-)Coulomb interaction ceases to provide a completely adequate description for the interaction between charges in different layers, while the Keldysh interaction ceases to provide a completely adequate description for the interaction between charges residing in the same atomic layer. A general derivation of the so-called “multilayer” interaction is given in Section 3.3, however, it is possible to analytically determine the Fourier components of the intralayer and interlayer charge-charge interaction potentials in the case of a bilayer.

Following a procedure similar to that in the preceding subsection, one can derive expressions for the electrostatic potential in space due to a point charge and the subsequent polarisation of a pair of two-dimensional polarisable planes. Letting the first of these layers reside at $z = 0$, with polarisability κ_1 , and the

second reside at $z = d$, with polarisability κ_2 ,⁸ we find (for a test charge density $\rho(\mathbf{s})\delta(z - d)$, isolated to layer 2) that

$$\phi_{\text{inter}}(\mathbf{q}) = \phi(\mathbf{q}, z = 0) = \frac{2\epsilon\rho(\mathbf{q})}{q [(2\epsilon + \kappa_1q)(2\epsilon + \kappa_2q)e^{q\alpha d} - \kappa_1\kappa_2q^2e^{-q\alpha d}]}, \quad (3.14)$$

$$\phi_{\text{intra}}(\mathbf{q}) = \phi(\mathbf{q}, z = d) = \frac{[(2\epsilon + \kappa_1q)e^{q\alpha d} - \kappa_1qe^{-q\alpha d}] \rho(\mathbf{q})}{q [(2\epsilon + \kappa_1q)(2\epsilon + \kappa_2q)e^{q\alpha d} - \kappa_1\kappa_2q^2e^{-q\alpha d}]}, \quad (3.15)$$

where we have identified $\phi_{\text{inter}}(\mathbf{q})$ and $\phi_{\text{intra}}(\mathbf{q})$ to reflect the fact that the interaction potential between different layers is, in this case,⁹ the value of the electrostatic potential in layer 1 due to the presence of a charge density isolated to layer 2, and that the interaction potential for charges in the same layer is the value of the electrostatic potential in layer 2 due to the charge density present in layer 2. The potential between charges which are both in layer 1 may be obtained by a re-arrangement of variables ($\kappa_1 \leftrightarrow \kappa_2$) in Eqn. (3.17). The factor $\alpha = \sqrt{\epsilon_{\parallel}/\epsilon_{\perp}}$ reflects the degree of anisotropy of the surrounding dielectric medium, and will be neglected in the remainder of this section: we consider encapsulation in anisotropic dielectric media only in the case of hexagonal boron nitride, however, for bulk hBN $0.71 < \alpha < 0.95$ [110–113], and uncertainties in experimental estimates of interlayer distances in bilayer heterostructures are overwhelmingly likely to be the dominant source of uncertainty in d anyway. The cumulative effect of this small perturbation *may* be important in the case of multilayer systems, and there (in Section 3.3) we include it.

As before, after defining $r_{*,i} = \kappa_i/2\epsilon$, and with α taken as unity, the bilayer potentials read

$$\phi_{\text{inter}}(\mathbf{q}) = \frac{1}{2\epsilon q [(1 + r_{*,1}q)(1 + r_{*,2}q)e^{q\alpha d} - r_{*,1}r_{*,2}q^2e^{-q\alpha d}]}, \quad (3.16)$$

$$\phi_{\text{intra}}(\mathbf{q}) = \frac{[(1 + r_{*,1}q)e^{q\alpha d} - r_{*,1}qe^{-q\alpha d}]}{2\epsilon q [(1 + r_{*,1}q)(1 + r_{*,2}q)e^{q\alpha d} - r_{*,1}r_{*,2}q^2e^{-q\alpha d}]}. \quad (3.17)$$

Whilst in the case of a polarisable monolayer, one can invert the Fourier compo-

⁸The consequence of this would be to add a term $-\kappa_2[\nabla_{\parallel}\phi(\mathbf{s}, z = d)]\delta(z - d)$ to \mathbf{D} .

⁹Up to a factor of particle charge.

nents to obtain $\phi(\mathbf{s})$, no such analytical expression seems procurable for either of ϕ_{inter} or ϕ_{intra} . The same is true of the multilayer potentials, and so we defer the discussion of obtaining real-space potentials numerically to Section 3.3.1.

3.2.2 Free and donor-bound complexes in hBN-encapsulated MoSe₂/WSe₂ heterobilayers

The van der Waals heterobilayer is the simplest possible van der Waals heterostructure – two different monolayers are exfoliated from their parent van der Waals crystals, and are stacked one atop the other. In analogy with the previous section on indirect trions in coupled quantum wells, but with no need of an applied external field, van der Waals heterobilayers of vertically stacked WSe₂ and MoSe₂ monolayers have been observed to feature a staggered type-II band alignment: it is energetically favourable for electrons and holes to become spatially separated over the two layers of the heterobilayer [114, 115], as depicted in Fig. 3.6. Type-II band alignment is also observed in the sulphides [116], and in the mixed MoS₂/WSe₂ system [117].

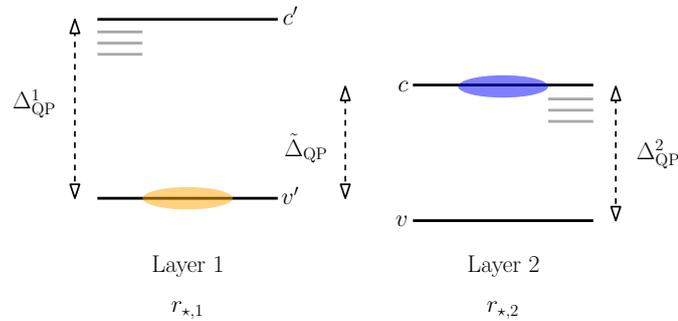


Figure 3.6: Schematic diagram depicting type-II (staggered) band alignment for a heterobilayer system. Orange (blue) shaded pockets highlight bands which host holes (electrons). In the MoSe₂/WSe₂ case, “Layer 1” is identified as MoSe₂ (where holes dominantly reside, in band v'), and “Layer 2” as WSe₂ (where electrons dominantly reside, in band c). We have also labelled the minority bands c' , v , where electrons and holes *may* reside, and do in some of our calculations.

We have performed VMC and DMC calculations for complexes of distinguishable charge carriers and stationary donor and acceptor charges in a model for hBN-encapsulated MoSe₂/WSe₂ heterobilayers. We have considered charges in-

teracting via the “full” bilayer potentials (Eqn. (3.16) and Eqn. (3.17)), and an approximate small- q Keldysh form of the potential (whose forms are given, and discussed, in Appendix B). Such a physical system was fabricated and characterised in Ref. [6], with theoretical contributions from some of the contents of the present section.

We have assumed the case of near-alignment, such that it is reasonable to ignore the effects of any superlattice perturbations. But for completeness, we note that collaborators have studied the twist-angle dependence of radiative and non-radiative decay rates in Ref. [118]. This said, it is generally expected that for heterobilayer systems, the small but inevitable misalignment has a particularly important consequence: it prevents direct recombination of excitons formed from constituent quasifermions which occupy different layers. This means that emphasis must be placed on other recombination mechanisms, *i.e.* those involving defect charges (scattering on which permits recombination without violation of momentum conservation) and those involving other sources of momentum and energy. One such source would be lattice phonons, phonon-assisted recombination was studied in Ref. [118]). In cases where there is no lattice mismatch (homomultilayers, for example), no such restriction applies, and recombination can in principle occur again by direct means.

Our naming scheme for complexes in heterobilayers is as follows. Generally, primes (') indicate that a charge-carrier is in the MoSe₂ layer; otherwise the charge carrier is in the WSe₂ layer. The subscripts c and v indicate whether charge carriers are electrons (c) or holes (v). Subscripts on a roman letter defining a charge-carrier complex indicate the charge-carrier content of that complex. For example, the interlayer exciton formed from an electron in the MoSe₂ layer and a hole in the WSe₂ layer is denoted $X_{vc'}$. Donor ions (D) are always assumed to be in the MoSe₂ layer (supplying excess electrons there), while acceptor ions (A) are always assumed to be in the WSe₂ layer (supplying excess holes there).

Table 3.4 displays the total energies of charge-carrier complexes in the hBN/

MoSe₂/WSe₂/hBN heterostructure. For completeness we include results in which the electrons are found in either layer; however, the results of immediate relevance to the present thesis are those for which the electrons are all found in the MoSe₂ layer. DMC results for two-particle complexes agree with calculations performed using Mathematica’s finite-element method [119], which were performed as a test of our implementation of the bilayer Keldysh interaction in the CASINO QMC code. The dominant decay routes for charge-carrier complexes in the heterobilayer system, and their associated binding energies, are given in Table 3.5.

It is clear from Table 3.5 that the approximate Keldysh interaction of Appendix B performs well when calculating binding energies, provided the decay does not involve significant changes to short-range pair distributions. In this case, any error related with the spurious singularity (and associated wave function cusp) may cancel when one calculates a binding energy. If this is not the case, larger errors may manifest. The worst-case is that of the interlayer exciton, whose binding energy is simply equal to its total energy, and hence does not benefit from any error cancellation. Here, the binding energy is overestimated by 23% when the approximate Keldysh interaction is used.

We are not aware of any published experimental results on donor-bound interlayer complexes in heterobilayers, but we discuss the validity of our results for intralayer complexes in Section 3.2.2.2.

An interesting conclusion to note is that various possible complexes in the hBN-encapsulated heterobilayer are very weakly bound, with binding energies often around or lower than 10 meV. Such complexes are unlikely to feature as sharp, isolated peaks in the PL spectrum of an encapsulated heterobilayer held under normal laboratory conditions. At 300K, the typical thermal energy scale $k_B T$ is around 25 meV, more than enough to completely smear any physics taking place at a 10 meV energy scale. At 4K (a common cryogenic temperature, reached in the laboratory with liquid He cooling), the thermal energy is a fraction of an meV, and the prospects for observing some of these bound charge-carrier complexes improve

Complex	DMC total energy (meV)	
	Approximate Keldysh	Full bilayer
$X_{vc'}$	-103.958669(5)	-84.232(1)
$X_{vc'e'}^-$	-108.1967(4)	-88.32(3)
$X_{vvc'}^+$		-88.12(2)
$X_{vc'}X_{vc'}$		unbound
$X_{vc'}X_{vc'e'}^-$		unbound
$D_{c'}^0$	-163.2478711(5)	-229.03306(1)
$D_{c'e'}^-$	-176.9426(3)	-249.60(2)
$D_{c'}^0h_v$	-163.4819(8)	unbound
$D_{c'}^0X_{vc'}$	-278.73(2)	-335.781(4)
$D_{c'}^0X_{vvc'}^+$		-340.891(6)
$D_{c'e'}^-X_{vc'}$	-292.83(1)	-343.26(3)
$D_{c'}^0X_{vc'}X_{vc'}$		unbound
$D_{c'e'}^-X_{vc'}X_{vc'}$		-430.9(1)
A_v^0		-205.24083(1)
A_{vv}^+		-223.56(1)
$A_v^0e_{c'}$		unbound
$A_v^0X_{vc'}$		-309.411(4)
$A_v^0X_{vc'e'}^-$		-315.021(8)
X_{vc}	-114.601814(1)	-140.4303329(4)
D_c^0	-124.890219(9)	-102.5996(7)
$X_{vcc'}^-$	-120.6018(5)	unbound
X_{vcc}^-	-123.7189(5)	-152.25(1)
$D_{cc'}^-$	-165.8499(5)	unbound
D_{cc}^-	-129.3199(9)	unbound
D^+X_{vc}	-133.758(2)	-141.716(8)
$D_{c'}^0X_{vc}$	-279.776(5)	unbound
$D_{c'e'}^-X_{vc}$	-301.81(1)	unbound
$D_{c'}^0X_{vcc}^-$	-295.00(1)	unbound

Table 3.4: DMC total energies of various charge-carrier complexes in the hBN/MoSe₂/WSe₂/hBN heterostructure calculated using the monolayer Keldysh approximation to the bilayer potential (Eqn. (B.3) and Eqn. (B.4)) and using the full bilayer interaction (Eqn. (3.16) and Eqn. (3.17)). Interlayer complexes in which all the electrons are in the MoSe₂ layer and all the holes are in the WSe₂ layer are listed in the upper section of the table; complexes in which some of the electrons are in the WSe₂ layer are listed in the lower section of the table.

Decay route	DMC binding energy (meV)	
	Approximate Keldysh	Full bilayer
$X_{vc'c'}^- \rightarrow X_{vc'} + e_{c'}$	4.2380(4)	4.09(3)
$X_{vvc'}^+ \rightarrow X_{vc'} + h_v$		3.89(2)
$X_{vc'} X_{vc'}$		unbound
$X_{vc'} X_{vc'}^-$		unbound
$D_{c'c'}^- \rightarrow D_{c'}^0 + e_{c'}$	13.6948(3)	20.57(1)
$D_{c'}^0 h_v \rightarrow D_{c'}^0 + h_v$	0.2340(8)	unbound
$D_{c'}^0 X_{vc'} \rightarrow X_{vc'} + D_{c'}^0$	11.52(2)	22.516(4)
$D_{c'}^0 X_{vvc'}^+ \rightarrow D_{c'}^0 + X_{vvc'}^+$		23.74(2)
$D_{c'c'}^- X_{vc'} \rightarrow X_{vc'} + D_{c'c'}^-$	11.93(1)	9.43(4)
$D_{c'}^0 X_{vc'} X_{vc'}$		unbound
$D_{c'c'}^- X_{vc'} X_{vc'}$		3.3(2)
$A_{vv}^+ \rightarrow A_v^0 + h_v$		18.32(1)
$A_v^0 e_{c'} \rightarrow A_v^0 + e_{c'}$		unbound
$A_v^0 X_{vc'} \rightarrow A_v^0 + X_{vc'}$		19.938(4)
$A_v^0 X_{vc'c'}^- \rightarrow A_v^0 + X_{vc'c'}^-$		21.46(3)
$X_{vcc'}^- \rightarrow X_{vc} + e_{c'}$	6.0000(5)	unbound
$X_{vcc}^- \rightarrow X_v + e_c$	9.1170(5)	11.83(1)
$D_{c'}^0 X_{vc} \rightarrow X_{vc} + D_{c'}^0$	1.926(5)	unbound
$D_{c'c'}^- X_{vc} \rightarrow X_{vc} + D_{c'c'}^-$	10.26(1)	unbound
$D_{cc'}^- \rightarrow D_{c'}^0 + e_c$	2.6020(5)	unbound
$D_{c'}^0 X_{vcc}^- \rightarrow D_{c'}^0 + X_{vcc}^-$	8.03(1)	unbound
$D_{cc}^- \rightarrow D_c^0 + e_c$	4.4297(9)	unbound
$D^+ X_{vc} \rightarrow D^+ + X_{vc}$	19.156(2)	1.286(8)

Table 3.5: Decays of complexes and their associated binding energies in hBN/MoSe₂/WSe₂/hBN.

enormously. Owing to the typically very small binding energies of the majority of the interlayer charge complexes, we have neglected to calculate, where even possible, their de-excitonisation energies.

3.2.2.1 Sensitivity to model parameters

We have performed test calculations to determine the sensitivity of the DMC-evaluated $D_{c'}^0 X_{vc'}$ binding energy to the model parameters $m_{c'}$, m_v , $r_{\star,1}$, $r_{\star,2}$, and d in the hBN-encapsulated heterobilayer. Note that $r_{\star,1}$ and $r_{\star,2}$ are here the screening-length parameters in vacuum, so that the screening lengths in a dielectric environment are $r_{\star,1}/\epsilon$ and $r_{\star,2}/\epsilon$. We find that, upon variation of each of the parameters in turn by $\pm 10\%$ from the values listed in Table 3.1, the $D_{c'}^0 X_{vc'}$ binding energy never varies by more than 8% (1.8 meV), as shown in Table 3.6. The derivatives of the binding energy with respect to the parameters were evaluated numerically by the central difference approximation. In excitonic units, it is easy to show that the derivative of a binding energy E^b with respect to the dielectric constant is

$$\frac{\partial E^b}{\partial \epsilon} = -\frac{1}{\epsilon} \left(2r_{\star,1} \frac{\partial E^b}{\partial r_{\star,1}} + 2r_{\star,2} \frac{\partial E^b}{\partial r_{\star,2}} + d \frac{\partial E^b}{\partial d} + 2E^b \right), \quad (3.18)$$

allowing us to evaluate the sensitivity of the binding energy with respect to the dielectric constant. We find that the binding energies are most sensitive to the screening parameter $r_{\star,2}$, followed by the dielectric constant ϵ , followed by the electron mass $m_{c'}$, and that the sensitivity to the layer separation d is relatively weak. The sensitivity to the screening parameter $r_{\star,1}$ and to the hole mass m_v are particularly weak in the present case, because only one hole resides in the WSe₂ layer.

We have also performed DMC calculations with $\epsilon = 4.5$ (instead of $\epsilon = 4$), finding that the $X_{vc'}$, $X_{vc'}^-$, $D_{c'}^0$, and $D_{c'}^0 X_{vc'}$ binding energies are 76, 3.8, 207, and 20.7 meV, respectively. This directly confirms that the sensitivity to the precise value of the dielectric constant of the environment is relatively weak. The value of

P	δP	$E_{D_c^0 X_{v_c'}}^b$ (meV)	$P \cdot \partial E_{D_c^0 X_{v_c'}}^b / \partial P$
$m_{c'}$	+10%	23.27(1)	7.866 meV
	-10%	21.70(1)	
m_v	+10%	22.71(1)	2.108 meV
	-10%	22.29(1)	
$r_{*,1}$	+10%	22.46(1)	-0.5196 meV
	-10%	22.52(1)	
$r_{*,2}$	+10%	20.96(1)	-35.89 meV
	-10%	24.31(1)	
d	+10%	22.93(1)	4.568 meV
	-10%	22.02(1)	
ϵ	+10%	~ 20.97	-15.48 meV
	-10%	~ 24.06	

Table 3.6: Changes in binding energies per fractional parameter changes for DMC binding energies $E_{D_c^0 X_{v_c'}}^b$ of the interlayer donor-bound trion under variations δP of different model parameters P . With the exception of the dielectric constant ϵ , the parameters are varied by $\pm 10\%$ about the values listed in Table 3.1 and the central difference approximation is used to estimate the derivative with respect to the parameter value. The binding energy when all the parameters take the values listed in Table 3.1 is $E_{D_c^0 X_{v_c'}}^b = 22.516(4)$ meV. Note that $r_{*,1}$ and $r_{*,2}$ are here the screening lengths for monolayers isolated in vacuum. The derivative of the binding energy with respect to the dielectric constant ϵ was evaluated by the chain rule, as described in the text.

$\partial E_{D_c^0 X_{vc}'}^b / \partial \epsilon$ found by the forward difference approximation is -3.63 meV, which is in reasonable agreement with the value obtained using the chain rule, reported in Table 3.6.

3.2.2.2 Intralayer binding energies: comparison with experiment

For TMD monolayers, experimental agreement with QMC calculations of the binding energies of charge-carrier complexes employing the Keldysh interaction has previously been addressed in Refs. [98] and [3]. There, trion binding energies in particular are found to be in exceedingly good agreement with experimental results.

Relatively few experimental studies of charge-carrier complexes in heterobilayers have been performed to date. Ceballos *et al.* studied a SiO₂/MoSe₂/MoS₂/vacuum sample [120], performing PL measurements on monolayer MoSe₂, monolayer MoS₂, and heterobilayer MoSe₂/MoS₂ regions of their sample (where the two monolayers overlap). Gong *et al.* studied a SiO₂/MoS₂/WS₂/vacuum sample [121], again performing PL measurements on each of the three distinct surface regions of their sample. Both experiments, although studying different TMD bilayers prepared by different means, observed rather small shifts in the dominant intralayer exciton lines on moving from monolayer regions to bilayer regions. Our heterobilayer results of Table 3.5, in conjunction with monolayer binding-energy fitting formulas presented in Ref. [3] provide further support for this claim. The intralayer exciton energy reported in Table 3.4 for an exciton X_{vc} in the WSe₂ layer of a hBN/MoSe₂/WSe₂/hBN heterostructure is -140.4 meV, whereas the exciton total energy in monolayer WSe₂ encapsulated in hBN is -159.7 meV, according to the monolayer fitting formula of Ref. [3]. The intralayer negative trion X_{vcc}^- binding energy reported in Table 3.5 is 11.8 meV, whereas the fitted negative-trion binding energy in monolayer WSe₂ encapsulated in hBN is 13.6 meV. The intralayer exciton energy differs by about 19 meV from the monolayer exciton energy, while the intralayer trion binding energy differs by only 2 meV from the monolayer re-

sult. We therefore conclude by noting that intralayer complex binding energies in a heterobilayer are very similar to the monolayer binding energies in the host layer, and also that the validity of our model may be judged at least in part by examining previously reported results [3, 98] for TMD monolayers.

3.3 Complexes in multilayers of two-dimensional semiconductors

The addition of further polarisable layers, into the tri-, quadri-, ..., n -layer domain, further complicates the expressions for the electrostatic potential within a given layer of a multilayer heterostructure.

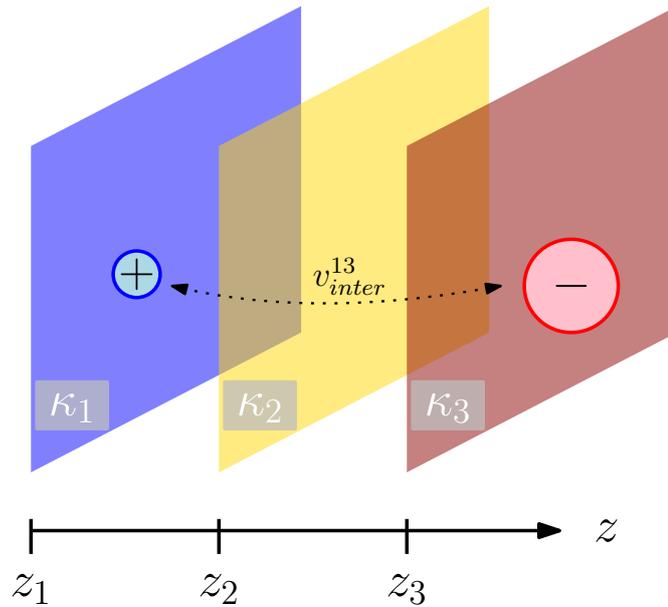


Figure 3.7: Schematic diagram for an $n = 3$ heterostructure of two-dimensional semiconductors. An interlayer exciton is also depicted.

Consider now a multilayer heterostructure comprised of n layers, each having polarisability κ_i and at fixed z -coordinate d_i ($i = 1, \dots, n$). A representation of the three-layer case is given in Fig. 3.7. Further, let us restrict to the case where $\epsilon = 1$.¹⁰ Let us denote a test charge density in the j -th layer of the heterostructure

¹⁰In the case of the bilayer, the inclusion of an all-permeating ϵ over-screens the interlayer interaction slightly, as the vacuum layer-spacing is replaced by that of a medium with ϵ which may be greater than unity. In the present case, this would in general lead to a much more severe error.

as $\rho_j(\mathbf{r}) = \rho_j(\mathbf{s}, d_j) = \rho_j(\mathbf{s})\delta(z - d_j)$. The electric displacement field in space due to this charge density, and due also to the combined contribution of all of the polarisation fields from layers in the heterostructure is given by

$$\mathbf{D}(\mathbf{r}) = -\nabla\phi(\mathbf{s}, z) - \sum_{i=1}^n \kappa_i [\nabla_{\parallel}\phi(\mathbf{s}, d_i)] \delta(z - d_i). \quad (3.19)$$

As earlier, application of Gauss' law yields

$$\rho_j(\mathbf{s})\delta(z - d_j) = -\nabla^2\phi(\mathbf{s}, z) - \sum_{i=1}^n \kappa_i \nabla_{\parallel}^2\phi(\mathbf{s}, d_i)\delta(z - d_i), \quad (3.20)$$

and subsequent Fourier transformation gives

$$\rho_j(\mathbf{q})e^{-q|z-d_j|} = 2q\phi(\mathbf{q}, z) + q^2 \sum_{i=1}^n \kappa_i \phi(\mathbf{q}, d_i)e^{-q|z-d_i|}. \quad (3.21)$$

At this stage, it is not generally possible to analytically determine the $\phi(\mathbf{q}, d_i)$. We have so far studied two cases where it is possible ($n = 1, 2$).¹¹ However, considering the evaluation of Eqn. (3.21) at $z = d_1, d_2, \dots, d_n$, we can form a matrix equation for the electrostatic potential within the layers as

$$\rho_{ij}(\mathbf{q}) = \sum_{k=1}^n M_{ik}(\mathbf{q})\phi_{kj}(\mathbf{q}), \quad (3.22)$$

where the matrices ρ and M are defined as

$$\rho_{ij}(\mathbf{q}) = e^{-q|z_i - z_j|}, \quad (3.23)$$

$$M_{ij}(\mathbf{q}) = 2q\delta_{ij} + q^2\kappa_j e^{-q|z_i - z_j|}, \quad (3.24)$$

and where the $\phi_{ij}(\mathbf{q})$ are the Fourier components of the electrostatic potential in layer i (j) due to the presence of a charge in layer j (i), and due to the polarisation of the remainder of the heterostructure. Notice that in the formation of this matrix

¹¹Such analytical solution also appears to be possible for tri-layers. For cases with $n > 4$ this need not be the case, as the solution would involve the solution of a quintic characteristic polynomial.

problem, we have specialised to the case of point charges in layers.

Such linear problems as in Eqn. (3.22) can be solved easily with existing computational routines. In LAPACK [122], for example, the relevant routine is `dgesv`.

3.3.1 Obtaining real-space interaction potentials

Once the Fourier components of some given interaction potential are known, the method of Ogata [123] is used to perform the zeroth-order Hankel transform (spherically-symmetric two-dimensional Fourier transform) necessary for the expression of the potentials in real-space. Firstly, however, we extract the dominant singular behaviour from the function which is to be transformed. Such a procedure increases the stability of the integration method, but also allows for the introduction of a cutoff radius on the corrective term.

As an explicit example, suppose we wish to evaluate the intralayer interaction in a few-layer heterostructure of two-dimensional materials. At sufficiently long-range in real-space, the interaction between charges in the layer is of Coulomb form.¹² The Fourier transform of $\phi_{\text{intra}}(\mathbf{q}) - \phi_{\text{Coulomb}}(\mathbf{q})$ is then finite-ranged in real-space, and has no singularity as $q \rightarrow 0$. If we take the Hankel transform of the difference, and add in $1/r$ (the analytical Hankel transform of $1/q$), we obtain the intralayer interaction in real-space, without numerical instabilities incurred by the inclusion of the singular term in our Hankel transform.

3.3.2 Excitons and Biexcitons in few-layer TMDs

We have performed VMC and DMC calculations to extract the binding energies of intra and interlayer excitons and biexcitons in few-layer TMDs, modelling a variety of dielectric media. We have considered the following treatments of dielectric screening:

- (v) A homostructure in vacuum.

¹²Indeed, *any* of the inter/intralayer potentials are of Coulomb form at long-range.

- (1) A homostructure in an isotropic dielectric, having $\epsilon_r = 4$.
- (2) A homostructure in an anisotropic dielectric material, with dielectric tensor $\tilde{\epsilon}_r = \text{diag}(4.69, 4.69, 2.65)$.¹³
- (3) A homostructure properly encapsulated with monolayer hBN, having $r_*(\text{hBN}) = 11.6$ a.u.¹⁴

A schematic diagram displaying the relative differences of these screening regimes is given in Fig. 3.8.

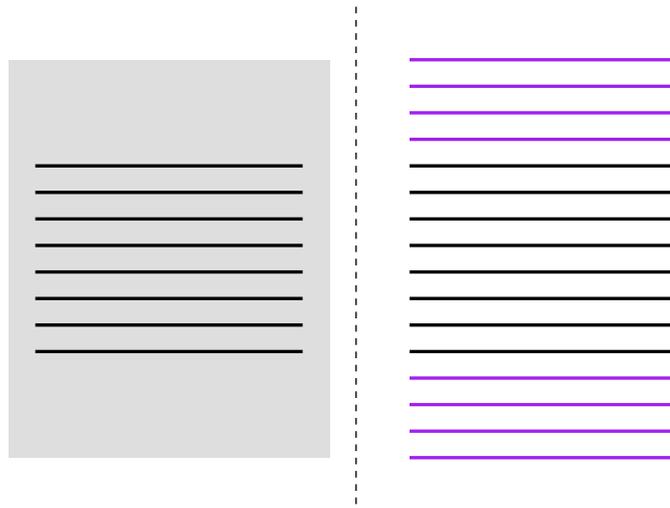


Figure 3.8: Encapsulation models for an $n = 8$ homostructure. (left) dielectric screening model in cases (1) and (2), grey space corresponds to dielectric medium. (right) dielectric screening model in case (3), with extra (purple) hBN layers for visibility.

The case of a TMD homostructure in vacuum is interesting for two reasons. Firstly, it is not inconceivable that experimentalists will succeed in replicating such a physical system with suspended multilayers of TMDs, fabricated directly by mechanical exfoliation in the first instance. In this case, the results of such a study could be useful in interpreting any photoluminescence experiment performed on such a sample. Secondly, without knowing *a priori* what families of charge-carrier complexes are likely to be bound in homostructures in general, the vacuum

¹³Taken as the electronic contribution to the dielectric tensor, as determined by density functional perturbation theory calculations on bulk hBN, with the PBE exchange-correlation functional.

¹⁴Determined by the prescription presented in the Supplementary Material of Ref. [99], with lattice constants taken from experiment [11].

case provides a useful baseline. If a charge carrier complex is unbound in the vacuum case, then the introduction of any dielectric medium is unlikely to lead to the formation of a bound state. The additional screening supplied by a dielectric medium of some kind can only weaken inter-particle interactions, and lower carrier complex binding energies.

The first non-trivial choice (1) is commonly made by other authors [124, 125] for nanomaterials encapsulated in hBN. The second choice is the more realistic extension, where the anisotropy of bulk hBN is taken into account. The final choice corresponds to a new treatment, wherein the multilayer Keldysh interaction is used to model the situation where individual hBN layers are placed above and below the heterostructure. Here, we do not consider the presence of charges inside the hBN layers, but include their effects in the formation of the multilayer interaction potentials between charges which reside inside the TMD homostructure. Note that in cases 1 and 2, for systems other than monolayers, we incur additional unphysical screening of interactions between charges due to the space between two-dimensional layers being treated as a dielectric medium (and not vacuum).

In the latter case, we have chosen to approximate the hBN-TMD interlayer distance as the average of the two bulk out-of-plane lattice constants, following Ref. [126]. An accurate treatment of such binding physics is, as we have previously stated, computationally challenging, but also unlikely to be important in forming qualitative conclusions about charge-carrier complexes in TMD homostructures.

The full results of each of our studies are given in Tables C.1, C.2, C.3, C.4 and C.5 of Appendix C.

3.3.2.1 Excitons

A graphical depiction of the binding energies of excitons in $n = 1, \dots, 5$ -layer homostructures (in vacuum) is given in Fig. 3.9.

As is evident from Fig. 3.9 and Tables C.1 to C.5, there is a clear separation of energy scales for excitons bound in few-layer homostructures. The most well-

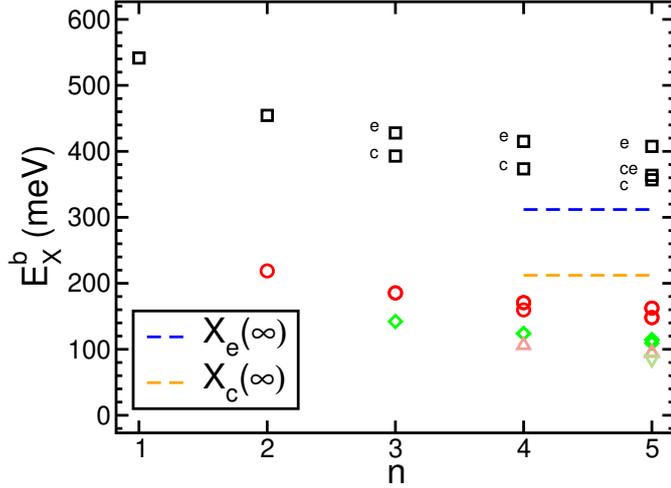


Figure 3.9: Binding energies of excitons in few-layer MoSe₂ homostructures. The upper branch of black squares represent intralayer excitons, with “c”, “e”, and “ce” labelling excitons in central, edge, and centre-edge layers, respectively. $X_{e(c)}(\infty)$ is defined in the text. Red circles, green diamonds, light-red upward-point triangles, and a single light-green downward-pointing triangle label interlayer complexes separated by 1, 2, and 3 empty MoSe₂ layers. The light-green downward-pointing triangle, for example, is the exciton formed in a five-layer system with electron and hole occupying opposite edge layers.

bound excitons are intralayer, and the least well-bound are interlayer. Interlayer excitons have, as one might expect, binding energies that decrease with the number of unoccupied medial layers (black to red to green symbols in Fig. 3.9). $X_{e(c)}(\infty)$ is the energy of an exciton at the edge (centre) of a bulk MoSe₂ system, with ϵ_{eff} as evaluated using a fitting formula from Ref. [3], and discussed in Section 3.3.2.3.

From the point of view of interpreting photoluminescence spectra of few-layer homostructures in vacuum, we expect somewhat generally to witness the formation of two (as opposed to one) bands of dominant excitonic emission. There are those attributed to intralayer complexes, and those attributed to interlayer complexes. The overall energetics of few-layer PL is then effectively similar to that in bilayers, which were studied in detail in Section 3.2.

There are however a few specific differences between bilayers and higher few-layer homostructures. The first is that one does not expect the need for consideration of the interlayer gap $\tilde{\Delta}_{\text{QP}}$, such that all excitonic lines appear from approximately the same reference point. This reference point, the quasiparticle gap of the homostructure, is however known to characteristically decrease with

n , and also to undergo a direct-to-indirect transition on increase of n [127, 128]. Whilst this does not affect the binding energetics of our model directly, the change in the nature of the gap and the bands from which it originates mean that the optical activity of any exciton will evolve with layer number (in order for recombination to be momentum-conserving), and also that the effective mass parameters featuring in our model may need adjustment. Ultimately, the dominant changes in exciton lines are likely to originate from the change in magnitude and nature of the gap. In our previous test for heterobilayers (see Table 3.6), effective mass changes of order 10% were shown to yield small changes in binding energies of around an meV.

3.3.2.2 Biexcitons

We will now examine closely the behaviour of the most relevant biexcitonic complexes from the point of few-layer systems, those which are entirely intralayer (originating from charges in the same layer). Our data for such complexes are given in Table 3.7, where biexcitons $XX_p(n)$ are indexed by their positions (p) in a homostructure, and the number of layers n in the homostructure. Our remaining biexciton data are presented in Appendix C.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
XX(1)	22.9(1)	15.356(5)	16.297(2)	22.26(9)
XX(2)	22.34(3)	14.574(3)	15.811(5)	21.94(2)
XX _c (3)	92.44(1)	14.491(3)	15.770(4)	21.94(1)
XX _e (3)	21.98(9)	14.494(4)	15.771(4)	21.84(8)
XX _c (4)	105.59(1)	14.487(3)	15.763(4)	21.92(1)
XX _e (4)	22.34(3)	14.491(4)	15.765(4)	21.94(3)
XX _c (5)	123.50(2)	14.469(4)	15.766(3)	21.93(1)
XX _{ce} (5)	110.12(2)	14.474(6)	15.766(3)	21.93(1)
XX _e (5)	22.3(3)	14.456(2)	15.765(4)	21.91(1)

Table 3.7: Binding energies in meV of intralayer biexcitons in few-layer TMD homostructures. Indices of c , e and ce denote biexcitons confined to central, edge, and centre-edge (that between the centre and the edge) layers.

For monolayer and bilayer MoSe₂ in vacuum firstly, we observe a non-interesting,

perhaps intuitive, behaviour. The binding energies $XX(1)$ and $XX(2)$ are very similar, but $XX(2)$ is lower by a statistically significant margin. This seems physically reasonable owing to the softening of the electron-hole interaction upon the addition of an empty polarisable layer to the homostructure. Interestingly, however, in the vacuum case the biexciton binding energies start to behave somewhat strangely for $n = 3$ and beyond, in cases where the biexciton resides in a centre or centre-edge layer. For instance, consider $XX_c(3)$, which has a binding energy of 92.44(1) meV, in contrast with $XX_e(3)$, which has a binding energy of 21.98(9) meV and appears to continue the smooth trend in binding energies from monolayer, to bilayer, to trilayer. The source of this discrepancy is evident from Fig. 3.9, and corresponds primarily to the loss in stability of the exciton daughter products. The biexciton total energy is less sensitive to the nature of the electron-hole interaction than the exciton total energy: what is lost in softening of the electron-hole attraction is largely made up for in the accompanied softening of the electron-electron and hole-hole repulsion. Whilst there are more attractive interactions in the biexciton from a combinatorics standpoint, screening disproportionately affects long-range interactions (which, for a stable biexciton, are the repulsive interactions). The exciton energy, however, is very sensitive to the nature of the dielectric screening and no cancellation effect occurs. The energy difference between the edge and centre intralayer excitons in the $n = 3$ homostructure is around 35 meV, and given that $E_{XX}^b = 2E_X - E_{XX}$, the 70 meV discrepancy in biexciton binding energies is clearly accounted for. Such an effect is less noticeable, even imperceptible, when the nature of the dielectric screening makes the particular layer in which a charge resides less important (in any of the other cases we have studied).

This behaviour forces a re-evaluation of the picture of the biexciton as a bound state of “two excitons”. Were this the case, the biexciton binding energy would not be so sensitive as it appears to the nature of the dielectric screening, and would monotonically decrease with the exciton binding energies (the screening of the dipole-dipole interaction between distinct excitons would be the driving force

for lowering of the binding energy). Instead, biexcitons would more aptly be described as the four-particle bound state that they are, with a character which does not simply originate from a favourable spatial orientation of two well-bound constituent excitons. This effect is a very sensitive one, however, as evidenced by the results we have presented for environment (3) in Table 3.7. Encapsulation in a single monolayers of hBN is enough to remove statistically significant changes in intralayer biexciton binding energies as one considers different numbers of layers, and placements of charges within those layers. The difference between the biexciton binding energies $XX(1)$ in environment (3) and $XX_c(3)$ in vacuum might be expected to be similar on the basis of the geometrical similarity of the two systems, however, the screening parameters of the two layers surrounding the central MoSe₂ layer are wildly different ($r_*(\text{hBN}) = 11.6$ a.u., $r_*(\text{MoSe}_2) = 75.19$ a.u.), and it is observed again that a large difference in biexciton binding energies exists between these similar trilayer systems. The arguments made before still apply – biexcitons are not weakly interacting excitons, and exciton energies are highly sensitive to the particulars of the dielectric environment. To further make this point, we note that the variation in the relevant intralayer exciton energies between the two systems is substantial (around 100 meV).

Complete insight into these exciton energies is given in Appendix C, but the same effect may also be observed graphically in Fig 3.9. Consider the move from the $n = 1$ exciton to the $n = 3$ (centre-layer; labelled “c” in Fig. 3.9) exciton. There, the addition of a single MoSe₂ layer (either side of the monolayer) corresponds to a sizeable decrease in the (central) intralayer exciton binding energy of almost 150 meV.

3.3.2.3 Bulk crossover

Theoretically, in the limit of large n , the multilayer interactions ϕ which solve Eqn. (3.22) must reduce at long-range to that which would be realised for a transversely isotropic dielectric medium. Let ϕ_n^{edge} denote the $\phi_{11}(\mathbf{q})$ (or $\phi_{nn}(\mathbf{q})$) interaction

in an n -layer heterostructure in vacuum, and ϕ_n^{bulk} denote the ϕ_{mm} interaction (m labels the central layer if n is odd, or either of the two central layers if n is even). In equations, our requirement that the large n limit of the multilayer interactions coincides with the description of a transversely isotropic dielectric medium becomes

$$2\pi \left/ \lim_{q \rightarrow 0} \left[\lim_{n \rightarrow \infty} q \phi_n^{\text{bulk}}(\mathbf{q}) \right] \right. = \epsilon_{\text{eff}}, \quad (3.25)$$

$$2\pi \left/ \lim_{q \rightarrow 0} \left[\lim_{n \rightarrow \infty} q \phi_n^{\text{edge}}(\mathbf{q}) \right] \right. = \epsilon_{\text{eff}}^{(\text{edge})} = \frac{\epsilon_{\text{eff}} + 1}{2}, \quad (3.26)$$

where $\epsilon_{\text{eff}} = \sqrt{\epsilon_{\parallel} \epsilon_{\perp}}$, and we have used the fact that the edge is terminated by a vacuum. For $\epsilon_{\text{eff}}^{(\text{edge})}$, we use the fact that the effective dielectric constant felt by a plane in-between two different dielectric media with constants ϵ_1 and ϵ_2 is equal to $\epsilon_{\text{eff}} = (\epsilon_1 + \epsilon_2)/2$.¹⁵ In-plane dielectric permittivity is expressed using the susceptibility parameter r_{\star} [99],

$$\epsilon_{\parallel} = \chi_{3\text{D}}^{\parallel} + 1 = \frac{\chi_{2\text{D}}}{d} + 1 = 2 \frac{r_{\star}}{d} + 1, \quad (3.27)$$

while the out-of-plane dielectric constant is assumed to be that of vacuum. $\chi_{3\text{D}}^{\parallel}$ is the bulk electric susceptibility of the material in the directions parallel to the plane, and $\chi_{2\text{D}}$ is the in-plane two-dimensional susceptibility analogue. This has shown to be true for the multilayer interactions to a high degree of numerical precision for very large n (by M. Szyniszewski) in Ref. [7], but an analytical connection has yet to be made.¹⁶ Here, we consider the application of this rule in predicting the trends of edge and centre exciton and biexciton binding energies in homostructures. Using fits given in Ref. [3], the binding energy of a exciton in monolayer MoSe₂ surrounded by a dielectric medium having $\epsilon_{\text{eff}} = 3.65$ (2.33) corresponding to the bulk (edge) configuration of the MoSe₂ homostructure is 212.2 (311.9) meV. Similarly, for the biexciton, these are 16.1 (19.1) meV. As seen from Fig. 3.9, and Table 3.7, it doesn't appear that we are close to reaching this limit for few-layer

¹⁵This was shown to be valid at first order in the difference $(\epsilon_1 - \epsilon_2)$ in Ref. [3].

¹⁶*I.e.* we cannot directly show that the multilayer interactions as derived earlier reduce exactly to a large- n limit yielding dielectric functions which have ϵ_{eff} as desired.

homostructures in vacuum. The intralayer excitons all have (for a given n) very similar binding energies, regardless of which layer they occupy, and the drop-off of these binding energies as n increases appears to be rather slow when compared with the binding energies of the fit (shown as dashed lines in Fig. 3.9).

3.3.2.4 General trends

Here we will give a brief summary of the extensive listings of binding energy results presented in Appendix C.

The first consideration we mention is that of energetically favourable decay mechanisms. In all cases, the dominant decay (with the least binding energy) in our biexciton calculations has been the one which leads to the formation of two excitons (never a trion and a lone charge). Furthermore, in cases of ambiguity, the decay which maximises the number of intralayer daughter products is always observed to be the most favourable.

The second consideration regards our modelling of hBN encapsulation. In these cases, we observe an almost total loss of structure in the binding energies of multilayer complexes. For example, consider the differences in binding energies of edge/centre/centre-edge intralayer biexcitons discussed in Section 3.3.2.2. These differences are effectively nullified in the cases where we have considered any treatment of encapsulation (even by a hBN monolayer). In these cases, the intralayer exciton energies depend very weakly on which layer the charges reside in, and therefore do not strongly affect intralayer biexciton binding energies. In environments (2) and (3), this could be argued to be a consequence of our inclusion of unphysical screening in the vacuum regions between layers, which leads to an over-screening of the multilayer interactions, and to an insensitivity to the particular layer number. Charges in central layers here experience a dielectric environment which is similar to that experienced by charges in edge layers, owing to the permeation of our “bulk hBN” dielectric environment into the homostructure. We note that alternative models in the literature have sought to exclude

this unphysical effect by explicit treatment of the realistic “piece-wise” dielectric environment [126, 129], but also that in those studies no attention was paid to complexes higher than excitons.

Our study on environment (3), however, shows that some of this insensitivity to layer number is perhaps reasonable. There, intralayer exciton binding energies are again largely insensitive to layer number, all without the unphysical interlayer screening supplied by an approximate treatment of bulk encapsulation. Interlayer exciton energies are only observed to be sensitive to the number of medial layers that separate charges, and not the absolute layer numbers (consider the similarity between the binding energies of 1X_2 and 2X_3 in the $n = 5$ case, for example). It is interesting that the observed diversity and sensitivity to layer number of binding energies of complexes in homostructures in vacuum is lost almost immediately on encapsulation with only a single monolayer of hBN.

Finally, again where we have modelled encapsulation, we note that there is a strong tendency for topographically similar biexcitonic complexes to group into bunches with characteristic binding energies. Consider for example the complexes which are of the form “intralayer exciton plus two charges in other, different layers” ($^{14}XX_{13}$, for example, in the notation of Appendix C). In $n = 4, 5$ homostructures, in environment (3) say, such complexes seem to collect in groups having binding energies which are approximately ~ 120 and ~ 115 meV, respectively. A similar trend is observed for complexes of the form “intralayer trion plus a charge in another layer” ($^{11}XX_{13}$, for example), where the binding energies in $n = 2, 3, 4$ homostructures are also grouped, with another splitting depending on how many layers separate the intralayer trion from the remaining charge. This physics might be expected in the case where the additional charges (occupying layers alone) are only weakly bound in the field of the intralayer complex, which itself dominates the binding energy. To directly test this, we have calculated the VMC binding energy of the intralayer negative trion at an edge of the monolayer hBN encapsulated $n = 5$ homostructure ($X_e^-(5)$) to be 27.891(1) meV. The grouped complexes, which

we might have hypothesised were only weakly bound decorations of a dominant intralayer trions, are therefore *not* simply bunched around the trion binding energy, but are significantly better bound than the trion. If anything, this system could be thought of as interlayer exciton-like, with a free carrier in a layer bound to a trion in another layer, so that it has a binding energy which is determined approximately by the internal energy of the trion, and the binding energy of an exciton comprised of the single carrier, and a heavier trion “quasi-carrier”. The bunching into groups with similar binding energies overall is then a consequence of the fact that interlayer exciton energies depend weakly on the number of medial layers in system (3), and also the presumed relative insensitivity to layer number on the binding energy of the trion (the intralayer exciton data in Appendix C support this hypothesis in part – the binding energies of the intralayer excitons only weakly depend on n).

Environment (3) is the most realistic model for dielectric encapsulation that we have studied, and does not feature any pathological over-screening of interactions within homostructures. We note that it would be interesting (although more computationally demanding) to consider the application of our model in a scenario such as this to the case of a more realistic hBN encapsulation scheme. Monolayers are not often chosen to encapsulate samples, but rather multilayers or even bulk-like hBN stacks are. A typical width of a hBN capping layer in an experimental sample might be of the order of tens of nm. Using the bulk hBN lattice constant of Ref. [11], 10nm of hBN corresponds to around 30 or so individual layers of hBN. It would therefore be possible to repeat the study presented here in the case of a more realistic encapsulation (30 hBN layers + n charge-bearing layers + 30 hBN layers, for example), should one wish to spend the computer time.

Chapter 4

Excited-states by QMC

Accurate determination of the excited-state properties of atoms, molecules, and solids is an outstanding goal of modern theoretical and computational physics. In this chapter, we review the use of continuum QMC methods for the calculation of energy gaps, a particularly important set of excited-state properties, from first principles. We present a broad set of excited-state calculations carried out with the VMC and (fixed-node) DMC methods on atoms, molecules, and solids. We compare our results with those of other theoretical methodologies, and with experiments. Our work in this chapter aims to set an important precedent for subsequent QMC studies, particularly when it comes to finite-size effects in calculations employing periodic boundary conditions.

4.1 Excited-state QMC

4.1.1 Quasiparticle and excitonic gaps

In order to perform a QMC supercell calculation with periodic boundary conditions, the trial wave function must satisfy the many-body Bloch conditions outlined in Ref. [130]. Specifically, the wave function should acquire a phase $\exp(i\mathbf{k}_s \cdot \mathbf{R}_s)$ whenever a single particle is translated through a supercell lattice point \mathbf{R}_s , where the constant vector \mathbf{k}_s is the *supercell Bloch vector* or *twist*. Furthermore, the

wave function should acquire a phase $\exp(i\mathbf{k}_p \cdot \mathbf{R}_p)$ when all the particles are together translated through a primitive lattice point \mathbf{R}_p , where \mathbf{k}_p lies in the first BZ of the primitive cell. This is usually achieved by requiring the Jastrow factor and backflow function to have the periodicity of the supercell under single-particle displacements and the periodicity of the primitive cell under all-particle displacements, while the Bloch orbitals in the Slater determinant lie on a regular grid of primitive-cell reciprocal lattice points offset by the supercell Bloch vector \mathbf{k}_s . *E.g.*, for an $l \times m \times n$ supercell, this grid would be an $l \times m \times n$ grid of \mathbf{k} -points in the primitive-cell BZ, centred on the supercell Bloch vector \mathbf{k}_s . Folding of these points into the supercell BZ results in all points being mapped to \mathbf{k}_s . The occupancies of the single-particle orbitals at each \mathbf{k} point in the primitive-cell BZ can then be used to define excitations.

The *quasiparticle gap* Δ_{QP} of a system is the energy required to create an unbound electron-hole pair in that system. It is given by the difference between a conduction-band minimum \mathcal{E}_{CBM} and a valence-band maximum \mathcal{E}_{VBM} ,¹ *i.e.*

$$\begin{aligned} \Delta_{\text{QP}}(\mathbf{k}_f, \mathbf{k}_t) &= \mathcal{E}_{\text{CBM}}(\mathbf{k}_t) - \mathcal{E}_{\text{VBM}}(\mathbf{k}_f) \\ &= [E_{N+1}(\mathbf{k}_t) - E_N(\mathbf{k}_t)] - [E_N(\mathbf{k}_f) - E_{N-1}(\mathbf{k}_f)] \\ &= E_{N+1}(\mathbf{k}_t) + E_{N-1}(\mathbf{k}_f) - E_N(\mathbf{k}_t) - E_N(\mathbf{k}_f), \end{aligned} \tag{4.1}$$

where E_N is the total ground-state energy of an N -electron system. The labels \mathbf{k}_f and \mathbf{k}_t denote the \mathbf{k} -points from which and to which excitations are made, and may be ignored in finite systems. The ground-state energies $E_N(\mathbf{k}_t)$ and $E_N(\mathbf{k}_f)$ are identical if the calculations used to evaluate the quasiparticle energies $E_{N\pm 1}$ are performed on the same grid of \mathbf{k} -vectors (*i.e.*, for cells with the same supercell Bloch vector \mathbf{k}_s we have $\Delta_{\text{QP}}(\mathbf{k}_f, \mathbf{k}_t) = E_{N+1}(\mathbf{k}_t) + E_{N-1}(\mathbf{k}_f) - 2E_N$); otherwise, they may differ. It is always possible to evaluate Δ_{QP} between any pair of \mathbf{k} -points \mathbf{k}_f and \mathbf{k}_t at any system size by appropriate choices of the supercell Bloch vector

¹In a finite system, the “conduction-band minimum” is $-A$, where A is the electron affinity and the “valence-band maximum” is $-I$, where I is the first ionisation potential.

\mathbf{k}_s (the offset of the \mathbf{k} -point grid) in the two cases.

The *excitonic gap* (or *optical gap*) of a system is the energy required to create a bound electron-hole pair in that system. It is given by the difference of total energies obtained with an electron promoted to an excited state of the system and the total energy of the ground state

$$\Delta_{\text{Ex}}(\mathbf{k}_f, \mathbf{k}_t) = E_N^+(\mathbf{k}_f, \mathbf{k}_t) - E_N, \quad (4.2)$$

with $E_N^+(\mathbf{k}_f, \mathbf{k}_t)$ the *excited-state* total energy of an N -electron system in which an electron has been promoted from an occupied valence-band orbital at \mathbf{k}_f to an unoccupied conduction-band orbital at \mathbf{k}_t (again, the \mathbf{k} -point labels may be ignored in the finite case). The ground-state energy E_N is in this case unambiguous, and has to be evaluated with the same \mathbf{k} -point grid as the excited-state energy $E_N^+(\mathbf{k}_f, \mathbf{k}_t)$. In the rest of this section, we will suppress the \mathbf{k} -point labels \mathbf{k}_f and \mathbf{k}_t . Note that, unlike the quasiparticle gap, the excitonic gap may only be evaluated between pairs of \mathbf{k} -points that are *simultaneously* included in the \mathbf{k} -point grid (*i.e.*, the set of \mathbf{k} points must contain both \mathbf{k}_f and \mathbf{k}_t). This is not generally possible for a given pair of \mathbf{k} -points at all system sizes. For example, it is possible to calculate a vertical excitonic gap ($\mathbf{k}_f = \mathbf{k}_t$) in any supercell by using an appropriate offset \mathbf{k}_s to the grid of \mathbf{k} vectors; however, it is only possible to calculate an excitonic gap from Γ to K in a two-dimensional hexagonal cell in supercells of $3l \times 3m$ primitive cells, where l and m are integers.

For our purposes, the total energy E_{N-1} (E_{N+1}) is evaluated by calculation of the QMC energy of a state with the removal (addition) of an electron from (into) an occupied (unoccupied) state in the Slater determinant. Similarly, the total energy E_N^+ is evaluated by calculation of the QMC energy of a state whose valence- and conduction-band occupancies have been switched for the particular orbitals of interest. This trial wave function describes a correlated state of an excited electron and remnant hole, *i.e.*, an exciton. The difference $E_X^b = \Delta_{\text{QP}} - \Delta_{\text{Ex}}$ is equal to the exciton binding energy for a particular configuration of electron and

hole, and is always greater than or equal to zero for a finite system or for an extended system in the thermodynamic limit, because the electron-hole Coulomb interaction is attractive. This may not be the case in QMC data obtained in a finite periodic cell, in which case finite-size effects may lead to the apparently unphysical scenario where $\Delta_{\text{QP}} < \Delta_{\text{Ex}}$. The origin of this behaviour is explained in Section 4.1.5. The exciton binding energy E_X^b can only be evaluated at system sizes for which calculation of Δ_{Ex} is permitted. It may be reexpressed as

$$\begin{aligned} E_X^b(\mathbf{k}_f, \mathbf{k}_t) &= \Delta_{\text{QP}}(\mathbf{k}_f, \mathbf{k}_t) - \Delta_{\text{Ex}}(\mathbf{k}_f, \mathbf{k}_t) \\ &= E_{N+1}(\mathbf{k}_t) + E_{N-1}(\mathbf{k}_f) - E_N^+(\mathbf{k}_f, \mathbf{k}_t) - E_N, \end{aligned} \quad (4.3)$$

with $E_N = E_N(\mathbf{k}_f) = E_N(\mathbf{k}_t)$. The four QMC total energies in Eqn. (4.3) are statistically independent, unlike Δ_{QP} and Δ_{Ex} , which both depend on the same ground-state energy E_N .

4.1.2 Singlet and triplet excitations

In the preceding section, we neglected to include information on the possible spin degree of freedom of the electrons involved in excitations. For Hamiltonians that include no spin-orbit coupling we can, with no added difficulty, define the quasi-particle and excitonic gaps including explicitly the spin $\sigma \in \{\uparrow, \downarrow\}$ of the electron, which is excited from (\mathbf{k}_f, σ_f) to (\mathbf{k}_t, σ_t) . Singlet excitations are those with $\sigma_f = \sigma_t$, while triplet excitations incur a spin flip, $\sigma_f \neq \sigma_t$. In QMC, the spin of any electrons involved in excitations can be controlled by specification of the (spin-dependent) orbital occupancies in the Slater part of the trial wave function. In most cases singlet excitations are more physically relevant, because triplet optical excitations are forbidden in first-order perturbation theory. The feasibility of calculating singlet-triplet splittings by QMC techniques depends on the magnitude of the singlet-triplet splitting; the resolution of a small energy difference requires small QMC statistical error bars. We have calculated the singlet-triplet splitting of

the lowest lying excitonic states of anthracene in Section 4.3.2.3 and of the ground state of O₂ in Section 4.3.2.2.

4.1.3 Wave-function nodes and variational principles

DMC gives the energy of any eigenstate of the Hamiltonian exactly if the nodal surface of the trial wave function is equal to that of the true eigenfunction, even if the trial wave function is approximate between the nodes [131]. In general, however, each of the total energies E_N , $E_{N\pm 1}$, and E_N^+ suffers a fixed-node error due to the inexact nodal surface of the trial wave function. Assuming the excitations are made into the lowest-energy quasiparticle bands, E_N and $E_{N\pm 1}$ are themselves ground-state total energies, and hence the fixed-node errors in E_N and $E_{N\pm 1}$ must be positive [131]. There is no rigorous variational principle on the quasiparticle gap $\Delta_{\text{QP}} = E_{N+1} + E_{N-1} - 2E_N$, although in practice gaps evaluated using total energies evaluated by a variational method usually provide upper bounds. In Hartree-Fock theory, the absence of electronic correlation has the consequence that electrons localise excessively to avoid one another, and hence quasiparticle energy gaps are overestimated significantly. For example, in Si the Hartree-Fock quasiparticle gap is an overestimate by around 4.5 eV [132, 133]. In QMC, as we recover more and more of the electronic correlation energy by optimising Jastrow factors and backflow functions, and performing DMC to project out the fixed-node ground state, we observe that quasiparticle gaps reduce substantially from their Hartree-Fock values towards their exact static-nucleus non-relativistic values. Apart from the unlikely case in which we recover significantly more correlation energy in the $(N \pm 1)$ -electron systems compared to the ground-state N -electron system, we therefore expect QMC quasiparticle gaps to be upper bounds on the exact gaps. Because individual contributions to the quasiparticle gap separately obey ground-state variational principles, one expects to obtain improved DMC estimates of quasiparticle gaps by reoptimising parameters that affect the nodal surfaces in the $(N \pm 1)$ -electron systems. Improving a Jastrow factor is expected

to improve VMC energy gaps, but not fixed-node DMC gaps, since the Jastrow factor does not affect the nodal surface. (Of course, improving the Jastrow factor reduces statistical error bars, finite-time-step bias, finite-population bias, and pseudopotential locality errors; furthermore, parameters that do affect the nodal surface should be optimised together with the Jastrow factor.)

Let us now consider the fixed-node error in the excitonic gap. Again, the ground-state energy can only be overestimated by the fixed-node variational principle. The excited-state energy E_N^+ , however, is not bounded by variational principles except in special circumstances. If the trial excited-state wave function transforms as a one-dimensional (1D) irreducible representation (irrep) of the full symmetry group of the many-body Hamiltonian, then the resultant fixed-node DMC energy provides an upper bound on the energy of the lowest-lying eigenstate that transforms as that 1D irrep. In that case, the error in the DMC energy is second order in the error in the nodal surface of the excited-state trial wave function, and there is a tendency for positive fixed-node errors to cancel in excitonic gaps. In the likely case that we recover more correlation energy in the ground state than in the excited-state calculation, QMC excitonic gaps act as upper bounds to their exact counterparts.

If, however, the trial excited-state wave function does not transform as a 1D irrep, or we are not studying the lowest-energy eigenstate that transforms as the same irrep as the trial wave function then the fixed-node error in the excited-state energy E_N^+ can be either positive or negative, and hence there could be cases in which the DMC excitonic gap is too small. As a consequence, reoptimisation of trial-wave-function parameters affecting the nodal surface can lead to absurd results, as the nodal surface becomes more like that of the ground state. We provide an example illustrating this behaviour in Section 4.3.1.

If the excited-state trial wave function transforms as a multidimensional irrep of the full symmetry group of the Hamiltonian, then weaker lower bounds on the estimate of the excited-state energy can be realised by forming trial wave

functions that transform as 1D irreps of subgroups of the full symmetry group of the Hamiltonian [134]. This is discussed in Section 4.2.2.2.

Importantly, for excitations made between different \mathbf{k} points, where complex Bloch states (having definite crystal momentum \mathbf{k}_T) can be chosen to populate the Slater part of the trial wave function, variational principles on the lowest energy excitations are always realised because of translational invariance (states of definite crystal momentum transform according to 1D irreps of the space group, in line with the many-body Bloch conditions) [134]. In the case where one wishes to form real linear combinations of complex Bloch states with crystal momenta \mathbf{k}_T and $-\mathbf{k}_T$, respectively, the subsequent real superposition does not generally transform as a 1D irrep of the space group, and hence excited-state variational principles are not in general realised. If \mathbf{k}_T happens to be on the edge of the BZ, however, \mathbf{k}_T and $-\mathbf{k}_T$ are equivalent, and an excited-state variational principle is realised once again. If one is not able to recover an excited-state variational principle in this way, then one should use complex Bloch orbitals (maintaining a variational principle, at the cost of added computational expense). The so-called *fixed-ray* method of Hipes has been developed specifically to ensure the existence of excited-state variational principles in cases of degeneracy such as this [135].

Variational bounds on excited-state energies may also be obtained by other means, *e.g.* via MacDonald’s theorem [136]. Zhao and Neuscamman have recently devised a method which allows for the realisation of a variational principle on selected excited-state energies, and also for practical optimisation of excited-state QMC trial wave functions [35]. Mussard *et al.* have extended the VMC method using the ideas of time-dependent linear-response theory to extract excited-state properties, and have presented example calculations within the Tamm-Dancoff approximation to the linear-response equations [137].

4.1.4 Nodal topology

Fixed-node DMC works by obtaining exact ground-state solutions to the Schrödinger equation within nodal pockets, *i.e.*, within the regions of configuration space bounded by the nodes of the trial wave function [131]. The boundary conditions on the Schrödinger equation in each nodal pocket are that the DMC wave function goes to zero at the edges of the pocket. If the nodes of an excited-state trial wave function are exact then the ground-state energy in each nodal pocket is equal to the excited-state energy corresponding to the trial wave function.

From the point of view of fixed-node DMC, the fundamental differences between the ground-state many-electron wave function and its excited-state counterparts are codified in the topology of their respective nodal surfaces, which completely determine the corresponding fixed-node DMC energies. The nodal surface of the many-electron ground state satisfies a tiling property (all nodal pockets are equivalent under permutations of identical fermions; this is also true of determinants of Kohn-Sham orbitals) [138], and it is conjectured that the presence of only two nodal pockets is a generic feature of the many-electron ground state [139, 140]. The nodal surfaces of excited states are less-well-understood; they do not satisfy a tiling property in general unless the trial state transforms as a 1D irrep of the group of the Hamiltonian, and in the general case the number of nodal pockets they possess can only be bounded: Hilbert and Courant [141] proved that the nodes of the n^{th} excited state divide configuration space into no more than $n + 1$ nodal pockets.² The fact that the number of inequivalent nodal pockets remains small in low-lying excited states means that, for a sufficiently large DMC target population, each set of equivalent nodal pockets will have a significant initial population of walkers; furthermore, the walker populations in high-energy sets of pockets are expected to die out on an imaginary-time scale given by the inverse of the difference between the energies of the different nodal pockets. Hence the

²In 1D, a rigorous analysis of the topology of the ground and excited-state nodal surfaces culminates in the Hilbert-Courant nodal line theorem. The ground state is nodeless, and the n^{th} (non-degenerate) excited-state has n nodes, dividing the 1D configuration space into $n + 1$ nodal pockets (saturating the earlier stated constraint, which applies in dimensions greater than one).

fixed-node DMC energy with an excited-state trial wave function is equal to the lowest of the pocket ground-state energies. An example of this behaviour is shown in Section 4.3.1.1.

It is not possible for a Jastrow factor to alter the nodal surface of a trial state, and nor is it possible for a smooth backflow function to alter the *topology* of a trial state. It is this fact that prevents variational collapse of excited-state energies in VMC calculations in the cases of electron addition, removal, or promotion where the trial state is a state of definite symmetry transforming as a 1D irrep. While nodal topology is an important factor in the description of excited states in QMC, and it is important that backflow functions preserve it, as we show in Section 4.3.1.1, the correct nodal topology does not guarantee that one will obtain reasonable results when optimising backflow functions in trial excited states. In cases where trial wave functions do not transform as 1D irreps of the symmetry group of the Hamiltonian, preserving the nodal topology can still lead to the formation of a pathological nodal surface and a DMC energy which is too low.

We note that, while we will not explicitly consider their use here, pfaffian and geminal pairing wave functions have recently been shown to be somewhat more efficient at accurately describing the nodes of a few systems where the exact nodes are known [142].

4.1.5 Finite-size effects

A major source of error in gap calculations for condensed matter using explicitly correlated wave-function methods such as QMC is the presence of finite-size (FS) effects. For calculations on solids, we are only able to simulate a finite supercell subject to periodic boundary conditions. This means that our raw DMC data contain unwanted contributions from the electrostatic interaction of added (or removed) charges with their periodic images, and we must either correct for this effect or extrapolate to infinite supercell size. A general simulation supercell in d dimensions is defined by a $d \times d$ integer “supercell matrix” S , which expresses the

supercell lattice vectors $\{\mathbf{a}_i^{\text{sc}}\}$ in terms of primitive-cell lattice vectors $\{\mathbf{a}_i^{\text{prim}}\}$:

$$\mathbf{a}_i^{\text{sc}} = \sum_j S_{ij} \mathbf{a}_j^{\text{prim}}. \quad (4.4)$$

A “diagonal supercell” is one for which the supercell matrix is diagonal; such a supercell consists of an $S_{11} \times S_{22} \times S_{33}$ array of primitive cells. In general a supercell contains $\det(S)$ primitive cells.

Various FS correction schemes exist for the total energies per primitive cell of solids calculated at fixed system size in DMC [143–145]. However, such FS errors cancel between ground and excited states and are of little relevance to the FS effects in excitation energies. Let us first consider the FS effects in Δ_{QP} . The leading-order FS error is due to the self-interaction of added quasielectrons or quasiholes. The energy of the resulting unwanted lattice of quasiparticles (each having charge $q = \pm 1$) is given by a screened Madelung sum over supercells, *i.e.*, $q^2 v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})/2$ with $v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$ being the screened Madelung constant for the supercell [146]. There are two separate terms of this type in a quasiparticle gap correction, one for $-\mathcal{E}_{\text{VBM}} = E_{N-1} - E_N$ and another for $\mathcal{E}_{\text{CBM}} = E_{N+1} - E_N$. A physically reasonable FS correction formula for Δ_{QP} therefore reads

$$\Delta_{\text{QP}}(\infty) \approx \Delta_{\text{QP}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}) - v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}), \quad (4.5)$$

where $\Delta_{\text{QP}}(\infty)$ is the infinite-system quasiparticle gap. A similar expression has previously been used at the DFT level to study FS effects in the formation energies of charged defects [147]. Assuming the separation of the neighbouring images of the quasiparticle is sufficiently large that linear response theory is valid, $v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$ can be evaluated using an appropriately screened Coulomb interaction. In QMC calculations with fixed ions, only the electronic contribution to the susceptibility is relevant to the FS effects in the quasiparticle gap, *i.e.*, the permittivity that should be used to evaluate the screened Madelung constant is the high-frequency permittivity. This can usually be evaluated with sufficient accuracy using density

functional perturbation theory [148], if experimental results are unavailable. In anisotropic materials, the Madelung constant must be evaluated using the permittivity tensor, as is done in DFT studies of charged defect formation energies [149]. A simple expression for the anisotropically screened Madelung constant can be obtained by a coordinate transformation to the principal axes of the permittivity tensor. If $\tilde{v}_M(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$ is the unscreened Madelung constant then the screened Madelung constant is

$$v_M(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}) = \frac{1}{\sqrt{\det(\epsilon)}} \tilde{v}_M(\epsilon^{-1/2} \mathbf{a}_1^{\text{sc}}, \epsilon^{-1/2} \mathbf{a}_2^{\text{sc}}, \epsilon^{-1/2} \mathbf{a}_3^{\text{sc}}), \quad (4.6)$$

where ϵ is the high-frequency permittivity tensor of the system. The properties of physical permittivity tensors mean that the square root of the inverse is always well-defined: positive-definite matrices have only one square root, also known as the principal square root. This expression can be obtained from an analysis of the Ewald interaction in the presence of an anisotropic medium (supplied in the present case by the rest of the system). Similar arguments were given by Fischerauer for the interaction between aperiodic point charges in anisotropic media [150]. In the case of an isotropic medium, Eqn. (4.6) reduces to division of the unscreened Madelung constant by the relative permittivity, *i.e.*, $v_M = \tilde{v}_M/\epsilon$.

In layered and two-dimensional materials, the in-plane polarisability of the layers modifies the form of the Coulomb interaction to the so-called Keldysh interaction [95–97]. Depending on the in-plane susceptibility and the spatial extent of the simulation cell, it may be necessary to employ this modified form of interaction in the evaluation of the screened “Madelung” constant. For supercells much larger than the length scale r_* defined by the ratio of the in-plane susceptibility to the permittivity of the surrounding medium, the Keldysh interaction between image charges reduces to Coulomb form, and the subtraction of the screened Coulomb Madelung constant is reasonable. On the other hand, if the supercell size is significantly less than r_* then the Keldysh interaction is of logarithmic form [3] and the resulting Madelung constant is roughly independent of system size, until the linear

size of the simulation cell reaches r_* . We discuss this further in Section 4.3.4.

If the leading-order FS error in the quasiparticle gap is removed by subtracting the screened Madelung constant, the remaining systematic FS errors are expected to be dominated by periodic charge-image quadrupole interactions, and to fall off rapidly as L^{-3} , where L is the linear size of the supercell. Depending on whether sufficient data are available, linear extrapolation in $1/L^3$ can be used to remove these errors. One could even attempt to eliminate these errors using the Makov-Payne expression for the correction to the formation energy of a charged defect [151]. For a two-dimensional material with a supercell size much less than r_* , the charge-image quadrupole Keldysh interaction falls off as $1/L^2$; when the linear size of the supercell exceeds r_* , a crossover to $1/L^3$ scaling takes place.

The corrected $\Delta_{\text{QP}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}) - v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$ data are also subject to additional, beyond-linear-response effects. These additional effects are quasirandom, scaling in no systematic way with system size; however, they do correlate with analogous charged-defect formation energies evaluated at the DFT level: see Section 4.3.3.1. We interpret these errors as commensurability effects: oscillations in the electron pair density arising from additional quasiparticles (in a metallic system these would be Friedel oscillations) are artificially made commensurate with the supercell.

Some earlier QMC studies have extrapolated gaps to infinite system size assuming FS errors in energy gaps scale as $1/L$ [152, 153]. For a fixed cell shape, $v_{\text{M}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$ itself scales like $1/L$, so this *Ansatz* is reasonable. However, this approach is invalid if the cell shape is varied. Furthermore, it is difficult to extrapolate reliably from a small number of data points suffering from unquantified quasirandom noise. In many cases, averaging corrected energy gaps is a more accurate way of removing systematic and quasirandom FS effects. As shown in Section 4.3.3.1 (Table 4.6, specifically), the magnitude of the quasirandom FS effects appears larger than any remnant systematic FS error after application of our proposed correction [Eqn. (4.5)] in three-dimensional Si; in two-dimensional phosphorene,

however, residual systematic FS errors are still present after the Madelung-constant correction has been applied, as shown in Section 4.3.4. Whether extrapolation in $1/L^3$ or simple averaging of corrected gaps is the most effective way of removing FS effects depends on the system and on the number of system sizes at which gap data are available. In either case, provided the quantified QMC statistical error bars are less than the unquantified quasirandom FS noise (typically around 0.1 eV), the data should not be weighted by the inverse square QMC error bars when extrapolating or averaging.

For a fixed supercell size N , one can choose a cell shape to maximise the distance between periodic images, thereby minimising remaining systematic FS effects not accounted for by Eqn. (4.5). For cubic materials, the cells that maximise the nearest-image distance are themselves cubic ($n \times n \times n$ arrays of unit cells). In other lattice systems, the supercells maximising the nearest-image distance need not be of the same shape as the primitive cell, or even be diagonal in their extent. Non-diagonal supercells have previously been used in studies of lattice dynamics at the DFT level [154], but purely as a means of reducing computational expense. The shape of the simulation supercell may also be of significance with regards to the quasirandom FS effects: see Section 4.3.3.1.

For the case of excitonic gaps, there is another FS effect to consider. The characteristic size of an exciton is usually the exciton Bohr radius $a_B^* = \epsilon/\mu$, where $\mu = m_e^* m_h^*/(m_e^* + m_h^*)$ is the electron-hole reduced mass, ϵ is the permittivity, and m_e^* and m_h^* are the electron and hole effective masses, respectively. (Note that the size of an exciton is different in two-dimensional materials where the screened interaction is of Keldysh form [3]; in that case the size of the exciton is $r_0 = \sqrt{r_*/(2\mu)}$.) If the simulation supercells used are of linear size much less than the characteristic exciton size then the exciton is artificially confined and the kinetic energy dominates the Coulomb interaction. The exciton consists of two weakly attracting, almost independent quasiparticles, and the FS behaviour of the resulting “excitonic” gap mimics that of the quasiparticle gap, with a FS error dominated by the

Madelung energies of the free electron and hole. If, on the other hand, the simulation supercell has a linear size exceeding the characteristic size of the exciton, the hydrogen-like bound state forms, and the leading-order systematic FS scaling in the excitonic energy gap is given by the energy of a lattice of self-image-interacting excitons. To investigate the binding energy of a lattice of exciton images, we have performed a series of two-particle DMC calculations in which an electron and a hole in the effective mass approximation and interacting by the Ewald interaction are confined to a face-centred cubic (FCC) cell of lattice parameter L . The results of this investigation are presented in Fig. 4.1, which clearly shows the crossover in the scaling of the FS error in the exciton binding energy from L^{-1} in small cells to L^{-3} in large cells when the linear size of the cell is about twice the exciton Bohr radius. The two-dimensional [155] and three-dimensional [151, 156] Ewald interactions, $v_{\text{Ew}}(r)$, may be expanded into the general form

$$v_{\text{Ew}}(r) - v_{\text{M}} = v_{\text{Coul}}(r) + a \frac{r^2}{L^3} + O(r^4), \quad (4.7)$$

where v_{M} is the Madelung constant and a is a geometrical factor, which is sensitive to dimensionality and the supercell shape. $v_{\text{Coul}}(r)$ is the aperiodic Coulomb interaction. This difference from the exact Coulomb interaction is the physical source of the L^{-3} FS error in the exciton binding energy as evaluated in calculations employing periodic boundary conditions.³ In a sufficiently large cell, the exciton wave function is nearly independent of linear system size L , and hence by first-order perturbation theory the effect of the ar^2/L^3 term goes as L^{-3} . Once again, the situation is different in two-dimensional materials when the simulation supercell is much smaller than r_* (but larger than the exciton size r_0); in that case the FS error in the exciton binding energy and hence excitonic gap scales as L^{-2} .

The approximate FS behaviour of the excitonic gap is determined by the FS behaviour of the exciton binding energy $E_{\text{X}}^{\text{b}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$. In particular, the excitonic

³In an anisotropic system, the quadratic term in Eqn. (4.7) would be replaced by a bilinear form $\mathbf{r}^T T \mathbf{r}$, with T a tensor depending on the lattice structure.

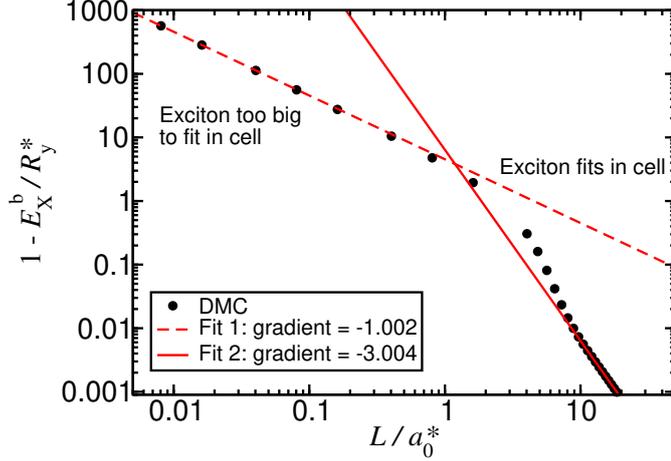


Figure 4.1: Scaled difference of exciton binding energy E_X^b and the exciton Rydberg against the lattice parameter L in an effective-mass model of a three-dimensional exciton confined in a periodic FCC cell. $R_y^* = m_e^* m_h^* / [2\epsilon^2(m_e^* + m_h^*)]$ and $a_0^* = \epsilon(m_e^* + m_h^*) / (m_e^* m_h^*)$ are the exciton Rydberg and the exciton Bohr radius, respectively, where m_e^* and m_h^* are the electron and hole masses and ϵ is the permittivity. The exciton Rydberg is the binding energy of a free exciton (or, one in a cell of infinite extent). The gradient on this log-log plot gives the scaling exponent of the FS error in the exciton binding energy.

gap in a finite supercell is approximately given by

$$\Delta_{\text{Ex}}(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}) \approx \Delta_{\text{QP}}(\infty) - E_X^b(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}). \quad (4.8)$$

If the exciton Bohr radius is large compared with the supercell then $E_X^b(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}}) \approx -v_M(\mathbf{a}_1^{\text{sc}}, \mathbf{a}_2^{\text{sc}}, \mathbf{a}_3^{\text{sc}})$, so that the FS behaviour of the quasiparticle and excitonic gaps is the same, and either can be used to estimate the infinite-system quasiparticle gap by subtracting the screened Madelung constant from the result obtained in a finite supercell. There is no point in attempting to calculate exciton binding energies using differences of quasiparticle and excitonic gaps in supercells smaller than the exciton Bohr radius suggested by the effective-mass approximation. On the other hand, if the simulation supercell is larger than the exciton Bohr radius then the FS errors in the exciton binding and hence excitonic gap are small and fall off rapidly as L^{-3} ; in this case it is possible to determine the exciton binding energy.

We have investigated whether single-particle FS effects (*i.e.*, momentum-quantisation

effects) are significant in DMC gaps by fitting $\Delta(N) = \Delta(\infty) + b/N^{1/3} + c[\Delta^{\text{DFT}}(N, \mathbf{k}_s) - \Delta^{\text{DFT}}(\infty)]$ to DMC gaps $\Delta(N)$ obtained in a series of cells of the same shape but different size N , where $\Delta^{\text{DFT}}(N, \mathbf{k}_s)$ is the DFT energy gap evaluated for a finite supercell containing N electrons, \mathbf{k}_s is the offset to the grid of \mathbf{k} -vectors used in the DFT calculation, and $\Delta^{\text{DFT}}(\infty)$ is the DFT gap converged with respect to \mathbf{k} -point sampling. However, we do not find the fitted values of c to be statistically significant. Nor do we find correlation between the ground-state DFT total energy and the QMC gaps. On the other hand, we do observe some correlation with FS effects in DFT-calculated defect-formation energies (see Fig. 4.7). Twist averaging [157] (TA) is a method for removing single-particle type FS effects from ground-state expectation values, and is discussed in detail in the context of Monte Carlo twist averaging (MCTA) in Chapter 5. Because of the requirement that \mathbf{k} points involved in the excitation are present in the grid of \mathbf{k} points used to define the many-body trial wave function TA, in the conventional sense, cannot be used in QMC excitation calculations. We note that a kind of TA might be used to define a band edge in QMC, by shifting \mathbf{k}_s and performing vertical gap calculations about some target point, one could in principle characterise properties like the effective mass. Such accurate determination of properties to feed into other methods (such as those discussed in Chapters 2, 3 and 5) are of critical importance to the accuracy of these methods, which are often able to address phenomena that are inaccessible to first-principles calculations (the binding energies and recombination dynamics of multi-particle excitonic complexes, for example).

4.1.6 QMC band structures

Quasiparticle energies are generally complex quantities, because quasiparticle excitations have finite lifetimes. The central quantity of interest in many spectroscopic experiments is the spectral function $A(\mathbf{k}, \omega)$, which characterises the electronic states of wave vector \mathbf{k} in a given material, having peaks centred on the quasiparticle energies ω whose widths relate to the lifetime of the quasiparticle excitation

in question. It would be possible to try to extract the energy-momentum spectral function from VMC calculations. As an example, one could calculate the squared matrix element

$$|\langle \Psi_N(\mathbf{r}_1, \dots, \mathbf{r}_N) \cdot \exp[i\mathbf{k} \cdot \mathbf{r}_{N+1}] | \Psi_{N+1}(\mathbf{r}_1, \dots, \mathbf{r}_{N+1}) \rangle|^2, \quad (4.9)$$

for the HEG at the VMC level, where Ψ_N is an optimised N-electron wave function. This would allow for determination of the broadening of the spectral peak at a particular momentum \mathbf{k} and extraction of the lifetime of quasiparticles in the quasielectron band at \mathbf{k} , complementing previous works. This would go some way to completing the first-principles description of the properties of the HEG from the point of view of Landau's Fermi liquid theory [158–160].

A similar possibility would be to try to calculate the radiative lifetime for an excitonic state. This relies on the evaluation of dipole matrix elements, which again is possible with VMC. This has already been performed for few-body systems in Ref. [118] (published alongside the work described in Section 3.2), and for the $2^2S \rightarrow 2^2P$ transition of the Li atom in Ref. [161].

One might think that a natural way to obtain improved estimates of quasiparticle lifetimes and radiative rates would be to evaluate the corresponding matrix elements at the DMC level. However, this is not immediately possible. The DMC method gives no direct information regarding many-electron wave functions (*i.e.*, produces no functional form for $\Psi_N(\mathbf{R})$).⁴

4.1.7 Excitations in metallic systems

Various studies have investigated, from a microscopic viewpoint, the excited-state properties of the two-dimensional HEG [162–165]. This involves the study of in-

⁴As is discussed in Chapter 1, the DMC algorithm normally samples the mixed distribution, and generates configurations distributed in configuration space as the mixed distribution. If a pure estimator is used, then the generated configurations are distributed as the square modulus of the DMC wave function. This does not change the salient point, which is that DMC *generates configurations* and does not supply a functional form for the fixed-node ground state, or even for the mixed distribution $\mathcal{M}(\mathbf{R}, \tau)$.

traband excitations, in which electrons are promoted or added into higher energy states on the free-electron-like band of the HEG in order to determine the quasi-particle effective mass and the Fermi liquid parameters. All of these studies have observed the presence of severe FS effects. In what remains of the present article, we will discuss only interband excitations to calculate energy gaps.

4.1.8 Computational expense

Methods developed to improve the scaling of QMC calculations [166, 167] may find use in excitation calculations. By localising low-lying states which are not directly involved in excitations, the number of non-zero orbitals to evaluate at a given point \mathbf{r} is reduced, and the Slater matrix is made sparse, improving the cost scaling of the Slater part of the wave function by a factor of N . An additional side effect of this is to reduce the computational expense of the inclusion of backflow correlations (whose dominant cost arises at the orbital-evaluation stage of a calculation). However, a major problem with the use of localised orbitals is that, in order to obtain efficiency increases, one sacrifices accuracy in individual total energies by truncating localised orbitals to zero at finite range. The extent to which this loss of accuracy will affect total-energy differences in solids is unclear, although early studies on molecules have provided positive results [168]. Given that other biases (single-particle FS effects, time-step bias, etc.) cancel so well in gap calculations in solids (see Section 4.3.3.1) we expect the loss in accuracy in energy gaps due to the truncation of localised low-lying electronic states to be very small. On the other hand, computational expense is often dominated by other factors such as the evaluations of two-body terms in the Jastrow factor and updates to the Slater matrix, limiting the scope for speedup.

Because highly precise total energies are required from the DMC calculations used in forming energy gaps, the most significant portion of computational time is spent in the statistics-accumulation phase; the equilibration phase is only a small fraction of the total computational expense. This means that QMC gap calcula-

tions are particularly suited to massively parallel computational architectures.

4.1.9 Nuclear relaxation and vibrational effects

The renormalisation of static-nucleus energy gaps by zero-point vibrational effects is important for any comparison of theoretical results with experiment [169]. In the extreme case of hexagonal ice, this effect contributes a correction in the range of 1.5–1.7 eV [170, 171]. Related work has also demonstrated a large renormalisation of the energy gap in the benzene molecule by more than 0.5 eV [153]. We investigate this issue in Section 4.3.2.1, where we present results for an H₂ molecule with a full quantum treatment of both protons and electrons.

A second issue is the equilibrium geometry of electronic excited states. In an adiabatic ionisation potential, electron affinity, or quasiparticle gap, the geometry of the molecule or crystal is allowed to relax after the addition or removal of an electron. By contrast, in a “vertical” ionisation potential, electron affinity, or quasiparticle gap, the atomic structure of the cation or anion is assumed to be the same as that of the ground state. An important point to note here is that, from the point of view of experiment, atomic relaxation may or may not be relevant. Experimental measurements that occur on timescales smaller than those associated with the structural relaxation of a molecule or a solid (for example, as with photoemission/inverse photoemission spectroscopy) are insensitive to any relaxation effects which are instigated by the measurement. On the other hand, in experimental measurements that occur on timescales greater than those associated with the structural relaxation (for example, as in zero electron kinetic energy spectroscopy [172]), one can expect that one will measure directly an adiabatic excitation energy, and that comparison to fully relaxed *ab initio* results is reasonable. The situation is less clear in the case that the experimental and structural relaxation timescales are comparable. Geometrical relaxation in excited states typically reduces quasiparticle gaps by 0.1–0.5 eV. We present many of our quasiparticle-gap results with and without relaxation in excited states, using DFT

to relax structures. A closely related issue is the Stokes shift, which is the difference between excitonic absorption and emission gaps. In an absorption gap, the geometry is that of the ground state; in an emission gap, the geometry is that of the excited state. QMC calculations have previously been performed to calculate Stokes shifts in diamondoids using DFT geometries [173].

Both of these issues complicate the detailed comparison of *ab initio* gaps with experimental measurements.

4.2 Computational details

4.2.1 DFT orbital generation

Our DFT calculations were carried out with the CASTEP plane-wave-basis code [174]. In the case of molecules and of phosphorene, prior to any wave-function generation calculation, we relaxed the ground-state (and, where explicitly stated, excited-state) geometries to within a force tolerance of at most 0.05 eV/Å, with ultrasoft pseudopotentials [175] representing the nuclei and core electronic states. All of our DFT calculations used the Perdew-Burke-Ernzerhof (PBE) parameterisation of the generalised gradient approximation to the exchange-correlation energy [176]. For our calculations on solids, we used experimentally obtained geometries (Si from Ref. [177], hexagonal boron nitride from Ref. [11], cubic boron nitride (BN) from Ref. [178], and α -SiO₂ from Ref. [179]).

We have used Trail-Needs Dirac-Fock averaged-relativistic-effect pseudopotentials [52, 53] for all wave-function generation calculations and subsequent QMC calculations, except in our all-electron calculations. We have chosen the local channels of our pseudopotentials such that no ghost states exist, and we have used plane-wave cutoff energies which lead to an estimated DFT basis-set error per atom of at most 10^{-4} a.u. (2.72 meV) [55].

After their generation, the DFT single-particle orbitals were rerepresented in a blip (B-spline) basis [180]. This allows for improved computational efficiency

of QMC calculations, and the removal of unphysical periodicity in calculations on zero-, one-, and two-dimensional systems.

4.2.2 QMC calculations

4.2.2.1 Slater-Jastrow(-backflow) wave functions

We have used Jastrow factors of the form outlined in Ref. [181] in all of our QMC calculations, with system-appropriate terms and with free parameters optimised by unweighted variance minimisation and subsequent energy minimisation [100, 182–184]. We have not (except where explicitly stated) reoptimised Jastrow-factor parameters in trial excited states. We have used backflow functions of the form outlined in Ref. [185], optimising free parameters by energy minimisation [100].

The results of our DMC calculations have been simultaneously extrapolated to infinite population size, and zero time step in an efficient manner [42]. We have used the “T-move” method of Casula to ensure that our DMC energies are variational in the presence of non-local pseudopotentials [186]. All of our QMC calculations have been carried out using the CASINO code [34].

4.2.2.2 Multideterminant trial wave functions

In a multideterminant wave function, the Slater part of the wave function of Eqn. (1.36) is replaced by

$$\mathcal{D}(\mathbf{R}) \rightarrow \mathcal{D}(\mathbf{R}) + \sum_j c_j \mathcal{D}_j(\mathbf{R}), \quad (4.10)$$

where the original determinant \mathcal{D} is chosen as the “dominant” determinant, and the excited determinants \mathcal{D}_j are populated with single-particle orbitals with substituted degenerate or near-degenerate orbitals of interest with respect to those appearing in \mathcal{D} . Unless one believes the single-particle theory used to generate the orbitals to be qualitatively incorrect, the order of the eigenvalues of the orbitals occupied in the Slater determinant of single-particle orbitals is preserved with respect to the interacting case: the states of the interacting and non-interacting

systems are assumed to be adiabatically connected. In the case of a failure of the single-particle theory, this is not guaranteed, and the state formed from the determinant of single-particle orbitals is not a reasonable trial state. *E.g.*, in a case where DFT metallizes an insulator, one might attempt to remedy the problem by inclusion of exact exchange (the use of a hybrid functional, or even Hartree-Fock theory itself) or artificial separation of the occupied and unoccupied manifolds (*i.e.*, the use of a scissor correction) in the orbital-generation calculation.

One is able to obtain better estimates of ground-state total energies by variation of the multideterminant expansion coefficients $\{c_j\}$. One might also be able to obtain better estimates of certain excited-state energies (see Section 4.1.1). However, general excited states do not obey variational principles, and so it is not obviously the case that one would always want to form a multideterminant expansion for the excited state.

There are cases where the formation of a (restricted) multideterminant expansion is desirable. Firstly, excited-state multideterminant expansions transforming as 1D irreps of the full symmetry group of the Hamiltonian of a system can be shown to obey variational principles in fixed-node DMC [134], as discussed in Section 4.1.3. Secondly, in cases of states with degeneracy or near-degeneracy, one might expect that the wave function should have some multireference character. Such degeneracies are much more likely to occur in the excited state than in the ground state. The inclusion of determinants characterising electron promotions (or additions, or removals) from the degenerate or near-degenerate energy levels might reduce excited-state energies, leading to lower QMC energy gaps. Towler *et al.* [187] paid a great deal of attention to the correct inclusion of degenerate determinants of specified symmetry classes in their study of diamond (which has the same symmetry properties as Si, with the same consequence that the valence-band maximum and conduction band at Γ are triply degenerate at the single-particle level). When choosing a multideterminant expansion to describe an excited state, one must apply a group theoretical projection operator to each of the possible

degenerate determinants in order to determine an excited-state trial wave function of definite symmetry. This “safe” trial wave function is then a few-determinant expansion in the space of degenerate determinants of single-particle orbitals, with a definite symmetry. However, this symmetry may only be maintained at the VMC level, and the fixed-node DMC algorithm may still break it if the trial wave function does not transform as a 1D irrep. The weaker variational principle for DMC excited states mentioned in Section 4.1.3 still applies in cases where trial functions have specific transformation properties, however.

We have explicitly tested the formation of multideterminant trial wave functions in some of our calculations in Si (see Section 4.3.3.1), where three bands at the Γ point are degenerate in the absence of spin-orbit coupling.

We note, but have not explicitly demonstrated, that multireference character is likely the source of some strange conclusions regarding hexagonal boron nitride in Section 4.3.3.2.

4.3 Results and discussion

4.3.1 Atoms

4.3.1.1 H atom: a model of excited-state fixed-node errors

An important class of fixed-node errors in excited-state DMC calculations is that which may arise due to the lack of a variational principle. Here we consider various modifications to the hydrogenic 2s orbital, whose exact energy is $-\frac{1}{8}$ a.u. The corresponding wave function is isotropic and hence transforms as the trivial 1D irrep of the $SO(3)$ geometric symmetry group of the H atom; however, it is not the lowest energy eigenfunction of this symmetry. The nodal surface of the 2s orbital is a sphere of radius 2 a.u. This example was previously investigated analytically in Ref. [134]; here we provide numerical results that corroborate the argument in Ref. [134], and we investigate the consequences for optimisation of backflow functions in excited states.

The two ways that a spheroid nodal surface can be inexact are that (a) the average positions of nodes is incorrect, and/or (b) the curvature of nodes is incorrect. We have studied two inexact nodal surfaces for the 2s state using the trial wave functions

$$\psi^\gamma(r) = C_\gamma (2\gamma - r) \exp\left(-\frac{r}{2}\right), \quad (4.11)$$

$$\psi_L^{\alpha,\beta}(r, \theta) = D_{\alpha,\beta} \{2\beta [1 + \alpha \mathcal{Y}_{L,0}(\theta)] - r\} \exp\left(-\frac{r}{2}\right), \quad (4.12)$$

which are exact (2s) eigenstates for $\gamma = 1$ and $\alpha = 0, \beta = 1$. The wave function $\psi^\gamma(r)$ encodes the scenario already explored in Ref. [134]. The normalisation constants C_γ and $D_{\alpha,\beta}$ are irrelevant in DMC, and \mathcal{Y}_{L,m_L} is a spherical harmonic. We have used ψ^γ as a DMC trial wave function with γ being a control parameter which varies the nodal volume, keeping the node spherical. This addresses point (a). We have also used $\psi_L^{\alpha,\beta}$ as a DMC trial wave function, with α a control parameter that sets the degree of non-spherical distortion of the nodal surface, this time with β chosen to fix the nodal volume to the exact value. This addresses point (b). The nodal topology of our trial wave function does not change as a function of γ and α ; there are always two nodal pockets. The results of varying γ and α are presented in Figs. 4.2 and 4.3.

Define the pocket eigenvalues $E_{\text{outside}}^{\text{pocket}}$ and $E_{\text{inside}}^{\text{pocket}}$ to be the energy eigenvalues associated with single electrons occupying the regions outside and inside the nodal surface of ψ^γ , respectively, where the boundary conditions are that the pocket eigenfunctions are zero outside of their respective pockets. For the first case, the pocket eigenvalues can be determined via numerical solution of a model eigenvalue problem. If the radial Schrödinger equation is integrated, but with a “nodal boundary condition” $\psi^\gamma(2\gamma) = 0$, then the lower of the corresponding eigenvalues $\min\{E_{\text{outside}}^{\text{pocket}}, E_{\text{inside}}^{\text{pocket}}\}$ matches very closely the DMC energy. Moreover, we can also find the pocket eigenvalues corresponding to solutions inside and outside the nodal surface for all γ (see extended dotted lines; only the lesser of these solutions

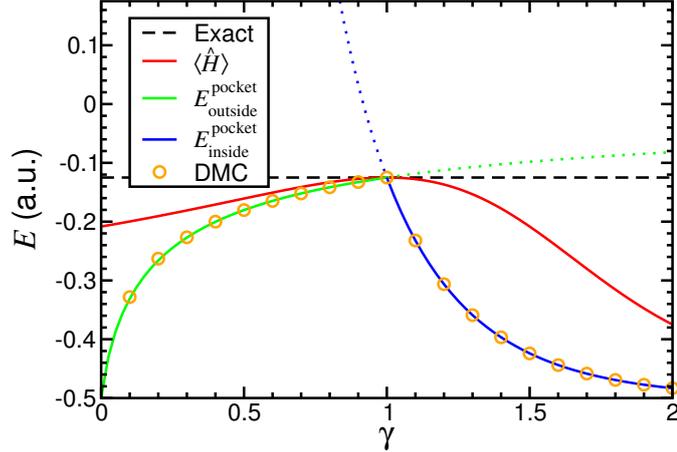


Figure 4.2: Approximations to the first-excited-state energy of an H atom using the ψ^γ excited-state trial wave function of Eqn. (4.11) as a function of γ , obtained by various means. DMC errors are smaller than the thickness of the lines. The pocket eigenvalues outside and inside the nodal surface, $E_{\text{outside}}^{\text{pocket}}$ and $E_{\text{inside}}^{\text{pocket}}$, were determined by numerical integration of the Schrödinger equation with fixed-node boundary conditions, and $\langle \hat{H} \rangle = \langle \psi^\gamma | \hat{H} | \psi^\gamma \rangle$, where \hat{H} is the Hamiltonian.

is sampled by the DMC algorithm). Even in the $\gamma \rightarrow 0$ and $\gamma \rightarrow \infty$ nodeless limits the ground-state variational principle is always obeyed, *i.e.*, $E \geq -\frac{1}{2}$ a.u.

There is a qualitative difference in the behaviour of the energy expectation value $\langle \hat{H} \rangle$ (which could be evaluated by VMC) versus the fixed-node DMC energy as a function of γ : the error in the DMC excited-state energy due to the use of an inexact nodal surface is more severe, and is first-order in the error in the nodal surface (as quantified by γ). Recall that the fixed-node error in the DMC ground-state energy is second order in the error in the trial nodal surface.

In the second case, as is shown in Fig. 4.3, the fixed-node error is always positive for $\alpha \neq 0$. This is not too surprising, given that if the wave function is to satisfy the nodal constraint, it must adopt additional curvature in both nodal pockets. Additional curvature in space corresponds to an increased kinetic energy of the wave function in both nodal pockets. The fixed-node DMC energy is second-order in the parameter α , because it is an even function of α .

This model serves as an illustrative example of the fact that excited-state fixed-node errors can be either positive or negative, depending on the nature of the inexactness of the nodal surface. This is important, in particular, if one is to

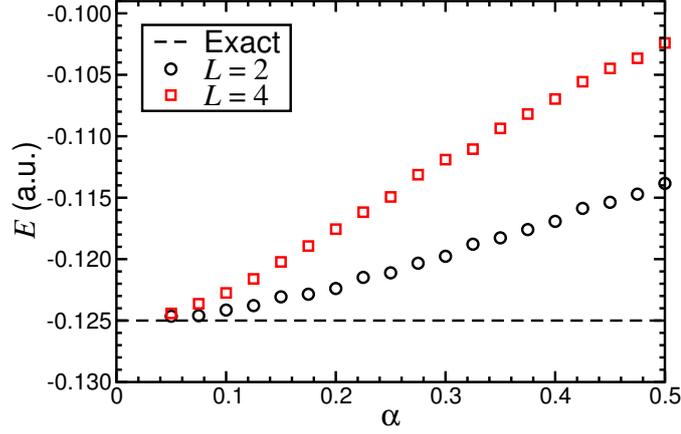


Figure 4.3: DMC first-excited-state energies of an H atom with the trial wave function $\psi_L^{\alpha,\beta}$ [see Eqn. (4.12)] for $L = 2$ and 4 at various amplitudes α of wrinkling of the nodal surface [see Eqn. (4.12)]. DMC error bars are of order the size of the symbols.

attempt to improve the nodal surface in a trial excited state. Even if the optimisable parameters of a trial excited-state wave function cannot change the nodal topology, optimisation by energy minimisation may result in the development of a pathological nodal surface that gives a DMC energy that is too low.

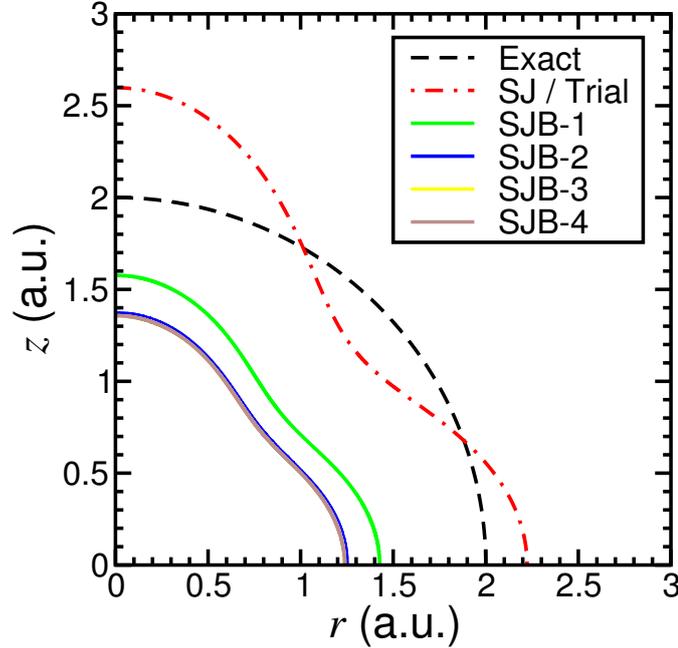


Figure 4.4: Nodal surface of the SJB trial wave function $\psi_{L=4}^{\alpha=0.1,\beta}$ [see Eqn. (4.12)] for the first excited state of an H atom. The wave function is optimised by VMC energy minimisation. SJB- n labels the nodal surface of the SJB wave function after the n^{th} cycle of energy minimisation. The $n = 3$ and 4 cases are indistinguishable from each other, and correspond to the termination of the optimisation process.

We have tested this explicitly for the case of a trial wave function $\psi_{L=4}^{\alpha=0.1,\beta}(r, \theta)$, with an electron-nucleus backflow function. Successive cycles of energy minimisation lower the VMC energy of this state from $-0.1180(2)$ a.u. ($> -\frac{1}{8}$ a.u., positive error) to $-0.1445(3)$ a.u. ($< -\frac{1}{8}$ a.u., negative error). This is exacerbated at the DMC level, where the energy of the state with the optimal backflow function drops further still to $-0.1562(3)$ a.u. Throughout VMC optimisation, the nodal surface alters significantly, as shown in Fig. 4.4. Interestingly, in our test calculations, the VMC nodal surface only ever became smaller as a result of the inclusion of a backflow function. However, the growth of the node (to a point at infinity) would also yield a state which mimics the ground state, and hence lead to a negative nodal error. This is prevented in practice by the inclusion of a cutoff function in backflow displacements.

This investigation of the hydrogen atom suggests that the lack of variational principle for excited-state energies is only a significant problem if one attempts to reoptimise a parameter that moves the nodal surface in an excited state.

4.3.1.2 Ne atom: VMC and backflow?

In terms of computational cost, VMC is several times cheaper than DMC. It would therefore be desirable to know whether or not energy gaps at the VMC level can be of comparable quality to their DMC counterparts. To this end we have calculated the n^{th} ionisation potential of all-electron Ne up to and including $n = 8$, at various levels of theory (SJ-VMC, SJB-VMC, SJ-DMC, and SJB-DMC). It has previously been shown that SJB-VMC is capable of retrieving large fractions (more than 99%) of the correlation energy (defined with respect to the then-best SJB-DMC energy) of the Ne and Ne⁺ species [188]; however, no attempt was made to evaluate the effectiveness of this approach beyond $n = 1$. Our results for the Ne atom are given in Table 4.1, alongside corrected non-relativistic literature values [14].

As can be seen, the DMC ionisation potentials match very closely the “exact” non-relativistic results. The general trend that more sophisticated levels of theory

n	Exact	Ionisation potential (eV)				Error in ionisation potential (eV)			
		SJ-VMC	SJB-VMC	SJ-DMC	SJB-DMC	SJ-VMC	SJB-VMC	SJ-DMC	SJB-DMC
1	21.61333	22.08(2)	21.96(2)	21.72(1)	21.72(1)	0.350	0.104	0.109	
2	40.99110	41.48(2)	41.39(2)	41.10(1)	41.06(1)	0.397	0.108	0.074	
3	63.39913	63.44(2)	63.23(1)	63.35(2)	63.39(1)	-0.173	-0.050	-0.010	
4	97.29312	97.91(2)	97.78(1)	97.75(2)	97.72(1)	0.489	0.458	0.424	
5	126.28846	126.85(2)	126.72(2)	126.85(1)	126.79(1)	0.436	0.564	0.504	
6	157.80001	158.43(2)	158.30(1)	158.25(2)	158.34(1)	0.630	0.453	0.545	
7	207.04137	204.48(2)	204.56(1)	205.04(2)	205.26(1)	-2.561	-2.005	-1.786	
8	238.78949	238.10(1)	238.49(1)	238.70(2)	238.79(1)	-0.687	-0.089	0.002	
MAE	0%	0.83%	0.67%	0.38%	0.34%				

Table 4.1: n^{th} ionisation potential of an all-electron Ne atom at various levels of QMC theory, together with corrected non-relativistic experimental values [14]. The mean absolute errors (MAEs) have been calculated over all ionisation potentials obtained within a given level of theory.

capture more of the correlation energy in excited states is observed, in that the MAE follows the expected trend: SJ-VMC does very well, SJB-VMC does better, SJ-DMC does better still, and SJB-DMC is our best method. In this case, the system is absent of vibrational effects and relativistic effects have been removed from the experimental data. Hence the major source of error in the DMC calculations is fixed-node effects. To test the impact of fixed-node error on our ionisation potentials, we have performed a test calculation with a SJB wave function which was reoptimised in the Ne^+ cationic state. Ionisation potentials are differences in ground-state energies for different numbers N of electrons, and hence fixed-node error is always positive in each of the two energies involved in forming the difference. We find that the SJB-VMC and SJB-DMC first ionisation potentials are 21.51(1) eV and 21.73(1) eV, respectively. The SJB-DMC first ionisation potentials with and without reoptimisation are consistent with each other. On the other hand, the SJB-VMC first ionisation potentials with and without reoptimisation are 21.51(1) eV and 21.96(2) eV respectively [with MAE values of 0.47(6)% and 1.62(8)%], and here we see the most improvement from reoptimisation. The MAE of the SJB-DMC result is 0.52(5)%, meaning that the results from SJB-VMC and SJB-DMC with reoptimised backflow functions are effectively as good as each other—although SJB-VMC underestimates and SJB-DMC overestimates the ionisation potential.

A recent coupled cluster [CCSD(T)] calculation determined the first and second ionisation potentials of Ne as 21.564 eV and 44.3 eV, respectively [absolute errors of 0.04930 eV (0.23%) and 3.30890 eV (8.1%) with respect to the “exact” non-relativistic results that we have compared against] [189]. A less recent configuration interaction calculation determined the eighth ionisation potential of Ne as 238.78440 eV [absolute error of 0.00509 eV (0.0021%)] [190].

4.3.2 Molecules

4.3.2.1 H₂ dimer

We have evaluated the SJ-DMC first ionisation potential of the H₂ dimer using orbitals expanded in plane-wave and Gaussian basis sets. Our plane-wave calculations employed Trail-Needs pseudopotentials, while our Gaussian basis set calculations were all-electron. In our all-electron calculations, we have used bond lengths matching the G2 values [191]. In the pseudopotential calculations, we have relaxed geometries in the ground (and excited, where specifically mentioned) states in DFT with the use of the PBE exchange-correlation functional.

We have also carried out plane-wave-basis all-electron calculations, where the full Coulomb interaction was used to evaluate the DFT total energy. Such calculations are prohibitively expensive for atoms beyond C, requiring very large plane-wave cutoff energies to achieve reasonable convergence of total energies. We have carried out total-energy convergence tests for this system, the results of which informed our choice of plane-wave cutoff in orbital-generation calculations (500 a.u.). We estimate the error in DFT total energies due to this choice of plane-wave cutoff energy to be $\sim 2 \times 10^{-3}$ a.u., and much smaller in DMC (where cusp corrections [31] act to correct the wave function behaviour at short range, which is the most difficult region to represent in a plane-wave basis). Our findings are displayed alongside experimental and other theoretical estimates in Table 4.2.

It is clear that the use of pseudopotentials has some bearing on the quality of the excitation results, but also that structural and vibrational effects are critically important, as evidenced by the strong reduction of the ionisation potentials upon relaxation of the excited-state geometry.

Experimental zero-point energies suggest that a reduction in the calculated ionisation potential of H₂ of around 0.02 eV is appropriate to properly allow for comparison with experiment [200]. This is not enough to fully bridge the gap between our best SJ-DMC results and the experimental ones. However, we have used DFT-derived geometries, and have already shown that the use of pseudopotentials

Method	Ionisation potential (eV)	
	H ₂	O ₂
SJ-DMC (AE-PW)	16.465(3)	–
SJ-DMC (AE-G)	16.462(6)	13.12(7)
SJ-DMC (PP-PW)	16.377(1)	12.84(2)
SJ-DMC (PP-PW-ER)	15.582(1)	12.33(2)
J-DMC (p ⁺ p ⁺ e ⁻ e ⁻)	15.4253(7)	–
QSGW	16.04 [192], 16.45 [193]	–
CC-EPT	–	12.34,12.43 [194]
MP2	–	11.72 [195]
CCSD	–	11.76,12.13 [196]
CCSD(T)	–	11.95 [196]
QCISD(T)	–	12.18 [195]
JCE	15.42580 [197]	–
Experiment	15.4258068(5) [198]	12.0697(2) [199]

Table 4.2: DMC ionisation potentials of the H₂ and O₂ dimers. All-electron (AE) and pseudopotential (PP) calculations have been performed with Gaussian (G) and plane-wave (PW) bases. Calculations employing relaxed excited-state geometries are denoted “ER.” The “J-DMC (p⁺p⁺e⁻e⁻)” calculations used a Jastrow wave function to describe the ground state of two distinguishable quantum protons and two distinguishable electrons for parahydrogen H₂, and the ground state of two distinguishable protons and one electron for the parahydrogen cation H₂⁺. Self-consistent quasiparticle *GW* results are denoted “QSGW,” coupled cluster results with single, double, and (triple) excitations “CCSD(T)” (“EPT” means electron propagator theory), second-order Møller-Plesset perturbation theory results “MP2,” quadratic configuration interaction “QCI” (with levels of excitations as with coupled cluster), and results obtained by means of the generalised James-Coolidge expansion “JCE.”

incurs an error of order the remaining difference between the (pseudopotential) SJ-DMC and experimental ionisation potential.

For the simple case of a parahydrogen H₂ molecule (*i.e.*, a molecule with opposite-spin protons) it is feasible to perform DMC calculations in which both the protons and electrons are treated as distinguishable quantum particles. Since the ground states of both the parahydrogen molecule H₂ and the parahydrogen cation H₂⁺ are nodeless, the fixed-node DMC calculations are exact non-relativistic calculations (in the limit of zero time step, etc.). We find the J-DMC total energies of parahydrogen H₂ and the parahydrogen cation H₂⁺ to be $-1.16401(2)$ and $-0.5971396(3)$ a.u., respectively.⁵ As shown in Table 4.2, the resulting ionisation

⁵To extrapolate the J-DMC H₂ energy $E(\tau)$ to zero time step we used nine different time

potential then agrees with experiment to within 0.01 eV. Another experimental study was able to resolve a para-ortho splitting of 19(9) μeV in the ionisation potential, and determined the first ionisation potential of parahydrogen specifically as 15.425808(6) eV [201], a value which is consistent with the averaged result of Ref. [198].

The results shown here demonstrate the critical importance of nuclear geometry and vibrational effects on energy gaps on a subelectronvolt scale. To obtain excellent agreement with the experimental ionisation potential of H_2 in *ab initio* DMC calculations it was necessary to treat both the electrons and the protons as quantum particles. Even for heavier atoms than hydrogen, it is unreasonable to expect quantitative agreement with experiment in the absence of vibrational corrections.

4.3.2.2 O_2 dimer

We have performed static-nucleus SJ-DMC ionisation-potential calculations for the O_2 molecule, similar to the calculations described in Section 4.3.2.1. Our results are shown in Table 4.2.

The triplet ground-state ${}^3\Sigma_g^-$ electronic configuration was used to obtain the results given in Table 4.2, with a geometry obtained from structural relaxation of the triplet state in spin-polarised DFT, and with explicitly spin-polarised single-particle orbitals populating the single Slater determinant of orbitals in the trial wave function.⁶ However, we have also evaluated the ${}^1\Delta_g$ singlet-state energy, evaluated with a geometry obtained from structural relaxation of the singlet state

steps τ , ranging from 0.0005 a.u. to 0.032 a.u., and we found the time-step bias to consist of a crossover between two different linear regimes. This is because there are two small length scales in the problem: the Bohr radius and the root-mean-square displacement of the protons in their vibrational ground state. We therefore performed the time-step extrapolation by fitting the Padé form $E(\tau) = [E(0) + a\tau + b\tau^2]/(1 + C\tau)$ to our data, where $E(0)$, a , b , and C are fitting parameters. We recommend this form of time-step extrapolation in other DMC calculations in which there is a separation of length scales that results in a crossover between two linear-bias regimes.

⁶We have calculated the energy of the triplet state with and without the use of spin-polarised DFT orbitals, finding that the spin-polarised orbitals provide a DMC total energy which is lower, but by a statistically insignificant amount [0.016(16) eV]. We have given results with the spin-polarised orbitals, owing to the physically reasonable nature of their use.

in DFT, finding that it is higher by 1.62(2) eV than the triplet ground-state energy. This is rather higher than the experimental splitting between these two spin configurations of 0.9773 eV [202].

There is an important way in which the single-determinant wave function we have thus far used to describe the singlet state of O₂ might be inadequate. The singlet state is degenerate at the single-particle level, and one could in principle find a significantly better singlet wave function by inclusion of all symmetry-allowed determinants in the subspace of these degenerate states: at the single-determinant level, the DMC energy of the singlet state is essentially arbitrary. We have performed multideterminant DMC calculations for the singlet state, forming a few-determinant expansion with spin-unpolarised DFT orbitals populating the Slater part of the trial wave function, and find that the multideterminant singlet ground state energy is lower in energy by 1.37(2) eV with respect to the single-determinant singlet state. The DMC singlet-triplet splitting of O₂ is then 0.20(3) eV, which is significantly lower than the previously quoted experimental value of 0.9773 eV [202].

The underestimate of the singlet-triplet splitting reflects the fact that the singlet trial wave function has more variational freedom via the use of multiple (degenerate) determinants. We could easily improve the triplet wave function by forming a multideterminant expansion using non-degenerate determinants. However, this illustrates a general difficulty with the use of multideterminant wave functions in QMC calculations of energy differences. Most QMC calculations rely on a cancellation of fixed-node errors and in general it is difficult to provide multideterminant wave functions of equivalent accuracy for two different systems.

A related difficulty is encountered in Section 4.3.3.2, where backflow functions play a role in revealing potential consequences of multireference character.

4.3.2.3 Non-dimer molecules

The aromatic compounds anthracene ($C_{14}H_{10}$) and benzothiazole (C_7H_5NS) are known to possess sizeable first ionisation potentials, as is boron trifluoride (BF_3). Tetracyanoethylene (C_6N_4), on the other hand, is a strong Lewis acid, with a large electron affinity. With this in mind, we have calculated the ionisation potentials and, where positive, the electron affinities of these molecules using SJ-DMC, with and without the effects of structural relaxation in the excited state at the DFT level. Our results for the first three of these molecules are displayed in Table 4.3. The structures of the molecules we have studied are shown in Fig. 4.5.

Molecule	Ionisation potential (eV)						Electron affinity (eV)					
	SJ-DMC	SJ-DMC (ER)	GW	TDDFT	CCSD(T)	Expt.	SJ-DMC	SJ-DMC (ER)	GW	TDDFT	CCSD(T)	Expt.
$C_{14}H_{10}$	7.35(3)	7.31(3)	7.06 [203]	7.02 _A [204] 7.09 _V [204]	7.52 [205]	7.439(6) _A [206]	0.33(3)	0.45(3)	0.32 [203]	0.53 _A [204] 0.43 _V [204]	0.33 [205]	0.530(5) _A [207]
C_7H_5NS	8.92(2)	8.80(2)	8.48 [203]			8.72(5) _A [208]	-	-	-	-	-	-
BF_3	16.226(6)	16.227(6)				15.96(1) _V [209]	-	-	-	-	-	-

Table 4.3: SJ-DMC ionisation potentials and electron affinities of various non-dimer molecules. Calculations employing relaxed excited-state geometries are designated with “(ER).” Adiabatic gaps are given the subscript “A” and vertical gaps the subscript “V.”

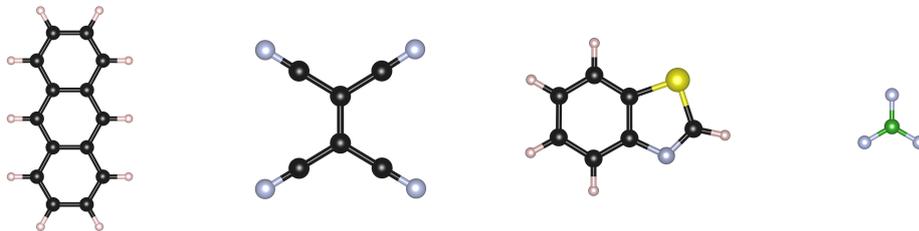


Figure 4.5: Non-dimer molecules whose energy gaps we have calculated. From left to right: anthracene ($C_{14}H_{10}$), tetracyanoethylene (C_6N_4), benzothiazole (C_7H_5NS), and boron trifluoride (BF_3).

As an example of an excitonic gap in a molecule, we have evaluated the first singlet and triplet excitation energies of anthracene at the SJ-DMC level. We find that the singlet excitation energy is 3.07(3) eV, while the corresponding triplet excitation energy is 2.36(3) eV. A recent QMC study obtained a significantly larger (vertical) singlet VMC excitation energy of 4.193(17) eV [4.00(4) eV at the DMC level] [210]; however, the form of trial wave function was qualitatively different, and various details of the underlying geometry-relaxation and orbital-generation

calculations differ from what we have reported here. Available experimental values for the singlet excitations are 3.38 [211] and 3.433 [212], while a single experiment (on molecules in a solvent) has claimed that the triplet excitation energy lies in the range 1.84–1.85 eV [213]. However, comparison is complicated due to the presence of vibrational effects, which generally differ for singlet and triplet excitations.

For the cases of C_6N_4 and BF_3 we have also performed some test SJB calculations. We find that the SJB-DMC ionisation potential of BF_3 is 16.221(4) eV [the difference from the SJ-DMC value of 16.226(6) eV being statistically insignificant], and present our C_6N_4 results in Table 4.4. Backflow correlations have little effect on the calculated ionisation potentials and electron affinities. Nor are the calculated energy differences significantly affected by the reoptimisation of excited-state geometries. We therefore expect that the dominant sources of error in these cases arise from the use of pseudopotentials and (in comparisons with experiment) vibrational renormalisation.

Method	IP (eV)	EA (eV)
SJ-DMC	11.87(1)	3.23(1)
SJ-DMC (ER)	11.85(1)	3.25(1)
SJB-DMC	11.88(1)	3.20(1)
SJB-DMC (ER)	11.86(1)	3.23(1)
SJB(R)-DMC	11.87(1)	–
SJB(R)-DMC (ER)	11.84(1)	–
<i>GW</i>	11.192–12.517 [205]	3.30–~3.9 [214] 2.732–3.804 [205]
CCSD(T)	11.99 [215]	3.05 [215]
Expt.	11.79(5) _V [216] 11.765(8) _A [218]	3.16(2) _A [217]

Table 4.4: DMC ionisation potentials (IPs) and electron affinities (EAs) of C_6N_4 at various levels of QMC theory, compared to experiment and other methods. Calculations employing relaxed excited-state geometries are designated “(ER),” and those employing reoptimised backflow functions “(R).” Adiabatic gaps are given the subscript “A,” vertical gaps the subscript “V.”

Excitation	SJ-DMC gap (eV)			
	$2 \times 2 \times 2$ supercell	$3 \times 3 \times 3$ supercell	$4 \times 4 \times 4$ supercell	FS corr. and av.
$\Delta_{\text{QP}}(\Gamma_v \rightarrow \Gamma_c)$	3.56(6)	3.9(2)	4.0(2)	4.18(6)
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow \Gamma_c)$	3.57(4)	3.82(9)	3.9(1)	4.14(3)
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow X_c)$	1.24(4)	–	1.8(1)	1.9(1)
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow L_c)$	2.39(4)	–	2.8(1)	2.97(7)
$\Delta_{\text{Ex}}(X_v \rightarrow X_c)$	4.55(4)	5.01(8)	5.1(1)	5.3(1)
$\Delta_{\text{Ex}}(L_v \rightarrow L_c)$	3.77(4)	4.00(8)	4.2(1)	4.35(4)

Table 4.5: Uncorrected quasiparticle and excitonic energy gaps Δ_{QP} and Δ_{Ex} of Si evaluated in SJ-DMC for different simulation supercells.

4.3.3 Three-dimensional solids

4.3.3.1 Diamond Si

Silicon in the diamond structure is an indirect-band-gap semiconductor with a valence-band maximum at the Γ point (Γ_v) in the FCC BZ and a conduction-band minimum at around 85% of the distance along the line $\overline{\Gamma X}$. Extensively studied over the past few decades by experimentalists and theorists alike, Si provides an ideal test-bed on which to benchmark QMC band-gap results. To this end, we have calculated the excitonic gaps of Si between various high-symmetry points in the BZ. Specifically, we have considered promotions from $\Gamma_v \rightarrow \Gamma_c$, $\Gamma_v \rightarrow L_c$, $\Gamma_v \rightarrow X_c$, $L_v \rightarrow L_c$, and $X_v \rightarrow X_c$. Calculations of the $\Gamma_v \rightarrow L_c$ and $\Gamma_v \rightarrow X_c$ excitonic gaps are forbidden in the $3 \times 3 \times 3$ supercell, where no choice of supercell reciprocal lattice vector \mathbf{k}_s can ensure that both L and X appear simultaneously with Γ in the $3 \times 3 \times 3$ grid of \mathbf{k} points used to generate our single-particle orbitals. In order to address the issue of FS effects in our energy gaps, we have used simulation supercells comprised of $2 \times 2 \times 2$, $3 \times 3 \times 3$, and $4 \times 4 \times 4$ arrays of primitive cells, and averaged the FS-corrected SJ-DMC results. The exciton binding energy of Si is very weak [15.01(6) meV [219]], and the exciton Bohr radius is much larger than the simulation cells available to QMC calculations. We therefore expect the excitonic and quasiparticle gaps to be very similar and to show the same FS scaling. Our energy gap results are given in Table 4.5 and Fig. 4.6.

As a further test of our method and our treatment of FS effects, we have calculated the quasiparticle energy gap at the Γ point. We have also calculated

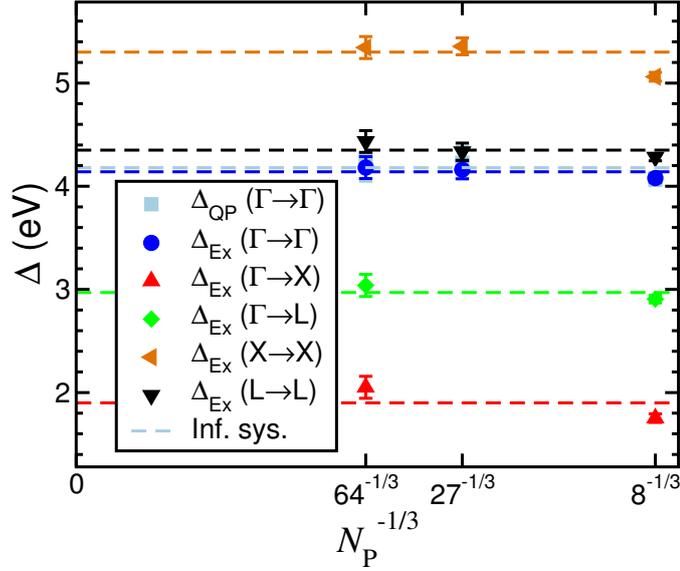


Figure 4.6: Finite-size errors in uncorrected SJ-DMC quasiparticle and excitonic gaps Δ_{QP} and Δ_{Ex} of Si as a function of the number of primitive cells N_P in the supercell. The dashed lines show the infinite-system gaps estimated by subtracting the supercell Madelung constant from the gaps obtained in finite cells and averaging over the different cells.

excitonic and quasiparticle gaps at the Γ point in various differently shaped (non-cubic, but diagonal) supercells.⁷ The results of this investigation are given in Table 4.6, showing the quasirandom variation with cell shape. We have found that the FS effects that exist in our SJ-DMC energy gap data correlate with those obtained from DFT calculations wherein charged defects have been introduced. Specifically, we have calculated the DFT total energies of supercells of intrinsic Si, Si with one P substitution, and Si with one Al substitution, with the total number of electrons fixed to that of the intrinsic Si calculation. This mimics the introduction of two point charges, and a DFT analogue quasiparticle gap can be defined as

$$\Delta_{AQP}^{DFT} = E_P + E_A - 2E_{Si}, \quad (4.13)$$

where E_X is the energy of the Si system with one substitution of atom type X. Our analogue DFT energies have been obtained with a fixed (dense) \mathbf{k} point sampling,

⁷Specifically, non-cubic cells comprised of: $2 \times 1 \times 1$, $3 \times 1 \times 1$, $2 \times 2 \times 1$, $3 \times 2 \times 1$, and $3 \times 3 \times 1$ arrays of primitive cells.

and with ultrasoft pseudopotentials generated on-the-fly in CASTEP.⁸ A plot of $\Delta_{\text{AQP}}^{\text{DFT}}$ against Δ_{QP} obtained from SJ-DMC simulations is given in Fig. 4.7. The correlation is statistically significant with or without the inclusion of the data point corresponding to the smallest cell size. This directly confirms that FS errors in QMC gap calculations are analogous to those in DFT defect-formation-energy calculations.

Supercell	Madelung const. (eV)	SJ-DMC gap (eV)	
		Δ_{QP}	Δ_{Ex}
$2 \times 1 \times 1$	-0.7364	4.00(7)	4.08(4)
$2 \times 2 \times 1$	-0.6009	3.8(1)	3.93(6)
$3 \times 1 \times 1$	-0.4116	4.3(1)	4.34(7)
$3 \times 2 \times 1$	-0.4370	3.8(1)	3.97(8)
$3 \times 3 \times 1$	-0.3342	4.5(2)	4.3(1)

Table 4.6: Finite-size-corrected SJ-DMC vertical quasiparticle gaps Δ_{QP} and SJ-DMC vertical excitonic gaps Δ_{Ex} at the Γ point in Si for various non-cubic supercells. After correction, the Δ_{QP} and Δ_{Ex} data sets have lower variances by a factor of more than two.

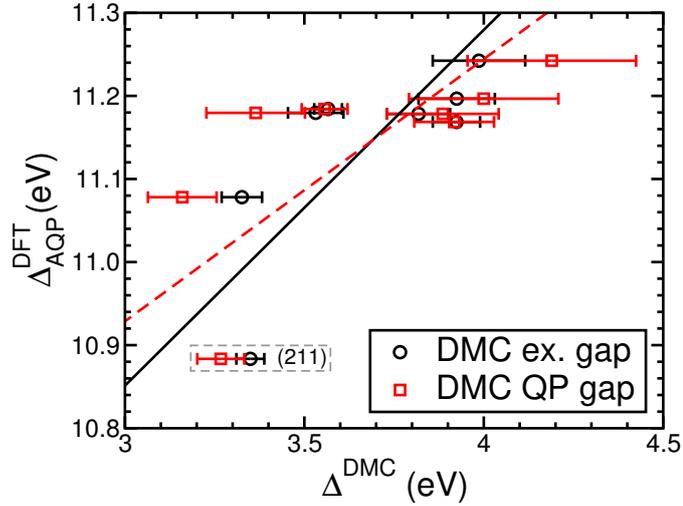


Figure 4.7: Uncorrected SJ-DMC quasiparticle and excitonic energy gaps of Si at Γ , plotted against DFT analogue “quasiparticle” (AQP) gaps, obtained using the defect formation energies for positive and negative charged defects. The results were obtained in different sizes and shapes of periodic cell. The straight lines are linear fits of SJ-DMC gap against DFT AQP gaps.

⁸Versions of CASTEP before 17.2 were subject to a bug which led to incorrect total energies in charged calculations. We have worked with version 17.2, avoiding the undesirable behaviour. We also note that tests with earlier versions of CASTEP indicate that the errors in individual total energies reported by CASTEP do *not* cancel when one calculates a defect formation energy. We thank S. Murphy for drawing our attention to this issue.

All of our DMC calculations for this system have employed time steps of 0.01 and 0.04 a.u., except for our tests in non-cubic cells, and our SJB tests, which employed larger time steps of 0.04 a.u. and 0.16 a.u (with a computational speed-up factor of four). However, we have observed in tests that, in conjunction with the T-move scheme [186], it is possible to use far larger time steps in SJ-DMC gap calculations. The results of these tests are displayed in Fig. 4.8. While time-step bias in total energies is significant at larger DMC time steps (of order a few eV), this bias cancels almost entirely in both excitonic and quasiparticle energy gaps at fixed system size and DMC population size. We expect that the use of even larger DMC time steps in other systems could allow for computational savings of at least an order of magnitude.

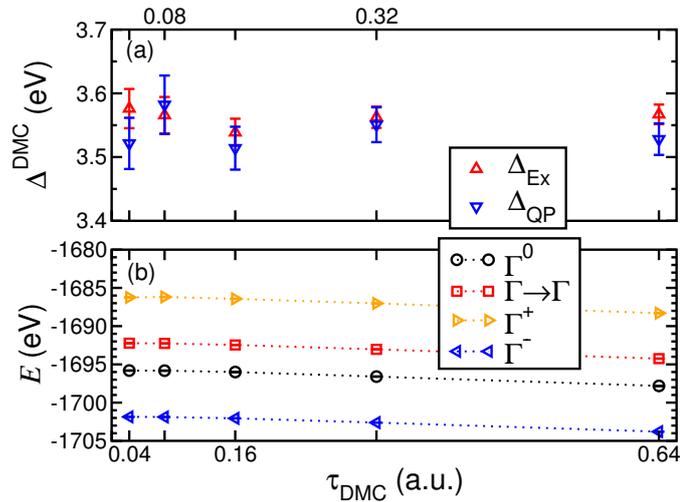


Figure 4.8: Time-step bias in (a) SJ-DMC energy gaps and (b) SJ-DMC total energies for ground (Γ^0), excitonic ($\Gamma \rightarrow \Gamma$), cationic (Γ^-), and anionic (Γ^+) states of Si. All calculations have been performed in a $2 \times 2 \times 2$ supercell with a target population of 256 walkers. The Madelung correction is not included (and would only offset the gaps by a constant).

Our largest family of calculations for Si, those in the $4 \times 4 \times 4$ supercell, required around 1.7 million core hours to complete. Had we opted to employ time steps of 0.04 and 0.16 a.u., which are still conservative choices in light of the information presented in Fig. 4.8, we would have required 0.5 million core hours of computer time.

To address the impact of fixed-node errors in our calculated energy gaps, we

have carried out tests including backflow correlations in our trial wave functions. We find that the inclusion of backflow lowers the (Madelung-corrected) DMC excitonic and quasiparticle gaps in a $2 \times 2 \times 2$ supercell of Si from the SJ-DMC values of 4.08(4) and 4.07(6) eV to the SJB-DMC values of 3.95(1) and 3.95(3) eV, respectively. This is an $O(0.1 \text{ eV})$ effect, which we expect to affect our results at larger system sizes to at least a similar extent. However, to explicitly verify this for the larger cells would incur significant further computational expense. Furthermore, we have considered the impact of reoptimisation of backflow functions in excited states. We find that in the case of the $\Gamma_v \rightarrow \Gamma_c$ quasiparticle gap, this reoptimisation lowers the (Madelung-corrected) SJB-DMC gap even further, from 3.95(3) to 3.77(3) eV in a $2 \times 2 \times 2$ cell. For the $\Gamma_v \rightarrow 0.85X_c$ quasiparticle gap in a $2 \times 2 \times 2$ supercell, reoptimisation lowers the (Madelung-corrected) SJB-DMC gap from 1.66(2) to 1.46(1) eV. In summary, the use of SJ trial wave functions leads to positive fixed-node errors in energy gaps of at least 0.2 eV for Si. In a material with a negligible exciton binding energy such as Si, one may choose to calculate “the gap” as either an excitonic gap or a quasiparticle gap; both exhibit the same FS errors. The quasiparticle gap allows the safe reoptimisation of backflow functions when electrons are added to or removed from a simulation supercell, and furthermore the quasiparticle gap can be calculated between any pair of wave vectors in any supercell. On the other hand, the excitonic gap requires just two QMC calculations to be performed in each simulation cell, rather than three or four for the quasiparticle gap.

A further potential source of fixed-node error at the Γ point arises from the three-fold degeneracy of the light-hole, heavy-hole, and “spin-orbit split-off” bands. Here, a DFT code will output three arbitrary linear combinations of the single-particle orbitals in question. To investigate the possible consequences of this, we have performed SJ-VMC test calculations with trial wave functions formed from three determinants including each of the three degenerate single-particle states at Γ . We find that the formation of a few-determinant expansion has, in this case,

no statistically significant effect on the resultant quasiparticle band energy. We have further investigated the potential impact of degeneracy by repeating these calculations on a grid with $\mathbf{k}_s \neq \mathbf{0}$. Here, the Γ point is not explicitly sampled, but instead the grid is centred on a wave vector of very small magnitude, $\mathbf{k}_s = (\epsilon, \epsilon', \epsilon'')$, so as to break the three-fold degeneracy of the orbitals at Γ . Here, we again find no change in the resultant quasiparticle band energy: if all three determinants are included in the expansion, the total energy of the cationic state at the SJ-VMC level is $-7.8179(1)$ a.u. The total energy of the single-determinant state is (again) $-7.8179(1)$ a.u., while the total energies corresponding to singlet excitations made from the two other (once degenerate) states are $-7.8177(1)$ and (again) $-7.8177(1)$ a.u. The differences are statistically insignificant, and we have therefore eliminated degeneracy as a source of error at the Γ point.

Early QMC studies on solids had claimed some success in the evaluation of band structures and energy gaps. The earliest examples of such calculations [diamond in Refs. [187, 220], Si in Ref. [221], solid atomic (I2₁3) N in Ref. [222], and manganese (II) oxide in Ref. [223]] considered direct calculation of the excitonic gap in small supercells [8 atoms for diamond, Si, and solid N, 16–20 atoms in manganese (II) oxide]. Quasiparticle energy gaps were evaluated, if at all, by means of an addition of an estimate of the exciton binding energy (in the Mott-Wannier model, for example). SJ trial wave functions were used exclusively, and no attempts were made to examine explicitly the nature of FS effects in energy gaps themselves, or to explore fixed-node errors. In common supercell shapes the Madelung constant is typically *negative*, so that a positive correction to quasiparticle gaps is required; this would have been partially offset by fixed-node errors. N.b., the cells used in QMC studies of Si are small compared with the exciton Bohr radius, so FS errors in the excitonic gap behave the same as FS errors in the quasiparticle gap (see Section 4.1.5).

Our QMC quasiparticle gaps in silicon are generally larger than those obtained from *GW* calculations. For example, a recent all-electron G_0W_0 calculation deter-

mined the $\Gamma_v \rightarrow \Gamma_c$, $\Gamma_v \rightarrow X_c$, $\Gamma_v \rightarrow L_c$, $X_v \rightarrow X_c$, and $L_v \rightarrow L_c$ quasiparticle gaps of silicon as 3.07, 0.95, 2.21, 3.46, and 4.09 eV, respectively [224]. A different study determined somewhat larger (pseudopotential) quasiparticle self-consistent *GW* quasiparticle gaps from $\Gamma_v \rightarrow \Gamma_c$, X_c , and L_c as 3.54, 1.60, and 2.41 eV, respectively [225].

4.3.3.2 Hexagonal boron nitride

Bulk hexagonal boron nitride (hBN) is a material similar in structure to graphite, consisting of a series of individual layers of boron and nitrogen atoms occupying the *A* and *B* hexagonal sublattice sites of a two-dimensional honeycomb lattice, where boron atoms lie above nitrogen atoms, and *vice versa* (*AA'* stacking). These constituent layers are weakly bonded, by van der Waals interactions, resulting in both the lubricating properties of hBN and the possibility of isolating monolayers by mechanical exfoliation, as with graphene from graphite. Bulk hBN is semiconducting, with experimental estimates of the band gap ranging from 5.2(2)–7.1(1) eV [226–230]. Watanabe *et al.* find the quasiparticle band gap to be direct and of value 5.971 eV in a single-crystal sample [228]. More recent experimental work by Cassabois *et al.* indicates that bulk hBN is in fact an indirect semiconductor with a quasiparticle band gap of 6.08 eV [230]. Many-body *GW* calculations also indicate that bulk hBN is an indirect-gap semiconductor, with a fundamental gap of 5.95–6.04 eV between the valence-band maximum (close to the K point, along the line $\overline{\Gamma K}$) and the conduction-band minimum at the M point [231–233].

Bulk hBN is believed to exhibit a large exciton binding energy, with values of 0.7–1.2 eV [232–234] predicted by *GW*-Bethe-Salpeter-equation (*GW*-BSE) calculations. Somewhat unexpectedly, experimental measurements find the exciton binding energy to be only 0.13–0.15 eV [228, 230].

We have calculated the quasiparticle and excitonic gaps of bulk hBN between various high-symmetry points in the hexagonal BZ. In performing our QMC calculations on bulk hBN, we used four optimal non-diagonal supercells, each chosen

to maximise the distance between nearest periodic images for a given number of primitive cells [235]. Our supercells contained $N_P = 9, 18, 27,$ and 36 primitive cells. Non-diagonal supercells are here necessitated by the inability to choose a large set of geometrically similar supercells that both maximise the distance between periodic images and have a tractable number of particles. As a consequence, the FS corrections made here employ the formula presented in Eqn. (4.6), with dielectric a tensor obtained from density functional perturbation theory calculations (identical to that presented in Section 3.3.2) [148]. The use of optimal non-diagonal supercells minimises this dominant FS contribution to quasiparticle gaps (the screened Madelung constant), but may leave a slightly non-systematic remnant FS effect of the charge-quadrupole variety. Furthermore, in our optimal cells containing $N_P = 18$ or 36 cells, we are unable to simultaneously include Γ and K in the grid of \mathbf{k} -vectors. In these cases, we have calculated the $\Gamma_v \rightarrow K_c$ excitonic gaps in supercells instead defined by diagonal supercell matrices $S(N_P = 18) = \text{diag}(3, 3, 2)$ and $S(N_P = 36) = \text{diag}(3, 3, 4)$, respectively. Our corrected QMC quasiparticle gaps versus system size are displayed in Fig. 4.9. A table of our final quasiparticle energy gap results, alongside literature estimates and those from density functional theory calculations, are presented in Table 4.7. Our excitonic gaps are given by $\Delta_{\text{Ex}}(\Gamma_v \rightarrow \Gamma_c) = 9.2(2)$, $\Delta_{\text{Ex}}(K_v \rightarrow \Gamma_c) = 8.2(5)$, and $\Delta_{\text{Ex}}(K_v \rightarrow K_c) = 8.3(1)$. Recent experiments give values for $\Delta_{\text{Ex}}(\text{VBM} \rightarrow \text{CBM})$ of 5.822 [228] and 5.955 [230].

The SJ-DMC $K_v \rightarrow K_c$ exciton binding energy of bulk hBN, which is corrected by the subtraction of the screened Madelung constant and then extrapolated against N_P^{-1} to infinite system size [8], is $0.8(1)$ eV. This is consistent with the range of *GW*-BSE values, and is significantly smaller than the monolayer exciton binding energy, as one would expect. The $K_v \rightarrow \Gamma_c$ exciton binding is $0.3(5)$ eV, which is smaller than the statistical error bars.

The SJ-DMC $K_v \rightarrow M_c$ quasiparticle gap is $7.96(9)$ eV. The VBM in bulk hBN is near the K point, while the CBM is at or near the M point [230]. Recent

Method	Δ_{QP} (eV)					
	$\Gamma_{\text{v}} \rightarrow \Gamma_{\text{c}}$	$M_{\text{v}} \rightarrow M_{\text{c}}$	$K_{\text{v}} \rightarrow \Gamma_{\text{c}}$	$K_{\text{v}} \rightarrow K_{\text{c}}$	$M_{\text{v}} \rightarrow \Gamma_{\text{c}}$	VBM \rightarrow CBM
DFT-LDA	6.09	4.54	4.93	4.84	5.28	4.05
DFT-PBE	6.65	4.76	5.42	4.94	5.78	4.28
DFT-HSE06	8.01	6.09	6.54	6.33	6.95	5.55
GW [232]	8.4	6.5	6.9	6.9	7.3	5.95
SJ-DMC	10.1(2)		8.5(2)	9.06(8)		7.96(9)*
Expt.						5.971, [228] 6.08 [230]

Table 4.7: Static-nucleus quasiparticle gaps for bulk hBN, determined by different methods, compared with experimental results. Recall that DFT vibrational-renormalisation calculations indicate that the static-nucleus gaps should be renormalised by -0.40 eV at 300 K. An asterisk (*) denotes the SJ-DMC energy gap from $K_{\text{v}} \rightarrow M_{\text{c}}$.

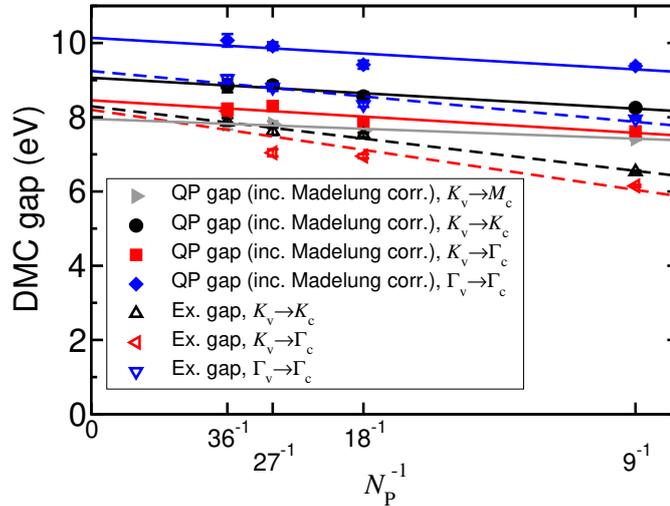


Figure 4.9: SJ-DMC quasiparticle gaps Δ_{QP} and excitonic gaps Δ_{Ex} of bulk hBN against $1/N_{\text{P}}$, where N_{P} is the number of primitive cells in the supercell. The quasiparticle gaps include the Madelung correction as described in the text. The statistical error bars show the random error in the SJ-DMC gap in a particular supercell; the noise due to quasirandom FS effects clearly exceeds the noise due to the Monte Carlo calculation.

DFT calculations yield a zero-temperature vibrational correction in bulk hBN of -0.35 eV (which increases to -0.40 eV at 300 K),⁹ even taking this into account the SJ-DMC quasiparticle gap appears to overestimate the experimental gap of around 6 eV significantly. As a probe of this discrepancy, we have considered (in the $N_P = 9$ supercell) the effects of a backflow transformation of the many-electron wave function. We have also investigated our use of high-symmetry points (K and M) in the BZ rather than the true positions of the VBM and CBM, as determined at the DFT-HSE06 level of theory.

We find that backflow lowers the DMC quasiparticle ($K_v \rightarrow$ CBM) gap of bulk hBN in the $N_P = 9$ supercell by $0.17(5)$ eV. By considering the exact VBM and CBM positions, we find a further energy lowering of $0.02(6)$ eV, which is not statistically significant. Further, we have also considered explicit re-optimisation of backflow functions in anionic and cationic states for the VBM \rightarrow CBM quasiparticle gap. This has recently been shown to lead to significant further lowering of SJB-DMC quasiparticle energy gaps [8]; however, in this case we find that re-optimisation of the backflow functions by minimising the VMC energy actually *raises* the SJB-DMC gap by $0.08(3)$ eV. Collaborators observe the same effect in the hBN monolayer (at the Γ point). Near-degeneracy of the bands at the M point is a possible cause of both the unusual behaviour of the DMC energy in the presence of backflow and the overestimate of the gap. Near-degeneracy can lead to multireference character and hence significant fixed-node errors with a single-determinant wave function.

As a further investigation of possible sources of error in our bulk gaps, we have recalculated the DMC excitonic and quasiparticle gaps at the Γ point (the direct gaps) in an $N_P = 9$ supercell using recently constructed energy-consistent correlated-electron pseudopotentials (eCEPPs) [237]. We find that the difference between the DMC results obtained with the Dirac-Fock pseudopotentials and the eCEPPs is statistically insignificant.

⁹Performed by B. Monserrat, and contained in Ref. [236].

We have also explored more sensitively the nature of the conduction band at the Γ point (Γ_c), which is nearly-free electron-like, and almost completely free in the case of monolayer hBN.¹⁰ In calculations performed by collaborators on monolayer hBN, this is found to be the conduction band minimum. The valence band maximum is at the K point. In practice, however, the energy difference between the direct and indirect excitations is small, and transitions into the conduction band at Γ are around three times weaker than those into the conduction band at K [236]. Monolayer hBN is then, from the point of view of experiment, effectively rendered a direct gap material at the K point. To investigate the effect of crossover from bulk to monolayer on Γ_c , we have calculated band structures and (normalised) band charge densities for a model of bulk hBN with increasingly larger c lattice vector magnitude. Our results are displayed in Fig. 4.10.

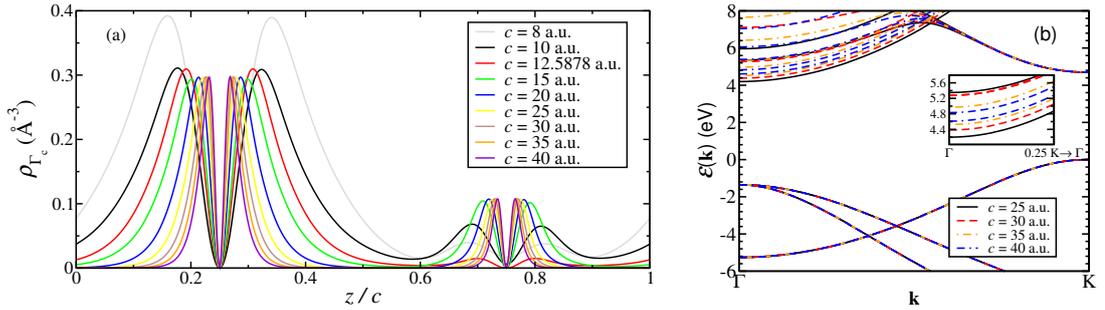


Figure 4.10: (a) DFT-PBE charge density of the state at Γ_c as a function of lattice parameter c for bulk hBN. $c = 12.5878$ a.u. is the experimental lattice parameter [11]. The charge density is plotted along a straight line in the z direction, passing through a boron atom at $z/c = 0.25$ and a nitrogen atom at $z/c = 0.75$. At large c the CBM at Γ_c is an arbitrary linear combination of the degenerate monolayer CBMs. (b) DFT-PBE bulk hBN band structure at three large values of the lattice parameter c . The inset to (b) displays a close-up of the two near-degenerate states at Γ_c .

Two noteworthy aspects of Fig. 4.10 are of immediate interest. Firstly, although in this case the DFT total energy is found to converge for layer separation much smaller than that we have considered (up to 40 a.u.), the band Γ_c is clearly not. The band density still changes significantly even at $c = 40$ a.u. Secondly, the band

¹⁰In monolayers, the band exists very close to the vacuum level in DFT calculations, and is particularly sensitive to the choice of artificial periodicity. It also has a rather isotropic band effective mass of almost unity in the in-plane directions, indicating that it is a nearly free state.

eigenvalues also do not converge for Γ_c . The Kohn-Sham eigenvalues of valence states, and conduction states away from Γ are converged, but there are pairs of eigenvalues near to the Γ point which are near degenerate, but separated, even at $c = 40$ a.u. The splitting of this state is effectively due to the formation of bonding and antibonding states of nearly-free-electrons originating from Γ_c , due to our use of a unit cell containing two layers in our investigation of the monolayer-to-bulk crossover. Regardless of this choice, it remains the case that the charge density, and the nature of the band at Γ_c , is more sensitive to the particular choice of artificial periodicity than one would expect from the convergence of the total energy alone. In conclusion; when performing excited-state calculations on two-dimensional materials, one must carefully check the convergence of conduction states, observing more than merely the total energy.¹¹

4.3.3.3 Cubic boron nitride

Cubic BN has the zincblende crystal structure, with diamond-structure sites alternately occupied by B and N atoms. It is an insulator with a large and indirect fundamental gap from $\Gamma_v \rightarrow X_c$. Experimental estimates of the indirect excitonic gap range from 5.5–7.0 eV [238, 239], and previous DFT investigations give a range for the indirect quasiparticle gap from 4.2–8.7 eV [240, 241]. Theoretical studies based on DFT [242] and on the Bethe-Salpeter equation [243], predict that many-body effects in the absorption spectra of cubic BN are significant, and that a Mott-Wannier exciton formed between the valence and conduction bands at Γ , with binding energy around 0.35 eV, should exist in the bulk material. We have calculated the excitonic energy gaps of cubic BN between the same high-symmetry points as for Si, and have also calculated the quasiparticle gap from $\Gamma_v \rightarrow \Gamma_c$. Our energy gap results for cubic BN are given in Table 4.8. We find that the quasiparticle gap from $\Gamma_v \rightarrow \Gamma_c$ is 12.8(2) eV, but are unable to resolve a statistically significant $\Gamma_v \rightarrow \Gamma_c$ exciton binding energy, because our SJ-DMC error bars are

¹¹We have observed no such similar effect in phosphorene, studied in Section 4.3.4.

~ 0.2 eV, compared to the expected exciton binding of around 0.35 eV. Our value of 7.5(3) eV for the indirect excitonic gap is consistent with the range of experimental estimates.

Excitation	SJ-DMC gap (eV)			
	$2 \times 2 \times 2$ supercell	$3 \times 3 \times 3$ supercell	$4 \times 4 \times 4$ supercell	FS corr. and av.
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow \Gamma_c)$	10.45(4)	11.60(9)	12.06(4)	12.9(2)
$\Delta_{\text{QP}}(\Gamma_v \rightarrow \Gamma_c)$	10.37(5)	11.7(2)	12.00(8)	12.8(2)
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow X_c)$	5.12(4)	–	6.76(5)	7.5(3)
$\Delta_{\text{Ex}}(\Gamma_v \rightarrow L_c)$	11.67(4)	–	13.16(4)	14.0(2)
$\Delta_{\text{Ex}}(X_v \rightarrow X_c)$	10.77(4)	11.85(8)	12.50(5)	13.2(2)
$\Delta_{\text{Ex}}(L_v \rightarrow L_c)$	13.60(4)	14.81(8)	15.37(5)	16.1(2)

Table 4.8: Uncorrected quasiparticle and excitonic energy gaps Δ_{QP} and Δ_{Ex} of cubic BN evaluated in SJ-DMC for different simulation supercells.

4.3.3.4 α -quartz: SiO_2

The α -quartz polymorph of SiO_2 is the most thermodynamically stable at ambient conditions, and hence common. Recent quasiparticle self-consistent GW (QSGW) calculations [244] corroborate earlier theoretical claims [245] that the system hosts a very-well-bound exciton formed at the Γ point in the BZ. The exciton binding energy obtained in Ref. [244] is 1.2 eV, compared with 1.7 eV in Ref. [245]. Experiment finds that the exciton binding is around 1 eV [246]. We have calculated the quasiparticle and excitonic gaps from $\Gamma_v \rightarrow \Gamma_c$, in $1 \times 1 \times 1$ and $2 \times 2 \times 2$ supercells in an attempt to explore this phenomenon. The crystal structure of α -quartz makes the study of larger supercells prohibitively expensive (the unit cell consists of three Si atoms and six O atoms, or 48 electrons when using Trail-Needs pseudopotentials to describe core electronic states). We find that the SJ-DMC quasiparticle and excitonic gaps of α - SiO_2 are 11.4(2) eV and 11.51(7) eV, respectively. We are hence unable to extract a statistically significant exciton binding in α - SiO_2 , perhaps due to the limited sizes of simulation cell that we can study in this case.

4.3.4 Two-dimensional phosphorene

Phosphorene (monolayer black phosphorus) is a two-dimensional material that exhibits a large exciton binding according to *GW*-BSE calculations [247–249], an effective-mass model parameterised by DFT [250], and experimental studies of few-layer black phosphorus on a substrate together with an effective-mass model [251]. Phosphorene consists of phosphorus atoms, four in each unit cell, in a two-dimensional armchair structure with a rectangular Bravais lattice. We used DFT-PBE to obtain a relaxed geometry with lattice parameters $a = 3.31$ Å and $b = 4.56$ Å. As a two-dimensional material, the screened interaction between charge carriers is of Keldysh form, and care is required in the treatment of FS effects (as outlined in Section 5.1.3). The electron and hole effective masses $m_e^* = 0.44 m_0$ and $m_h^* = 0.98 m_0$ may be roughly estimated as geometrical means of the masses in the zig-zag and armchair directions [252]; the vacuum in-plane susceptibility parameter is estimated to be $r_* = 24.24$ Å [250]. The physical size of the exciton in the effective-mass approximation is therefore $r_0 = \sqrt{r_*/(2\mu)} = 4.6$ Å for free-standing phosphorene in vacuum.

Due to the anisotropic nature of phosphorene we studied simulation supercells comprised of 2×2 , 3×2 , 4×3 , 5×4 , and 7×5 primitive cells. Each supercell was chosen to be as square as possible, maximising the nearest-image distance in the space of diagonal supercells. The radii of the largest spheres that can be inscribed in the Wigner-Seitz cells of the simulation supercells are 3.3, 4.6, 6.6, 8.3, 11.4 Å, respectively. Thus we are in the regime in which the Keldysh interaction must be used to evaluate the Madelung correction to the quasiparticle gap (see 5.1.3), with the correction being roughly independent of system size, at least for the smaller cells. We exclude the 2×2 supercell from our extrapolation of the excitonic gap to the thermodynamic limit, since it is too small to contain the exciton. Residual FS errors in the Madelung-corrected quasiparticle gap and in the excitonic gap are expected to scale as $1/L^2$, *i.e.* as $1/N_P$, where N_P is the number of primitive cells, over our range of supercell sizes (this would cross over to $1/L^3$ behaviour if

the supercell size exceeded r_*). We have also studied one non-diagonal supercell containing six primitive cells, which has a slightly larger Wigner-Seitz cell radius (4.9 Å) than the 3×2 supercell. We find that the energy gaps in the non-diagonal cell differ from those obtained in the 3×2 supercell by amounts which are not statistically significant.

Our results for the excitonic gap Δ_{EX} , the quasiparticle gap Δ_{QP} , and the exciton binding energy E_{X}^{b} are shown in Fig. 4.11. After subtraction of the Keldysh-screened Madelung constant, we extrapolate the excitonic gap and Madelung-corrected quasiparticle gap to the thermodynamic limit assuming the error scales as $1/L^2$ (*i.e.* we neglect the effects of the crossover to $1/L^3$ scaling at $L \sim r_*$). For phosphorene specifically, M. Szyniszewski performed the raw (uncorrected) energy gap calculations.

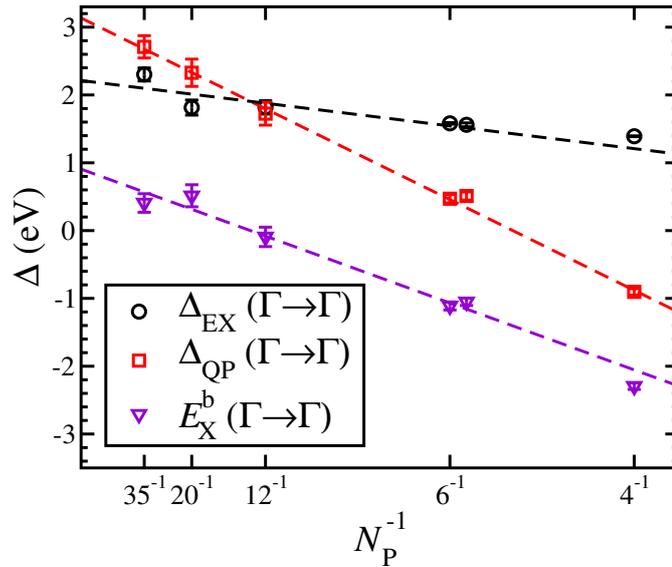


Figure 4.11: DMC quasiparticle gaps Δ_{QP} , excitonic gaps Δ_{EX} , and exciton binding energies E_{X}^{b} at Γ against the inverse of the number N_{P} of primitive cells in the supercell for a free-standing phosphorene monolayer. The Keldysh Madelung constant correction has been applied to the quasiparticle gaps; no FS correction has been applied to the excitonic gaps. The non-diagonal supercell results (filled symbols) have been slightly shifted relative to the 3×2 supercell result for readability.

The resulting energy gaps are slightly larger than previous estimates [247–251] for a free-standing phosphorene monolayer, but our exciton binding energy is consistent with these estimates, as shown in Table 4.9.

Environment	Method	Δ_{QP} (eV)	Δ_{Ex} (eV)	E_{X}^{b} (eV)
Vacuum	(SJ-DMC, extrap. in N_{P}^{-1})	3.13(4)	2.2(2)	0.9(1)
Vacuum	EMA [247–250], EMA+expt. [251]	2.0–2.26	1.2–1.41	0.762–0.85
SiO ₂ subst.	Theory [253, 254]	2.15	1.77	0.38–0.396
SiO ₂ subst.	Expt. [255, 256]	2.05	1.75	0.3

Table 4.9: Comparison of the SJ-DMC energy gaps and exciton binding of monolayer phosphorene with results available in the literature for a free-standing monolayer and a monolayer on an SiO₂ substrate. EMA labels results of an effective mass approximation analysis.

We have explicitly tested the effect of a backflow transformation in the optimal non-diagonal $N_{\text{P}} = 6$ supercell of phosphorene, finding that the inclusion of a backflow transformation (optimised in the ground state) has no statistically significant effect on the DMC energy gaps. The SJB-DMC quasiparticle gap is 0.03(9) eV lower in energy than the SJ-DMC quasiparticle gap, and the SJB-DMC excitonic gap is 0.04(5) eV lower in energy than the SJ-DMC excitonic gap.

Comparison with experiment is complicated by the fact that the exciton binding energy is strongly dependent on the dielectric environment of the monolayer sample. For example, available theoretical [253, 254] and experimental [255, 256] results for phosphorene on a SiO₂ substrate show a decrease in the exciton binding and a larger excitonic gap, as compared to vacuum results. Generally speaking, the larger the relative permittivity of an encapsulating material, the lower the exciton binding. This fact is corroborated by the fitting formulae proposed in Ref. [3], which predicts that exciton (and trion, biexciton...) binding energies decrease monotonically for charges confined to polarising layers which are surrounded by dielectric media of increasing relative permittivity.

A very recent QMC study of phosphorene has explored the use of “hard-wall” boundary conditions for the evaluation of energy gaps, by studying hydrogen-terminated molecular flakes of phosphorene [257]. In this case, the dominant FS effect appears as an $O(L^{-2})$ confinement effect in the kinetic energy of the added or removed charge rather than the slowly decaying image-interaction effect that occurs in a periodic supercell. Interestingly, the authors of Ref. [257] appear to have calculated what we refer to as an excitonic gap (but which they refer to, mistak-

only we believe, as a quasiparticle gap), the energy difference between the ground state and one where a singlet promotion has been made to the single-determinant reference. They subsequently extrapolate data taken at a range of system sizes to the thermodynamic limit, assuming an $\mathcal{O}(N_{\text{P}}^{-1})$ scaling of the FS effect. If we opt to re-interpret their calculations in the way that we see as proper, we conclude that they have calculated the excitonic gap of phosphorene, and correctly *guessed* the FS scaling exponent which is relevant for excitonic gaps, which we have shown to be a consequence of aforementioned exciton image interactions. The result of Frank *et al.* is that the “quasiparticle gap” of phosphorene is around 2.4 eV¹². Our value for the excitonic gap is 2.2(2) eV. We conclude by remarking that we believe that Frank *et al.* have correctly calculated the excitonic gap of phosphorene, and that there is no sound reason for claiming that an excitonic promotion calculation should yield a fundamental gap, or that the fixed-node constraint can in some way prevent the binding interaction of quasiparticles in a continuum quantum Monte Carlo calculation.

For three-dimensional crystals, it is relatively straightforward to remove the $\mathcal{O}(L^{-1})$ FS error in the quasiparticle gap under periodic boundary conditions by using the Madelung correction. The use of finite clusters to approximate the bulk introduces other non-systematic FS errors, such as edge-termination effects. Indeed, the nature of the electronic states involved in the excitation are not necessarily even qualitatively similar to the relevant electronic states in the infinite system. For example, the lowest unoccupied molecular orbital in a diamondoid molecule is a delocalised surface state that does not correspond to the bulk diamond conduction-band minimum [258], and were one to attempt to calculate the band gap of bulk diamond by consideration of larger and larger diamondoids one would have to address this difficulty.

For two-dimensional materials, however, hard-wall boundary conditions pro-

¹²A few values are given, which differ in that the underlying orbitals come from different DFT wave function generation calculations. They are 2.68(10), 2.54(12) and 2.41(17) (the same within error bars).

vide a relatively attractive method for *ab initio* calculations of quasiparticle gaps and the energies of charged excitations. As shown here, calculations in periodic supercells smaller than r_* are absolutely dependent on a Madelung correction evaluated using the Keldysh interaction; since this is roughly constant in cells with $L < r_*$, it is not possible even in principle to extrapolate gaps to the thermodynamic limit. By contrast, gaps obtained in hydrogen-terminated flakes can be extrapolated to infinite size without relying on model interactions. For excitonic gaps the FS errors go as $1/L^2$ under periodic boundary conditions on supercells with $L < r_*$, and hence can be extrapolated if the crossover to $1/L^3$ behaviour is neglected. In this case calculations using periodic boundary conditions maybe preferable, as they are less affected by surface effects.

We emphasise that the need for large periodic cells to describe charged quasiparticles in two-dimensional materials is not an artefact of QMC calculations, but an inevitable consequence of the physics of two-dimensional materials and the Keldysh interaction, which must affect all attempts at *ab initio* gap calculations in these materials. Similar considerations must arise in calculations of charged defect formation energies in layered and two-dimensional materials.

4.3.5 Summary

We have reviewed the use of QMC methods for the calculation of energy gaps in atoms, molecules, and crystalline solids. Although the quasiparticle gap does not formally satisfy a variational principle, we have argued that in practice the fixed-node error in the quasiparticle energy gap is overwhelmingly likely to be positive. Reoptimisation of trial wave functions for systems in which electrons have been added or removed can be expected to improve the calculated quasiparticle energy gaps. For neutral excitations this is not necessarily the case, as was shown in Sec. 4.1.1, and reoptimisation may result in the formation of a pathological excited-state trial nodal surface. Unless the neutral excitation results in a trial wave function that transforms as a 1D irrep of the full symmetry group of the

system and the target state is the lowest-energy eigenstate transforming as that irrep, reoptimisation of the free parameters in the excited-state wave function should not be attempted. Since Jastrow factors do not affect the nodal surface and hence DMC energy, there is little to be gained by reoptimising Jastrow factors in excited states; on the other hand, reoptimising backflow functions in states in which electrons have been added to or removed from the neutral ground state can significantly improve DMC quasiparticle energy gaps.

The use of larger-than-typical DMC time steps for excitation calculations has been shown to be a major source of possible computational savings in DMC energy gap calculations. Time-step bias appears to cancel extraordinarily well in energy gaps. In Si we have made computational savings of a factor of four by using larger time steps in backflow calculations.

Our calculations employing multideterminant trial wave functions for Si at the Γ point show that, even where bands are exactly degenerate, it is not necessarily the case that a few-determinant excited-state wave function comprised of contributions from all possible combinations of degenerate single-particle orbitals performs any better than the single-determinant alternative. On the other hand, such a multideterminant wave function significantly lowers the energy of the singlet first-excited state of O_2 . The need for multideterminant wave functions appears to be more of an issue in studies of excitations in molecules than those in crystals.

We have evaluated energy gaps in atomic, molecular, and crystalline systems using the VMC and DMC methods with single-determinant SJ and SJB trial wave functions. In atomic Ne, where vibrational and finite-size effects are not present, we have achieved highly accurate ionization potentials in comparison with experimental data from which relativistic effects have been removed. The MAE across all of our SJB-DMC calculated ionization potentials for Ne is 0.34%, demonstrating the intrinsic high accuracy achieved by the SJB-DMC method.

In various molecules, where vibrational effects may be present, but finite-size effects never are, we have repeatedly achieved energies which are in reasonable

agreement with their experimental counterparts, with differences attributable to vibrational corrections. We have investigated using DFT to relax excited-state geometries. It too is important, having the largest impact in the H₂ (~ 0.8 eV) and O₂ (~ 0.5 eV) dimers, of all the molecules we have studied. For the parahydrogen molecule we performed DMC calculations of the ionization potential with the protons treated as distinguishable quantum particles, demonstrating excellent agreement with highly-accurate experimental results. This makes clear the fundamental importance of geometrical and vibrational effects when comparing *ab initio* gaps with experiment.

We have probed the effects of fixed-node errors in SJ-DMC energy gap calculations for atoms, molecules, and crystalline solids, finding that the inclusion of backflow functions generally improves DMC energy gaps in these systems (especially in solids, where backflow lowers gaps by 0.1–0.2 eV). We have shown that, in the case of Si, the use of backflow functions reoptimised in anionic and cationic states is crucial in order to achieve reasonable agreement with experiment. Residual overestimates ($\mathcal{O}(0.5$ eV) for first-row atoms) are expected in solids due to the presence of vibrational effects, which are the dominant remaining source of uncertainty when it comes to comparison with experiment. We have also performed gap calculations for free-standing monolayer phosphorene, showing that systematic finite-size effects are qualitatively different in 2D materials, and that an explicit treatment of two-dimensional screening is important in forming corrections (see also Sec. 5.1.3).

Chapter 5

Electron gases in doped two-dimensional semiconductors

There are a number of ways in which the two-dimensional electron gases in electron-doped semiconducting TMDs and in metallic two-dimensional materials differ from those which have been previously realised, *e.g.* in Si MOSFET inversion layers, at III-V semiconductor heterostructure interfaces, at ZnO/ZnMgO interfaces [259,260], and possibly that at the LaAlO₃/SrTiO₃ interface [261].¹ The key difference is that in the atomically flat limit, charge carriers in two-dimensional materials experience much less electrostatic screening, as “the bulk” no longer exists in the same sense. In two-dimensional semiconductors, this lack of out-of-plane screening coupled with the fact that the layers themselves are polarisable leads to the realisation of the already discussed Keldysh interaction (See 3.1.1).

The focus of this chapter will be on the study of 2DEGs wherein charge-carriers interact via a screened interaction which (in the low-density limit) is of Keldysh form. The phase-diagram of the two-dimensional homogeneous electron gas has previously been studied by a variety of authors [262,263], but none of those cases are directly comparable to the physical situation realised in (electron-doped) two-dimensional semiconductors or two-dimensional metals. Modelling such a scenario requires the development of a version of the Keldysh interaction which is compat-

¹Although the mechanism underlying the formation of this 2DEG is still unknown.

ible with periodic boundary conditions, which is where this chapter begins.

5.1 Periodic Keldysh interaction

In order to study electron gases in which the polarisation field of the remainder of a two-dimensional material acts to screen the interaction of the conduction electrons, under the Keldysh model, one must find means of dealing with lattice sums, for particles at \mathbf{r}_i and \mathbf{r}_j , of the form

$$v(\mathbf{r}_i, \mathbf{r}_j) = \sum_{\mathbf{R}} v_{\text{K}}(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|), \quad (5.1)$$

where v_{K} is the Keldysh interaction, and \mathbf{R} denotes the set of two-dimensional lattice vectors of the periodic system in question. Such sums are conditionally convergent by the same arguments as those for the Coulomb interaction. Ultimately, it is the behaviour of the sum at long-range which is problematic, but at distances much greater than the Keldysh interaction screening parameter, r_{\star} , the Keldysh interaction reduces to the Coulomb interaction (inheriting the same problematic long-range behaviour). The Ewald method [264] is one way of overcoming the difficulty posed by such conditionally convergent lattice sums, and is routinely used in electronic structure codes to evaluate lattice sums over the Coulomb potential.

Here, we will present an original derivation of an Ewald-like version of the Keldysh interaction, whose lattice sums are absolutely convergent.

5.1.1 Lattice sum

We may re-write the lattice sum Eqn. (5.1) as follows

$$\sum_{\mathbf{R}} v_{\text{K}}(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|) = \sum_{\mathbf{R}} v_{\text{C}}(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|) + \sum_{\mathbf{R}} [v_{\text{K}}(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|) - v_{\text{C}}(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|)], \quad (5.2)$$

wherein the latter sum is absolutely convergent. Convergence is guaranteed, by the fact that the asymptotic behaviour of the summand is $\mathcal{O}(r^{-3})$ at long-range. Using an asymptotic expansion of the Struve function,² we find

$$v_K - v_C \sim \frac{1}{2r_\star} \sum_{k=1}^{\infty} \frac{\Gamma(\frac{1}{2} + k)}{\Gamma(\frac{1}{2} - k)} \left(\frac{2r_\star}{r}\right)^{1+2k}, \quad (5.3)$$

with Γ the Gamma function.³ Power law summands which go like r^{-2} and weaker lead to conditionally convergent lattice sums in two-dimensions. Incidentally, various power-law interactions which have stronger decays ($r^{-(2+\delta)}$, $\delta > 0$) have two-dimensional lattice sums which are known analytically from their connection to θ -functions [266–268].

This separation made, we may replace the Coulomb lattice sum with the well-known Ewald form, and need only correct the so-called “ $\mathbf{q} = \mathbf{0}$ ” term. This is the term associated with the average value of the electrostatic potential within the supercell, which is fixed to zero. Specifically, in order for the Ewald-Keldysh interaction to have the same property, we must subtract from the Ewald-Coulomb part of the interaction the quantity

$$\begin{aligned} \frac{1}{A} \lim_{\mathbf{q} \rightarrow 0} [v_K(\mathbf{q}) - v_C(\mathbf{q})] &= \frac{2\pi}{A} \lim_{\mathbf{q} \rightarrow 0} \left[\frac{1}{q(1 + r_\star q)} - \frac{1}{q} \right], \\ &= \frac{2\pi}{A} \lim_{\mathbf{q} \rightarrow 0} \left[\frac{-r_\star}{(1 + r_\star q)} \right], \\ &= -\frac{2\pi r_\star}{A}. \end{aligned} \quad (5.4)$$

which is non-zero and finite. This is the average value of the term we are adding to the Ewald-Coulomb interaction, and its subtraction satisfies our need. In the limit of no polarisability, $r_\star \rightarrow 0$, the corrective constant is zero, and we recover the Ewald-Coulomb interaction. An absolutely convergent expression for the lattice

²See DLMF entry: <https://dlmf.nist.gov/11.6.E1>.

³One might worry that the coefficients in the expansion of Eqn. (5.3), and their oscillations for larger and larger k , give rise to a diverging sequence. In fact, if the expansion is truncated at the m^{th} order it is known [265] that the remainder term of the sum, R_m , is of the same sign as the first neglected term, having numerical value *lesser* than the first neglected term.

sum Eqn. (5.1) therefore reads

$$v(\mathbf{r}_i, \mathbf{r}_j) = v_C^{\text{Ew}}(\mathbf{r}_i, \mathbf{r}_j) + \sum_{\mathbf{R}} [v_K(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|) - v_C(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|)] + \frac{2\pi r_\star}{A}, \quad (5.5)$$

where $v_C^{\text{Ew}}(\mathbf{r}_i, \mathbf{r}_j)$ is the Ewald expression for a pair of charges at \mathbf{r}_i and \mathbf{r}_j interacting via the Coulomb interaction in a periodic cell (the Ewald-Coulomb interaction).

Physically speaking, this form of the extended Keldysh interaction is one which has so-called “tin-foil” boundary conditions at infinity. Any polarisation charges that would be present, if the simulation cell had a dipole moment, are then screened by the presence of a fictitious, perfectly metallic ($\epsilon = \infty$), “tin-foil”.

5.1.2 Computational implementation

Two important aspects of the Ewald-Keldysh interaction mean that its computational implementation differs from the conventional Ewald-Coulomb interaction. Firstly, the cusp conditions on identical and non-identical fermion pairs are logarithmic.⁴ Secondly, the lattice sum over $v_K - v_C$, whilst absolutely convergent, contains terms which are asymptotically $\mathcal{O}(r^{-3})$. This contrasts with terms in the Ewald-Coulomb interaction, which are modulated by rapidly-decaying complementary error functions. This means that convergence requires the inclusion of a much greater number of terms in the sum, and/or the inclusion of a correction to account for missing terms.

5.1.2.1 Cusp conditions: \log^{-u} Jastrow term

The cusp conditions on coalescing particles which interact via the logarithmic or Keldysh interaction were derived in [99], and such pairwise Jastrow terms are normally dealt with for few-particle systems via use of the “ex2D” Jastrow term within the CASINO code. Whilst the form of this term is reasonable for bound complexes of charges in the effective mass approximation (exponential localisation

⁴Only in monolayer systems - one can imagine extending the Ewald-Keldysh interaction to the case of bi/multilayers.

is enforced), we do not use it here; instead we add the minimal cusp-satisfying term to the Jastrow u -term in calculations where the Ewald-Keldysh interaction is to be used. Such a minimal term has form (between distinguishable particles i and j)

$$u_{\text{K}}(r_{ij}) = - \left(1 - \frac{r_{ij}}{L_{u\text{K}}}\right)^C \Theta(L_{u\text{K}} - r_{ij}) \frac{q_i q_j m_i m_j}{\underbrace{2r_{\star}(m_i + m_j)}_{\Gamma_{ij}}} r_{ij}^2 \log(r_{ij}), \quad (5.6)$$

where $L_{u\text{K}}$ is the cutoff of the u_{K} -term, C is the truncation order ($C = 3$ ensures continuous first and second derivatives of u_{K}), and m_i, q_i are the mass and charge of particle i . For indistinguishable pairs i and j , the value of the pairing coefficient Γ_{ij} is halved.⁵ In principle, $L_{u\text{K}}$ is an optimisable parameter, whose value ought to be less than unity (else the logarithm permits for the formation of a pathological “repulsive Jastrow factor” for unlike charges). Here, we will consider it fixed at its maximal value (unity).

5.1.2.2 Precomputation and runtime evaluation

The long-range behaviour of the summand $v_{\text{K}} - v_{\text{C}}$ means that we require the inclusion of a large number of starts of lattice vectors in order to achieve acceptable convergence of the lattice sum. Whilst beyond a certain point a continuum approximation is to be made (see 5.1.2.4), it is still necessary to sum a large quantity of terms to ensure the validity of such a continuum approximation (the larger the cutoff on the corrective sum, the more valid the continuum approximation, as stars of lattice vectors are then separated by smaller and smaller distances). We therefore seek a means of pre-computing the parts of the sum associated with this long-ranged behaviour.

We start by splitting the corrective term summand (where we have used the

⁵Similarly with the Coulomb case, as was examined in 1.4.4.

shorthand $\mathcal{S}_{ij}(R)$ to represent $[v_K(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|) - v_C(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|)]$ as follows

$$\mathcal{S}_{ij}(R) = f(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|)\mathcal{S}_{ij}(R) + (1 - f(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|))\mathcal{S}_{ij}(R), \quad (5.7)$$

where f is a function that tends to 1 as $|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}| \rightarrow 0$ and tends to 0 as $|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}| \rightarrow \infty$. The rationale for doing this is so that the first term contains the singular parts of the summand which result from the $\mathbf{R} = \mathbf{0}$ part of the sum (specifically when $|\mathbf{r}_i - \mathbf{r}_j| \rightarrow 0$), and that the second term contains only the non-singular long-range part of the summand. When summed, the first term is a rapidly convergent real-space sum (convergence determined by f) whilst the second is the sum of a long-ranged function. The sum of the second summand is finite, having no singularities, and can therefore be pre-computed over a grid of points within the simulation cell in a setup calculation. We can interpolate on this grid to find the value of the sum at points not on the grid, and add the value obtained to the sum of the first term (which is not computationally expensive, and can be calculated at runtime). We have chosen f to be of gaussian form

$$f(x) = \exp[-\alpha x^2], \quad (5.8)$$

with a parameter α which is tuned to ensure that the FWHM of the gaussian occurs at some fixed distance within the simulation cell. This makes the pre-computed part of the sum as smooth as possible.

5.1.2.3 Blip representation of pre-computed sum

We have chosen to use a B-spline (blip) grid to store the information evaluated in our pre-calculation. A single blip basis function has the form

$$\Phi_\gamma(x, y) = \phi\left(\frac{x - x_\gamma}{a_x}\right) \phi\left(\frac{y - y_\gamma}{a_y}\right), \quad (5.9)$$

with

$$\phi(z) = \begin{cases} 1 - 3\left(\frac{z^2}{2} - \frac{|z|^3}{4}\right), & 0 \leq |z| \leq 1 \\ \frac{1}{4}(2 - |z|)^3, & 1 \leq |z| \leq 2, \\ 0, & \text{else} \end{cases} \quad (5.10)$$

and the value of the pre-computed sum in the simulation cell may be expanded as

$$\sum_{\mathbf{R}} [1 - f(|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|)] \mathcal{S}_{ij}(R) = \sum_{\gamma} b_{\gamma} \Phi_{\gamma}(x, y), \quad (5.11)$$

with γ labelling a point $\mathbf{r}_{\gamma} = (x_{\gamma}, y_{\gamma})$ on the blip grid, a_x, a_y the grid spacings,⁶ and b_{γ} the blip coefficients. We have chosen to define our grid in (fractional) units of supercell lattice vectors.⁷ The value of the blip coefficients may be determined by requiring that the value of the interpolation function at a point on the blip grid (labelled by γ' , say) is exactly equal to the numerically evaluated sum at that point

$$\sum_{\gamma} b_{\gamma} \Phi_{\gamma}(x_{\gamma'}, y_{\gamma'}) = \sum_{\mathbf{R}} [1 - f(|\mathbf{r}_{\gamma'} - \mathbf{R}|)] \mathcal{S}_{\gamma'}(R), \quad (5.12)$$

where $\mathbf{r}_{\gamma'} = (x_{\gamma'}, y_{\gamma'})$ is a vector on the grid, and $\mathcal{S}_{\gamma'}(R)$ is just $\mathcal{S}_{ij}(R)$, evaluated with $\mathbf{r}_i - \mathbf{r}_j = \mathbf{r}_{\gamma'}$. This, with the localisation of the basis functions ϕ taken into account, defines a system of linear equations for the $\{b_{\gamma}\}$

$$\begin{aligned} \sum_{\mathbf{R}} [1 - f(|\mathbf{r}_{\gamma'} - \mathbf{R}|)] \mathcal{S}_{\gamma'}(R) &= \sum_{\langle \gamma, \gamma' \rangle} b_{\gamma} \Phi_{\gamma}(x_{\gamma}, y_{\gamma}), \\ &= \underbrace{b_{\gamma'}}_{\text{Star 0}} + \frac{1}{4} \underbrace{(b_{\gamma'\uparrow} + b_{\gamma'\downarrow} + b_{\gamma'\leftarrow} + b_{\gamma'\rightarrow})}_{\text{Star 1}} \\ &\quad + \frac{1}{16} \underbrace{(b_{\gamma'\nearrow} + b_{\gamma'\nwarrow} + b_{\gamma'\searrow} + b_{\gamma'\swarrow})}_{\text{Star 2}}. \end{aligned} \quad (5.13)$$

where $\langle \gamma, \gamma' \rangle$ means that γ and γ' are pairs of points separated by no more than two grid spacings (contributions from individual ϕ functions are non-zero). We have

⁶Although in our case we will work in fractional coordinates, on uniformly spaced grids.

⁷This is to allow for ease of coding, and also to allow for a clean tessellation of real space. Our sum is periodic in the supercell lattice vectors, and the cell formed by spanning both supercell lattice vectors is guaranteed to tile the plane in a simple fashion.

here assumed that our grid is evenly spaced in both directions, and have labelled a neighbouring grid point with an arrow signalling the direction of the neighbour. Neighbours in “Star 1”, for example, are the points directly above or below the current lattice point (for which $\Phi = \phi(0)\phi(1) = \frac{1}{4}$; hence the pre-factor of $\frac{1}{4}$). At the edges of the cell, the neighbouring grid points are taken modulo 1 in fractional coordinates (we impose periodic boundaries, as our lattice sum is periodic in the supercell lattice vectors). Any points not close enough to an interpolation site to matter (*i.e.* within two grid spacings) are not taken account of in our evaluation routine (only potential neighbours of the kind listed above have their contributions $b_\gamma\Phi_\gamma$ calculated). The solution to the system defined by Eqn. (5.13) is obtained with use of the LAPACK routine `dgesv` [122].

A schematic of the blip grid for a two-dimensional rectangular system with $n = 4$ is displayed in Fig. 5.1.

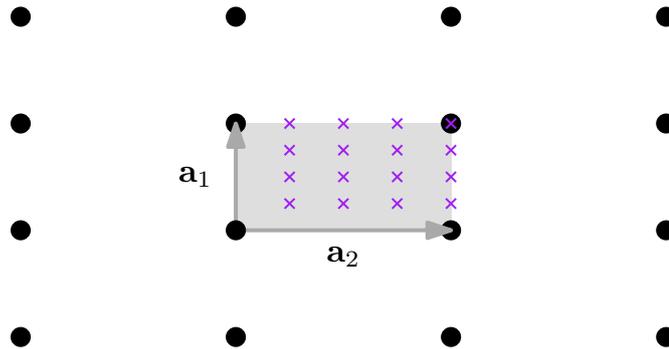


Figure 5.1: Blip grid points (purple crosses) within the simulation cell (shaded grey), for a two-dimensional rectangular lattice with simulation cell lattice vectors \mathbf{a}_1 and \mathbf{a}_2 .

5.1.2.4 Integral correction

The sum in Eqn. (5.11) can only, in practice, be taken over a finite number of two-dimensional lattice vectors. Supposing we sum from the origin until a star of lattice vectors whose magnitudes are strictly less than some cutoff r_c . Beyond the cutoff, we can approximate the remainder of the sum by integrating over a continuous series of contributions weighted by the density of lattice points. A schematic illustrating this idea is given in Fig. 5.2.

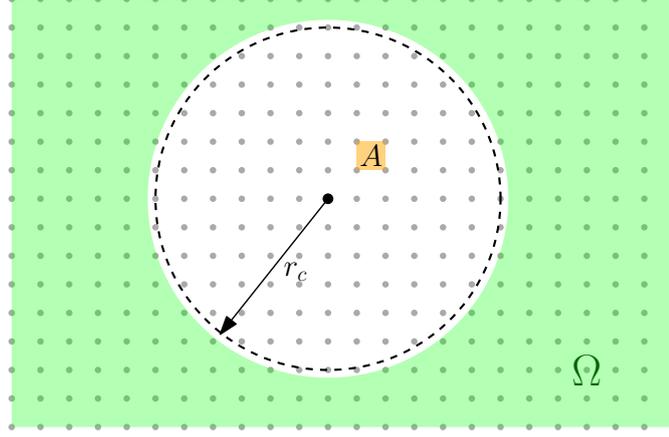


Figure 5.2: A schematic representation of the integral correction scheme for lattice sums appearing in the Ewald-Keldysh interaction. The green region Ω defines the integration domain.

Let the integral corrective term be denoted I , then

$$\begin{aligned}
 I &= \int_{\Omega} \int d^2\mathbf{r} \rho_{\text{LV}} \Delta v(r), \\
 &= 2\pi \int_{r_c}^{\infty} dr r \frac{1}{A} \Delta v(r), \\
 &= \frac{2\pi}{A} \int_{r_c}^{\infty} dr r [v_{\text{K}}(r) - v_{\text{C}}(r)], \tag{5.14}
 \end{aligned}$$

where we have used the fact that the density of lattice points $\rho_{\text{LV}} = \frac{1}{A}$, and assumed that at and beyond r_c our smoothing function f is zero. Further simplification is possible. Making the substitution $\alpha = r/r_*$,

$$\begin{aligned}
 \frac{2\pi}{A} \int_{r_c}^{\infty} dr r [v_{\text{K}}(r) - v_{\text{C}}(r)] &= \frac{2\pi r_*^2}{A} \int_{r_c/r_*}^{\infty} d\alpha \alpha [v_{\text{K}}(\alpha r_*) - v_{\text{C}}(\alpha r_*)], \\
 &= \frac{2\pi r_*}{A} \int_{r_c/r_*}^{\infty} d\alpha \alpha \left[\frac{\pi}{2} \left(H_0(\alpha) - Y_0(\alpha) \right) - \frac{1}{\alpha} \right], \\
 &= \frac{2\pi r_*}{A} \int_{r_c/r_*}^{\infty} d\alpha \left[\frac{\pi}{2} \alpha \left(H_0(\alpha) - Y_0(\alpha) \right) - 1 \right], \\
 &= \frac{2\pi r_*}{A} \left[\alpha \frac{\pi}{2} \left(H_1(\alpha) - Y_1(\alpha) \right) - \alpha \right]_{\alpha=r_c/r_*}^{\alpha \rightarrow \infty}, \\
 &= \frac{2\pi r_c}{A} \left[1 - \frac{\pi}{2} \left(H_1 \left(\frac{r_c}{r_*} \right) - Y_1 \left(\frac{r_c}{r_*} \right) \right) \right]. \tag{5.15}
 \end{aligned}$$

In the last equality, we have used the limit

$$\lim_{x \rightarrow \infty} x \left[\frac{\pi}{2} (H_1(x) - Y_1(x)) - \frac{1}{x} \right] = 0, \quad (5.16)$$

and in the penultimate equality, we have used the standard integral (C an integration constant)

$$\int dx x(H_0(x) - Y_0(x)) = x(H_1(x) - Y_1(x)) + C. \quad (5.17)$$

Our computational routine, as implemented in CASINO, uses the parameterisation of $H_1 - Y_1$ proposed in Ref. [269], which has a maximal absolute error of 2.5×10^{-8} on the entirety of the real line.

5.1.3 Finite-size errors in first-principles calculations on two-dimensional materials

An important first-use of the periodic Keldysh interaction involves the FS correction scheme for first-principles energy gap calculations on two-dimensional materials (see 4.3.4, and Ref. [9]). There, it has been shown that the leading order FS error in quasiparticle energies is given by the (screened) Madelung energy for particles interacting with their periodic images and neutralising background charge. The Madelung energies in those studies were evaluated under the formalism presented here.

Under the earlier prescription, the Madelung energy for a series of self-interacting charges confined in a polarisable two-dimensional layer is given by

$$v_M = v_M^C + \frac{2\pi r_\star}{A} + \sum_{\mathbf{R} \neq \mathbf{0}} [v_K - v_C], \quad (5.18)$$

which, again, is easily computed.

5.2 Crystallisation of the electron gas in doped two-dimensional semiconductors

In metallic systems at low density the phase transition from a fluid (de-localised) to a crystalline (localised) phase is known as the Wigner crystallisation transition, after Eugene Wigner, who offered the first description of it in 1934 [270]. For two-dimensional electron gas systems, the density at which this transition occurs (the *crystallisation density*) has been studied by various authors [262, 271, 272]. Various more recent studies have addressed the possibility for the formation of a Wigner crystal in two-dimensional materials [273, 274].

Here, we describe a study of the Wigner crystallisation transition in a model of electrostatically doped MoSe₂, encapsulated in hBN, with a view towards answering a specific question: is the crystallisation density in such a system so low that it is likely to be irrelevant for experimental study?

Our model is predicated on the use of the periodic Keldysh interaction to describe carrier interactions for conduction electrons in a weakly-to-moderately doped MoSe₂ charge-bearing layer. The effects of the valence electrons on the conduction electrons are to provide carrier effective masses, and background screening of the Keldysh variety (a finite r_* value). We consider a population of electrons which are distributed equally amongst the K and K' valleys, occupying the lowest of the two spin-split bands (*i.e.* a spin-valley polarised two-component gas). One expects that for the ground-state, four-fermion-component physics should only arise either at a density sufficient to lead to complete filling of the lower spin-split bands, or at a temperature high enough such that a significant portion of the population can exist (at thermal equilibrium) in the upper of the two spin-split conduction bands. We will work on a model for the electron gas in thermal equilibrium, at low temperature and density (with no population of the upper spin-split band whatsoever). Estimates for the conduction (valence) band spin splitting are $1 \text{ meV} < \Delta_{\sigma\text{C}} < 50 \text{ meV}$ ($150 \text{ meV} < \Delta_{\sigma\text{V}} < 460 \text{ meV}$) [275–279]. Estimates for

the density n_{cb} at which the upper of the two spin-split conduction bands start to fill exist in the literature. Kormanyos *et al.* [85] studied a series of TMDs, and find (on the basis of hybrid DFT calculations) $n_{\text{cb}}(\text{MoSe}_2) = 4.5 \times 10^{12} \text{ cm}^{-2}$. Pisoni *et al.* [280] inferred $n_{\text{cb}}(\text{MoS}_2) \sim 4 \times 10^{12} \text{ cm}^{-2}$ by measurements of Shubnikov–de Haas oscillations in hBN-encapsulated monolayers of MoS_2 (the corresponding DFT estimate of Kormanyos *et al.* for MoS_2 is notably lower, at $0.54 \times 10^{12} \text{ cm}^{-2}$). The upper range of reasonably obtainable charge densities by electrostatic gating of monolayer samples is $\sim 4 \times 10^{13} \text{ cm}^{-2}$ [281]. Lower estimates on reasonably obtainable electron densities are harder to come by, as the limiting factor here relates strongly with the nature (*i.e.* the electronic structure) and concentration of defects in a given sample [282].

We will study two phases of the system, our de-localised phase is a non-polarised (in the spin or valley sense) Fermi fluid. Our localised phase is a non-polarised Wigner crystal of electrons occupying *stripes* on the sites of a triangular (potential energy minimising) lattice. Up(down)-spin, K(')-valley electrons (say) occupy alternate stripes on the triangular lattice (see Fig. 5.3). This configuration is unlikely to be the true spin-texture of the ground state, which is in fact likely to be a spin liquid, owing to geometrical frustration [283–285], but the energy difference associated with phase competition between spin states is on a much lower scale than that associated with localisation physics, and so the phase boundary for the striped crystal-to-fluid transition is likely a reasonable estimate for the true transition from spin liquid-to-fluid.⁸ In our fluid-phase calculations, we have chosen to work with numbers of electrons which correspond to closed-shell configurations in momentum space (so-called “magic numbers”) under periodic boundary conditions.⁹ In our crystal-phase calculations, we have chosen to work with numbers of

⁸In fact, the difference in energy between different antiferromagnetic spin textures is likely to be much smaller than the difference in energy between an antiferromagnetic crystal phase and a ferromagnetic one. Previously, this (in principle much larger) energy difference was shown to be very difficult to resolve in QMC calculations [262].

⁹A magic number is the number of electrons needed to fill a star of reciprocal lattice vectors in momentum space, where “filling” a \mathbf{k} -point means including plane-wave orbitals at that \mathbf{k} -point in the many-particle trial wave function.

electrons which correspond to exact multiples of the minimal electron content of a single cell for our stripe configuration, but which are as close as possible given this constraint to the corresponding fluid phase electron numbers, and again under periodic boundary conditions. In both cases, the extent of our periodic simulation cells is chosen to exactly reproduce a target electron density. All of our calculations are performed in excitonic units (Eqn. (2.6)), and with material parameters from Table 3.1. Note that there are no holes in these calculations, so in fact we use excitonic units with $\mu = m_e$.

5.2.1 Trial wave functions

5.2.1.1 Crystal phase

Our crystal phase many-body trial wave functions are Slater determinants of Gaussian orbitals, arranged on stripes of a triangular lattice, multiplied by Jastrow factors. A particular lattice point is associated with a single-particle orbital of the form

$$\psi_{\mathbf{P}_\sigma}(\mathbf{r}_i) = \exp[-g|\mathbf{r}_i - \mathbf{P}_\sigma|^2], \quad (5.19)$$

where we denote the Wigner crystal site associated with the i^{th} electron (having spin σ) as \mathbf{P}_σ , and the (optimisable) parameter controlling the width of the Gaussian orbital g . This removes any possible confusion over the underlying lattice vectors for our TMD host material ($\{\mathbf{R}\}$), which now play no role in our continuum model for conduction charges. The \mathbf{P}_σ are distributed as shown in Fig. 5.3.

The Jastrow factors in our crystal calculations are initially comprised of isotropic two-body, and plane-wave expansions in simulation cell reciprocal lattice vectors (u, p terms respectively), which are optimised by variance minimisation [36,38,182]. We then include a backflow function comprised solely of a two-body (η) term, and re-optimize the Jastrow factor and backflow function in tandem with energy minimisation [39,40]. Throughout the optimisation process, we allow the Gaussian parameter g to vary, setting its value initially to equal the ‘‘DMC-optimised’’ value

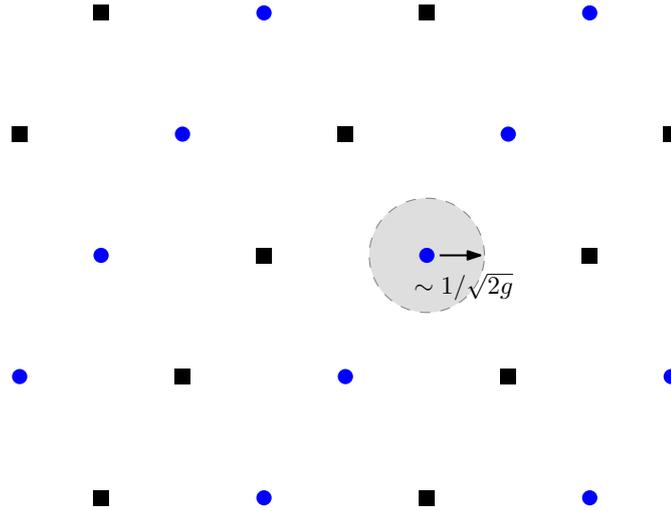


Figure 5.3: Lattice points for the striped configuration of electrons in our crystal phase trial function. Black squares (blue circles) denote up(down)-spin, K' -valley electrons (*i.e.* the $\{P_{\uparrow}\}$ ($\{P_{\downarrow}\}$) set). The grey shaded region denotes one standard deviation of an individual electron orbital, as related to the orbital parameter g .

reported for the two-dimensional HEG in Ref. [262]. Having tested the strategy of optimising g at the DMC level, then optimising a backflow on top of a fixed g value, we find that we are able to obtain better VMC energies by utility of our aforementioned strategy: let all parameters vary, allow backflow parameters and the Gaussian parameter to (ultimately) vary together.

The success of this approach is likely related with the additional freedom it allows in the description of the Gaussian orbital itself. On its own, the backflow function enters the r -dependence of the individual Gaussian orbitals, and hence acts as a means to control the width of the Gaussian, in a somewhat restrictive way. By also allowing the parameter that directly controls the width to vary (g), we effectively create variational freedom in the long-range part of the Gaussian orbital. Such freedom is clearly important in lowering the energy of the crystal phase, at least near the phase boundary, where the de-localised fluid phase competes with the crystal.

In principle, g can be spin-dependent, but we have neglected to include this dependence here, as we find no significant lowering of the VMC energy (no symmetry breaking) when allowing this additional freedom.

5.2.1.2 Fluid phase

Our fluid phase many-body trial wave functions are simply Slater determinants of plane-wave orbitals, $\psi_{\mathbf{k}}(\mathbf{r}_i)$. The grid of \mathbf{k} points on which we perform our calculations is formed as described in 4.1.1. An important difference in this case, however, is that the twist \mathbf{k}_s is of critical importance. Finite-size effects in fluid-phase calculations are extreme, and linked directly with the varying occupation of single-particle states as one changes system size, and/or twist angle \mathbf{k}_s . We describe our strategy for twist averaging our fluid-phase total energies later in this chapter.

The Jastrow factors in our fluid calculations are initially comprised of isotropic two-body, and plane-wave expansions in simulation cell reciprocal lattice vectors (u, p terms respectively), which are optimised by variance minimisation [36,38,182]. We then include a backflow function comprised solely of a two-body (η) term, and again re-optimize the Jastrow factor and backflow function in tandem by energy minimisation [39,40].

Finally, we note that we only consider the paramagnetic case where an equal number of electrons of each spin and valley degree of freedom are present. In the case of the two-dimensional HEG, it has been observed [262] that there is phase competition between the paramagnetic and ferromagnetic fluids for an intermediate density regime around the crystallisation density, but that the ferromagnetic fluid is never stable at any density. At high density, the paramagnetic fluid is always the most stable. However, in the present case, the full inclusion of physics of this variety would also require an additional treatment of inter-valley scattering, and a description of the spin splitting of the conduction bands. Whilst desirable in general, such a study is beyond the scope of the present work.

We can make some basic comments on the likely crystallisation prospects, and those of itinerant ferromagnetism in metallic two-dimensional systems. The Keldysh interaction suppresses the strong Coulomb repulsion, especially at short range. Short-range interactions are most important for the paramagnetic fluid

(unlike spins don't experience an exchange interaction), followed by the ferromagnetic fluid, followed by the Wigner crystal phase, which is largely insensitive to short-range interactions where it is relevant (low density). We expect, then, that starting from the two-dimensional HEG, switching on a finite polarisability leads to, predominantly, an energy lowering of the paramagnetic fluid phase. This acts to move the crystallisation density lower (to a higher r_s value). Assuming the effects of the softened interaction stabilise the paramagnetic fluid the most, as stated above, the transition from paramagnetic to ferromagnetic fluid in a HEG with polarisability (not necessarily under the restricted model we describe) becomes even less likely, and it is therefore unreasonable to expect itinerant ferromagnetism to be any more likely in doped two-dimensional semiconductor systems as compared to the two-dimensional HEG.

5.2.2 DMC calculations

We have performed a series of DMC total energy calculations for both phases over a range of densities. We represent densities by the Wigner-Seitz radius, r_s , which is the radius of a circle which contains one electron (on average). A formula relating r_s in excitonic units to the number density of electrons (in experimentally relevant units of cm^{-2}) reads

$$r_s = \frac{\mu}{\epsilon} \frac{1.06632 \times 10^8}{\sqrt{n}} = \frac{1.01286 \times 10^7}{\sqrt{n}}, \quad (5.20)$$

where the middle equality applies generally, and the right-most term is specific to our model system. This means that, for example, an electron density of $n = 10^{10} \text{ cm}^{-2}$ corresponds to an r_s value of 101.28, and $n = 10^{11} \text{ cm}^{-2}$ to $r_s = 32.02$, in our system. In the context of aforementioned experimentally accessible densities for TMDs, these are low densities.

For our crystal phase calculations, we have used DMC time steps of 2.0 and 8.0 (in excitonic units). Reasonable crystal phase time steps are those for which the

typical diffusion distance $\sqrt{\alpha d \tau}$ (symbols as in Chapter 1) is much smaller than the typical width of an individual Gaussian orbital (which is $\sigma = 1/\sqrt{2g}$). This implies $\tau \ll 1/(4g)$, since the DMC acceptance probability $\alpha \sim 1$ and $d = 2$. Taking g from our optimised VMC wave functions,¹⁰ we find generically that $\tau \ll \mathcal{O}(10^3)$, which is clearly satisfied by our stated choice. We have simultaneously extrapolated time step and population control bias [42], and used population sizes of at least 512 configurations.

In our fluid phase calculations, we have used identical DMC time steps, and population sizes that are smaller by a factor of two. Fluid phase calculations are significantly more expensive, and we do not expect this change to have a significant bearing on the quality of our results. If anything, we expect that our crystal phase calculations are performed for excessively large population sizes. More importantly for fluid phase calculations, however, are single-particle FS effects. Here, we have performed Monte Carlo twist-averaging (MCTA) throughout the course of a DMC calculation. Here, the twist \mathbf{k}_s is changed on-the-fly after a designated number of steps, and the trial wave function re-constructed to match. A period of re-equilibration is undertaken, followed by a further period of statistics accumulation. This is repeated for the desired number of twists. In our calculations, we have averaged over at least 15 twists. We have also applied a statistical efficiency improvement afforded by post-processing of our DMC data [286]. We have used the Hartree-Fock kinetic and potential energy of the Keldysh-screened electron gas as control variates, in contrast to previous works.¹¹

¹⁰After having checked that the backflow function does not significantly couple with g in such a way that g alone is a meaningless measure of spread for the crystal orbitals.

¹¹As an example, in a test calculation at $r_s = 40$ with $N = 16$, and 100 MCTA twists, the statistical improvement factor in control-variate-extrapolated DMC energy is a factor of 88.36289 with the Coulomb exchange energy as the control variate, and slightly higher at 88.48953 with Keldysh exchange energy taken into account. *I.e.* the control variate method for suppressing fluctuations reduces the statistical uncertainty in DMC energies by around two orders of magnitude, with the particulars of screening being negligible for low densities (as expected).

5.2.3 Finite-size effects

We have studied crystal (fluid) systems comprised of $N = 64, 100$ ($N = 62, 110$) electrons. Extrapolation to the thermodynamic limit is therefore a necessity.

In previous studies of two-dimensional electron phases, various authors have presented, and used, FS extrapolation formulae based on physical intuition or assumption. For crystal phases, it was often assumed that the dominant FS error was due to image-interactions (as we have argued is the case for energy gaps in Chapter 4). In two dimensions, such effects go like $\mathcal{O}(N^{-3/2})$. Extrapolation of total energies to the thermodynamic limit for crystal phases under this assumption occurs in Refs. [262, 271, 287]. For fluid phases, arguments first presented in Ref. [271] set the precedent for the assumption that the dominant size effects were $\mathcal{O}(N^{-1})$, and that they were rooted in the size effects incurred predominantly in the potential energy when using finite simulation cells. Such scaling was used to extrapolate total energies to the thermodynamic limit in Refs. [271, 288, 289].

It is now understood that the dominant FS effect in total energies for *both* of these systems originates from the potential energy contribution, and is in fact $\mathcal{O}(N^{-5/4})$ [290]. Such a scaling emerges from the fact that the exchange-correlation hole – the region of suppressed electron density surrounding an electron in a two-dimensional system as a consequence of exchange and correlation effects – is artificially compressed, and distorted, from the use of periodic boundary conditions (and the associated Ewald interaction). We have therefore extrapolated all of our fluid¹² and crystal phase data to the thermodynamic limit assuming a systematic FS effect which is $\mathcal{O}(N^{-5/4})$.

5.2.4 Phase diagram

For the purpose of extracting the crystallisation density, we have fit our fluid and crystal total energy data in the thermodynamic limit to parameterised functions of r_s .

¹²After twist-averaging.

For the fluid, we have fit the difference between the Keldysh total energy at r_s and the Coulomb total energy at r_s (as determined by evaluation of Eqn. (28) of Ref. [289]; with the fitting parameters of Ref. [262]) to a polynomial in r_s . The form of the Keldysh-screened total energy of the fluid at Hartree-Fock level is only modulated by changes to the interaction energy. Such changes may be expressed as providing an additive contribution to the Coulomb interaction energy, and so it is not so unreasonable to simply parameterise the difference as we have done. We have added a polynomial term of the form ar_s^b to the aforementioned Coulomb total energy fits, and allowed a , b to *both* vary. More concrete motivation for our approach is presented at the Hartree-Fock level in Appendix D. We have checked that the parameter b does not vary too much from its optimal low-density HF value of -2 . Allowing slight variation of b does allow for more variational freedom in the total energy fit, and seems like a fair compromise between the full restriction of enforcing a HF density scaling exponent and adding arbitrary, unmotivated, terms to the expansion of the total energy – especially with such a limited data set. Our optimal parameters for the fluid fit are $a = 1.267(2)$ and $b = -1.9197(6)$. Happily, the parameter b does not stray too greatly from the HF scaling exponent of -2 .

In the case of the crystal, we have fit our total energy directly to a polynomial in $1/\sqrt{r_s}$. The first term being $\mathcal{O}(r_s^{-1})$, and with a fixed coefficient so as to correctly reproduce the Coulomb-interaction Madelung constant in the limit of extremely low density. We have additionally fixed the $\mathcal{O}(r_s^{-3/2})$ coefficient so as to reproduce the quasi-harmonic zero-point phonon energy of the triangular lattice. At low density, the Madelung energy of the crystal must approach the Coulomb limit, as for distances much larger than r_* , the Keldysh interaction reduces to the Coulomb interaction. Our remaining crystal phase fitting function is $c/r_s^2 + d/r_s^{5/2} + e/r_s^3$. Our optimal fitting parameters are $c = 2.11(2)$, $d = -4.4(2)$ and $e = 11.5(6)$.

In both cases, and for extraction of the value and uncertainty of the transition density, we have used a Monte Carlo bootstrapping procedure as described in Appendix A.

Our results for the phase diagram of our model for doped, hBN-encapsulated MoSe₂ are displayed in Fig. 5.4. We have determined the crystallisation density to be at $r_s^c = 34.2(1)$, or (propagating errors) $n^c = 8.80(6) \times 10^{10} \text{ cm}^{-2}$.

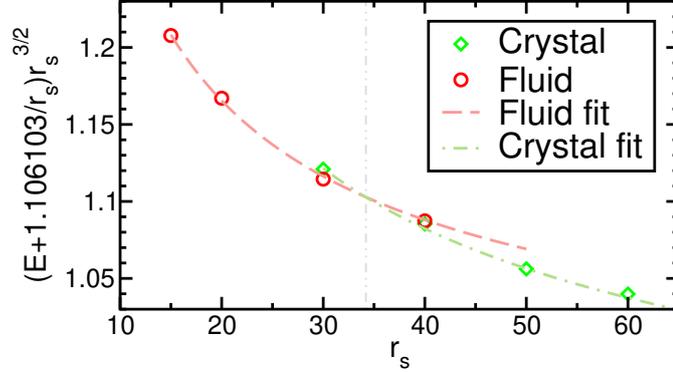


Figure 5.4: DMC energy as a function of density parameter r_s for paramagnetic fluid and Wigner crystal phases of our model.

In comparison with previous studies in the Coulomb limit, our model (in our rescaled units) has only one difference: the form of the interaction term. The screening of repulsive interactions leads, as we have now seen, to a lowering of the crystallisation density for electrons in two dimensions: we witness a slight renormalisation of the Coulomb crystallisation transition point from $r_s^c = 31(1)$ of Drummond *et al.* to $r_s^c = 34.2(1)$. This is *not* primarily from an energy raising of the fluid phase, however. Our crystal energies are also significantly higher than those from studies under the Coulomb interaction.

The crystallisation transition we have established corresponds to an electron density in real units of $n^c = 8.80(6) \times 10^{10} \text{ cm}^{-2}$. This is exceedingly low, and corresponds to a very slight filling of the conduction band in MoSe₂.

In comparison with the results of Zarenia *et al.* [273], who studied crystallisation in TMDs in vacuum in a DFT-based model, we predict a similar crystallisation density. For MoSe₂, Zarenia *et al.* predict crystallisation at a density of $1.5 \times 10^{11} \text{ cm}^{-2}$, a density slightly higher than ours, but they also neglect encapsulation. The electron-electron interaction is stronger in a monolayer in vacuum, and therefore by the same argument we made earlier, it is perhaps not surprising that in this case a higher transition density would be realised. Plainly put, we can

generally conclude that more screening leads to lower transition densities: from the Coulomb two-dimensional HEG of Drummond *et al.*, to the screened MoSe₂ in vacuum of Zarenia *et al.*, to the screened, hBN-encapsulated MoSe₂ that we have studied.

The outlook for Wigner crystallisation in MoSe₂ (and, presumably, other TMDs) is then rather bleak. As an example, consider the experiment of Qiu *et al.* [291] on MoS₂. There, an electron density of $1 \times 10^{11} \text{cm}^{-2}$ was realised (giving hope of perhaps achieving lower densities), but all the while alongside an estimated (S vacancy) defect density of order 10^{13}cm^{-2} . For every electron in such gated samples there can be of order a hundred defects, which would likely cast doubt on any similar measurement claiming a crystallisation transition, for reasons we will now discuss.

5.2.5 Some comments on experimental detection

The unambiguous detection of a Wigner crystal phase of electrons or holes is a challenge which is yet to be fully overcome. Such a state has an obvious hallmark, in that it is insulating, but there are many other insulating phases that may occur in two dimensions, especially at low density. Consider a fluid at low density in a real sample. Whilst theoretically such a state would be conducting, even a modest amount of disorder (in the form of crystalline defects, say) could gap the system, pinning charge carriers to defect sites, and creating an insulating phase.

Disorder aside, other complications involving using homogeneous theory to describe real systems exist. The purported observation of a hole Wigner crystal in GaAs at $r_s \sim 35$ by Yoon *et al.* [292] closely matches the Coulomb interaction QMC predictions [262,271]. However, the applicability of the theory studied in the QMC simulations is questionable in the real-world case, where holes in GaAs are known to exist in non-parabolic bands, and also to experience a strong spin-orbit coupling. Both of these effects act to sow doubt in the comparison between theory and experiment, and whether or not such a state is actually a Wigner crystal of

holes, or a gapped phase borne of defect-trapped conductors is still a matter of debate. That said, the sample of Yoon *et al.* was, in the same study, measured to have an exceptionally large peak mobility at higher carrier density, and so the likelihood of a gapped state emerging solely from the presence of defects does seem lessened.

Essentially, the claimed experimental detection of a Wigner crystal in GaAs amounts to the discovery of a metal-insulator transition at sufficiently low density. The worries over “other” insulating phases are addressed by a supplementary measurement of the (exceedingly large) mobility of the same sample at higher carrier density.

Chapter 6

Future Work

6.1 Superfluidity in electron-hole bilayers: revisited

It has recently been suggested that a superfluid phase of indirect excitons could be observed in van der Waals heterostructures of spatially separated TMD monolayers [293–295]. Existing QMC [296, 297] calculations, aimed at describing the same situation in CQW heterostructures, had an important role in characterising the various interesting pieces of physics at play there.

In the last decade or so, with some contributions originating in the present thesis, it has become apparent that the nature of dielectric screening is perhaps the most important difference between the old class of III-V CQW heterostructures and those based on van der Waals stacks of two-dimensional materials. That said, it would be possible to make the bilayer Keldysh interaction of Chapter 3 “periodic”, in the sense of Chapter 5. This would allow for a more realistic (at least in terms of the interaction potential) revisit to the electron-hole bilayer, in the new and interesting context of van der Waals heterostructures. There is an intense research effort in this avenue, and various open questions about the nature of the condensate exist.

QMC calculations are necessary in the first place for such systems because

the superfluid phases which appear therein are mediated by a complex interplay between long-ranged Coulomb interlayer attraction and intralayer repulsion. The relative differences introduced by the inclusion of realistic screening (namely, the inequivalent softening of these two interactions) mean that significant changes to the phase diagram of the electron-hole bilayer may be realised. Intralayer repulsion suffers the greatest softening, and so it is at least reasonable to expect that the onset of superfluidity in such systems may significantly depart from that in the Coulomb regime.

6.2 Holes immersed in a doped two-dimensional semiconductor

An outstanding topic of research which this thesis enables is the study of excitonic complexes in two-dimensional semiconductors at *finite density*. In Chapters 2 and 3, charge-carrier complexes were studied in the (very) dilute limit. Here, the constituent quasiparticles in a charge-carrier complex interact only with one another, and any effects of a finite concentration of carriers in the conduction band are neglected entirely.

It is possible to consider the case where an appreciable density of conduction electrons exists within the two-dimensional layer, as was considered in Chapter 5, but also where a small population of holes exists. This would enable one to mimic the effects of doping on physical quantities such as the charge-carrier complex binding energy – bridging the gap between unrealistic models which entirely neglect the conduction Fermi sea, and experiment, where a small but nevertheless finite concentration of electrons will always exist.

The use of the periodic Keldysh interaction of Chapter 5 to study, for example, the Mahan exciton in a two-dimensional semiconductor (see *e.g.* [77]), is perfectly possible. In fact, the study presented in Chapter 5 is well suited to this continuation study: a parameterisation of the Fermi energy of the fluid phase for a range

of electronic densities is a necessity in calculating the immersion energy of a hole in an electron gas, and for identifying the crossover from the dilute limit to that in which the hole and the Fermi sea of electrons start to behave collectively. Recent (analytical) theoretical work indicates that this crossover begins to happen almost immediately [298].

Chapter 7

Conclusions

Despite having access to our condensed matter physics “theory of everything”, and the means to access its predictions as presented in Chapter 1, the boundaries of our understanding of real materials are still finite in extent. Here, we have worked to nudge these boundaries outward, in particular with respect to the excited state properties of two-dimensional semiconductors.

In Chapter 2 we uncovered some interesting facts about semiconductor systems comprised of III-V elements, and which are of continued interest to experimentalists. The coupled quantum well supports the binding of a trion for system geometries which are realistic and achievable, and such a state is significantly more stable than the previously considered biexciton. The quantum ring hosts several bound states of electrons and holes, at least theoretically, although we have shown that such states are fragile, and unlikely to play a role in the room-temperature physics of rings.

In Chapter 3 we turned our attention to two-dimensional semiconductors and their heterostructures, deriving various extensions to the so-called Keldysh interaction, and solving the associated few-body problems for charge-carrier complexes. In this respect, Chapters 2 and 3 are special, as they both contain studies of distinguishable particle ensembles, and quantum Monte Carlo methods are hence numerically exact. Here, QMC methods act as a means of obtaining benchmarks of our models themselves – any disagreement with experiment or alternative theoret-

ical methodology is unlikely to be rooted in our solutions to these model problems, but to the models themselves. We have probed one possible limiting feature of our models in our study of *ab initio* dielectric functions, in Section 3.1.4. In modelling van der Waals homostructures, we have shown that the traditional picture of higher charge carrier complexes being formed from effective subsystems (a biexciton as a pair of bound excitons, for example) does not hold ground, and cannot be used to interpret or predict charge-carrier complex binding energies.

Numerically exact QMC was abandoned in Chapter 4, were we sought to describe the excited state properties of atoms, molecules, and solids in a truly *ab initio* manner, by direct Monte Carlo solution of the many-body Schrödinger equation. Exactness is abandoned for two reasons: we here study systems of indistinguishable fermions (hence require the fixed-node approximation), and, where necessary, subject to periodic boundary conditions (hence inherit finite-size errors). There, amongst a slew of other findings, we demonstrated that the fixed-node approximation itself *does* place practical limits on QMC methods in the context of excited states. This is evidenced with the direct illustration of so-called “negative nodal error” in an excited state calculation with a backflow function (Section 4.3.1.1), and with the demonstration of significant lowering (or, even absurd raising) of energy gaps in solid systems where backflow correlations are used in an attempt to alleviate the fixed-node error. In a series of calculations with a happier conclusion, we have performed the first published QMC calculations for the energy gaps and exciton binding of a two-dimensional material, phosphorene. We envision that the results uncovered in this chapter will inform future QMC studies of excitation in real materials, in particular with respect to the treatments of finite-size errors in energy gap calculations.

Finally, in Chapter 5, we derived a variant of the Keldysh interaction which is safely applicable to the study of continuum phases of matter (*i.e.* which has absolutely convergent lattice sums). We then used this “periodic Keldysh interaction” to study the possible formation of an electronic Wigner crystal in a

model for n -doped MoSe₂, establishing a transition density of $r_s^c = 34.2(1)$ (or $8.80(6) \times 10^{11} \text{cm}^{-2}$ in real units). This work also enables further interesting studies concerning excitonic effects in two-dimensional semiconductors at finite density. We further outlined this idea in Section 6.2.

Bibliography

- [1] 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, Feb 2018. [4](#), [67](#), [68](#), [71](#)
- [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, Mar 2019. [4](#), [72](#), [75](#)
- [3] Mostaani, E. and Szyniszewski, M. and Price, C. H. and Maezono, R. and Danovich, M. and Hunt, R. J. and Drummond, N. D. and Fal’ko, V. I., “Diffusion quantum Monte Carlo study of excitonic complexes in two-dimensional transition-metal dichalcogenides,” *Phys. Rev. B*, vol. 96, p. 075431, Aug 2017. [4](#), [20](#), [68](#), [82](#), [86](#), [87](#), [90](#), [91](#), [100](#), [101](#), [106](#), [110](#), [124](#), [126](#), [164](#)
- [4] E. Mostaani, R. J. Hunt, D. M. Thomas, M. Szyniszewski, and N. D. Drummond, “Charge-carrier complexes in monolayer transition metal dichalcogenides: quintons, external fields, and the accuracy of the Keldysh interaction potential.” In preparation. [4](#)
- [5] M. Danovich, D. A. Ruiz-Tijerina, R. J. Hunt, M. Szyniszewski, N. D. Drummond, and V. I. Fal’ko, “Localized interlayer complexes in heterobilayer transition metal dichalcogenides,” *Phys. Rev. B*, vol. 97, p. 195452, May 2018. [4](#)

- [6] F. Vialla, M. Danovich, D. A. Ruiz-Tijerina, M. Massicotte, P. Schmidt, T. Taniguchi, K. Watanabe, R. J. Hunt, M. Szyniszewski, N. D. Drummond, T. G. Pedersen, V. I. Fal’ko, and F. H. L. Koppens, “Tuning of impurity-bound interlayer complexes in a van der Waals heterobilayer,” *2D Materials*, vol. 6, p. 035032, May 2019. [5](#), [94](#)
- [7] R. J. Hunt, M. Szyniszewski, D. M. Thomas, E. Mostaani, N. D. Drummond, and V. I. Fal’ko, “Diffusion Monte Carlo study of charge-carrier complexes in multilayer van der Waals heterostructures.” In preparation. [5](#), [110](#)
- [8] R. J. Hunt, M. Szyniszewski, G. I. Prayogo, R. Maezono, and N. D. Drummond, “Quantum Monte Carlo calculations of energy gaps from first principles,” *Phys. Rev. B*, vol. 98, p. 075122, Aug 2018. [5](#), [156](#), [158](#)
- [9] R. J. Hunt, N. D. Drummond, B. Monserrat, V. Zólyomi, and V. I. Fal’ko, “Diffusion quantum Monte Carlo and *GW* study of the electronic properties of monolayer and bulk hexagonal boron nitride.” In preparation. [5](#), [178](#)
- [10] M. Aghajanian, A. A. Mostofi, and J. Lischner, “Tuning electronic properties of transition-metal dichalcogenides via defect charge,” *Sci. Rep.*, vol. 8, 2018. [14](#), [89](#)
- [11] R. W. Lynch and H. G. Drickamer, “Effect of High Pressure on the Lattice Parameters of Diamond, Graphite, and Hexagonal Boron Nitride,” *J. Chem. Phys.*, vol. 44, no. 1, pp. 181–184, 1966. [18](#), [104](#), [113](#), [133](#), [159](#)
- [12] R. M. Lee, N. D. Drummond, and R. J. Needs, “Exciton-exciton interaction and biexciton formation in bilayer systems,” *Phys. Rev. B*, vol. 79, p. 125308, Mar 2009. [20](#), [67](#), [69](#), [70](#), [71](#)
- [13] R. J. Young, E. P. Smakman, A. M. Sanchez, P. Hodgson, P. M. Koenraad, and M. Hayne, “Optical observation of single-carrier charging in type-II quantum ring ensembles,” *Appl. Phys. Lett.*, vol. 100, no. 8, p. 082104, 2012. [20](#), [71](#), [75](#), [76](#)

- [14] S. J. Chakravorty, S. R. Gwaltney, E. R. Davidson, F. A. Parpia, and C. F. Fischer, “Ground-state correlation energies for atomic ions with 3 to 18 electrons,” *Phys. Rev. A*, vol. 47, no. 5, p. 3649, 1993. [22](#), [140](#), [141](#)
- [15] P. A. M. Dirac, “Quantum Mechanics of Many-Electron Systems,” *P. Roy. Soc. A: Math. Phys.*, vol. 123, pp. 714–733, Apr 1929. [24](#), [28](#)
- [16] 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, no. 5696, pp. 666–669, 2004. [25](#), [26](#)
- [17] J. H. van Vleck, “Nonorthogonality and Ferromagnetism,” *Phys. Rev.*, vol. 49, pp. 232–240, Feb 1936. [28](#)
- [18] W. Kohn, “Nobel Lecture: Electronic structure of matter—wave functions and density functionals,” *Rev. Mod. Phys.*, vol. 71, no. 5, p. 1253, 1999. [28](#)
- [19] O. Perron, “Zur Theorie der Matrices,” *Math. Ann.*, vol. 64, pp. 248–263, Jun 1907. [32](#)
- [20] G. Frobenius, “Über Matrizen aus nicht negativen Elementen,” *Ber. Akad. Wiss. Berlin*, pp. 456–477, Jun 1912. [32](#)
- [21] 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, no. 6, pp. 1087–1092, 1953. [33](#), [43](#)
- [22] W. K. Hastings, “Monte Carlo sampling methods using Markov chains and their applications,” *BioMetrika*, vol. 57, no. 1, pp. 97–109, 1970. [33](#)
- [23] H. Flyvbjerg and H. G. Petersen, “Error estimates on averages of correlated data,” *J. Chem. Phys.*, vol. 91, no. 1, pp. 461–466, 1989. [36](#)
- [24] U. Wolff, “Monte Carlo errors with less errors,” *Comput. Phys. Commun.*, vol. 156, no. 2, pp. 143–153, 2004. [37](#)

- [25] M. Born, “Quantenmechanik der stoßvorgänge,” *Z. Phys.*, vol. 38, no. 11-12, pp. 803–827, 1926. [38](#)
- [26] M. Born and R. Oppenheimer, “Zur Quantentheorie der Molekeln,” *Ann. Phys. (Berl.)*, vol. 389, no. 20, pp. 457–484, 1927. [39](#)
- [27] M. Born and K. Huang, *Dynamical theory of crystal lattices*. Clarendon Press, 1954. [39](#)
- [28] 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, Dec 2004. [46](#), [82](#)
- [29] 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, Dec 2006. [46](#), [49](#), [50](#)
- [30] T. M. Whitehead, M. H. Michael, and G. J. Conduit, “Jastrow correlation factor for periodic systems,” *Phys. Rev. B*, vol. 94, p. 035157, Jul 2016. [46](#), [47](#)
- [31] T. Kato, “On the eigenfunctions of many-particle systems in quantum mechanics,” *Commun. Pure Appl. Math.*, vol. 10, no. 2, pp. 151–177, 1957. [48](#), [82](#), [143](#)
- [32] R. P. Feynman and M. Cohen, “Energy Spectrum of the Excitations in Liquid Helium,” *Phys. Rev.*, vol. 102, pp. 1189–1204, Jun 1956. [49](#)
- [33] M. Holzmann, D. M. Ceperley, C. Pierleoni, and K. Esler, “Backflow correlations for the electron gas and metallic hydrogen,” *Phys. Rev. E*, vol. 68, p. 046707, Oct 2003. [49](#), [50](#)
- [34] R. J. Needs, M. D. Towler, N. D. Drummond, and P. López Ríos, “Continuum variational and diffusion quantum Monte Carlo calculations,” *J. Phys. Condens. Matter*, vol. 22, no. 2, p. 023201, 2009. [51](#), [57](#), [82](#), [134](#)

- [35] L. Zhao and E. Neuscamman, “An Efficient Variational Principle for the Direct Optimization of Excited States,” *J. Chem. Theory Comput.*, vol. 12, no. 8, pp. 3436–3440, 2016. [51](#), [120](#)
- [36] N. D. Drummond and R. J. Needs, “Variance-minimization scheme for optimizing Jastrow factors,” *Physical Review B*, vol. 72, no. 8, p. 085124, 2005. [51](#), [82](#), [181](#), [183](#)
- [37] C. J. Umrigar, K. G. Wilson, and J. W. Wilkins, “Optimized trial wave functions for quantum Monte Carlo calculations,” *Physical Review Letters*, vol. 60, no. 17, p. 1719, 1988. [51](#), [82](#)
- [38] P. R. C. Kent, R. J. Needs, and G. Rajagopal, “Monte Carlo energy and variance-minimization techniques for optimizing many-body wave functions,” *Physical Review B*, vol. 59, no. 19, p. 12344, 1999. [51](#), [181](#), [183](#)
- [39] J. Toulouse and C. J. Umrigar, “Optimization of quantum Monte Carlo wave functions by energy minimization,” *The Journal of Chemical Physics*, vol. 126, no. 8, p. 084102, 2007. [52](#), [181](#), [183](#)
- [40] M. P. Nightingale and V. Melik-Alaverdian, “Optimization of ground-and excited-state wave functions and van der Waals clusters,” *Physical Review Letters*, vol. 87, no. 4, p. 043401, 2001. [52](#), [181](#), [183](#)
- [41] 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, no. 2, pp. 183–190, 1976. [55](#)
- [42] R. M. Lee, G. J. Conduit, N. Nemeč, P. López Ríos, and N. D. Drummond, “Strategies for improving the efficiency of quantum Monte Carlo calculations,” *Phys. Rev. E*, vol. 83, p. 066706, Jun 2011. [56](#), [57](#), [59](#), [134](#), [185](#)
- [43] B. L. Hammond, W. A. Lester, and P. J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry*. WORLD SCIENTIFIC, 1994. [58](#)

- [44] S. Baroni and S. Moroni, “Reptation Quantum Monte Carlo: A Method for Unbiased Ground-State Averages and Imaginary-Time Correlations,” *Phys. Rev. Lett.*, vol. 82, pp. 4745–4748, Jun 1999. 58
- [45] R. Gaudoin and J. M. Pitarke, “Hellman-Feynman Operator Sampling in Diffusion Monte Carlo Calculations,” *Phys. Rev. Lett.*, vol. 99, p. 126406, Sep 2007. 58
- [46] R. Needs, M. Towler, N. Drummond, and P. López Ríos, “User’s guide version 2.13 (2015),” 2019. 59
- [47] 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, no. 4, pp. 2865–2890, 1993. 59
- [48] D. M. Ceperley, “The statistical error of Green’s function Monte Carlo,” *J. Stat. Phys.*, vol. 43, no. 5-6, pp. 815–826, 1986. 59
- [49] 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, no. 2, pp. 1130–1136, 1987. 59
- [50] 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, Jun 2005. 59
- [51] J. H. Lloyd-Williams, R. J. Needs, and G. J. Conduit, “Pseudopotential for the electron-electron interaction,” *Phys. Rev. B*, vol. 92, p. 075106, Aug 2015. 60
- [52] J. R. Trail and R. J. Needs, “Norm-conserving Hartree-Fock pseudopotentials and their asymptotic behavior,” *J. Chem. Phys.*, vol. 122, no. 1, p. 014112, 2005. 60, 133

- [53] 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, no. 17, p. 174109, 2005. [60](#), [133](#)
- [54] M. Casula, S. Moroni, S. Sorella, and C. Filippi, “Size-consistent variational approaches to nonlocal pseudopotentials: Standard and lattice regularized diffusion Monte Carlo methods revisited,” *J. Chem. Phys.*, vol. 132, no. 15, p. 154113, 2010. [60](#)
- [55] 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, Oct 2016. [60](#), [133](#)
- [56] F. Bloch, “Über die quantenmechanik der elektronen in kristallgittern,” *Z. Phys.*, vol. 52, no. 7-8, pp. 555–600, 1929. [61](#)
- [57] N. Ashcroft and N. Mermin, *Solid State Physics*. Philadelphia: Saunders College, 1976. [62](#)
- [58] C. Kittel, *Introduction to Solid State Physics*, vol. 8. Wiley New York, 1976. [62](#)
- [59] R. M. Martin, *Electronic structure: basic theory and practical methods*. Cambridge University Press, 2004. [62](#)
- [60] G. Dresselhaus, “Effective mass approximation for excitons,” *J. Phys. Chem. Solids*, vol. 1, no. 1, pp. 14–22, 1956. [63](#)
- [61] G. H. Wannier, “The Structure of Electronic Excitation Levels in Insulating Crystals,” *Phys. Rev.*, vol. 52, pp. 191–197, Aug 1937. [64](#)
- [62] M. Combescot, O. Betbeder-Matibet, and R. Combescot, “Bose-Einstein Condensation in Semiconductors: The Key Role of Dark Excitons,” *Phys. Rev. Lett.*, vol. 99, p. 176403, Oct 2007. [66](#)

- [63] L. V. Butov, A. C. Gossard, and D. S. Chemla, “Macroscopically ordered state in an exciton system,” *Nature*, vol. 418, pp. 751–754, 2002. [66](#), [71](#)
- [64] S. Yang, A. T. Hammack, M. M. Fogler, L. V. Butov, and A. C. Gossard, “Coherence Length of Cold Exciton Gases in Coupled Quantum Wells,” *Phys. Rev. Lett.*, vol. 97, p. 187402, Oct 2006. [66](#)
- [65] V. B. Timofeev and A. V. Gorbunov, “Bose-Einstein condensation of dipolar excitons in double and single quantum wells,” *Phys. Status Solidi C*, vol. 5, no. 7, pp. 2379–2386, 2008. [66](#)
- [66] Y. Shilo, K. Cohen, B. Laikhtman, K. West, L. Pfeiffer, and R. Rapaport, “Particle correlations and evidence for dark state condensation in a cold dipolar exciton fluid,” *Nat. Commun.*, vol. 4, p. 2335, Aug 2013. [66](#)
- [67] M. Alloing, M. Beian, M. Lewenstein, D. Fuster, Y. González, L. González, R. Combescot, M. Combescot, and F. Dubin, “Evidence for a Bose-Einstein condensate of excitons,” *Europhys. Lett.*, vol. 107, no. 1, p. 10012, 2014. [66](#)
- [68] M. Combescot, R. Combescot, and F. Dubin, “Bose-Einstein condensation and indirect excitons: a review,” *Rep. Prog. Phys.*, vol. 80, no. 6, p. 066501, 2017. [66](#)
- [69] Y. E. Lozovik and V. I. Yudson, “Feasibility of superfluidity of paired spatially separated electrons and holes; a new superconductivity mechanism,” *J. Exp. Theor. Phys. Lett.*, vol. 22, p. 274, 1975. [66](#)
- [70] X. Zhu, P. B. Littlewood, M. S. Hybertsen, and T. M. Rice, “Exciton condensate in semiconductor quantum well structures,” *Phys. Rev. Lett.*, vol. 74, no. 9, pp. 1633–1636, 1995. [66](#)
- [71] M. Gutowski, P. Skurski, A. I. Boldyrev, J. Simons, and K. D. Jordan, “Contribution of electron correlation to the stability of dipole-bound anionic states,” *Phys. Rev. A*, vol. 54, pp. 1906–1909, Sep 1996. [67](#)

- [72] K. D. Jordan and F. Wang, “Theory of Dipole-Bound Anions,” *Annu. Rev. Phys. Chem.*, vol. 54, no. 1, pp. 367–396, 2003. [67](#)
- [73] H. Hao, J. Shee, S. Upadhyay, C. Ataca, K. D. Jordan, and B. M. Rubenstein, “Accurate Predictions of Electron Binding Energies of Dipole-Bound Anions via Quantum Monte Carlo Methods,” *J. Phys. Chem. Lett.*, vol. 9, no. 21, pp. 6185–6190, 2018. [67](#)
- [74] E. Fermi and E. Teller, “The capture of negative mesotrons in matter,” *Phys. Rev.*, vol. 72, no. 5, p. 399, 1947. [67](#)
- [75] J. E. Turner, “Minimum dipole moment required to bind an electron—molecular theorists rediscover phenomenon mentioned in fermi-teller paper twenty years earlier,” *Am. J. Phys.*, vol. 45, no. 8, pp. 758–766, 1977. [68](#)
- [76] R. A. Sergeev and R. A. Suris, “The X^+ trion in a system with spatial separation of the charge carriers,” *Semiconductors*, vol. 37, pp. 1205–1210, Oct 2003. [68](#)
- [77] G. G. Spink, P. López Ríos, N. D. Drummond, and R. J. Needs, “Trion formation in a two-dimensional hole-doped electron gas,” *Phys. Rev. B*, vol. 94, p. 041410, Jul 2016. [68](#), [192](#)
- [78] 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, Jan 2005. [71](#)
- [79] R. Timm, A. Lenz, H. Eisele, L. Ivanova, M. Dähne, G. Balakrishnan, D. L. Huffaker, I. Farrer, and D. A. Ritchie, “Quantum ring formation and antimony segregation in GaSb/GaAs nanostructures,” *J. Vac. Sci. Technol. B*, vol. 26, no. 4, pp. 1492–1503, 2008. [71](#)
- [80] 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,” *Crys. Growth Des.*, vol. 13, no. 3, pp. 1226–1230, 2013. 71
- [81] 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, no. 11, p. 117127, 2014. 71
- [82] 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, no. 4, p. 044305, 2016. 71
- [83] A. Marent, T. Nowozin, M. Geller, and D. Bimberg, “The QD-flash: a quantum dot-based memory device,” *Semicond. Sci. Technol.*, vol. 26, p. 014026, Dec 2010. 71
- [84] 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, Jul 2013. 77, 79
- [85] A. Kormányos, G. Burkard, M. Gmitra, J. Fabian, V. Zólyomi, N. D. Drummond, and V. Fal’ko, “k·p-theory for two-dimensional transition metal dichalcogenide semiconductors,” *2D Materials*, vol. 2, p. 022001, Apr 2015. 77, 180
- [86] Q. H. Wang, K. Kalantar-Zadeh, A. Kis, J. N. Coleman, and M. S. Strano, “Electronics and optoelectronics of two-dimensional transition metal dichalcogenides,” *Nat. Nanotechnol.*, vol. 7, pp. 699–712, Nov 2012. 77
- [87] O. Lopez-Sanchez, D. Lembke, M. Kayci, A. Radenovic, and A. Kis, “Ultra-sensitive photodetectors based on monolayer MoS₂,” *Nat. Nanotechnol.*, vol. 8, pp. 497–501, Jun 2013. 77

- [88] 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 Res.*, vol. 9, pp. 1543–1560, Apr 2016. 77
- [89] E. Mostaani, N. D. Drummond, and V. I. Fal’ko, “Quantum Monte Carlo Calculation of the Binding Energy of Bilayer Graphene,” *Phys. Rev. Lett.*, vol. 115, p. 115501, Sep 2015. 78
- [90] T. Cheiwchanchamnangij and W. R. L. Lambrecht, “Quasiparticle band structure calculation of monolayer, bilayer, and bulk MoS_2 ,” *Phys. Rev. B*, vol. 85, p. 205302, May 2012. 79
- [91] R. Coehoorn, C. Haas, J. Dijkstra, C. J. F. Flipse, R. A. de Groot, and A. Wold, “Electronic structure of MoSe_2 , MoS_2 , and WSe_2 . I. Band-structure calculations and photoelectron spectroscopy,” *Phys. Rev. B*, vol. 35, pp. 6195–6202, Apr 1987. 79
- [92] A. Kumar and P. K. Ahluwalia, “Tunable dielectric response of transition metals dichalcogenides MX_2 (M=Mo, W; X=S, Se, Te): Effect of quantum confinement,” *Physica B*, vol. 407, no. 24, pp. 4627–4634, 2012. 79
- [93] M. K. Agarwal, P. D. Patel, and R. M. Joshi, “Growth conditions and structural characterization of $\text{MoSe}_x\text{Te}_{2-x}$ ($0 \leq x \leq 2$) single crystals,” *J. Mater. Sci. Lett.*, vol. 5, pp. 66–68, Jan 1986. 79
- [94] W. J. Schutte, J. L. De Boer, and F. Jellinek, “Crystal structures of tungsten disulfide and diselenide,” *J. Solid State Chem.*, vol. 70, no. 2, pp. 207–209, 1987. 79
- [95] N. S. Rytova, “Coulomb interaction of electrons in a thin film,” *Dokl. Akad. Nauk. SSSR*, vol. 163, pp. 1118–1120, 1965. 81, 124
- [96] N. S. Rytova, “Screened potential of a point charge in the thin film,” *Vestn. Mosk. Univ. Fiz. Astron.*, vol. 3, pp. 30–37, 1967. 81, 124

- [97] L. V. Keldysh, “Coulomb interaction in thin semiconductor and semimetal films,” *J. Exp. Theor. Phys.*, vol. 29, no. 11, p. 658, 1979. [81](#), [124](#)
- [98] M. Szytniszewski, 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, Feb 2017. [81](#), [82](#), [87](#), [90](#), [100](#), [101](#)
- [99] B. Ganchev, N. Drummond, I. Aleiner, and V. Fal’ko, “Three-Particle Complexes in Two-Dimensional Semiconductors,” *Phys. Rev. Lett.*, vol. 114, p. 107401, Mar 2015. [82](#), [104](#), [110](#), [172](#)
- [100] 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, Mar 2007. [82](#), [134](#)
- [101] J. R. Schaibley, H. Yu, G. Clark, P. Rivera, J. S. Ross, K. L. Seyler, W. Yao, and X. Xu, “Valleytronics in 2D materials,” *Nat. Rev. Mater.*, vol. 1, Aug 2016. [85](#)
- [102] J. Huang, T. B. Hoang, and M. H. Mikkelsen, “Probing the origin of excitonic states in monolayer WSe₂,” *Sci. Rep.*, vol. 6, Mar 2016. [85](#)
- [103] A. Neumann, J. Lindlau, L. Colombier, M. Nutz, S. Najmaei, J. Lou, A. D. Mohite, H. Yamaguchi, and A. Högele, “Opto-valleytronic imaging of atomically thin semiconductors,” *Nat. Nanotechnol.*, vol. 12, pp. 329–334, Jan 2017. [85](#)
- [104] M. Barbone, A. R.-P. 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, Sep 2018. [88](#)

- [105] M. L. Trolle, T. G. Pedersen, and V. Véniard, “Model dielectric function for 2D semiconductors including substrate screening,” *Sci. Rep.*, vol. 7, Jan 2017. [88](#)
- [106] K. S. Thygesen, “Calculating excitons, plasmons, and quasiparticles in 2D materials and van der Waals heterostructures,” *2D Mater.*, vol. 4, p. 022004, Jun 2017. [88](#)
- [107] F. Hüser, T. Olsen, and K. S. Thygesen, “How dielectric screening in two-dimensional crystals affects the convergence of excited-state calculations: Monolayer MoS₂,” *Phys. Rev. B*, vol. 88, p. 245309, Dec 2013. [88](#)
- [108] J. P. Echeverry, B. Urbaszek, T. Amand, X. Marie, and I. C. Gerber, “Splitting between bright and dark excitons in transition metal dichalcogenide monolayers,” *Phys. Rev. B*, vol. 93, p. 121107, Mar 2016. [90](#)
- [109] C. Chakraborty, A. Mukherjee, L. Qiu, and A. N. Vamivakas, “Electrically tunable valley polarization and valley coherence in monolayer WSe₂ embedded in a van der Waals heterostructure,” *Opt. Mater. Express*, vol. 9, pp. 1479–1487, Mar 2019. [90](#)
- [110] R. Geick, C. H. Perry, and G. Rupprecht, “Normal modes in hexagonal boron nitride,” *Phys. Rev.*, vol. 146, no. 2, p. 543, 1966. [92](#)
- [111] M. F. Plass, W. Fukarek, A. Kolitsch, N. Schell, and W. Möller, “Layered growth of boron nitride thin films,” *Thin Solid Films*, vol. 305, no. 1-2, pp. 172–184, 1997. [92](#)
- [112] K.-L. Barth, W. Fukarek, H.-P. Maucher, M. F. Plass, and A. Lunk, “In situ characterization of cubic boron nitride film growth in the IR spectral region,” *Thin Solid Films*, vol. 313, pp. 697–703, 1998. [92](#)
- [113] S. L. Rumyantsev, M. E. Levinshtein, A. D. Jackson, S. N. Mohammad, G. L. Harris, M. G. Spencer, and M. Shur, “Boron nitride (bn),” in *Properties of*

- Advanced Semiconductor Materials GaN, AlN, InN, BN, SiC, SiGe* (M. E. Levinshtein, S. L. Rumyantsev, and M. S. Shur, eds.), ch. 4, pp. 67–92, John Wiley & Sons New York, NY, USA, 2001. [92](#)
- [114] C. Gong, H. Zhang, W. Wang, L. Colombo, R. M. Wallace, and K. Cho, “Band alignment of two-dimensional transition metal dichalcogenides: Application in tunnel field effect transistors,” *Appl. Phys. Lett.*, vol. 103, no. 5, p. 053513, 2013. [93](#)
- [115] J. Kang, S. Tongay, J. Zhou, J. Li, and J. Wu, “Band offsets and heterostructures of two-dimensional semiconductors,” *Appl. Phys. Lett.*, vol. 102, no. 1, p. 012111, 2013. [93](#)
- [116] H. M. Hill, A. F. Rigosi, K. T. Rim, G. W. Flynn, and T. F. Heinz, “Band Alignment in MoS₂/WS₂ Transition Metal Dichalcogenide Heterostructures Probed by Scanning Tunneling Microscopy and Spectroscopy,” *Nano Lett.*, vol. 16, no. 8, pp. 4831–4837, 2016. [93](#)
- [117] M.-H. Chiu, C. Zhang, H.-W. Shiu, C.-P. Chuu, C.-H. Chen, C.-Y. S. Chang, C.-H. Chen, M.-Y. Chou, C.-K. Shih, and L.-J. Li, “Determination of band alignment in the single-layer MoS₂/WSe₂ heterojunction,” *Nat. Commun.*, vol. 6, Jul 2015. [93](#)
- [118] M. Danovich, D. A. Ruiz-Tijerina, R. J. Hunt, M. Szytniszewski, N. D. Drummond, and V. I. Fal’ko, “Localized interlayer complexes in heterobilayer transition metal dichalcogenides,” *Phys. Rev. B*, vol. 97, p. 195452, May 2018. [94](#), [130](#), [234](#)
- [119] *Mathematica, Version 11.0.1*. Champaign, Illinois, U.S.A.: Wolfram Research, Inc., 2016. [95](#)
- [120] F. Ceballos, M. Z. Bellus, H.-Y. Chiu, and H. Zhao, “Ultrafast charge separation and indirect exciton formation in a MoS₂–MoSe₂ van der Waals heterostructure,” *ACS Nano*, vol. 8, no. 12, pp. 12717–12724, 2014. [100](#)

- [121] Y. Gong, J. Lin, X. Wang, G. Shi, S. Lei, Z. Lin, X. Zou, G. Ye, R. Vajtai, B. I. Yakobson, H. Terrones, M. Terrones, B. K. Tay, J. Lou, S. T. Pantelides, Z. Liu, W. Zhou, and P. M. Ajayan, “Vertical and in-plane heterostructures from ws_2/mos_2 monolayers,” *Nat. Mater.*, vol. 13, no. 12, p. 1135, 2014. [100](#)
- [122] E. Anderson, Z. Bai, J. Dongarra, A. Greenbaum, A. McKenney, J. Du Croz, S. Hammarling, J. Demmel, C. Bischof, and D. Sorensen, “LAPACK: A Portable Linear Algebra Library for High-performance Computers,” in *Proceedings of the 1990 ACM/IEEE Conference on Supercomputing*, Supercomputing '90, (Los Alamitos, CA, USA), pp. 2–11, IEEE Computer Society Press, 1990. [103](#), [176](#)
- [123] H. Ogata, “A numerical integration formula based on the Bessel functions,” *Publ. Res. Inst. Math. Sci.*, vol. 41, no. 4, pp. 949–970, 2005. [103](#)
- [124] J. Xue, J. Sanchez-Yamagishi, D. Bulmash, P. Jacquod, A. Deshpande, K. Watanabe, T. Taniguchi, P. Jarillo-Herrero, and B. J. LeRoy, “Scanning tunnelling microscopy and spectroscopy of ultra-flat graphene on hexagonal boron nitride,” *Nat. Mater.*, vol. 10, pp. 282–285, Feb 2011. [105](#)
- [125] A. Laturia, M. L. Van de Put, and W. G. Vandenberghe, “Dielectric properties of hexagonal boron nitride and transition metal dichalcogenides: from monolayer to bulk,” *npj 2D Mater. Appl.*, vol. 2, Mar 2018. [105](#)
- [126] S. Latini, K. T. Winther, T. Olsen, and K. S. Thygesen, “Interlayer Excitons and Band Alignment in $\text{MoS}_2/\text{hBN}/\text{WSe}_2$ van der Waals Heterostructures,” *Nano Lett.*, vol. 17, no. 2, pp. 938–945, 2017. [105](#), [112](#)
- [127] Y. Sun, D. Wang, and Z. Shuai, “Indirect-to-Direct Band Gap Crossover in Few-Layer Transition Metal Dichalcogenides: A Theoretical Prediction,” *J. Phys. Chem. C*, vol. 120, no. 38, pp. 21866–21870, 2016. [107](#)
- [128] M. H. Naik and M. Jain, “Origin of layer dependence in band structures of two-dimensional materials,” *Phys. Rev. B*, vol. 95, p. 165125, Apr 2017. [107](#)

- [129] K. Andersen, S. Latini, and K. S. Thygesen, “Dielectric Genome of van der Waals Heterostructures,” *Nano Lett.*, vol. 15, no. 7, pp. 4616–4621, 2015. [112](#)
- [130] 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, no. 16, pp. 10591–10600, 1995. [114](#)
- [131] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Rev. Mod. Phys.*, vol. 73, pp. 33–83, Jan 2001. [118](#), [121](#)
- [132] R. Dovesi, M. Causa, and G. Angonoa, “Exact-exchange Hartree-Fock calculations for periodic systems. V. Ground-state properties of silicon,” *Phys. Rev. B*, vol. 24, no. 8, p. 4177, 1981. [118](#)
- [133] I. Ohkoshi, “An ab initio Hartree-Fock pseudopotential approach to the ground-state properties of Si,” *J. Phys. C*, vol. 18, no. 28, p. 5415, 1985. [118](#)
- [134] 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, no. 7, p. 4558, 1999. [120](#), [135](#), [136](#), [137](#)
- [135] P. G. Hipes, “Diffusion Monte Carlo method with symmetry-constrained variational principle for degenerate excited states,” *Phys. Rev. B*, vol. 83, p. 195118, May 2011. [120](#)
- [136] J. K. L. MacDonald, “Successive Approximations by the Rayleigh-Ritz Variation Method,” *Phys. Rev.*, vol. 43, pp. 830–833, May 1933. [120](#)
- [137] B. Mussard, E. Coccia, R. Assaraf, M. Otten, C. J. Umrigar, and J. Toulouse, “Chapter Fourteen - Time-Dependent Linear-Response Variational Monte

- Carlo,” in *Novel Electronic Structure Theory: General Innovations and Strongly Correlated Systems* (P. E. Hoggan, ed.), vol. 76 of *Advances in Quantum Chemistry*, pp. 255–270, Academic Press, 2018. 120
- [138] D. M. Ceperley, “Fermion nodes,” *J. Stat. Phys.*, vol. 63, pp. 1237–1267, Jun 1991. 121
- [139] L. Mitas, “Structure of Fermion Nodes and Nodal Cells,” *Phys. Rev. Lett.*, vol. 96, p. 240402, Jun 2006. 121
- [140] D. Bressanini, “Implications of the two nodal domains conjecture for ground state fermionic wave functions,” *Phys. Rev. B*, vol. 86, p. 115120, Sep 2012. 121
- [141] R. Courant and D. Hilbert, *Methods of Mathematical Physics*, vol. I. New York: John Wiley & Sons, 1953. 121
- [142] M. Bajdich, L. Mitas, L. K. Wagner, and K. E. Schmidt, “Pfaffian pairing and backflow wavefunctions for electronic structure quantum Monte Carlo methods,” *Phys. Rev. B*, vol. 77, p. 115112, Mar 2008. 122
- [143] 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, Aug 2006. 123
- [144] 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, no. 12, p. 125106, 2008. 123
- [145] H. Kwee, S. Zhang, and H. Krakauer, “Finite-Size Correction in Many-Body Electronic Structure Calculations,” *Phys. Rev. Lett.*, vol. 100, p. 126404, Mar 2008. 123
- [146] E. Madelung, “Das elektrische Feld in Systemen von regelmäßig angeordneten Punktladungen,” *Phys. Z.*, no. XIX, pp. 524–533, 1918. 123

- [147] N. D. M. Hine, K. Frensch, W. M. C. Foulkes, and M. W. Finnis, “Super-cell size scaling of density functional theory formation energies of charged defects,” *Phys. Rev. B*, vol. 79, p. 024112, Jan 2009. [123](#)
- [148] S. Baroni, S. De Gironcoli, A. Dal Corso, and P. Giannozzi, “Phonons and related crystal properties from density-functional perturbation theory,” *Rev. Mod. Phys.*, vol. 73, no. 2, p. 515, 2001. [124](#), [156](#)
- [149] S. T. Murphy and N. D. M. Hine, “Anisotropic charge screening and super-cell size convergence of defect formation energies,” *Phys. Rev. B*, vol. 87, p. 094111, Mar 2013. [124](#)
- [150] G. Fischerauer, “Comments on real-space Green’s function of an isolated point-charge in an unbounded anisotropic medium,” *IEEE Trans. Ultras. Ferroelectr. Freq. Control*, vol. 44, pp. 1179–1180, Nov 1997. [124](#)
- [151] G. Makov and M. C. Payne, “Periodic boundary conditions in ab initio calculations,” *Phys. Rev. B*, vol. 51, pp. 4014–4022, Feb 1995. [125](#), [127](#)
- [152] E. Ertekin, L. K. Wagner, and J. C. Grossman, “Point-defect optical transitions and thermal ionization energies from quantum Monte Carlo methods: Application to the F -center defect in MgO,” *Phys. Rev. B*, vol. 87, p. 155210, Apr 2013. [125](#)
- [153] E. Mostaani, B. Monserrat, N. D. Drummond, and C. J. Lambert, “Quasi-particle and excitonic gaps of one-dimensional carbon chains,” *Phys. Chem. Chem. Phys.*, vol. 18, pp. 14810–14821, 2016. [125](#), [132](#)
- [154] J. H. Lloyd-Williams and B. Monserrat, “Lattice dynamics and electron-phonon coupling calculations using nondiagonal supercells,” *Phys. Rev. B*, vol. 92, p. 184301, Nov 2015. [126](#)
- [155] B. Wood, W. M. C. Foulkes, M. D. Towler, and N. D. Drummond, “Coulomb finite-size effects in quasi-two-dimensional systems,” *J. Phys. Condens. Matter*, vol. 16, no. 6, p. 891, 2004. [127](#)

- [156] 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, pp. 1814–1832, Jan 1996. [127](#)
- [157] C. Lin, F. H. Zong, and D. M. Ceperley, “Twist-averaged boundary conditions in continuum quantum Monte Carlo algorithms,” *Phys. Rev. E*, vol. 64, no. 1, p. 016702, 2001. [129](#)
- [158] L. D. Landau, “The theory of a Fermi liquid,” *JETP*, vol. 3, no. 6, pp. 920–925, 1957. [130](#)
- [159] L. D. Landau, “Oscillations in a Fermi liquid,” *JETP*, vol. 5, no. 1, pp. 101–108, 1957. [130](#)
- [160] L. D. Landau, “On the theory of the Fermi liquid,” *JETP*, vol. 8, no. 1, p. 70, 1959. [130](#)
- [161] R. N. Barnett, P. J. Reynolds, and W. A. Lester, “Monte Carlo determination of the oscillator strength and excited state lifetime for the $\text{Li}_2^2\text{S} \rightarrow 2^2\text{P}$ transition,” *Int. J. Quantum Chem.*, vol. 42, no. 4, pp. 837–847, 1992. [130](#)
- [162] Y. Kwon, D. M. Ceperley, and R. M. Martin, “Quantum Monte Carlo calculation of the Fermi-liquid parameters in the two-dimensional electron gas,” *Phys. Rev. B*, vol. 50, pp. 1684–1694, Jul 1994. [130](#)
- [163] M. Holzmann, B. Bernu, V. Olevano, R. M. Martin, and D. M. Ceperley, “Renormalization factor and effective mass of the two-dimensional electron gas,” *Phys. Rev. B*, vol. 79, p. 041308, Jan 2009. [130](#)
- [164] N. D. Drummond and R. J. Needs, “Diffusion quantum Monte Carlo calculation of the quasiparticle effective mass of the two-dimensional homogeneous electron gas,” *Phys. Rev. B*, vol. 87, p. 045131, Jan 2013. [130](#)

- [165] N. D. Drummond and R. J. Needs, “Quantum Monte Carlo calculation of the Fermi liquid parameters of the two-dimensional homogeneous electron gas,” *Phys. Rev. B*, vol. 88, p. 035133, Jul 2013. [130](#)
- [166] D. Alfè and M. J. Gillan, “Linear-scaling quantum Monte Carlo technique with non-orthogonal localized orbitals,” *J. Phys. Condens. Matter*, vol. 16, no. 25, p. L305, 2004. [131](#)
- [167] A. J. Williamson, R. Q. Hood, and J. C. Grossman, “Linear-Scaling Quantum Monte Carlo Calculations,” *Phys. Rev. Lett.*, vol. 87, p. 246406, Nov 2001. [131](#)
- [168] A. J. Williamson, J. C. Grossman, R. Q. Hood, A. Puzder, and G. Galli, “Quantum Monte Carlo Calculations of Nanostructure Optical Gaps: Application to Silicon Quantum Dots,” *Phys. Rev. Lett.*, vol. 89, p. 196803, Oct 2002. [131](#)
- [169] B. Monserrat, G. J. Conduit, and R. J. Needs, “Extracting semiconductor band gap zero-point corrections from experimental data,” *Phys. Rev. B*, vol. 90, p. 184302, Nov 2014. [132](#)
- [170] E. A. Engel, B. Monserrat, and R. J. Needs, “Vibrational renormalisation of the electronic band gap in hexagonal and cubic ice,” *J. Chem. Phys.*, vol. 143, no. 24, p. 244708, 2015. [132](#)
- [171] E. A. Engel, B. Monserrat, and R. J. Needs, “Vibrational effects on surface energies and band gaps in hexagonal and cubic ice,” *J. Chem. Phys.*, vol. 145, no. 4, p. 044703, 2016. [132](#)
- [172] K. Muller-Dethlefs and E. W. Schlag, “High-resolution zero kinetic energy (ZEKE) photoelectron spectroscopy of molecular systems,” *Annu. Rev. Phys. Chem.*, vol. 42, no. 1, pp. 109–136, 1991. [132](#)
- [173] F. Marsusi, J. Sabbaghzadeh, and N. D. Drummond, “Comparison of quantum Monte Carlo with time-dependent and static density-functional theory

- calculations of diamondoid excitation energies and Stokes shifts,” *Phys. Rev. B*, vol. 84, p. 245315, Dec 2011. [133](#)
- [174] 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. Cryst. Mater.*, vol. 220, no. 5/6, pp. 567–570, 2005. [133](#)
- [175] D. Vanderbilt, “Soft self-consistent pseudopotentials in a generalized eigenvalue formalism,” *Phys. Rev. B*, vol. 41, no. 11, p. 7892, 1990. [133](#)
- [176] J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized Gradient Approximation Made Simple,” *Phys. Rev. Lett.*, vol. 77, no. 18, pp. 3865–3868, 1996. [133](#)
- [177] P. J. Mohr, D. B. Newell, and B. N. Taylor, “CODATA recommended values of the fundamental physical constants: 2014,” *Rev. Mod. Phys.*, vol. 88, p. 035009, Sep 2016. [133](#)
- [178] A. F. Goncharov, J. C. Crowhurst, J. K. Dewhurst, S. Sharma, C. Sanloup, E. Gregoryanz, N. Guignot, and M. Mezouar, “Thermal equation of state of cubic boron nitride: Implications for a high-temperature pressure scale,” *Phys. Rev. B*, vol. 75, p. 224114, Jun 2007. [133](#)
- [179] L. Levien, C. T. Prewitt, and D. J. Weidner, “Structure and elastic properties of quartz at pressure,” *Am. Mineral.*, vol. 65, no. 9-10, pp. 920–930, 1980. [133](#)
- [180] 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, Oct 2004. [133](#)
- [181] N. D. Drummond, M. D. Towler, and R. J. Needs, “Jastrow correlation factor for atoms, molecules, and solids,” *Phys. Rev. B*, vol. 70, no. 23, p. 235119, 2004. [134](#)

- [182] C. J. Umrigar, K. G. Wilson, and J. W. Wilkins, “Optimized trial wave functions for quantum Monte Carlo calculations,” *Phys. Rev. Lett.*, vol. 60, pp. 1719–1722, Apr 1988. [134](#), [181](#), [183](#)
- [183] N. D. Drummond and R. J. Needs, “Variance-minimization scheme for optimizing Jastrow factors,” *Phys. Rev. B*, vol. 72, p. 085124, Aug 2005. [134](#)
- [184] J. Toulouse and C. J. Umrigar, “Optimization of quantum Monte Carlo wave functions by energy minimization,” *J. Chem. Phys.*, vol. 126, no. 8, p. 084102, 2007. [134](#)
- [185] 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, Dec 2006. [134](#)
- [186] M. Casula, “Beyond the locality approximation in the standard diffusion Monte Carlo method,” *Phys. Rev. B*, vol. 74, no. 16, p. 161102, 2006. [134](#), [152](#)
- [187] M. D. Towler, R. Q. Hood, and R. J. Needs, “Minimum principles and level splitting in quantum Monte Carlo excitation energies: Application to diamond,” *Phys. Rev. B*, vol. 62, no. 4, p. 2330, 2000. [135](#), [154](#)
- [188] N. D. Drummond, P. López Ríos, A. Ma, J. R. Trail, G. G. Spink, M. D. Towler, and R. J. Needs, “Quantum Monte Carlo study of the Ne atom and the Ne⁺ ion,” *J. Chem. Phys.*, vol. 124, no. 22, p. 224104, 2006. [140](#)
- [189] J. White and E. Ackad, “A validation of a simple model for the calculation of the ionization energies in X-ray laser-cluster interactions,” *Phys. Plasmas*, vol. 22, no. 2, p. 022123, 2015. [142](#)
- [190] K. T. Chung, “Ionization potential of the lithiumlike $1s^22s$ states from lithium to neon,” *Phys. Rev. A*, vol. 44, pp. 5421–5433, Nov 1991. [142](#)

- [191] L. A. Curtiss, K. Raghavachari, G. W. Trucks, and J. A. Pople, “Gaussian-2 theory for molecular energies of first- and second-row compounds,” *J. Chem. Phys.*, vol. 94, no. 11, pp. 7221–7230, 1991. [143](#)
- [192] F. Kaplan, M. E. Harding, C. Seiler, F. Weigend, F. Evers, and M. J. van Setten, “Quasi-Particle Self-Consistent *GW* for Molecules,” *J. Chem. Theory Comput.*, vol. 12, no. 6, pp. 2528–2541, 2016. [144](#)
- [193] P. Koval, D. Foerster, and D. Sánchez-Portal, “Fully self-consistent *GW* and quasiparticle self-consistent *GW* for molecules,” *Phys. Rev. B*, vol. 89, no. 15, p. 155417, 2014. [144](#)
- [194] J. V. Ortiz, “Coupled-cluster reference electron propagator calculations on the ionization energies of O_2 ,” *Chem. Phys. Lett.*, vol. 199, no. 6, pp. 530–534, 1992. [144](#)
- [195] N. Q. Su, I. Y. Zhang, J. Wu, and X. Xu, “Calculations of ionization energies and electron affinities for atoms and molecules: A comparative study with different methods,” *Front. Chem. China*, vol. 6, no. 4, pp. 269–279, 2011. [144](#)
- [196] J. F. Stanton, R. J. Bartlett, and C. M. L. Rittby, “Fock space multireference coupled-cluster theory for general single determinant reference functions,” *J. Chem. Phys.*, vol. 97, no. 8, pp. 5560–5567, 1992. [144](#)
- [197] W. Kolos, “Extrapolated Born–Oppenheimer energy for the ground state of the hydrogen molecule,” *J. Chem. Phys.*, vol. 101, no. 2, pp. 1330–1332, 1994. [144](#)
- [198] E. McCormack, J. M. Gilligan, C. Cornaggia, and E. E. Eyler, “Measurement of high Rydberg states and the ionization potential of H_2 ,” *Phys. Rev. A*, vol. 39, no. 4, p. 2260, 1989. [144](#), [145](#)

- [199] R. G. Tonkyn, J. W. Winniczek, and M. G. White, “Rotationally resolved photoionization of O^{2+} near threshold,” *Chem. Phys. Lett.*, vol. 164, no. 2-3, pp. 137–142, 1989. [144](#)
- [200] K. K. Irikura, “Experimental Vibrational Zero-Point Energies: Diatomic Molecules,” *J. Phys. Chem. Ref. Data*, vol. 36, no. 2, pp. 389–397, 2007. [143](#)
- [201] W. L. Glab and J. P. Hessler, “Multiphoton excitation of high singlet np Rydberg states of molecular hydrogen: Spectroscopy and dynamics,” *Phys. Rev. A*, vol. 35, pp. 2102–2110, Mar 1987. [145](#)
- [202] C. Schweitzer and R. Schmidt, “Physical Mechanisms of Generation and Deactivation of Singlet Oxygen,” *Chem. Rev.*, vol. 103, no. 5, pp. 1685–1758, 2003. [146](#)
- [203] X. Blase, C. Attaccalite, and V. Olevano, “First-principles *GW* calculations for fullerenes, porphyrins, phtalocyanine, and other molecules of interest for organic photovoltaic applications,” *Phys. Rev. B*, vol. 83, no. 11, p. 115103, 2011. [147](#)
- [204] G. Mallocci, G. Cappellini, G. Mulas, and A. Mattoni, “Electronic and optical properties of families of polycyclic aromatic hydrocarbons: A systematic (time-dependent) density functional theory study,” *Chem. Phys.*, vol. 384, no. 1-3, pp. 19–27, 2011. [147](#)
- [205] J. W. Knight, X. Wang, L. Gallandi, O. Dolgounitcheva, X. Ren, J. V. Ortiz, P. Rinke, T. Körzdörfer, and N. Marom, “Accurate ionization potentials and electron affinities of acceptor molecules III: a benchmark of *GW* methods,” *J. Chem. Theory Comput.*, vol. 12, no. 2, pp. 615–626, 2016. [147](#), [148](#)
- [206] J. W. Hager and S. C. Wallace, “Two-laser photoionization supersonic jet mass spectrometry of aromatic molecules,” *Anal. Chem.*, vol. 60, no. 1, pp. 5–10, 1988. [147](#)

- [207] J. Schiedt and R. Weinkauff, “Photodetachment photoelectron spectroscopy of mass selected anions: Anthracene and the anthracene-H₂O cluster,” *Chem. Phys. Lett.*, vol. 266, no. 1-2, pp. 201–205, 1997. 147
- [208] J. H. D. Eland, “Photoelectron spectra of conjugated hydrocarbons and heteromolecules,” *Int. J. Mass Spectrom.*, vol. 2, no. 6, pp. 471–484, 1969. 147
- [209] C. F. Batten, J. A. Taylor, B. P. Tsai, and G. G. Meisels, “Photoionization processes at threshold. II. Threshold photoelectron, photoionization, and coincidence ion-threshold photoelectron spectra of BF₃,” *J. Chem. Phys.*, vol. 69, no. 6, pp. 2547–2552, 1978. 147
- [210] N. Dupuy, S. Bouaouli, F. Mauri, S. Sorella, and M. Casula, “Vertical and adiabatic excitations in anthracene from quantum Monte Carlo: Constrained energy minimization for structural and electronic excited-state properties in the JAGP ansatz,” *J. Chem. Phys.*, vol. 142, no. 21, p. 214109, 2015. 147
- [211] D. Biermann and W. Schmidt, “Diels-Alder reactivity of polycyclic aromatic hydrocarbons. 1. Acenes and benzologs,” *J. Am. Chem. Soc.*, vol. 102, no. 9, pp. 3163–3173, 1980. 148
- [212] M. Baba, M. Saitoh, K. Taguma, K. Shinohara, K. Yoshida, Y. Semba, S. Kasahara, N. Nakayama, H. Goto, T. Ishimoto, and U. Nagashima, “Structure and excited-state dynamics of anthracene: Ultrahigh-resolution spectroscopy and theoretical calculation,” *J. Chem. Phys.*, vol. 130, no. 13, p. 134315, 2009. 148
- [213] M. R. Padhye, S. P. McGlynn, and M. Kasha, “Lowest triplet state of anthracene,” *J. Chem. Phys.*, vol. 24, no. 3, pp. 588–594, 1956. 148
- [214] X. Ren, N. Marom, F. Caruso, M. Scheffler, and P. Rinke, “Beyond the *GW* approximation: A second-order screened exchange correction,” *Phys. Rev. B*, vol. 92, no. 8, p. 081104, 2015. 148

- [215] R. M. Richard, M. S. Marshall, O. Dolgounitcheva, J. V. Ortiz, J.-L. Brédas, N. Marom, and C. D. Sherrill, “Accurate Ionization Potentials and Electron Affinities of Acceptor Molecules I. Reference Data at the CCSD(T) Complete Basis Set Limit,” *J. Chem. Theory Comput.*, vol. 12, no. 2, pp. 595–604, 2016. [148](#)
- [216] K. N. Houk and L. L. Munchausen, “Ionization potentials, electron affinities, and reactivities of cyanoalkenes and related electron-deficient alkenes. A frontier molecular orbital treatment of cyanoalkene reactivities in cycloaddition, electrophilic, nucleophilic, and radical reactions,” *J. Am. Chem. Soc.*, vol. 98, no. 4, pp. 937–946, 1976. [148](#)
- [217] D. Khuseynov, M. T. Fontana, and A. Sanov, “Photoelectron spectroscopy and photochemistry of tetracyanoethylene radical anion in the gas phase,” *Chem. Phys. Lett.*, vol. 550, pp. 15–18, 2012. [148](#)
- [218] D. J. Knowles and A. J. C. Nicholson, “Ionization energies of formic and acetic acid monomers,” *J. Chem. Phys.*, vol. 60, no. 3, pp. 1180–1181, 1974. [148](#)
- [219] M. A. Green, “Improved value for the silicon free exciton binding energy,” *AIP Adv.*, vol. 3, no. 11, p. 112104, 2013. [149](#)
- [220] L. Mitáš, “Quantum Monte Carlo for Electronic Structure of Solids,” in *Electronic Properties of Solids Using Cluster Methods* (T. A. Kaplan and S. D. Mahanti, eds.), ch. 9, p. 155, New York: Plenum, 1995. [154](#)
- [221] A. J. Williamson, R. Q. Hood, R. J. Needs, and G. Rajagopal, “Diffusion quantum Monte Carlo calculations of the excited states of silicon,” *Phys. Rev. B*, vol. 57, no. 19, p. 12140, 1998. [154](#)
- [222] L. Mitáš and R. M. Martin, “Quantum Monte Carlo of nitrogen: atom, dimer, atomic, and molecular solids,” *Phys. Rev. Lett.*, vol. 72, no. 15, p. 2438, 1994. [154](#)

- [223] J.-W. Lee, L. Mitáš, and L. K. Wagner, “Quantum Monte Carlo study of MnO solid,” *arXiv:0411247*, Nov. 2004. [154](#)
- [224] S. Ishii, S. Iwata, and K. Ohno, “All-Electron *GW* Calculations of Silicon, Diamond, and Silicon Carbide,” *Mater. Trans.*, vol. 51, no. 12, pp. 2150–2156, 2010. [155](#)
- [225] F. Bruneval, N. Vast, and L. Reining, “Effect of self-consistency on quasiparticles in solids,” *Phys. Rev. B*, vol. 74, p. 045102, Jul 2006. [155](#)
- [226] D. M. Hoffman, G. L. Doll, and P. C. Eklund, “Optical properties of pyrolytic boron nitride in the energy range 0.05–10 eV,” *Phys. Rev. B*, vol. 30, pp. 6051–6056, Nov 1984. [155](#)
- [227] L. G. Carpenter and P. J. Kirby, “The electrical resistivity of boron nitride over the temperature range 700 degrees C to 1400 degrees C,” *J. Phys. D*, vol. 15, pp. 1143–1151, Jul 1982. [155](#)
- [228] K. Watanabe, T. Taniguchi, and H. Kanda, “Direct-bandgap properties and evidence for ultraviolet lasing of hexagonal boron nitride single crystal,” *Nat. Mater.*, vol. 3, pp. 404–409, May 2004. [155](#), [156](#), [157](#)
- [229] Y. Shi, C. Hamsen, X. Jia, K. K. Kim, A. Reina, M. Hofmann, A. L. Hsu, K. Zhang, H. Li, Z.-Y. Juang, M. S. Dresselhaus, L.-J. Li, and J. Kong, “Synthesis of Few-Layer Hexagonal Boron Nitride Thin Film by Chemical Vapor Deposition,” *Nano Lett.*, vol. 10, pp. 4134–4139, Oct 2010. [155](#)
- [230] G. Cassabois, P. Valvin, and B. Gil, “Hexagonal boron nitride is an indirect bandgap semiconductor,” *Nat. Photonics*, vol. 10, pp. 262–266, Jan 2016. [155](#), [156](#), [157](#)
- [231] G. Cappellini, V. Fiorentini, K. Tenelsen, and F. Bechstedt, “Structure and quasiparticle energies of cubic, wurtzite and hexagonal BN,” Materials Research Society, Pittsburgh, PA (United States), 1996. [155](#)

- [232] B. Arnaud, S. Lebègue, P. Rabiller, and M. Alouani, “Huge Excitonic Effects in Layered Hexagonal Boron Nitride,” *Phys. Rev. Lett.*, vol. 96, p. 026402, Jan 2006. [155](#), [157](#)
- [233] L. Wirtz, A. Marini, and A. Rubio, “Excitons in Boron Nitride Nanotubes: Dimensionality Effects,” *Phys. Rev. Lett.*, vol. 96, p. 126104, Mar 2006. [155](#)
- [234] B. Cunningham, M. Grüning, P. Azarhoosh, D. Pashov, and M. van Schilfgaarde, “Effect of ladder diagrams on optical absorption spectra in a quasiparticle self-consistent *GW* framework,” *Phys. Rev. Materials*, vol. 2, p. 034603, Mar 2018. [155](#)
- [235] J. H. Lloyd-Williams and B. Monserrat, “Lattice dynamics and electron-phonon coupling calculations using nondiagonal supercells,” *Phys. Rev. B*, vol. 92, p. 184301, Nov 2015. [156](#)
- [236] R. J. Hunt, N. D. Drummond, B. Monserrat, V. Zólyomi, and V. I. Fal’ko, “Diffusion quantum Monte Carlo and *GW* study of the electronic properties of single-layer and bulk hexagonal boron nitride,” *Manuscript in preparation*, 2018. [158](#), [159](#)
- [237] J. R. Trail and R. J. Needs, “Shape and energy consistent pseudopotentials for correlated electron systems,” *J. Chem. Phys.*, vol. 146, no. 20, p. 204107, 2017. [158](#)
- [238] V. A. Fomichev and M. A. Rumsh, “Investigation of X-ray spectra of hexagonal and cubic boron nitride,” *J. Phys. Chem. Solids*, vol. 29, no. 6, pp. 1015–1024, 1968. [160](#)
- [239] R. M. Chrenko, “Ultraviolet and infrared spectra of cubic boron nitride,” *Solid State Commun.*, vol. 14, no. 6, pp. 511–515, 1974. [160](#)
- [240] Y.-N. Xu and W. Y. Ching, “Calculation of ground-state and optical properties of boron nitrides in the hexagonal, cubic, and wurtzite structures,” *Phys. Rev. B*, vol. 44, pp. 7787–7798, Oct 1991. [160](#)

- [241] P. Rodríguez-Hernández, M. González-Díaz, and A. Muñoz, “Electronic and structural properties of cubic BN and BP,” *Phys. Rev. B*, vol. 51, pp. 14705–14708, May 1995. [160](#)
- [242] G. Cappellini, G. Satta, M. Palummo, and G. Onida, “Optical properties of BN in cubic and layered hexagonal phases,” *Phys. Rev. B*, vol. 64, p. 035104, Jun 2001. [160](#)
- [243] G. Satta, G. Cappellini, V. Olevano, and L. Reining, “Many-body effects in the electronic spectra of cubic boron nitride,” *Phys. Rev. B*, vol. 70, no. 19, p. 195212, 2004. [160](#)
- [244] G. Kresse, M. Marsman, L. E. Hintzsche, and E. Flage-Larsen, “Optical and electronic properties of Si_3N_4 and $\alpha\text{-SiO}_2$,” *Phys. Rev. B*, vol. 85, p. 045205, Jan 2012. [161](#)
- [245] E. K. Chang, M. Rohlfing, and S. G. Louie, “Excitons and Optical Properties of α -Quartz,” *Phys. Rev. Lett.*, vol. 85, pp. 2613–2616, Sep 2000. [161](#)
- [246] H. R. Philipp, “Optical transitions in crystalline and fused quartz,” *Solid State Commun.*, vol. 4, no. 1, pp. 73–75, 1966. [161](#)
- [247] V. Tran, R. Soklaski, Y. Liang, and L. Yang, “Layer-controlled band gap and anisotropic excitons in few-layer black phosphorus,” *Phys. Rev. B*, vol. 89, p. 235319, Jun 2014. [162](#), [163](#), [164](#)
- [248] V. Tran, R. Fei, and L. Yang, “Quasiparticle energies, excitons, and optical spectra of few-layer black phosphorus,” *2D Mater.*, vol. 2, no. 4, p. 044014, 2015. [162](#), [163](#), [164](#)
- [249] J.-H. Choi, P. Cui, H. Lan, and Z. Zhang, “Linear Scaling of the Exciton Binding Energy versus the Band Gap of Two-Dimensional Materials,” *Phys. Rev. Lett.*, vol. 115, p. 066403, Aug 2015. [162](#), [163](#), [164](#)

- [250] L. Seixas, A. S. Rodin, A. Carvalho, and A. H. Castro Neto, “Exciton binding energies and luminescence of phosphorene under pressure,” *Phys. Rev. B*, vol. 91, p. 115437, Mar 2015. [162](#), [163](#), [164](#)
- [251] G. Zhang, A. Chaves, S. Huang, F. Wang, Q. Xing, T. Low, and H. Yan, “Determination of layer-dependent exciton binding energies in few-layer black phosphorus,” *Sci. Adv.*, vol. 4, no. 3, 2018. [162](#), [163](#), [164](#)
- [252] J. Qiao, X. Kong, Z.-X. Hu, F. Yang, and W. Ji, “High-mobility transport anisotropy and linear dichroism in few-layer black phosphorus,” *Nat. Commun.*, vol. 5, p. 4475, Jul 2014. [162](#)
- [253] A. Castellanos-Gomez, L. Vicarelli, E. Prada, J. O. Island, K. L. Narasimha-Acharya, S. I. Blanter, D. J. Groenendijk, M. Buscema, G. A. Steele, J. V. Alvarez, H. W. Zandbergen, J. J. Palacios, and H. S. J. van der Zant, “Isolation and characterization of few-layer black phosphorus,” *2D Mater.*, vol. 1, no. 2, p. 025001, 2014. [164](#)
- [254] A. Chaves, T. Low, P. Avouris, D. Çakır, and F. M. Peeters, “Anisotropic exciton Stark shift in black phosphorus,” *Phys. Rev. B*, vol. 91, p. 155311, Apr 2015. [164](#)
- [255] L. Liang, J. Wang, W. Lin, B. G. Sumpter, V. Meunier, and M. Pan, “Electronic Bandgap and Edge Reconstruction in Phosphorene Materials,” *Nano Lett.*, vol. 14, no. 11, pp. 6400–6406, 2014. [164](#)
- [256] J. Yang, R. Xu, J. Pei, Y. W. Myint, F. Wang, Z. Wang, S. Zhang, Z. Yu, and Y. Lu, “Optical tuning of exciton and trion emissions in monolayer phosphorene,” *Light Sci. Appl.*, vol. 4, p. e312, Jul 2015. [164](#)
- [257] T. Frank, R. Derian, K. Tokár, L. Mitas, J. Fabian, and I. Stich, “Many-Body Quantum Monte Carlo Study of 2D Materials: Cohesion and Band Gap in Single-Layer Phosphorene,” *Phys. Rev. X*, vol. 9, p. 011018, Jan 2019. [164](#)

- [258] N. D. Drummond, A. J. Williamson, R. J. Needs, and G. Galli, “Electron emission from diamondoids: a diffusion quantum Monte Carlo study,” *Phys. Rev. Lett.*, vol. 95, no. 9, p. 096801, 2005. [165](#)
- [259] 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, no. 13, p. 132113, 2006. [169](#)
- [260] 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, no. 11, p. 111908, 2010. [169](#)
- [261] A. Ohtomo and H. Y. Hwang, “A high-mobility electron gas at the LaAlO₃/SrTiO₃ heterointerface,” *Nature*, vol. 427, no. 6973, p. 423, 2004. [169](#)
- [262] 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, Mar 2009. [169](#), [179](#), [180](#), [182](#), [183](#), [186](#), [187](#), [189](#)
- [263] B. Bernu, F. Delyon, M. Holzmann, and L. Baguet, “Hartree-Fock phase diagram of the two-dimensional electron gas,” *Phys. Rev. B*, vol. 84, p. 115115, Sep 2011. [169](#)
- [264] P. P. Ewald, “Die Berechnung optischer und elektrostatischer Gitterpotentiale,” *Ann. Phys. (Berl.)*, vol. 369, no. 3, pp. 253–287, 1921. [170](#)
- [265] R. B. Dingle, *Asymptotic Expansions: Their Derivation and Interpretation*. London-New York: Academic Press, 1973. [171](#)
- [266] M. L. Glasser and I. J. Zucker, “Lattice sums,” in *Theoretical Chemistry*, pp. 67–139, Elsevier, 1980. [171](#)

- [267] J. M. Borwein and P. B. Borwein, *Pi and the AGM*. Wiley, New York, 1987. [171](#)
- [268] J. M. Borwein, M. L. Glasser, R. C. McPhedran, J. G. Wan, and I. J. Zucker, *Lattice sums then and now*, vol. 150. Cambridge University Press, 2013. [171](#)
- [269] J. N. Newman, “Approximations for the Bessel and Struve functions,” *Math. Comput.*, vol. 43, no. 168, pp. 551–556, 1984. [178](#)
- [270] E. Wigner, “On the Interaction of Electrons in Metals,” *Phys. Rev.*, vol. 46, pp. 1002–1011, Dec 1934. [179](#)
- [271] B. Tanatar and D. M. Ceperley, “Ground state of the two-dimensional electron gas,” *Phys. Rev. B*, vol. 39, pp. 5005–5016, Mar 1989. [179](#), [186](#), [189](#)
- [272] N. Choudhury and S. K. Ghosh, “Density-functional theory of Wigner crystallization in two dimensions,” *Phys. Rev. B*, vol. 51, pp. 2588–2591, Jan 1995. [179](#)
- [273] 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, Mar 2017. [179](#), [188](#)
- [274] P. Karpov and S. Brazovskii, “Modeling of networks and globules of charged domain walls observed in pump and pulse induced states,” *Sci. Rep.*, vol. 8, Mar 2018. [179](#)
- [275] N. Zibouche, A. Kuc, J. Musfeldt, and T. Heine, “Transition-metal dichalcogenides for spintronic applications,” *Ann. Phys. (Berl.)*, vol. 526, no. 9-10, pp. 395–401, 2014. [179](#)
- [276] A. Kormányos, V. Zólyomi, N. D. Drummond, P. Rakyta, G. Burkard, and V. I. Fal’ko, “Monolayer MoS₂: Trigonal warping, the Γ valley, and spin-orbit coupling effects,” *Phys. Rev. B*, vol. 88, p. 045416, July 2013. [179](#)

- [277] A. Kormányos, V. Zólyomi, N. D. Drummond, and G. Burkard, “Spin-Orbit Coupling, Quantum Dots, and Qubits in Monolayer Transition Metal Dichalcogenides,” *Phys. Rev. X*, vol. 4, p. 011034, March 2014. 179
- [278] K. Kośmider, J. W. González, and J. Fernández-Rossier, “Large spin splitting in the conduction band of transition metal dichalcogenide monolayers,” *Phys. Rev. B*, vol. 88, p. 245436, December 2013. 179
- [279] Z. Y. Zhu, Y. C. Cheng, and U. Schwingenschlögl, “Giant spin-orbit-induced spin splitting in two-dimensional transition-metal dichalcogenide semiconductors,” *Phys. Rev. B*, vol. 84, p. 153402, October 2011. 179
- [280] R. Pisoni, A. Kormányos, M. Brooks, Z. Lei, P. Back, M. Eich, H. Overweg, Y. Lee, P. Rickhaus, K. Watanabe, T. Taniguchi, A. Imamoglu, G. Burkard, T. Ihn, and K. Ensslin, “Interactions and Magnetotransport through Spin-Valley Coupled Landau Levels in Monolayer MoS₂,” *Phys. Rev. Lett.*, vol. 121, p. 247701, December 2018. 180
- [281] B. Radisavljevic and A. Kis, “Mobility engineering and a metal–insulator transition in monolayer MoS₂,” *Nat. Mater.*, vol. 12, pp. 815–820, Jun 2013. 180
- [282] K. Sotthewes, R. van Bremen, E. Dollekamp, T. Boulogne, K. Nowakowski, D. Kas, H. J. W. Zandvliet, and P. Bampoulis, “Universal Fermi-Level Pinning in Transition-Metal Dichalcogenides,” *J. Phys. Chem. C*, vol. 123, no. 9, pp. 5411–5420, 2019. 180
- [283] P. Anderson, “Resonating valence bonds: A new kind of insulator?,” *Mater. Res. Bull.*, vol. 8, pp. 153–160, Feb 1973. 180
- [284] L. Savary and L. Balents, “Quantum spin liquids: a review,” *Rep. Progr. Phys.*, vol. 80, p. 016502, Nov 2016. 180
- [285] J. Wen, S.-L. Yu, S. Li, W. Yu, and J.-X. Li, “Experimental identification of quantum spin liquids,” *npj Quantum Mater.*, vol. 4, Apr 2019. 180

- [286] G. G. Spink, R. J. Needs, and N. D. Drummond, “Quantum Monte Carlo study of the three-dimensional spin-polarized homogeneous electron gas,” *Phys. Rev. B*, vol. 88, p. 085121, Aug 2013. [185](#)
- [287] D. Ceperley, “Ground state of the fermion one-component plasma: A Monte Carlo study in two and three dimensions,” *Phys. Rev. B*, vol. 18, pp. 3126–3138, Oct 1978. [186](#)
- [288] C. Attaccalite, S. Moroni, P. Gori-Giorgi, and G. B. Bachelet, “Correlation Energy and Spin Polarization in the 2D Electron Gas,” *Phys. Rev. Lett.*, vol. 88, p. 256601, Jun 2002. [186](#)
- [289] F. Rapisarda and G. Senatore, “Diffusion Monte Carlo Study of Electrons in Two-dimensional Layers,” *Aust. J. Phys.*, vol. 49, no. 1, p. 161, 1996. [186](#), [187](#)
- [290] 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, Sep 2008. [186](#)
- [291] 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, Oct 2013. [189](#)
- [292] J. Yoon, C. C. Li, D. Shahar, D. C. Tsui, and M. Shayegan, “Wigner Crystallization and Metal-Insulator Transition of Two Dimensional Holes in GaAs at $B = 0$,” *Phys. Rev. Lett.*, vol. 82, pp. 1744–1747, Feb 1999. [189](#)
- [293] M. M. Fogler, L. V. Butov, and K. S. Novoselov, “High-temperature superfluidity with indirect excitons in van der Waals heterostructures,” *Nat. Commun.*, vol. 5, Jul 2014. [191](#)

- [294] O. L. Berman and R. Y. Kezerashvili, “High-temperature superfluidity of the two-component Bose gas in a transition metal dichalcogenide bilayer,” *Phys. Rev. B*, vol. 93, p. 245410, Jun 2016. [191](#)
- [295] E. V. Calman, C. J. Dorow, M. M. Fogler, L. V. Butov, S. Hu, A. Mishchenko, and A. K. Geim, “Control of excitons in multi-layer van der Waals heterostructures,” *Appl. Phys. Lett.*, vol. 108, no. 10, p. 101901, 2016. [191](#)
- [296] R. Maezono, P. López Ríos, T. Ogawa, and R. J. Needs, “Excitons and biexcitons in symmetric electron-hole bilayers,” *Phys. Rev. Lett.*, vol. 110, p. 216407, May 2013. [191](#)
- [297] P. López Ríos, A. Perali, R. J. Needs, and D. Neilson, “Evidence from Quantum Monte Carlo Simulations of Large-Gap Superfluidity and BCS-BEC Crossover in Double Electron-Hole Layers,” *Phys. Rev. Lett.*, vol. 120, p. 177701, Apr 2018. [191](#)
- [298] D. K. Efimkin and A. H. MacDonald, “Many-body theory of trion absorption features in two-dimensional semiconductors,” *Phys. Rev. B*, vol. 95, p. 035417, Jan 2017. [193](#)
- [299] O. L. Berman and R. Y. Kezerashvili, “Superfluidity of dipolar excitons in a transition metal dichalcogenide double layer,” *Phys. Rev. B*, vol. 96, p. 094502, Sep 2017. [234](#)

Appendix A

Monte Carlo bootstrapping

Suppose one performs a least squares fit (which need not be linear) of a function $y = F(x)$ to some multidimensional data. Suppose further that F is parameterised by a series of optimal fit parameters $\{\tilde{f}_i\}$, and an associated set of uncertainties $\{\sigma_{f_i}\}$. Let $F(x|f_1, \dots, f_{N_f})$ denote the value of the fitting function given the fitting parameters $\{f_1, \dots, f_{N_f}\}$ (which may not necessarily be the optimal ones).

The procedure of Monte Carlo bootstrapping is as follows. We re-sample our fit parameters, $f_i \rightarrow \xi_i$, where ξ_i is a Gaussian distributed random variate, with mean \tilde{f}_i , and standard deviation σ_{f_i} . We form $y = F(x|\xi_1, \dots, \xi_{N_f})$, and evaluate some derived quantity $z(y)$. This could be a root of the fit function (as in Ch. 2), or, more generally, *any* arbitrary quantity which may be derived from the fit function. We then store the value of $z(y)$ in a table, and repeat the procedure. For an ensemble of resamplings, we now have a series of $z(y)$ realisations, distributed in some way. If we assume that the $z(y)$ are normally distributed, for example, we may then extract the mean (\bar{z}) and standard deviation (σ_z) of our derived quantity z .

Appendix B

Approximate bilayer interactions

At long-range (small q), the bilayer interaction potentials adopt limiting forms. The use of these limiting forms *for all* q constitutes the “approximate bilayer Keldysh interaction”. From the full bilayer potentials (ϕ), given in Eqn. (3.16) and Eqn. (3.17), and for point charge at the origin, we can extract small q limits that define the approximate bilayer potentials ($\tilde{\phi}$) as

$$\tilde{\phi}_{\text{inter}}(\mathbf{q}) = \frac{1}{2\epsilon q(1 + (r_{\star,1} + r_{\star,2} + d)q)}. \quad (\text{B.1})$$

$$\tilde{\phi}_{\text{intra}}(\mathbf{q}) = \frac{1}{2\epsilon q(1 + (r_{\star,1} + r_{\star,2})q)}. \quad (\text{B.2})$$

In both cases, the leading order error is $\mathcal{O}(q)$, which translates to a leading order error which is $\mathcal{O}(r^{-3})$ in real space. Both $\tilde{\phi}_{\text{inter}}$ and $\tilde{\phi}_{\text{intra}}$ can be evaluated in real space exactly. They are both of Keldysh form, with effective r_{\star} values $r_{\star,\text{eff}}^{\text{inter}} = r_{\star,1} + r_{\star,2} + d$ and $r_{\star,\text{eff}}^{\text{intra}} = r_{\star,1} + r_{\star,2}$. *I.e.*

$$\tilde{\phi}_{\text{inter}}(\mathbf{s}) = \frac{\pi}{2\epsilon(r_{\star,1} + r_{\star,2} + d)} \left[H_0 \left(\frac{s}{r_{\star,1} + r_{\star,2} + d} \right) - Y_0 \left(\frac{s}{r_{\star,1} + r_{\star,2} + d} \right) \right] \quad (\text{B.3})$$

$$\tilde{\phi}_{\text{intra}}(\mathbf{s}) = \frac{\pi}{2\epsilon(r_{\star,1} + r_{\star,2})} \left[H_0 \left(\frac{s}{r_{\star,1} + r_{\star,2}} \right) - Y_0 \left(\frac{s}{r_{\star,1} + r_{\star,2}} \right) \right] \quad (\text{B.4})$$

We note that, in a recent study, a hybrid of these two forms was guessed as appropriate for the treatment of interlayer interactions in TMD bilayers [299], and used to obtain a harmonic approximation to the interlayer interaction. There, the Keldysh interaction with argument $s \rightarrow \sqrt{(s^2 + d^2)}$ is used to model interlayer interactions. This is clearly not the same as the full bilayer interaction defined by Eqn. (3.16), but it does avoid a pathological feature encountered in our current approximation to the interlayer interaction, in Eqn. (B.3). Our approximation features a logarithmic singularity at $s = 0$, which is not physical, and not present in the exact bilayer interaction. As a consequence, our approximate Keldysh interaction requires unphysical cusps in the trial wave functions of interlayer complexes, and we may expect it to behave poorly where interlayer physics dominates the binding of a complex.

As noted in Ref. [118], another non-pathological approximation to $\phi_{\text{inter}}(s)$ is possible, which is of displaced Coulomb form, and does not feature the divergence at $s = 0$. Such an approximation does, however, fail to describe adequately the effects of two-dimensional screening at intermediate range.

Appendix C

Biexcitons in TMD

homostructures

Here, we present our complete data set for calculations on excitons and biexcitons in few-layer TMD homostructures. Note that, because of the z -inversion symmetry of our homostructure, some possible complexes are not listed (they are equivalent to ones that are listed). Complexes (or similar families of complexes) that are not explicitly listed in a homostructure of n layers are omitted because their analogues were identified as unbound in a lesser m layer homostructure ($m < n$).

Our naming convention indexes the layer number of electrons as prefix superscript numbers, and holes as standard subscripts. For example, the intralayer exciton formed from an electron in layer 1, and a hole in layer 2 of a homostructure is labelled “ 1X_2 ”. An interlayer biexciton formed from electrons in layers 1 and 4, with holes in layers 1 and 3 is labelled “ $^{14}XX_{13}$ ”.

For the case of excitons, we will quote binding energies to five significant figures.¹ In all cases, the error in the exciton binding energy is less than a millionth of an meV. In calculations of biexciton binding energies, full precision in exciton energies has been maintained, and individual errors propagated accordingly.

¹So as to save horizontal space on the page.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
1X_1	541.46	194.82	219.51	492.20
$^{11}XX_{11}$	22.9(1)	15.356(5)	16.297(3)	22.26(9)

Table C.1: Charge-carrier complex binding energies (in meV) in a MoSe₂ monolayer, in various dielectric environments.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
1X_1	454.44	170.97	197.17	430.51
1X_2	218.75	63.992	66.421	85.813
$^{11}XX_{11}$	22.34(3)	14.574(3)	15.811(5)	21.94(2)
$^{11}XX_{12}$	29.30(1)	39.839(8)	41.997(7)	146.21(1)
$^{12}XX_{11}$	28.767(9)	39.807(3)	41.923(5)	145.788(9)
$^{11}XX_{22}$	–	–	–	–

Table C.2: Charge-carrier complex binding energies (in meV) in an $n = 2$ MoSe₂ homostructure, in various dielectric environments.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
1X_1	427.91	162.86	189.99	408.69
1X_2	185.44	57.921	60.754	72.298
1X_3	142.14	48.178	48.947	67.794
2X_1	185.44	57.921	60.754	72.298
2X_2	392.86	162.86	189.99	408.69
$^{12}XX_{11}$	28.97(1)	36.502(4)	38.873(5)	132.33(1)
$^{13}XX_{13}$	132.37(9)	24.86(3)	21.49(3)	143.30(6)
$^{11}XX_{11}$	21.98(9)	14.494(4)	15.771(4)	21.84(8)
$^{13}XX_{12}$	–	–	–	–
$^{22}XX_{22}$	92.44(1)	14.491(3)	15.770(4)	21.94(1)
$^{11}XX_{13}$	62.559(8)	33.67(1)	36.013(4)	127.14(2)
$^{22}XX_{23}$	64.512(9)	36.518(4)	38.951(5)	132.787(9)
$^{12}XX_{13}$	–	–	–	–
$^{12}XX_{12}$	168.05(4)	24.52(3)	22.57(2)	143.06(6)
$^{11}XX_{12}$	29.47(1)	36.515(4)	38.951(5)	132.765(9)
$^{12}XX_{22}$	64.015(9)	36.483(8)	38.866(7)	132.30(1)
$^{13}XX_{11}$	163.410(6)	64.505(4)	64.461(3)	242.785(7)

Table C.3: Charge-carrier complex binding energies (in meV) in an $n = 3$ MoSe₂ homostructure, in various dielectric environments.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
1X_1	415.11	158.85	186.52	397.57
1X_2	170.74	54.577	57.730	64.363
1X_3	123.97	45.065	46.115	60.299
1X_4	106.18	38.556	38.536	56.924
2X_1	170.74	54.577	57.730	64.363
2X_2	373.47	158.85	186.52	397.57
2X_3	159.70	54.577	57.730	64.363
2X_4	123.97	45.065	46.111	60.299
$^{12}XX_{24}$	211.8(1)	37.25(4)	34.70(3)	224.62(6)
$^{12}XX_{23}$	15.51(5)	18.86(3)	16.94(3)	120.40(4)
$^{14}XX_{23}$	–	18.86(3)	18.14(3)	–
$^{13}XX_{13}$	166.91(2)	22.13(3)	19.98(2)	137.33(3)
$^{11}XX_{11}$	22.34(3)	14.491(4)	15.765(4)	21.94(3)
$^{14}XX_{14}$	125.21(4)	21.71(4)	19.63(3)	137.4(1)
$^{22}XX_{23}$	180.888(4)	63.135(3)	63.353(3)	240.800(4)
$^{11}XX_{12}$	128.209(4)	63.142(3)	63.348(3)	240.804(4)
$^{13}XX_{23}$	–	10.39(3)	9.80(3)	–
$^{12}XX_{14}$	46.74(8)	18.89(2)	16.71(3)	120.22(2)
$^{23}XX_{22}$	180.547(4)	63.017(3)	63.162(3)	240.478(4)
$^{13}XX_{14}$	4.16(8)	18.60(3)	16.84(2)	120.2(1)
$^{13}XX_{12}$	12.4(4)	18.72(2)	16.60(4)	119.8(1)
$^{13}XX_{11}$	170.11(6)	63.244(2)	63.364(3)	240.46(2)
$^{12}XX_{12}$	166.64(5)	21.62(2)	19.41(2)	137.6(1)
$^{23}XX_{14}$	–	18.130(8)	18.15(2)	–
$^{11}XX_{13}$	170.454(3)	63.315(2)	63.474(3)	240.769(3)
$^{12}XX_{11}$	127.869(3)	63.025(3)	63.170(3)	240.471(3)
$^{23}XX_{23}$	208.08(7)	21.45(4)	19.40(3)	137.4(1)
$^{22}XX_{12}$	71.13(1)	35.404(7)	37.967(4)	127.980(9)
$^{14}XX_{13}$	4.33(6)	18.63(4)	16.64(3)	120.4(1)
$^{13}XX_{24}$	–	28.8931(5)	29.2724(8)	192.0(1)
$^{11}XX_{14}$	183.2(2)	63.249(2)	63.308(2)	240.5(1)
$^{14}XX_{12}$	46.73(7)	19.02(2)	16.77(3)	120.6(1)
$^{12}XX_{22}$	70.603(9)	35.382(4)	37.886(5)	127.56(4)
$^{22}XX_{22}$	105.59(1)	14.487(3)	15.763(4)	21.92(1)
$^{23}XX_{13}$	–	10.51(3)	9.83(3)	53.1(1)
$^{13}XX_{33}$	117.42(1)	44.895(5)	49.498(5)	131.59(1)
$^{14}XX_{11}$	184.216(4)	63.229(2)	63.276(2)	240.429(4)
$^{22}XX_{24}$	212.090(4)	63.316(2)	63.472(3)	240.75(3)

Table C.4: Charge-carrier complex binding energies (in meV) in an $n = 4$ MoSe₂ homostructure, in various dielectric environments.

Complex	Dielectric environment			
	Vac.	Iso. hBN	Aniso. hBN	mono. hBN
¹ X ₁	407.58	156.48	184.48	390.83
¹ X ₂	162.43	52.475	55.866	59.081
¹ X ₃	114.39	43.071	44.336	55.255
¹ X ₄	94.598	32.550	36.850	52.103
¹ X ₅	85.035	32.081	31.667	49.447
² X ₁	162.43	52.475	55.866	59.081
² X ₂	363.68	156.48	184.48	390.83
² X ₃	148.02	52.474	55.866	59.081
² X ₄	109.30	43.071	44.336	55.255
² X ₅	94.598	36.672	36.851	52.103
³ X ₁	114.39	43.071	44.336	55.255
³ X ₂	148.02	52.475	55.866	59.081
³ X ₃	356.97	156.48	184.48	390.83
²² XX ₁₂	73.39(1)	34.900(3)	37.539(4)	125.83(1)
¹³ XX ₃₃	117.46(1)	31.900(3)	34.397(3)	119.08(1)
³³ XX ₁₃	117.99(1)	31.954(3)	34.500(4)	119.513(9)
²² XX ₂₂	110.12(2)	14.474(6)	15.766(3)	21.93(1)
¹¹ XX ₁₄	189.682(4)	66.8(3)	62.853(3)	–
¹³ XX ₁₂	8(6)	16.45(2)	14.30(4)	–
²³ XX ₂₄	87.9(1)	16.35(3)	14.92(2)	115.32(7)
²⁴ XX ₂₅	–	8.54(3)	8.14(2)	48.11(7)
²⁴ XX ₂₂	221.94(7)	62.733(2)	62.919(2)	239.43(1)
¹⁴ XX ₃₄	31.1(1)	7.56(2)	6.60(4)	43.04(5)
¹⁴ XX ₁₃	15.2(1)	16.22(3)	14.77(3)	115.53(9)
¹⁴ XX ₄₄	121.60(2)	33.5(3)	31.766(4)	113.290(8)
²³ XX ₂₃	216.83(4)	21.22(2)	18.86(3)	137.79(9)
³³ XX ₃₃	123.50(2)	14.469(4)	15.766(3)	21.93(1)
¹⁴ XX ₄₅	152.9(1)	33.900(4)	30.58(4)	–
²³ XX ₂₄	58.2(2)	16.32(3)	14.79(3)	115.1(1)
¹² XX ₁₃	15.3(1)	16.47(2)	14.39(4)	115.41(4)
¹⁵ XX ₁₄	0.8(1)	16.31(3)	14.46(3)	115.25(8)
¹⁵ XX ₁₁	196.075(3)	62.590(2)	62.664(2)	239.487(6)
³³ XX ₃₄	94.42(8)	34.897(3)	37.540(4)	125.72(3)
¹³ XX ₁₁	172.989(4)	62.732(2)	62.919(3)	239.519(4)
²² XX ₂₅	233.49(8)	56.316(2)	55.365(2)	239.801(4)
²³ XX ₂₂	187.478(4)	62.526(2)	62.737(3)	239.553(4)
¹² XX ₂₂	72.88(1)	34.888(3)	37.461(4)	125.39(1)
¹⁵ XX ₁₅	122.6(2)	21.33(2)	18.32(3)	135.01(1)
¹⁴ XX ₁₁	189.365(4)	62.709(2)	62.833(3)	239.59(8)
¹⁴ XX ₂₄	–	8.45(3)	7.75(3)	47.96(6)
¹¹ XX ₁₅	196.384(3)	62.542(2)	62.616(2)	239.789(3)
¹³ XX ₃₃	94.01(1)	34.892(3)	37.459(4)	125.368(9)
¹¹ XX ₁₂	129.6(1)	62.640(2)	62.918(3)	239.862(3)
¹² XX ₁₂	166.51(4)	21.08(2)	18.51(3)	135.12(5)
²² XX ₂₃	187.811(4)	62.644(2)	62.920(3)	239.865(4)
¹¹ XX ₁₃	173.311(5)	62.802(2)	63.031(3)	239.835(3)
²² XX ₂₄	222.293(4)	62.801(2)	63.029(3)	239.838(4)
¹³ XX ₂₃	–	8.72(2)	7.66(4)	48.3(2)
¹⁵ XX ₁₂	61.1(1)	20.8(3)	14.70(3)	115.49(2)

$^{13}\text{XX}_{13}$	173.35(6)	21.42(2)	18.76(2)	134.70(7)
$^{24}\text{XX}_{24}$	210.25(3)	21.44(2)	18.96(2)	–
$^{13}\text{XX}_{15}$	44.80(9)	16.30(4)	14.53(3)	115.3(1)
$^{12}\text{XX}_{15}$	60.73(9)	10.31(2)	14.66(5)	112.23(2)
$^{13}\text{XX}_{14}$	15.4(1)	16.14(3)	14.76(2)	–
$^{24}\text{XX}_{23}$	59.1(1)	16.61(2)	14.61(3)	115.31(9)
$^{15}\text{XX}_{13}$	44.7(1)	16.11(2)	14.70(3)	115.2(1)
$^{14}\text{XX}_{12}$	49.7(2)	16.82(3)	14.68(4)	115.28(4)
$^{23}\text{XX}_{13}$	–	8.49(3)	7.95(3)	48.0(1)
$^{12}\text{XX}_{11}$	129.177(4)	62.627(2)	62.745(3)	239.546(3)

Table C.5: Charge-carrier complex binding energies (in meV) in an $n = 5$ MoSe₂ homostructure, in various dielectric environments.

Appendix D

Hartree-Fock theory of Keldysh-screened electron gases

Here we give an overview of the differences in the Hartree-Fock treatment of electron gases which are subject to Keldysh screening, and those which are not.

The single relevant change is that to the exchange contribution of the single-particle energy band and hence the HF total energy. The exchange contribution to the renormalised single-particle energy is given by

$$\begin{aligned}\varepsilon_{\mathbf{k}\sigma}^{(\text{ex})} &= - \int_{q \leq k_{F\sigma}} \frac{d^2\mathbf{q}}{(2\pi)^2} \frac{2\pi e^2}{|\mathbf{k} - \mathbf{q}|(1 + r_\star|\mathbf{k} - \mathbf{q}|)}, \\ &= - \frac{e^2}{2\pi} \int_{q \leq k_{F\sigma}} d^2\mathbf{q} \left[\frac{1}{|\mathbf{k} - \mathbf{q}|} - \frac{r_\star}{1 + r_\star|\mathbf{k} - \mathbf{q}|} \right].\end{aligned}\quad (\text{D.1})$$

where we've employed a partial fraction decomposition. Recognising the first contribution as that which accompanies the Coulomb interaction (which we will label $\varepsilon_{\mathbf{k}\sigma}^{(\text{ex};C)}$), we then have

$$\begin{aligned}\varepsilon_{\mathbf{k}\sigma}^{(\text{ex})} &= \varepsilon_{\mathbf{k}\sigma}^{(\text{ex};C)} + \frac{e^2 k_{F\sigma}}{2\pi} \int_0^1 dx \, x \int_0^{2\pi} \frac{d\theta}{(r_\star k_{F\sigma})^{-1} + \sqrt{x^2 + y^2 - 2xy \cos(\theta)}}, \\ &= \varepsilon_{\mathbf{k}\sigma}^{(\text{ex};C)} + \varepsilon_{\mathbf{k}\sigma}^{(\text{ex};K)},\end{aligned}\quad (\text{D.2})$$

where $x = q/k_{F\sigma}$ and $y = k/k_{F\sigma}$. Notice that when r_\star tends to zero the contribu-

tion from the right-most term also tends to zero.

The excess Hartree-Fock total energy, per particle, is given by

$$\begin{aligned}
\delta e_{\text{HF}} &= \frac{1}{2N} \sum_{\mathbf{k}\sigma} n_{\mathbf{k},\sigma} \varepsilon_{\mathbf{k}\sigma}^{(\text{ex};K)}, \\
&= \frac{A}{2N} \int_0^{k_{F\sigma}} dk \, 2\pi k \, \varepsilon_{\mathbf{k}\sigma}^{(\text{ex};K)}, \\
&= \frac{e^2 k_{F\sigma}^3}{2} \frac{A}{N} \int_0^1 dx \, x \int_0^1 dy \, y \int_0^{2\pi} \frac{d\theta}{(r_\star k_{F\sigma})^{-1} + \sqrt{x^2 + y^2 - 2xy \cos(\theta)}}, \\
&= \frac{e^2 k_{F\sigma}^3}{2} \pi r_s^2 F(r_\star k_{F\sigma}) \sim r_s^{-1} F(r_\star k_{F\sigma}). \tag{D.3}
\end{aligned}$$

This is to be added to the Coulomb Hartree-Fock energy (per particle). The function F appears to admit no analytical form, though it may be studied numerically with relative ease. Let $\eta = r_\star k_{F\sigma} \sim r_\star/r_s$, then the generic behaviour of $F(\eta)$ tells us how the HF total energy depends on both r_s and r_\star (kinetic term aside). For $\eta \gg 1$, F appears to approach a constant value. For $\eta \ll 1$, F scales as η . Importantly, at low densities, or low values of the Keldysh screening parameter, F is a small factor which scales linearly with η . At fixed r_\star , F then goes like r_s^{-1} , giving the excess HF total energy a density scaling which is $\mathcal{O}(r_s^{-2})$.