

Durham E-Theses

*A computational study of Rydberg excitons and
defects in cuprous oxide*

ALISTAIR BREWIN

How to cite:

BREWIN, ALISTAIR (2025) A computational study of Rydberg excitons and defects in cuprous oxide. Doctoral thesis, Durham University.

Use policy

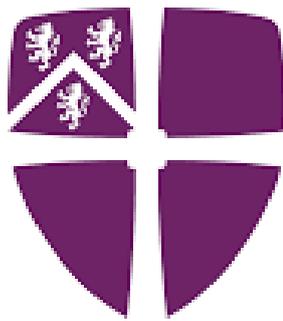


This work is licensed under a [Creative Commons Attribution Share Alike 3.0 \(CC BY-SA\)](https://creativecommons.org/licenses/by-sa/3.0/)

A computational study of Rydberg excitons and defects in cuprous oxide

Alistair Brewin

A thesis presented for the degree of
Doctor of Philosophy



Durham
University

Physics Department
The University of Durham
United Kingdom
March 2025

A computational study of Rydberg excitons and defects in cuprous oxide

Alistair Brewin

Abstract

Various theoretical and computational techniques were applied to Rydberg excitons in cuprous oxide. The techniques fall into two categories: top down, and bottom up. In the first chapter, Floquet theory is applied to a hydrogen-like model of an exciton to model the ultrastrong driving of the exciton by a microwave field. This is the top down approach, where the exciton is considered as an atom and the effect of the crystal is considered only insofar as it alters the energy levels and dipole transition moments of the atom. This method is very successful, quantitatively reproducing the experimental results up to the highest field strengths achievable. We demonstrate how the high energy exciton states hybridise into a quasi-continuum of Floquet states, and we can also predict qualitatively the intensity of sidebands on the laser, opening avenues for microwave-to-optical conversion. However, there are places the model fails, such as in correctly predicting the dependence on the polarisation of the microwave field. This is the limit of what can be done with models that neglect the effects of the crystal. Beginning the ab initio 'bottom up' approach, in chapter 2 we study the effect of point defects on the electronic structure of Cu_2O using density functional theory, with the aim of determining the origin of experimentally observed photoluminescence peaks. The method is very successful in determining which defects are not involved in photoemission, having accounted for many possible sources of erroneous results. However, it is difficult to make positive assignments as to exactly which peak is caused by which defect. Finally, in chapter 3, we begin to bridge the two approaches, studying ab initio the excitations of Cu_2O with time-dependent density functional theory. We demonstrate a novel method for estimating the radius of the 1S exciton, and show the effect that a local defect has on the localisation of the exciton and on the excitation energies. These results demonstrate new theoretical approaches to understanding of Rydberg excitons, especially in applying ab initio techniques, which is not currently seen in the field.

Supervisors: Matthew Jones and Stewart Clark

Acknowledgements

Before I started, I had this romantic idea of a PhD as the ultimate solo project. I am so glad to have discovered this is not the case. Firstly, I am exceedingly lucky to have received the support from many smart, generous people with all manner of technical aspects of this work, giants who enthusiastically offer me their shoulders to stand on. Secondly, and I think more importantly, I am so grateful for the people who have made me someone capable of doing something like this.

Two people who fall into both of these categories are my supervisors, Matt Jones and Stewart Clark. The circumstances of my project gave me two equal supervisors, and I have benefitted greatly from it. Matt Jones is a man of incredible insight, something I have tried my best to emulate. He has taught me not only to be a careful, attentive scientist but also to give good talks, ask better questions and to write well (as I hope the reader will agree). The only thing he cares more about than science are the people he is responsible for. Stewart Clark is a man of outstanding generosity. His door is quite literally always wide open. I've benefitted from many hours of conversation with him about any DFT problem I could possibly have, and he always has the answer and the patience for me to slowly understand it.

In the camp of people who have made me, I want to thank my family, especially my mum Liz, my dad David, my brother Jamie, and my aunt Sue. They've had endless time to talk to me about my journey through this project, both to help me through all the mental challenges, and to share in the joys of success. The people I have lived with and spent all my time with have also become a second family to me. I could write another thesis on all these people, but I'm tired from this one, so I'll just say: to Ben, Cam, and Adil, thanks for keeping me from the dangers of working too hard in my first year; to Eoin, Mike, Adarsh, Camille, Tim, Alex, Big Sean, Josh, Joe, even on the hardest days I knew I'd be coming home to some of my best friends; to Ally and Bex, for the great and terrible influences they have had

on me; to Sam for keeping me off the ground; and to my friends Jack, Pietro, Rose, Keerthi, Emma, Paul, Harish, Nathan, Mitch, Lolly, and all the other members of QLM and CMP; I love you all.

Special special thanks from the QLM side to my friends and colleagues Dr Liam Gallagher (no relation) and DFT wizard Vis Ravindran, who in all fairness have as much claim to this PhD as I do.

Contents

Declaration	vii
List of Figures	viii
List of Tables	xi
Nomenclature	xii
1 Introduction	1
1.1 Background on Rydberg excitons in Cu_2O	3
2 Ultrastrong driving with microwaves	5
2.1 Background	5
2.2 The experiment	8
2.3 Theory	9
2.3.1 Floquet Theory	9
2.3.2 The Hamiltonian	11
2.3.3 Schrödinger equation	14
2.3.4 Absorption	16
2.3.5 Sidebands	17
2.3.6 Floquet State Contribution	18
2.3.7 Polarisation Dependence	20
2.3.8 Background absorption	22
2.3.9 Model parameters	23
2.3.10 Microwave field calibration	24
2.4 Results	26
2.4.1 Absorption Spectra	26
2.4.2 Floquet states	28

2.4.3	Phases of the dipole matrix elements	34
2.4.4	Sidebands	36
2.5	Extensions to the model	37
2.5.1	A possible broadening of the excitonic states	39
2.5.2	Future experiments	41
2.6	Conclusions	41
3	DFT study of defects in Cu₂O	43
3.1	Background	43
3.2	Theory	46
3.2.1	Born-Oppenheimer Approximation	47
3.2.2	Density Functional Theory	48
3.2.3	Hohenberg-Kohn Theorems	48
3.2.4	The Kohn-Sham auxiliary system	50
3.2.5	Periodic systems and Bloch's theorem	53
3.2.6	Exchange-Correlation Functionals	55
3.2.7	Pseudopotentials	58
3.2.8	Spin polarisation and spin-orbit coupling	59
3.3	Convergence with parameters	61
3.4	DFT on pure Cu ₂ O	63
3.4.1	Band Structures	63
3.4.2	Perdew–Burke–Ernzerhof	64
3.4.3	HSE06	65
3.5	Supercells	65
3.6	The Defects	67
3.6.1	Searching for defect states	67
3.6.2	Oxygen interstitials	71
3.6.3	Copper interstitials	74
3.6.4	Copper vacancy	77
3.6.5	Split copper vacancies	77
3.6.6	Oxygen vacancy	81
3.6.7	Anti-sites	83
3.6.7.1	Copper replacing oxygen	83
3.6.7.2	Oxygen replacing copper	83
3.6.8	The effect of charge	86
3.6.9	Spin-orbit coupling	86
3.6.10	Formation energies	89
3.7	Discussion and conclusions	90
4	Excitons in DFT	94

4.1	Background	94
4.2	TDDFT Theory	95
4.2.1	Runge-Gross theorem	95
4.2.2	Linear response	96
4.2.3	KS linear response	97
4.2.4	XC kernel	99
4.2.5	Computational details	100
4.3	Results and discussion	100
4.3.1	Pure Cu ₂ O	100
4.3.2	Tetrahedral oxygen interstitial	104
4.4	Conclusions	108
5	Conclusion	111
	Appendix A Transforming the Dyson-type equation into frequency space	114
	Bibliography	117

Declaration

The work in this thesis is based on research carried out at the Department of Physics, University of Durham, England. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

Some of the work presented in this thesis has been published in journals and conference proceedings - the relevant publications are listed below.

Publications

Alistair Brewin, Liam A P Gallagher, Jon D. Pritchett, Horatio Q. X. Wong, Robert M. Potvliege, Stewart J. Clark and Matthew P. A. Jones. “Microwave-optical spectroscopy of Rydberg excitons in the ultrastrong driving regime”. *New Journal of Physics*, Volume 26, 113018, November 2024.

Jon D. Pritchett, Liam A. P. Gallagher, **Alistair Brewin**, Horatio Q. X. Wong, Wolfgang Langbein, Stephen A. Lynch, C. Stuart Adams, Matthew P. A. Jones. “Giant microwave–optical Kerr nonlinearity via Rydberg excitons in cuprous oxide”. *APL Photonics* 9, 031303 (2024)

Copyright © 2025 by Alistair Brewin.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

List of Figures

1.1	An example absorption spectrum of a laser by Cu_2O	3
2.1	A diagram of an optical cavity made of two mirrors coupled to an atom (or artificial analogue)	5
2.2	A diagram of the critical parts of the experimental setup	8
2.3	An example of a scan over the spectral range of the Fabry-Pérot etalon	9
2.4	A diagram of the Hamiltonian of the model	12
2.5	A diagram of the Floquet Hamiltonian	18
2.6	The dependence of the change in laser absorption on the relative polarisation angle between the microwave and laser fields	22
2.7	Experimental exciton absorption spectrum, with the fitted phonon background function shaded in red	25
2.8	The best-fitting microwave electric field strength for $\Delta\alpha L$ (eq. 2.58) at small input microwave powers is plotted against the square root of the power	26
2.9	Plots of the (a) experimental and (b) theoretical absorption spectra at different microwave field strengths, \mathcal{E}_m	27
2.10	Comparison between the experimental (filled circles) and theoretical (solid lines) values of the change in absorption $\Delta\alpha L$ as a function of laser energy	29
2.11	The real part of the eigenenergies of all the Floquet states in the region near $\hbar\omega_p = E_{10P}$	30
2.12	The real part of the eigenenergies of the Floquet states are plotted to show how the underlying structure of the exciton changes with microwave field strength	32
2.13	A heuristic for how much of the wavefunction oscillates at less than frequency $ N \omega_m$, when $\hbar\omega_p = E_{12P}$	33

2.14	Plots of the theoretical absorption spectra at different microwave field strengths, \mathcal{E}_m , for two different sets of randomly chosen phases of $\langle 0 \hat{D} \cdot \epsilon_p nP \rangle$	35
2.15	The intensity, expressed in counts per second, of second (colour) and fourth (black) order red- and blue-detuned sidebands for theory (lines) and experiment (points) at four different probe laser detunings	37
2.16	The intensity, expressed in counts per second, of second (colour) and fourth (black) order red- and blue-detuned sidebands for theory (lines) and experiment (points) at four different probe laser detunings	38
2.17	The probe absorption spectra at different microwave field strengths for (a) the experiment, and (b) the model with artificially large widths for the S, D, and F states	40
2.18	Comparison of the change in absorption $\Delta\alpha L$ as a function of laser energy between the experiment (filled circles) and the model with artificially large widths for the S, D, and F states (solid lines)	41
2.19	The theoretical laser absorption spectra for microwave field strengths above those of the experiment	42
3.1	A sketch of an example PL spectrum	45
3.2	An example of a pseudopotential and its pseudo-wavefunction	60
3.3	The convergence of the final energy of the Cu ₂ O unit cell	62
3.4	The high symmetry point labels in the first Brillouin zone for the primitive cubic lattice (the lattice for cuprous oxide). The Γ point is at $\mathbf{k} = (0, 0, 0)$. Band structures in this thesis follow paths along the red lines. The points R, M, and X represent $\mathbf{k} = (1, 1, 1) \pi/a$, $\mathbf{k} = (1, 1, 0) \pi/a$, and $\mathbf{k} = (1, 0, 0) \pi/a$, respectively.	63
3.5	The band structure of a single Cu ₂ O unit cell under the PBE xc-functional. Occupied bulk states are shown in green, while unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The band gap, $E_g = 0.499$ eV, is at the Γ point, making Cu ₂ O a direct-gap semiconductor. Note that E_g as given by PBE-DFT is less than the experimental gap.	64
3.6	Cu ₂ O unit cell band structure under HSE06	66
3.7	Supercell band structures for pure Cu ₂ O	69
3.8	Band structures for the octahedral form of the oxygen interstitial	70
3.9	Band structures for the tetrahedral form of the oxygen interstitial	72
3.10	Band structures for the octahedral form of the copper interstitial	73
3.11	Band structures for the tetrahedral form of the copper interstitial	75
3.12	Band structures for the simple copper vacancy	76

3.13	Band structures for the first form of the split copper vacancy	79
3.14	Band structures for the second form of the split copper vacancy	80
3.15	Band structures for the oxygen vacancy	82
3.16	Band structures for the copper anti-site	84
3.17	Band structures for the oxygen anti-site	85
3.18	The effect of spin-orbit coupling on band structures	88
3.19	A sketch of all the commonly observed peaks in photoluminescence experiments on Cu_2O , and our assignment of the defects which cause them	92
4.1	The first excited state of the pure Cu_2O crystal under PBE in the $2 \times 2 \times 2$ supercell	101
4.2	Isosurfaces of the electron density of the VBM and CBM of Cu_2O in the $2 \times 2 \times 2$ supercell.	102
4.3	The band structure of the tetrahedral oxygen interstitial under PBE in the $2 \times 2 \times 2$ supercell	103
4.4	The first excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell	105
4.5	Isosurfaces of the electron density of relevant bands of O_{tet}^i in the $2 \times 2 \times 2$ supercell.	106
4.6	The fifth excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell	107
4.7	The sixth excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell	109

List of Tables

3.1	The enthalpy of formation of the defects studied	89
4.1	Convergence of excitation energies with supercell size	100
4.2	The first 10 excited states of O_{tet}^i in the $2 \times 2 \times 2$ supercell	104

Nomenclature

FWHM full width half maximum

HWHM half width half maximum

RWA rotating wave approximation

SC strong coupling

USC ultrastrong coupling

DSC deep strong coupling

SD strong driving

USD ultrastrong driving

DSD deep strong driving

DFT density functional theory

PL photoluminescence

VBM valence band maximum

CBM conduction band minimum

KS Kohn-Sham

LDA local density approximation

GGA generalised gradient approximation

HEG homogeneous electron gas

HF Hartree-Fock

NCP Norm-conserving pseudopotentials

MP Monkhorst-Pack

TDDFT time-dependent density functional theory

Introduction

Rydberg atoms are atoms with an electron excited to a high principle quantum number, n . In these states, the excited electron orbits far from the nucleus and core electrons, allowing Rydberg atoms to be treated analogous to hydrogen. Much work has been done in the last few decades with Rydberg atoms in pursuit of quantum technologies. The separation between the electron and the nucleus in Rydberg states gives pairs of Rydberg atoms dipole-dipole interactions 10 orders of magnitude larger than that of ground state atoms [1, 2]. These tunable interactions make them one of the centres of the current quantum computing boom, where they are used to make quantum logic gates between Rydberg atom qubits [3, 4], or quantum simulators of many-body quantum systems [5, 6].

The same dipoles that make them sensitive to other Rydberg atoms also make them sensitive to electromagnetic fields. This makes them useful for various non-linear quantum optical effects, like single photon emitters [7, 8], receivers [9] and transistors [10, 11]. Parallel to their use as qubits in quantum computing, there are also ongoing efforts to use their sensitivity to ‘read out’ microwave photons in superconducting microwave qubits [12, 13, 14]. However, the marriage of these fields is difficult. Rydberg atoms require ultra-high vacuum chambers with laser access to cool and trap them, a complex and lengthy process to set up in and of itself. Superconducting qubits operate in dilution fridges in the dark to maintain sub-kelvin temperatures in solid state, which is hard to experimentally realise in the same location.

The fundamental excitation of semiconductors is the exciton, a bound state of an electron and a hole also resembling a hydrogen atom, and in 2014 Kazimierczuk et al. [15] observed Rydberg states of excitons up to $n = 25$ in cuprous oxide, opening the field of Rydberg excitons. As an alternative platform for Rydberg physics, Rydberg excitons have a lot of promise, exhibiting similarly strong interactions [16]

and sensitivity to fields [17]. Like superconducting resonators, they are observed in materials at cryogenic temperatures, with the highest n observed in dilution fridges [18], making them naturally more compatible.

Excitons were first predicted in the 1930s under two different schemes by Frenkel [19] and Wannier [20]. The Frenkel exciton is tightly bound and is typically seen in the context of molecules as the bound state of occupied and unoccupied molecular orbitals [21]. In contrast, the Mott-Wannier excitons seen in semiconductors can have extended wavefunctions due to the relative mobility of conduction electrons and valence holes in crystals. This latter form of exciton exhibits energy levels analogous to a hydrogen atom, and they were first observed in the 1950s [22] in Cu_2O . Cuprous oxide is particularly suited to supporting Mott-Wannier excitons because of the exceptionally free particle-like behaviour of the electron and hole (discussed in chapter 3). Whereas in most semiconductors exciton energy levels are rarely seen above $n = 3$ [23], excitons in Cu_2O have been seen up to $n = 30$, with radii in excess of $1 \mu\text{m}$ [18].

The theory of Rydberg excitons is usually approached by analogy to Rydberg atoms, studying their wavefunctions from Mott-Wannier theory like atoms in a vacuum chamber and forgetting about the crystal environment. This has proven successful, and can be used to understand effects like large optical non-linearities [24, 25] and inter-exciton interactions [16]. In chapter 2 of this thesis, I explore the limits of this ‘top down’ approach by applying it to the ultrastrong driving of Rydberg excitons by a microwave field. It performs remarkably, reproducing well many experimental results up to the strongest fields applied, and provides useful insight into the experiment. However, as discussed in the chapter, we find corners where the lack of a microscopic treatment of the crystal environment limits the performance of this type of model.

In chapter 3, we take the ‘bottom up’ approach, where we study *ab initio* the physics of native defects in Cu_2O , long-known to have a significant effect on the Rydberg exciton series [26, 27, 28]. There are many methods in condensed matter physics for studying the electronic properties of solids, but by far the most popular is density functional theory (DFT) [29], which is the approach taken in this work as well. There have been several previous studies on defects in Cu_2O with DFT [30, 31], as it is of interest to the solar cell community, however despite seeming a simple material, it has proven resistant to naive applications of DFT. Here we present the most thorough DFT study of defects in Cu_2O to date, combining fragmentary knowledge from previous works with new insights to identify which features caused by the defects are robust and which are spurious effects of theoretical limitations, an attention to detail not usually seen in the study of defects in general [32, 33,

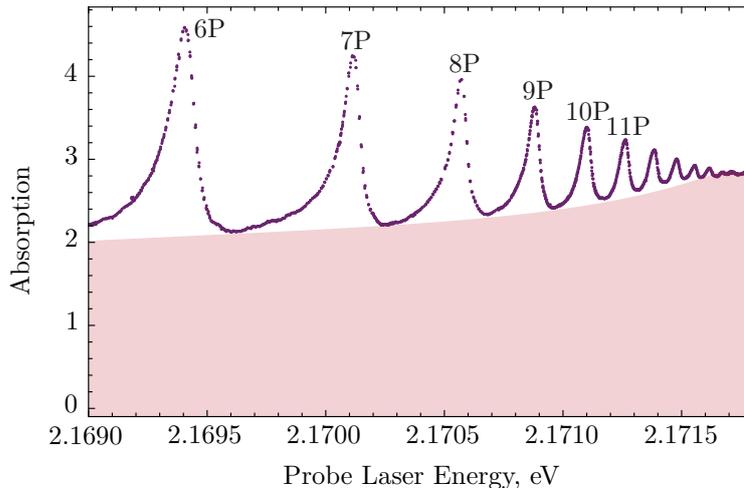


Figure 1.1: An example absorption spectrum of a laser by Cu_2O , showing some of the peaks of the $n\text{P}$ Rydberg exciton series. The phonon-mediated absorption background is shaded in red.

34, 35, 36].

Chapter 4 begins to build the bridge between these two regimes. There are several methods that can be used to extend DFT (a ground state theory) to excited states, and so study excitons from a purely ‘bottom up’ perspective. Due to its exploratory nature, this study uses time-dependent DFT (TDDFT), the simplest approximation beyond DFT, to make the first steps in this direction. These methods have previously been applied to excitons [37], and to defects [38], but in this work we apply them for the first time to the defect-exciton interaction in Cu_2O .

1.1 Background on Rydberg excitons in Cu_2O

Cuprous oxide has a band gap of $E_g = 2.18$ eV between the uppermost valence band and the lowermost conduction band. The exciton formed by the excitation between these two bands is called the yellow exciton series after the colour of laser required to excite it. There are other colours of exciton that are studied in Cu_2O [39], but in this work we only consider the yellow exciton series.

In the Mott-Wannier theory of excitons, each state can be labelled, like hydrogen, by its principal quantum n , angular momentum l , and z-component of orbital angular momentum m . Aside from the low lying states ($n < 4$), the states’ energies follow a Rydberg series with a quantum defect,

$$E_{nlm} = E_g - \frac{R_X}{(n - \delta_l)^2} \quad (1.1)$$

where R_X is the Rydberg constant, and δ_l is called the *quantum defect*. Whereas in hydrogen the l states are (almost) degenerate, for the exciton in Cu_2O (as well as Rydberg states in other atoms) the Rydberg energies are shifted by a different δ_l for each l , which are determined by fitting to experimental spectra [40], an effect arising from the non-parabolicity of the valence band [41].

The valence band in cuprous oxide has a character dominated by copper D orbitals, and so has even parity. As such, it is dipole-forbidden to excite from the valence band directly to an even parity exciton, such as nS or nD states. Single photon absorption experiments show no excitations of nS or nD states, demonstrating that l is a good quantum number for this system. The same experiments show excitations of nP and nF states, with the coupling to nF states several orders of magnitude smaller than for nP states. Therefore, to a good approximation, single photon absorption experiments only excite the nP exciton series, an example of which is shown in figure 1.1. Even though the excitation of the $1S$ exciton is dipole forbidden, the excitation of the $1S$ exciton and an even parity phonon is allowed. This absorption process is of similar strength to that of the P series excitons, and so the nP Rydberg series sits on a large absorption background.

Ultrastrong driving with microwaves

2.1 Background

At the forefront of modern quantum physics is the ever stronger interaction between light and matter. The coupling of the two marries the two biggest fields in quantum physics, with far reaching applications in quantum communication [42, 43, 44], simulation [45, 46], computation [47], and chemistry [48, 49]. To maximise the interaction, one can either increase the strength of the electromagnetic field, so that it drives the matter harder, or increase the electric dipole of the matter, so it responds more to the light.

One approach to the first strategy, known as cavity QED, is to place the matter in what is called an optical cavity (figure 2.1) by sandwiching it between two (typ-

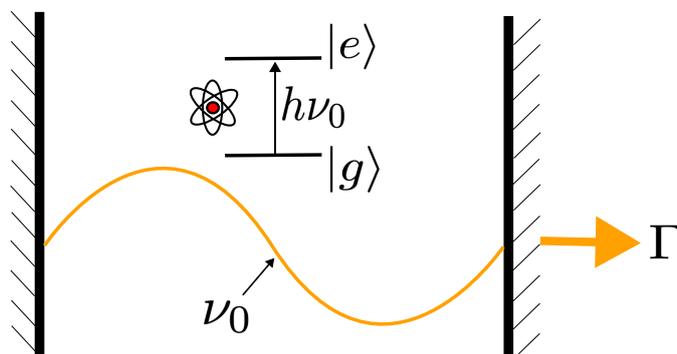


Figure 2.1: A diagram of an optical cavity made of two mirrors coupled to an atom (or artificial analogue). The resonant frequency of the cavity, ν_0 , matches the transition frequency between the ground, $|g\rangle$, and excited, $|e\rangle$, states of the ‘atom’. Photons are lost from the cavity at a rate Γ .

ically Bragg) mirrors. Alternatively, in circuit QED, artificial atoms are coupled electrically to a superconducting resonator [50]. Either way, the field is confined to a small mode volume, which greatly increases the strength of the field at the resonant frequency, ν_0 , and can result in quantum light interacting with quantum matter. Outside of resonators, strong classical electromagnetic fields can also be used to drive quantum systems, which have been shown to exhibit optical nonlinearities like high-harmonic generation [51].

To address the second strategy, common quantum systems with large electric dipoles include superconducting ‘flux qubits’[50], optomechanical devices [52], and, of particular note to this work, Rydberg atoms [53, 6]. Rydberg atoms provide an excellent platform for light-matter coupling because the transition dipole moments between neighbouring states in Rydberg atoms scales as n^2 for principal quantum number n . Since this scaling derives from the n^2 scaling of the Bohr radius, the scaling also exists for Rydberg excitons [54, 25].

It is also much easier to fabricate and cool semiconductor-Bragg mirror stacks than to work with individual trapped atoms. Because of this, much work is being done at the moment on semiconductors in optical microcavities, where the exciton and photon hybridise to form quasi-particles called exciton-polaritons. Exciton-polaritons have been observed to form Bose-Einstein condensates [55], with interesting applications including quantum computing [56], and novel kinds of lasers [57], and can exhibit topological effects [58]. Many groups are now studying cuprous oxide in microcavities, as part of the emerging field of Rydberg exciton-polaritons [59, 60, 24, 61], which can exhibit large optical nonlinearities useful for things like quantum information processing.

To categorise the level of coupling between the light and the matter, several thresholds are used as milestones. They are usually defined (see fig. 2.1) in terms of the bare energy gap of the atom (or analogue), $h\nu_0$, the rate of loss of photons from the cavity, Γ , and the Rabi frequency Ω of the atomic transition. The Rabi frequency is a common quantity in atomic physics quantifying the coupling strength between a transition and a field, given by

$$\Omega = \langle g | \hat{\mathbf{D}} \cdot \mathbf{E} | e \rangle / \hbar \quad (2.1)$$

where $|g\rangle$ and $|e\rangle$ are the ground and excited states, $\hat{\mathbf{D}}$ is the dipole operator, and \mathbf{E} is the electric field*. The strong coupling (SC) regime is said to be reached in cavities when the Rabi frequency for the electromagnetic field of zero photons, called the vacuum Rabi frequency, Ω_0 , exceeds the decay rate but not (yet) the

*This definition is for a classical electric field, an approximation which is always valid for the strong microwave fields considered in this thesis.

bare transition frequency, or $\nu_0 > \Omega_0 > \Gamma$ [49, 62, 63]. In this regime, the atom begins to coherently exchange a photon many times with the cavity before it is lost spontaneously to the environment [64]. Analogous to the SC regime, the strong driving (SD) regime for classical driving fields occurs when $\nu > \Omega > \Gamma$, where $h\nu$ is the energy gap of the driven transition, Ω is the Rabi frequency, and Γ is the decay rate of the system.

When the coupling to the electromagnetic field approaches or even exceeds the transition frequency, one enters the regimes of ultrastrong ($\Omega > 0.1\nu$) and deep-strong ($\Omega > \nu$) coupling/driving [49, 62, 63], where the commonly used rotating wave approximation (RWA) breaks down. This results in extreme quantum effects [65, 62] and complex dynamics instead of simple Rabi oscillations [66], with consequences for quantum computation [67, 68], materials [48, 49, 69], and microwave-optical conversion [70, 71]. Reaching these thresholds requires a large electric dipole and also small energy spacings. Conveniently, the energy gaps between Rydberg (and Rydberg exciton) states scales as n^{-3} , while the electric dipole moments scale as n^2 , so the threshold field strength decreases as n^{-5} with increasing n . Moreover, because the state widths are comparable to the state separations, many transitions can be ultrastrongly driven at the same time with a single microwave field, something that is impossible in atomic systems [53].

In previous works [54, 25], Rydberg excitons were coupled to a weak microwave field and the results could be understood with a perturbative model for the change in susceptibility. Nevertheless, giant non-linear effects were observed [25], with implications for microwave-to-optical conversion. This chapter follows the work done in Brewin et al. [72], where the ultrastrong driving of Rydberg excitons by a microwave field was experimentally demonstrated and theoretically modelled. The focus will be on the modelling I did based on Floquet theory [73], a common method for going beyond the RWA, with just enough experimental details to understand what we are comparing the model to. In that work, we studied the effect of a strong microwave field on the exciton absorption spectrum and the intensity of generated sidebands on the laser, and achieved excellent agreement between the theoretical model and experimental data. The success of the model allows us an understanding of the system that cannot be gained from direct measurement, which is discussed in detail in that work and here. For example, it shows parts of the exciton spectrum reach $\Omega/\nu \sim 4$, well into the deep strong driving regime.

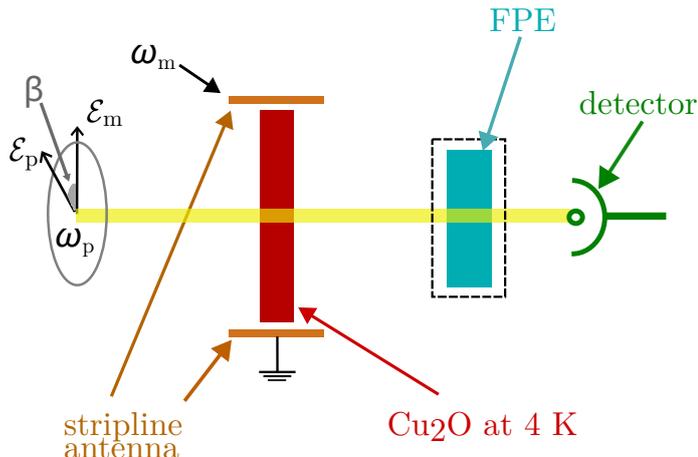


Figure 2.2: A diagram of the critical parts of the experimental setup. The probe laser at optical frequency ω_p passes through a thin film of Cu_2O , kept at 4K in a cryostat, creating excitons. The amount of transmitted light is measured on a detector. A stripline antenna driven with a microwave signal ω_m applies a microwave field across the Cu_2O . The polarisation of the probe laser field \mathcal{E}_p is at an angle β to the polarisation of the microwave field \mathcal{E}_m . A Fabry-Pérot etalon (FPE) is inserted to resolve the sidebands on the transmitted light.

2.2 The experiment

The experiments were run by Liam Gallagher, Jon Pritchett, and Horatio Q. X. Wong and their mechanics have been detailed in various experiment-focussed works [25, 74]. Here I will give the basic overview required for a reader of this thesis, to give the idea of the system we are trying to model.

Figure 2.2 shows a diagram of the experiment. The probe laser is monochromatic at a wavelength of $\lambda \approx 571$ nm. The intensity in the sample is $\sim 20 \mu\text{W mm}^{-2}$, which when accounting for the spot size, is less power than used for the probe laser in [75], allowing us to neglect the effect of inter-exciton interactions. It is incident on a thin slice of Cu_2O , $55 \pm 10 \mu\text{m}$ thick, inside a cryostat at $\sim 4\text{K}$. The stripline antenna (detailed in [54]) is driven at $\omega_m = 7$ GHz by a microwave signal generator, which both alters the probe laser absorption spectrum, and generates new sidebands on the laser light at even multiples of ω_m . The absorption spectrum is measured by comparing the light transmitted through the sample, $T(\omega_p)$, to a reference taken without the sample, $T_0(\omega_p)$, to get the optical depth $\alpha L = -\ln(T(\omega_p)/T_0(\omega_p))$. The Fabry-Pérot etalon, used to resolve the sidebands, has a finesse of 44.5 ± 0.7 and a free spectral range of 60.1 ± 0.2 GHz. An example of the output from scanning the spectral range of the FPE with the microwave field on and off is shown in figure 2.3. The intensity of the carrier signal ($\omega_p/2\pi$, zero detuning) is seen in the broad

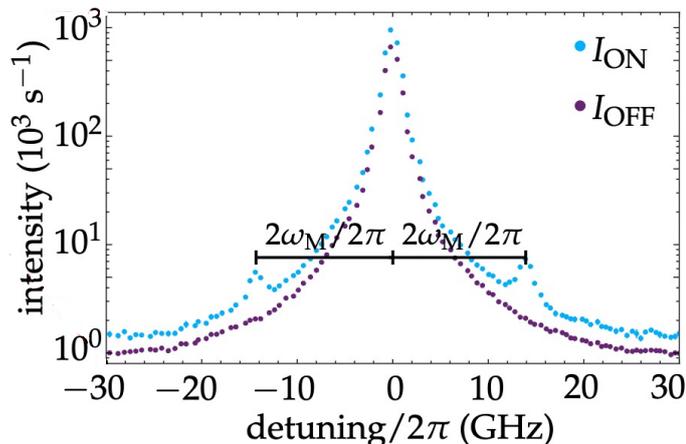


Figure 2.3: An example of a scan over the spectral range of the Fabry-Pérot etalon, with zero detuning corresponding to $\omega_p/2\pi$. The data labelled I_{ON} is the intensity of the light for $\mathcal{E}_m = 170 \text{ V m}^{-1}$, and the data labelled I_{OFF} was taken under no microwave field. At zero detuning the intensity of the carrier signal is measured. The two peaks either side of the central peak are the $\pm 2\omega_m$ ‘second order’ sidebands, only present when the microwave field is switched on. At higher microwave field strengths, fourth order sidebands can be seen.

central peak, and the second order sidebands ($\omega_p \pm 2\omega_m$) are seen either side. The data labelled I_{ON} is the intensity of the light for $\mathcal{E}_m = 170 \text{ V m}^{-1}$, and the data labelled I_{OFF} was taken under no microwave field. At higher microwave field strengths, 4th order sidebands can also be observed. While the features are broad, the frequency width of the transmitted light is actually very narrow. By measuring the transmission of a laser through the etalon, it was determined that the width of these features is introduced almost entirely by the etalon itself. Note that the higher frequency ‘blue’ sideband has a different intensity to the lower frequency ‘red’ sideband; in general, the intensity of the sidebands is not symmetric. The intensity of all the peaks was extracted by fitting Lorentzians to the etalon output spectrum.

2.3 Theory

2.3.1 Floquet Theory

Before beginning the theory proper, I want to give a short introduction to Floquet theory, on which the theory heavily relies. I will give an explanation here that should suffice for this work, but more in depth guides can be found elsewhere if needed [76, 73].

It is very common in quantum physics to encounter time-periodic Hamiltonians, $\hat{\mathcal{H}}(t)$, most notably for oscillating electric fields, as in this work. We can separate the Hamiltonian into a static part and a time periodic ‘interaction’ part, as

$$\hat{\mathcal{H}}(t) = H_0 + H_I(t) = H_0 + H_I(t + T), \quad (2.2)$$

where the interaction part has periodicity $T = 2\pi/\omega$. The static part H_0 we take to have a complete set of eigenstates, $\{|j\rangle\}$, that span the Hilbert space of possible solutions. The Schrödinger equation for these Hamiltonians,

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{\mathcal{H}}(t) |\psi(t)\rangle, \quad (2.3)$$

is known to have solutions of the form

$$|\psi(t)\rangle = \exp(-iEt/\hbar) |\phi(t)\rangle \quad (2.4)$$

where $|\phi(t)\rangle = |\phi(t + T)\rangle$ is periodic with the same period as the Hamiltonian. This result is known as Floquet’s Theorem, the states $|\psi(t)\rangle$ we will call Floquet states*, and the vectors $|\phi(t)\rangle$ we will call the Floquet mode. Although E has the units of energy, the time dependent Hamiltonian does not conserve energy, and so E is called the quasi-energy. Just like all differential equations, the solution we will find for the model will be a linear combination of these particular solutions. Note that the number of independent Floquet modes is not limited to the dimension of the Hilbert space in which we are working (spanned by the exciton states). For integer N ,

$$|\psi(t)\rangle = \exp(i(E + N\hbar\omega)t/\hbar) \exp(-iN\omega t) |\phi(t)\rangle = \exp(iE't/\hbar) |\phi'(t)\rangle \quad (2.5)$$

defines a new Floquet mode $|\phi'(t)\rangle$ for the same Floquet state solution, with quasi-energy $E' = E + N\hbar\omega$. Therefore each solution comes with an infinite set of modes with energies spaced by $\hbar\omega$.

We can use the periodicity to write $|\phi(t)\rangle$ as a Fourier series,

$$|\phi(t)\rangle = \sum_{N=-\infty}^{\infty} f_N \exp(iN\omega t) |\phi_N\rangle \quad (2.6)$$

where the vectors $|\phi_N\rangle$ do not depend on time. We can expand $|\phi_N\rangle$ in the basis $\{|j\rangle\}$ such that

$$|\phi_N\rangle = \sum_j f_{j;N} |j\rangle, \quad (2.7)$$

$$|\phi(t)\rangle = \sum_{N=-\infty}^{\infty} \sum_j c_{j;N} \exp(iN\omega t) |j\rangle, \quad (2.8)$$

*Different sources call them different things. Shirley [73], something of a father of this kind of physics, calls something else a Floquet state, but sadly no one else agrees with him [77, 76, 78].

where $c_{j;N} = f_N f_{j;N}$. Similarly, the components of the Hamiltonian,

$$\mathcal{H}_{jk}(t) = \langle j | \hat{\mathcal{H}}(t) | k \rangle, \quad (2.9)$$

we can expand as

$$\mathcal{H}_{jk}(t) = \sum_{N=-\infty}^{\infty} \mathcal{H}_{jk}^N \exp(iN\omega t). \quad (2.10)$$

Substituting these results into the Schrödinger equation and pre-multiplying by $\langle j |$, we obtain

$$\sum_{N=-\infty}^{\infty} c_{j;N} (E - N\hbar\omega) \exp(iN\omega t) = \sum_{N'=-\infty}^{\infty} \sum_{N''=-\infty}^{\infty} \sum_k \mathcal{H}_{jk}^{N''} c_{k;N'} \exp(i[N' + N'']\omega t). \quad (2.11)$$

Because we require the equality to hold at all times t , and exponentials of different frequencies are orthogonal, we can individually equate the terms in $\exp(iN\omega t)$. These terms appear on the left when $N'' = N - N'$, thus, after some simplifying, we find the set of coupled equations

$$\sum_{N'=-\infty}^{\infty} \sum_k [\mathcal{H}_{jk}^{N-N'} + N'\hbar\omega \delta_{NN'} \delta_{jk}] c_{k;N'} = E c_{j;N}. \quad (2.12)$$

This defines an eigenvalue problem with a new, infinite-dimensional, time-independent Hermitian operator which we will call the Floquet Hamiltonian,

$$\langle j; N | \hat{F} | k; N' \rangle = F_{(j;N),(k;N')} = \mathcal{H}_{jk}^{N-N'} + N'\hbar\omega \delta_{NN'} \delta_{jk} \quad (2.13)$$

with rows indexed by the double indices j and N , and columns by k and N' , written in the basis of product states $|j; N\rangle = |j\rangle \exp(iN\omega t)$. Its eigenstates and eigenenergies are the Floquet modes and their quasi-energies. Its infinite size is an echo of the non-uniqueness of the Floquet modes. While it has an infinite number of eigenvalues, simply changing where we start counting (changing the index $N \rightarrow N + M$) yields for the same state ($c_{j;N+M}$) a shifted eigenvalue $E \rightarrow E + M\omega$. This symmetry is broken when, in doing the actual calculation, we truncate N at some finite value. However, for large enough N_{\max} , it is a good approximation.

2.3.2 The Hamiltonian

The theory in this section was developed with great help from Robert Potvliege, who has since published a paper on the general concept [79]. In this section I will distil the key points of the theory and apply them to Rydberg excitons in cuprous

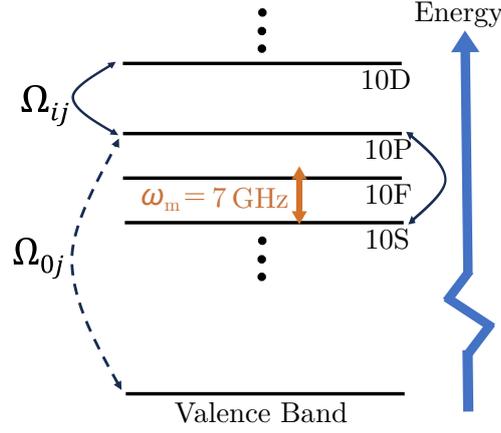


Figure 2.4: A diagram of the Hamiltonian of the model. A ladder of bare excitonic states are included, whose energies and widths are derived from experimental measurements, along with a ground state representing the valence band. The $|n\text{P}\rangle$ states are coupled to the valence band by a weak, optical frequency laser field (dotted arrow), and excitonic states that differ in angular momentum by $\Delta l = \pm 1$ are coupled by a strong microwave field of frequency $\omega_m = 7 \text{ GHz}$ (solid arrows).

oxide. The goal is to model the case of an ensemble of excitons coupled to a strong, homogenous, microwave field,

$$\mathcal{E}_m(t) = \frac{1}{2}\epsilon_m \left[\mathcal{E}_m \exp(-i\omega_m t) + \mathcal{E}_m^* \exp(i\omega_m t) \right], \quad (2.14)$$

and probed by a weak laser field,

$$\mathcal{E}_p(x, t) = \frac{1}{2}\epsilon_p \left[\mathcal{E}_p(x) \exp(-i\omega_p t) + \mathcal{E}_p^*(x) \exp(i\omega_p t) \right], \quad (2.15)$$

where the probe (microwave) electric field has amplitude $\mathcal{E}_p(x)$ (\mathcal{E}_m), polarisation vector ϵ_p (ϵ_m), and angular frequency ω_p (ω_m). Both fields are linearly polarised, and at an arbitrary angle β relative to one another. Because, in our case, the probe field is sufficiently weak (as discussed in section 2.2), the exciton density is low enough that we will neglect inter-exciton interactions and consider just one exciton. The appropriate Hamiltonian for this exciton, as described in figure 2.4, can be thought of as the sum of a set of field-free energy levels (the horizontal lines), the weak probe field coupling the valence band to the $n\text{P}$ states (dashed arrow), and the strong microwave field coupling states of opposite parity (solid arrows),

written respectively as

$$\begin{aligned}
 \hat{H}(t) &= \hbar \sum_{j=1}^{\mathcal{N}} \left(\omega^{(j)} - i \frac{\Gamma^{(j)}}{2} \right) |j\rangle \langle j| \\
 &\quad - \frac{\hbar}{2} \sum_{k=1}^{\mathcal{N}} \left[\left(\Omega_{0k} \exp(-i\omega_p t) + \Omega_{k0}^* \exp(i\omega_p t) \right) |0\rangle \langle k| + \text{h.c.} \right] \\
 &\quad - \frac{\hbar}{2} \sum_{j,k=1}^{\mathcal{N}} \left[\left(\Omega_{jk} \exp(-i\omega_m t) + \Omega_{kj}^* \exp(i\omega_m t) \right) |j\rangle \langle k| \right]. \quad (2.16)
 \end{aligned}$$

The \mathcal{N} excitonic states, labelled $|j\rangle$, have energies $\hbar\omega^{(j)}$ and full width half maximum (FWHM) $\hbar\Gamma^{(j)}$, where $|0\rangle$ labels the ground state (representing the valence band) with $\omega^{(0)} = \Gamma^{(0)} = 0$. The abbreviation ‘h.c.’ stands for Hermitian conjugate.

Keen readers will notice that the complex energies $\omega^{(j)} - i\Gamma^{(j)}/2$ on the diagonal of the Hamiltonian makes it non-Hermitian. Normally, to account for the decay of the exciton, we would have to solve the Lindblad master equation for the above Hamiltonian with real energies, and introduce the decays as jump operators. However, as seen in [79], under the weak probe approximation the same result can be obtained by solving the von Neumann equation for the above non-Hermitian Hamiltonian.

The Rabi frequencies between the states $|j\rangle$ and $|k\rangle$ are given by

$$\Omega_{0k} = \mathcal{E}_p \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | k \rangle / \hbar, \quad (2.17)$$

$$\Omega_{jk>0} = \mathcal{E}_m \langle j | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_m | k \rangle / \hbar, \quad (2.18)$$

with $\hat{\mathbf{D}}$ denoting the dipole operator. The dipole approximation is easily justified for the microwave field since the wavelength of the microwave field (~ 1 cm) is much larger than a Rydberg exciton ($\lesssim 1$ μm). Since the scale over which the laser field varies is much larger than that of the valence band, we also drop the spatial dependence of \mathcal{E}_p . While, in general, the field amplitudes can be complex, in this work we take them to be real, so moving forward we will use $\Omega_{jk} = \Omega_{kj}^*$.

The energies and widths of the $|n\text{P}\rangle$ states and the Rabi frequencies Ω_{0k} coupling the excitons to the valence band were obtained from one-photon excitation experiments on our source, described in section 2.3.8, and are zero if $|k\rangle \notin \{|n\text{P}\rangle\}$. Though F-state resonances were observed in single-photon experiments [80], since the crystal is not perfectly spherically symmetric, they were very weak and the coupling was hard to measure accurately, so we assume $\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | n\text{F} \rangle = 0$. The energies of the S, D and F states were obtained from two-photon excitation experiments [40]. Both even- and odd-parity states could be measured in two-photon experiments since the process for even parity excitation involves dipole absorption

and quadrupole emission, and the odd parity excitation process involves quadrupole absorption and dipole emission, resulting in similar strength signals [81]. The Rabi frequencies Ω_{jk} between excitonic states were calculated from the wavefunctions of a Wannier-Mott model [41, 16] by Valentin Walther, the details of which and of their dependence on β (or lack thereof) presented a significant challenge and are described in section 2.3.7.

The first simplification we can make to the Hamiltonian in equation 2.16 is to transform to rotating coordinates. Under a unitary transformation $\hat{U}(t)$, a Hamiltonian $\hat{H}(t)$ transforms as $\hat{H}_{\text{tr}} = \hat{U}\hat{H}\hat{U}^\dagger + i\hbar\dot{\hat{U}}\hat{U}^\dagger$. If we choose to transform to coordinates where all basis vectors except $|0\rangle$ rotate with the probe laser, using $\hat{U} = |0\rangle\langle 0| + \exp(i\omega_{\text{p}}t)\sum_{j>0}|j\rangle\langle j|$, our transformed Hamiltonian becomes

$$\begin{aligned}\hat{H}_{\text{tr}}(t) &= -\hbar\sum_{j=1}^{\mathcal{N}}(\omega_{\text{p}} - \omega^{(j)} + i\frac{\Gamma^{(j)}}{2})|j\rangle\langle j| \\ &\quad - \frac{\hbar}{2}\sum_{k=1}^{\mathcal{N}}\left[\Omega_{0k}\left(\exp(-2i\omega_{\text{p}}t) + 1\right)|0\rangle\langle k| + \text{h.c.}\right] \\ &\quad - \frac{\hbar}{2}\sum_{j,k=1}^{\mathcal{N}}\left[\Omega_{jk}\left(\exp(-i\omega_{\text{m}}t) + \exp(i\omega_{\text{m}}t)\right)|j\rangle\langle k|\right].\end{aligned}\quad (2.19)$$

Now, since ω_{p} is several orders of magnitude larger than ω_{m} , we can make the rotating wave approximation, neglecting the rapidly oscillating terms in $\exp(\pm 2i\omega_{\text{p}}t)$, so that

$$\begin{aligned}\hat{H}_{\text{tr}}(t) &= -\hbar\sum_{j=1}^{\mathcal{N}}(\omega_{\text{p}} - \omega^{(j)} + i\frac{\Gamma^{(j)}}{2})|j\rangle\langle j| \\ &\quad - \frac{\hbar}{2}\sum_{k=1}^{\mathcal{N}}\left[\Omega_{0k}|0\rangle\langle k| + \Omega_{k0}|k\rangle\langle 0|\right] \\ &\quad - \frac{\hbar}{2}\sum_{j,k=1}^{\mathcal{N}}\left[\Omega_{jk}\left(\exp(-i\omega_{\text{m}}t) + \exp(i\omega_{\text{m}}t)\right)|j\rangle\langle k|\right].\end{aligned}\quad (2.20)$$

2.3.3 Schrödinger equation

We will now solve the Schrödinger equation for the state $|\Psi_{\text{tr}}\rangle = \hat{U}|\Psi\rangle$ of the system,

$$i\hbar\frac{d}{dt}|\Psi_{\text{tr}}\rangle = \hat{H}_{\text{tr}}|\Psi_{\text{tr}}\rangle. \quad (2.21)$$

Decomposing the wavefunction into the basis of excitonic states, $|\Psi_{\text{tr}}\rangle = \sum_{j=0}^{\mathcal{N}}c_j(t)|j\rangle$, we can write equations of motion for each of the coefficients c_j . Initially, we take the population of the valence band to be 1, and the population of all other states as 0. Under the weak probe approximation, we can assume the population of the

where

$$V = -\frac{\hbar}{2} \begin{pmatrix} \Omega_{\text{VB};\text{VB}} & \Omega_{\text{VB};5\text{S}} & \Omega_{\text{VB};5\text{P}} & \dots & \Omega_{\text{VB};20\text{F}} \\ \Omega_{5\text{S};\text{VB}} & \Omega_{5\text{S};5\text{S}} & & & \\ \Omega_{5\text{P};\text{VB}} & & \Omega_{5\text{P};5\text{P}} & & \vdots \\ \vdots & & & \ddots & \\ \Omega_{20\text{F};\text{VB}} & & \dots & & \Omega_{20\text{F};20\text{F}} \end{pmatrix}, \quad (2.27)$$

$$A_N = \hbar \begin{pmatrix} N\omega_{\text{m}} & 0 & \dots & 0 \\ 0 & \omega^{5\text{S}} - i\Gamma^{5\text{S}} + N\omega_{\text{m}} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & \omega^{20\text{F}} - i\Gamma^{20\text{F}} + N\omega_{\text{m}} \end{pmatrix}, \quad (2.28)$$

and I is the (in this case infinite-dimensional) identity matrix. The matrix F is indeed the Floquet Hamiltonian. To solve this matrix equation, we must truncate it to some finite $N \in \{-N_{\text{max}}, N_{\text{max}}\}$, chosen such that the desired observables are converged to an appropriate tolerance. In this work, $N_{\text{max}} = 15$ was considered converged because any increase in N_{max} changed the resulting absorption of the probe laser by less than 0.01 in optical depth. After setting it up correctly, the matrix equation is then straight-forwardly solved for the coefficients $c_{j;N}$ by a sparse matrix solver. The python code that was written for this is available at [82], and uses the `sparse.linalg` package from the python library Scipy.

2.3.4 Absorption

To calculate the absorption of the laser by the excitons [83], we need the susceptibility, $\chi(\omega)$. Taking the medium to be homogeneous and isotropic, near the probe frequency it is approximately given by

$$\mathbf{P}(\omega) \approx \epsilon_0 \chi(\omega) \mathcal{E}_{\text{p}} \quad (2.29)$$

where $\mathbf{P}(\omega)$ is the polarisation field near the probe frequency, and ϵ_0 is the permittivity of free space. We now make the assumption that, near ω_{p} , and after sufficient time that the system is in a steady state, the polarisation is dominated by just a few frequency terms. We therefore make the time-domain ansatz

$$\mathbf{P}(t) = \epsilon_0 \epsilon_{\text{p}} \left[P_{\pm\text{p}} \exp(\pm i\omega_{\text{p}} t) + P_{\pm\text{p}}^{\pm\text{m}} \exp(\pm i[\omega_{\text{p}} \pm \omega_{\text{m}}] t) \right. \\ \left. P_{\pm\text{p}}^{\pm 2\text{m}} \exp(\pm i[\omega_{\text{p}} \pm 2\omega_{\text{m}}] t) + \dots \right] \quad (2.30)$$

where $P_{\text{p}}^{N\text{m}}$ ($P_{-\text{p}}^{-N\text{m}}$) is the amplitude of the polarisation field at the frequency $(-)\omega_{\text{p}} + (-)N\omega_{\text{m}}$. All of these terms are near ω_{p} since $\omega_{\text{m}} \ll \omega_{\text{p}}$.

We can calculate the polarisation field near the probe frequency from the wavefunction of the exciton as

$$\mathbf{P}(t) = n_d \langle \Psi | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | \Psi \rangle \boldsymbol{\epsilon}_p, \quad (2.31)$$

where n_d stands for some effective number density of excitons. As discussed in section 2.3.8, however, since we can only extract the quantity $n_d |\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle|^2$ from measurements, in practice we do not need to calculate n_d . Expanding equation 2.31 with $|\Psi\rangle = \hat{U}^\dagger |\Psi_{\text{tr}}\rangle$ and equation 2.23, we find

$$\begin{aligned} \mathbf{P}(t) = & \sum_N \left(c_{0;N}^* \exp(-iN\omega_m t) \langle 0 | + \sum_{j>0} c_{j;N}^* \exp(i[\omega_p - N\omega_m]t) \langle j | \right) \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p \\ & \sum_{N'} \left(c_{0;N'} \exp(iN'\omega_m t) | 0 \rangle + \sum_{j'>0} c_{j';N'} \exp(-i[\omega_p - N'\omega_m]t) | j' \rangle \right) \boldsymbol{\epsilon}_p. \end{aligned} \quad (2.32)$$

Since $c_{0;N} = \delta_{N,0}$, and most of the dipole matrix elements are zero, this simplifies to

$$\mathbf{P}(t) = \boldsymbol{\epsilon}_p \sum_N \sum_{j>0} \left[c_{j;N} \exp(-i[\omega_p - N\omega_m]t) \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle + \text{c.c.} \right]. \quad (2.33)$$

If we equate terms in equal frequencies in equation 2.29, we find that, rather neatly, at the probe frequency,

$$\chi(\omega_p) = \frac{n_d}{\epsilon_0 \mathcal{E}_p} \sum_j \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle c_{j;0}, \quad (2.34)$$

where we get $c_{j;0}$ from solving equation 2.25. This allows us to calculate the absorption per unit length of the probe laser,

$$\alpha(\omega_p) = 2 \frac{\omega_p}{c} \text{Im} \left[\sqrt{1 + \chi(\omega_p)} \right], \quad (2.35)$$

where c is the speed of light in vacuum.

2.3.5 Sidebands

Now, notice that in equation 2.33, $\mathbf{P}(t)$ also oscillates at frequencies $-N\omega_m$ away from ω_p . One might suspect that the excitons will generate additional fields at these frequencies that we can detect in experiments. In fact, careful analysis of equation 2.24 will reveal that $c_{nP;N} = 0$ for odd N , so the polarisation field only oscillates at even multiples of ω_m away, just as we only observe even sidebands on the probe laser in the experiment! This led us to identify the amplitude of the polarisation field of the excitons at the sideband frequencies with that of the probe

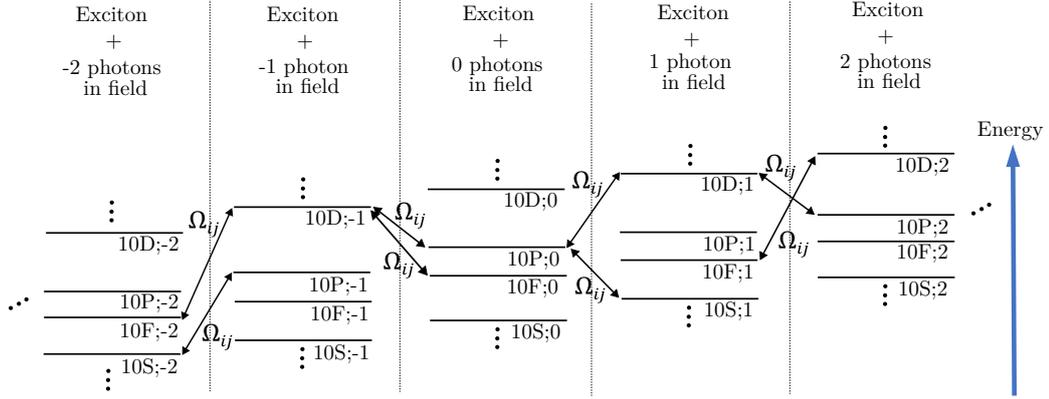


Figure 2.5: A diagram of the Floquet Hamiltonian. It describes different total states of the exciton and the microwave field. States of the exciton with N photons in the microwave field couple to the dipole-allowed states of the exciton with $N - 1$ ($N + 1$) photons in the field by absorbing (emitting) a photon from (to) the field. Negative photon numbers are allowed because the photon number is understood to be relative to some large total number [73].

sidebands. Since we only measure the intensity of the sidebands, $\mathcal{I}_{\text{SB}}(\omega_p, \omega_m, N)$, we write

$$\mathcal{I}_{\text{SB}}(\omega_p, \omega_m, N) = \eta \left| \frac{n_d}{\epsilon_0 \mathcal{E}_p} \sum_j \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle c_{j, -N} \right|^2, \quad (2.36)$$

where η is a constant of proportionality to account for the fact that we do not know the collection efficiency of the sideband light (see section 2.2). Here, N selects the order of the sideband, for example $N = 2$ would be the second-order blue sideband, or $N = -4$ would be the fourth-order red sideband. We take the coefficient of $-N$ in equations 2.36 so that $\mathcal{I}_{\text{SB}}(\omega_p, \omega_m, N)$ is the intensity of the sideband $N\omega_m$ greater than the probe frequency. We could absorb the other constants into η , however the quantity like the one in equation 2.34 is more convenient to calculate as it is ultimately independent of \mathcal{E}_p and of n_d (see equations 2.44 and 2.57). As seen in section 2.4.4, this assignment gives good agreement with the data. While 2.31 is only strictly true at the probe frequency, we measure the sidebands in the experiment to follow the polarisation of the carrier laser, lending further credibility to this assignment.

2.3.6 Floquet State Contribution

A useful consequence of casting equations 2.24 into block-matrix form is it reveals the Floquet Hamiltonian, F . The Floquet Hamiltonian defines a new, infinite-dimensional, time-independent system, related to the original time-dependent Hamilto-

nian. The A_N matrices on the diagonal, with energies $\hbar([\omega^{(j)} - i\Gamma^{(j)}/2] + N\omega_m)$, describe a ladder of field-free excitons plus N photons in the microwave field, as seen in figure 2.5. We have not formally quantised the electric field, however, it can be shown [73] that the full second-quantised theory reduces to the system described by a Floquet Hamiltonian in the limit of large photon number, where N is relative to the total number of photons. We will label these product states $|j; N\rangle = |j\rangle \exp(iN\omega_m t)$, and they form the larger, time-independent basis which describes the Floquet interpretation of the system. They are then coupled, via V , to the appropriate state $|k, N \pm 1\rangle$ on an adjacent rung of the ladder, so that

$$\langle j; N | F | k; N \pm 1 \rangle = -\frac{\hbar\Omega_{jk}}{2}, \quad (2.37)$$

where the $N+1$ (or $N-1$) state can be read as the exciton emitting (or absorbing) one microwave photon. This also gives clarity to the minus sign in equation 2.36, as the fewer photons are in the field, the more the exciton must have absorbed and so oscillates faster. The eigenstates of F , which we call Floquet states (also called dressed states), can be thought of as the underlying structure which the probe laser is coupling to. Using this concept, we will consider the contributions from the different Floquet states, which we will do by diagonalising F in equation 2.26. Since F has complex energies on the diagonal, it is non-Hermitian, so it will have different left and right eigenvectors for each quasi-energy E_q . We will denote them $|L_q\rangle$ and $|R_q\rangle$, respectively, so that $\langle L_q | F = E_q \langle L_q |$ and $F | R_q \rangle = E_q | R_q \rangle$, and such that their inner product $\langle L_q | R_{q'} \rangle = \delta_{qq'}$. Aside from at possible exceptional points (which I did not encounter in any of my calculations) we can decompose F as

$$F = \sum_q E_q |R_q\rangle \langle L_q|. \quad (2.38)$$

We can also write their completeness relation,

$$I = \sum_q |R_q\rangle \langle L_q|. \quad (2.39)$$

Substituting these into equation 2.26, and left-multiplying by $\langle L_{q'} |$, we find

$$\langle L_{q'} | c \rangle = \frac{\langle L_{q'} | d \rangle}{(\hbar\omega_p - E_{q'})}. \quad (2.40)$$

Left multiplying by $|R_{q'}\rangle$ and summing over q' , we find that we can write the coefficients of the state of the exciton, $c_{j;N}$, in terms of the Floquet states as

$$|c\rangle = \sum_q |R_q\rangle \frac{\hbar \langle L_q | d \rangle}{(\omega_p - E_q)}, \quad (2.41)$$

where we relabel $q' \rightarrow q$. If we write the components of the eigenvectors in the $|j; N\rangle$ basis as $\langle L_q | j; N \rangle = L_{j;N}^{(q)}$ and $\langle j; N | R_q \rangle = R_{j;N}^{(q)}$, equation 2.34 becomes

$$\chi(\omega_p) = \frac{n_d}{\epsilon_0 \mathcal{E}_p} \sum_j \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle \sum_q R_{j;0}^{(q)} \frac{\sum_k L_{k;0}^{(q)} \hbar \Omega_{k0} / 2}{(\hbar \omega_p - E_q)}. \quad (2.42)$$

Rearranging the sums into the most convenient order and rewriting Ω_{k0} in terms of the dipole for symmetry,

$$\chi(\omega_p) = \frac{-n_d}{2\epsilon_0} \sum_q \frac{1}{(\hbar \omega_p - E_q)} \sum_j R_{j;0}^{(q)} \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle \sum_k L_{k;0}^{(q)} \langle k | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | 0 \rangle, \quad (2.43)$$

and similarly for the sidebands,

$$\mathcal{I}_{\text{SB}}(\omega_p, \omega_m, N) = \eta \left| \frac{n_d}{2\epsilon_0} \sum_q \frac{1}{(\hbar \omega_p - E_q)} \sum_j R_{j;N}^{(q)} \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle \sum_k L_{k;0}^{(q)} \langle k | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | 0 \rangle \right|^2. \quad (2.44)$$

This form of the susceptibility gives a clear way to think about the workings of the system: the interaction of the exciton with the strong microwave field creates an underlying structure of broad Floquet states with Lorentzian line shapes which the laser couples to (remember E_q is complex); the coupling from the valence band to each state is a sum of its overlap with the excitonic states ($R_{j;0}^{(q)}$) weighted by their couplings to the valence band (which in this case is only non-zero for $|nP\rangle$); and the coupling from each state back to the valence band is a sum of its overlap with the excitonic states ($L_{j;0}^{(q)}$) weighted by the valence band coupling to them. Figure 2.12 demonstrates this way of understanding the experiment. The python code available in [82] also has the option to solve the system from this Floquet eigenstate approach.

2.3.7 Polarisation Dependence

While the model used here does not account for the effect of the crystal structure directly on the exciton, the Rabi frequencies (defined in equation 2.18) remain dependent on the angle β between the probe and microwave electric fields. Dependence on β is an important point of comparison between the model and the experiment, and in this section I outline how it was treated theoretically.

The Wannier-Mott theory of excitons is a common way to describe Rydberg excitons [84, 16]. Indeed, in this model, it was used to calculate the dipole matrix elements between pairs of excitonic states, $\langle j | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_m | k \rangle$. It results in spherically symmetric wavefunctions, $|\Psi_{nlm}\rangle_p$, which can therefore be separated into radial and angular parts, so that $|\Psi_{nlm}\rangle_p = |R_l^n\rangle |Y_m^l\rangle_p$. The subscript p indicates that

the quantum number m , the z -component of the orbital angular momentum l , is defined with respect to the polarisation direction of the probe field ϵ_p . In general, the choice of z direction is arbitrary, and we could also define the wavefunctions with respect to the microwave field direction as $|\Psi_{nlm}\rangle_m = |R_l^n\rangle |Y_m^l\rangle_m$. However, assuming that the laser excites from the valence band into only the $m = 0$ state, in experiments we measure the value of the dipole matrix elements $\langle 0 | \hat{D} \cdot \epsilon_p | \Psi_{nP0} \rangle_p$, so we are constrained to work in the probe basis. If we denote $\{x_m, y_m, z_m\}$ the Cartesian coordinates in the microwave system, where z_m is parallel to ϵ_m , then the microwave Rabi frequency between two states $|\Psi_{n_j l_j m_j}\rangle_p$ and $|\Psi_{n_k l_k m_k}\rangle_p$ becomes

$$\Omega_{jk>0} = \mathcal{E}_m \langle \Psi_{n_j l_j m_j} | z_m | \Psi_{n_k l_k m_k} \rangle_p / \hbar. \quad (2.45)$$

The coordinate inside the inner product $\langle \Psi | z | \Psi \rangle$ indicates that the coordinate is integrated over. This cannot be evaluated directly, since the wavefunctions are in a different coordinate system to the operator z_m . We express the states in the probe coordinate system in terms of the states in the microwave coordinate system as

$$|R_l^n\rangle |Y_m^l\rangle_p = |R_l^n\rangle \sum_{m'} d_{m,m'}^l(\beta) |Y_{m'}^l\rangle_m, \quad (2.46)$$

where $d_{m_j,m'}^{l_j}(\beta)$ is the Wigner small-d matrix [85]. Transforming to spherical coordinates,

$$\Omega_{jk>0} = -e\mathcal{E}_m \langle R_{l_j}^{n_j} | r | R_{l_k}^{n_k} \rangle \times \sum_{m',m''} [d_{m_j,m'}^{l_j}(\beta)]^* d_{m_k,m''}^{l_k}(\beta) \langle Y_{m'}^{l_j} | \cos \theta_m | Y_{m''}^{l_k} \rangle_m, \quad (2.47)$$

where $\{r, \phi_m, \theta_m\}$ are the spherical coordinates in the microwave system, aligned such that θ_m is parallel to ϵ_m . In essence, the states are initially aligned with the probe field, we rotate them to align with the microwave field, we evaluate the microwave dipole, and then we rotate back.

Figure 2.6 compares the β -dependence of the effect of the microwave field on the absorption spectrum, $\Delta\alpha L$ (defined in equation 2.58), between this theory (red line) and experiment (red dots). As seen in the figure, equation 2.47 greatly overestimates the dependence on β . A scheme where the valence band couples equally to all the $|nP, m = \{1, 0, -1\}\rangle$ states was tested, but gave similarly poor angle dependence, so the most likely reason is that the spherical symmetry of the Wannier-Mott wavefunction does not properly capture the symmetry of the crystal. In Jon Pritchett's thesis [74] a tensor treatment of the susceptibility is carried out, and it is shown to be compatible with no angle dependence. We use this fact to motivate the rather crude approximation

$$\Omega_{jk>0} = \begin{cases} -e\mathcal{E}_m \langle R_{l_j}^{n_j} | r | R_{l_k}^{n_k} \rangle, & \text{if } |l_j - l_k| = 1, \\ 0, & \text{otherwise,} \end{cases} \quad (2.48)$$

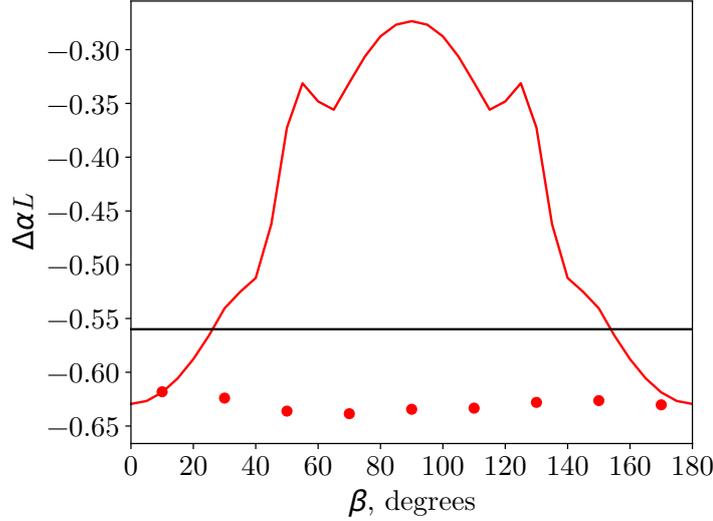


Figure 2.6: The dependence of the change in laser absorption on the relative polarisation angle between the microwave and laser fields. This data is taken at laser detuning $\hbar\omega_p = E_{11P}$ and microwave field strength $\mathcal{E}_m = 688 \pm 5 \text{ V m}^{-1}$. The red points show the experimental data, whose errorbars were too small to see. The solid red curve shows the model with Rabi frequencies defined in 2.47, and the horizontal black line shows the model with Rabi frequencies defined in 2.48.

removing the righthand side of equation 2.47 altogether, but maintaining the selection rule that $\Delta l = \pm 1$. The result is shown as the black line in figure 2.6.

2.3.8 Background absorption

In the experiment, the absorption due to the excitons sits on top of a large background due to excitations of a low-lying nS state and a phonon. Though the nS exciton is even parity, the parity of the phonon makes this process dipole allowed. The precise way to model this effect is an ongoing area of research [86, 87] and beyond the scope of this thesis, so for this project it was removed from experimental data. In [88], a function is used to fit the background absorption coefficient

$$\alpha_{\text{bg}}(E) = \begin{cases} \alpha_{\text{Ph}}(E) + \alpha_{\text{U}}(E) & E < \tilde{E}_g \\ \alpha_{\text{cont}} & E \geq \tilde{E}_g \end{cases}, \quad (2.49)$$

where \tilde{E}_g is the effective band gap energy, the point at which the exciton series smooths out into a continuum, and $E = \hbar\omega_p$ is the energy of the laser. The contribution due to the relevant $nS + \text{phonon}$ processes is given by

$$\alpha_{\text{Ph}}(E) = \alpha_{1S_y}^{\Gamma_3^-}(E) + \alpha_{1S_y}^{\Gamma_4^-}(E) + \alpha_{2S_y}^{\Gamma_3^-}(E) + \alpha_{1S_g}^{\Gamma_3^-}(E), \quad (2.50)$$

where the terms $\alpha_x^y(E)$ give the absorption from each exciton x and phonon y pair. The labels y and g are for the yellow and green exciton series, respectively. Each process only switches on when energetically allowed, so are modelled by

$$\alpha_x^y(E) = \begin{cases} A_x^y \sqrt{E - (E_x + E_y)} & E > E_x + E_y \\ 0 & E \leq E_x + E_y \end{cases}, \quad (2.51)$$

with A_x^y the amplitude of each process. The square-root dependence comes from the density of states of the phonons. The second term in the below-band gap contribution is the Urbach tail

$$\alpha_U(E) = c_U \exp\{(E - E_g)/E_U\}, \quad (2.52)$$

a phenomenological term that gives an exponential smearing of the continuum into the band gap, with amplitude c_U and width E_U . This smearing has been shown to originate in charged impurities in the crystal, the small, inhomogeneous electric fields of which ionise the high n excitons [89]. The continuum term, $\alpha_{\text{cont}}(E)$, is described functionally in [88], however in this project the experimental data did not extend far above the continuum, so α_{cont} was treated as constant.

In the microwave-field free case, the Rydberg exciton part of the absorption, α_{Ex} , was modelled with a sum of asymmetric Lorentzians

$$\alpha_{\text{Ex}}(E) = \frac{1}{L} \sum_{n\text{P}} A_{n\text{P}} \frac{\gamma_{n\text{P}} + q_{n\text{P}}(E - E_{n\text{P}})}{\gamma_{n\text{P}}^2 + (E - E_{n\text{P}})^2}. \quad (2.53)$$

Here, $A_{n\text{P}}$ is the amplitude of the Lorentzian, $\gamma_{n\text{P}} = \hbar\Gamma^{(n\text{P})}/2$ is the half width half maximum in energy, $q_{n\text{P}}$ is the asymmetry, $E_{n\text{P}} = \hbar\omega^{(n\text{P})}$ is the energy of the exciton resonance, and L is the depth of the sample. The asymmetric lineshape of the absorption peaks comes from the non-radiative decay of the excitons via phonon emission. Taking the phonon energies and amplitudes from [88], the rest of the parameters were fit all at once to the single photon excitation data. The full absorption spectrum, with the background function obtained from this fit shaded in red, is shown in figure 2.7.

2.3.9 Model parameters

The $n = 6$ to $n = 15$ P exciton energies and widths were taken directly from a fit to the data in figure 2.7. The S, D, and F energies and widths from $n = 6$ to $n = 11$ were fitted from 2-photon absorption spectroscopy [40]. The energies were extended to $n \in \{5, \dots, 20\}$ by fitting them to Rydberg series, $E_g - R_y/(n - \delta_l)^2$, with a shared Rydberg constant R_y and an l -dependent defect δ_l . The P widths

were extended by fitting them to a $1/n^3$ scaling law plus a constant, c_p . Because there is additional broadening in the 2-photon spectrum due to the SHG pump laser, the S, D, and F widths were fitted to a $1/n^3$ scaling law but with the same constant c_p . Details of these fits are in [40].

The dipole moment per unit volume, $n_d \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle$, can also be obtained from fitting. From [90] we know that, absent of the microwave field, the linear susceptibility due each nP exciton can be modelled as

$$\chi_{nP}^{(1)} = \frac{1}{2\epsilon_0} \frac{|n_d \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle|^2}{E - E_{nP} - i\gamma^{(nP)}}. \quad (2.54)$$

Assuming $\chi_{nP}^{(1)} \ll 1$, a reasonable assumption given the small effect of the excitons on the absorption, we find from 2.35 that

$$\alpha_{nP} = \frac{\omega_p}{2\epsilon_0 c} \frac{|n_d \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle|^2 \gamma_{nP}}{(E - E_{nP})^2 + \gamma_{nP}^2} \quad (2.55)$$

is the contribution to the absorption from each P state, or, with the laser on resonance with the P state,

$$\alpha_{nP} = \frac{\omega_p}{2\epsilon_0 c} \frac{|n_d \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle|^2}{\gamma_{nP}}. \quad (2.56)$$

If we compare this to the on-resonance Lorentzian contribution from each P state in equation 2.53 then we find that the square modulus of the dipole matrix elements per unit volume are given by

$$|n_d \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle|^2 = \frac{2A_{nP}\epsilon_0 c}{\omega_p L}. \quad (2.57)$$

Helpfully, the fact that we calculate these quantities per unit volume means we do not need to estimate a value for n_d in equations 2.43 or 2.41. All the parameters used in the model are available to download at [82].

Unfortunately, equation 2.57 does not give us any information about the phase of the matrix elements. This problem has arisen before, in fitting the 2-photon absorption data to a sum of complex poles [40]. To explore the significance of these phases, sets of phases were generated from a uniform random distribution, and the variation in the results of the model are discussed in section 2.4.3. In short, because they had little meaningful effect on the absorption spectra (fig. 2.14), we make the arbitrary decision that all such matrix elements are positive real wherever the absorption spectra are discussed in this work.

2.3.10 Microwave field calibration

All but one of the parameters that enter the model are measured in experiment or derived from some more fundamental theory. Unfortunately, there is no way

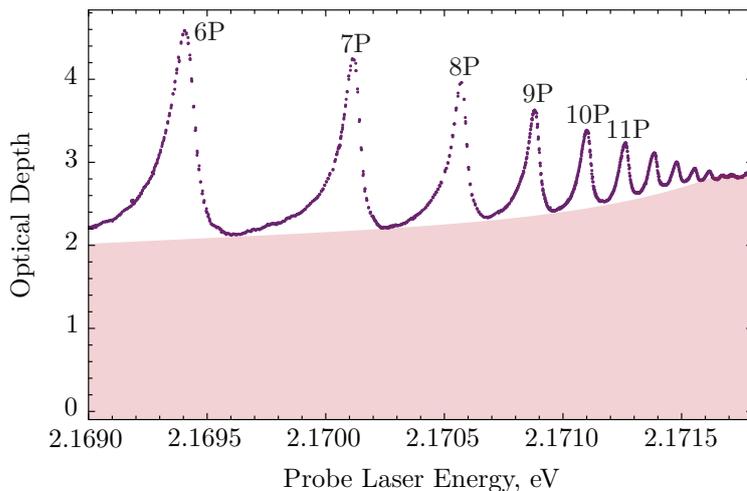


Figure 2.7: Experimental exciton absorption spectrum, with the fitted phonon background function shaded in red. For comparison with theory, this background is subtracted from the total absorption.

to measure the amplitude of the microwave field at the medium without affecting the experiment. Finite element electromagnetic simulations of the antenna found the field strength was very sensitive to the presence of metallic components inside the cryostat [54, 91]. It was also found that the transmission and reflection of the microwave signal through the antenna was poorly correlated with the field strength, so it could not be measured that way either. Therefore, the microwave field strength in the sample was calibrated by fitting the model for small values of \mathcal{E}_m , where the model reduces to the perturbative model in [25]. By fitting only to small values of input power, we can test the model at strong fields without any free parameters. This was done by fitting across the energy spectrum the change in absorption

$$\Delta\alpha L = \alpha(\mathcal{E}_m)L - \alpha(\mathcal{E}_m = 0)L, \quad (2.58)$$

where L is the thickness of the cuprite sample. In total, eight $\Delta\alpha L$ data sets were used to fit the model. One of these fits is seen in figure 2.10a. The best fit values of the microwave field strength are plotted against the square root of the input power to the microwave antenna in figure 2.8. The limits of the error bars displayed in fig. 2.8 represent the value of the electric field for which the chi-squared value of the fit doubled over the chi-squared value at the best fit electric field strength. The best fit line through figure 2.8 gives a calibration factor of $43.5 \pm 0.3 \text{ V m}^{-1}(\text{mW})^{-1/2}$, within the error bar of the value calculated in [25], and in reasonable agreement with the simulation in [54]. The slight difference from [25] and smaller error bar comes from using more points in the fit. It was measured that the output power from the antenna was linear with input power, so this calibration was used for all

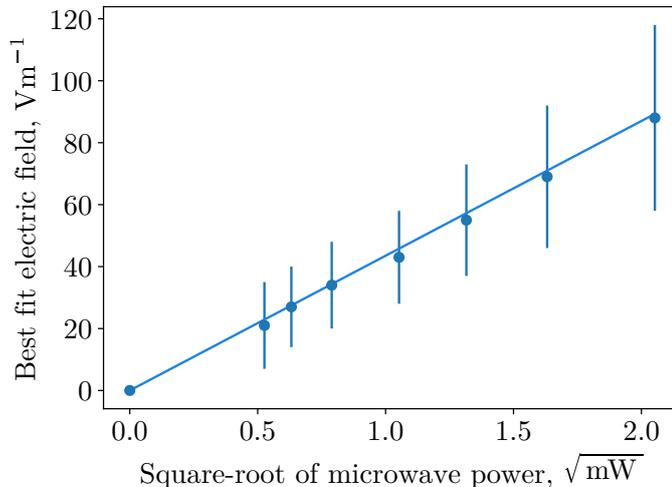


Figure 2.8: The best-fitting microwave electric field strength for $\Delta\alpha L$ (eq. 2.58) at small input microwave powers is plotted against the square root of the power. The errorbars on the best fit field are where the chi-squared for that fit doubles (the errorbars on the input power are too small to see). The linear fit allows us to extract the calibration factor of $43.5 \pm 0.3 \text{ V m}^{-1}(\text{mW})^{-1/2}$ between the power input to the antenna and the peak field strength at the sample.

input powers.

2.4 Results

2.4.1 Absorption Spectra

The absorption of the probe laser by the Rydberg excitons for varying microwave field strength is compared to predictions of equation 2.35 in figure 2.9. The absorption spectra are offset in optical depth by $\mathcal{E}_m/100$ to show six different microwave field strengths, as indicated on the righthand y-axis. As the microwave field strength increases, absorption at the peaks reduces and absorption between peaks increases, turning the high n exciton peaks into a broad absorption continuum. Qualitatively, this effect is reproduced well in the theory curves. Across the range of field strengths, the same peaks disappear in both the experimental and theoretical spectra, with the continuum in both growing to include all peaks above $9P$ at $\mathcal{E}_p = 688 \pm 5 \text{ V m}^{-1}$, the maximum field strength applied in the experiment. Unlike similar experiments on atoms [92], we do not observe obvious Stark shifts or Autler-Townes splittings of the peaks. In those systems, the states are narrow compared to their energy spacing, so there is a significant perturbative

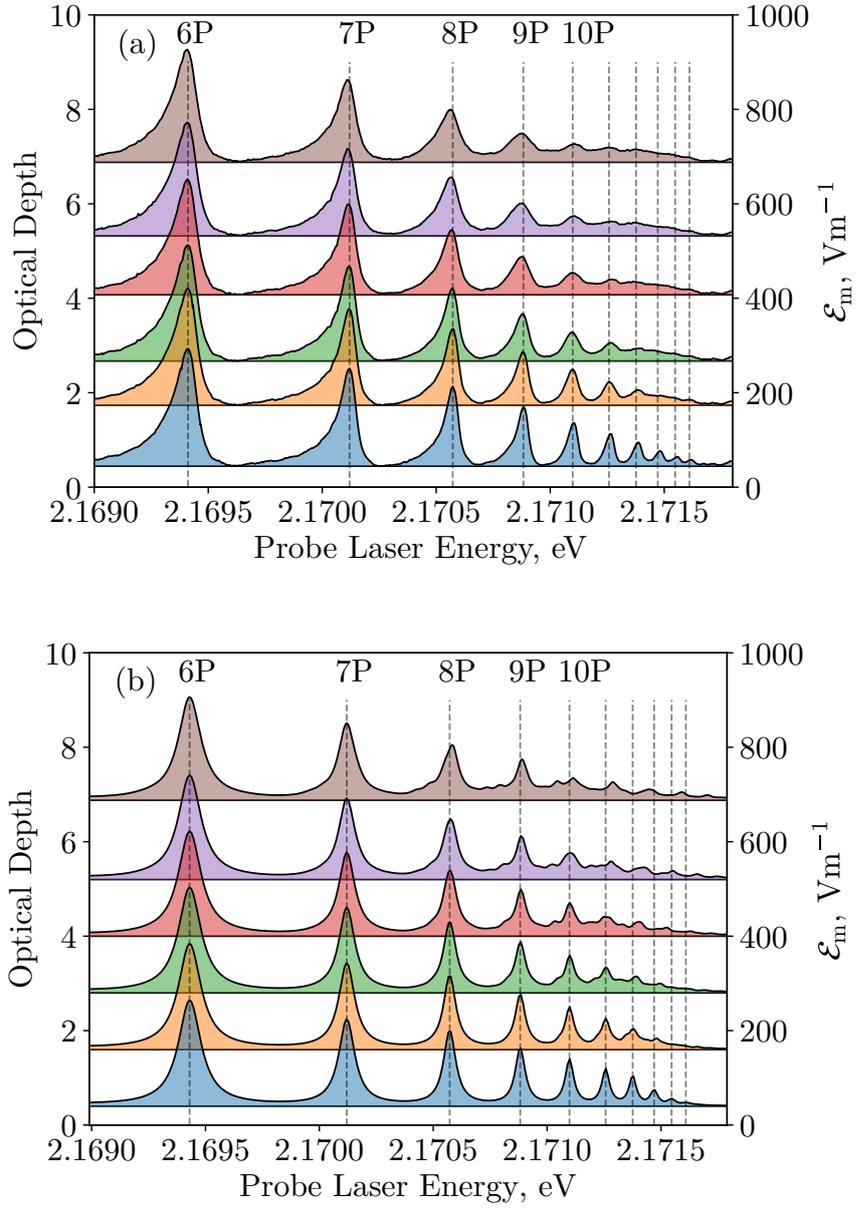


Figure 2.9: Plots of the (a) experimental and (b) theoretical absorption spectra at different microwave field strengths, \mathcal{E}_m . Each spectrum has been offset in optical depth by $\mathcal{E}_m/100$ to indicate the microwave field strength it was taken at, shown on the righthand y-axis. The dashed vertical lines show the positions of the zero-field nP resonances.

regime in which a state is coupled to exactly one other, and where the AT splitting (i.e. the Rabi frequency) is bigger than the linewidth. In our system, with large linewidths comparable to the energy separation between states, the (ultra)strong driving threshold is crossed even at low field strengths, where many states are coupled together at once, resulting in more complicated behaviour.

To compare quantitatively the experimental and theoretical absorption spectra, figure 2.10 shows the effect on the absorption by switching on the microwave field, defined in equation 2.58. The subfigures (a), (b) and (c) show, respectively, a small value of \mathcal{E}_m where $\Omega \ll \Gamma$, the threshold of strong driving ($\Omega \sim \Gamma$), and the maximum field strength achieved in the experiment ($\Omega \gg \Gamma$). In the low field regime, where the model reduces to the perturbative model in [54, 25], the agreement is excellent. While this is a useful result, it is somewhat expected, since this is where the microwave field strength was calibrated. Remarkably, the model maintains near quantitative agreement across all applied microwave fields. The agreement is best at the nP peaks, and reproduces well the expanding envelope of $\Delta\alpha L$. It should be stressed here that once the microwave field strength is calibrated against the data at low powers ($\mathcal{E}_m < 80 \text{ V m}^{-1}$), there are no free parameters in the model. For comparison, the perturbative model in [25] predicts $\Delta\alpha L = -6$ for the maximum field strength at $\hbar\omega_p = E_{10P}$, much larger than the actual size of the absorption peak, indicating we are well beyond the perturbative regime.

There are some discrepancies between the model and the experiment. Asymmetry can be seen in the nP peaks in the experimental data even with no microwave field applied. This asymmetry is well-studied and known to be a result of the non-radiative decay via phonons [86, 87]. However, modelling this effect is complicated and not easily compatible with this model. There are ways to introduce asymmetry into the peaks parametrically, for example adding an additional imaginary part to the matrix elements appearing in equation 2.41, however it is not obvious that this will have the desired effect on the physics of the system once the microwave field is switched on. The lack of asymmetry can explain some of the weaker agreement between nP peaks. However, there are still spurious features in the theoretical spectra not present in experiment, for example the two peaks that emerge between 8P and 9P, or the splitting of 10P into two peaks. These discrepancies are likely caused by inaccurate state structure and Rabi couplings in the model. Indeed, it is possible that the unusual smoothness of the experimental spectra when compared to the theoretical spectra is caused by additional broadenings of the S, D, and F states that are not accounted for in this model, as discussed in section 2.5.1. Nevertheless, the level of agreement gives strong evidence that the model is correctly capturing the essential physics of the system.

2.4.2 Floquet states

The Floquet state picture of the model provides powerful ways to understand the results seen in the experiment. Figure 2.11 shows a Stark map of the real part of the eigenenergies of the Floquet states, $\text{Re}[E_q]$, and how they change with microwave

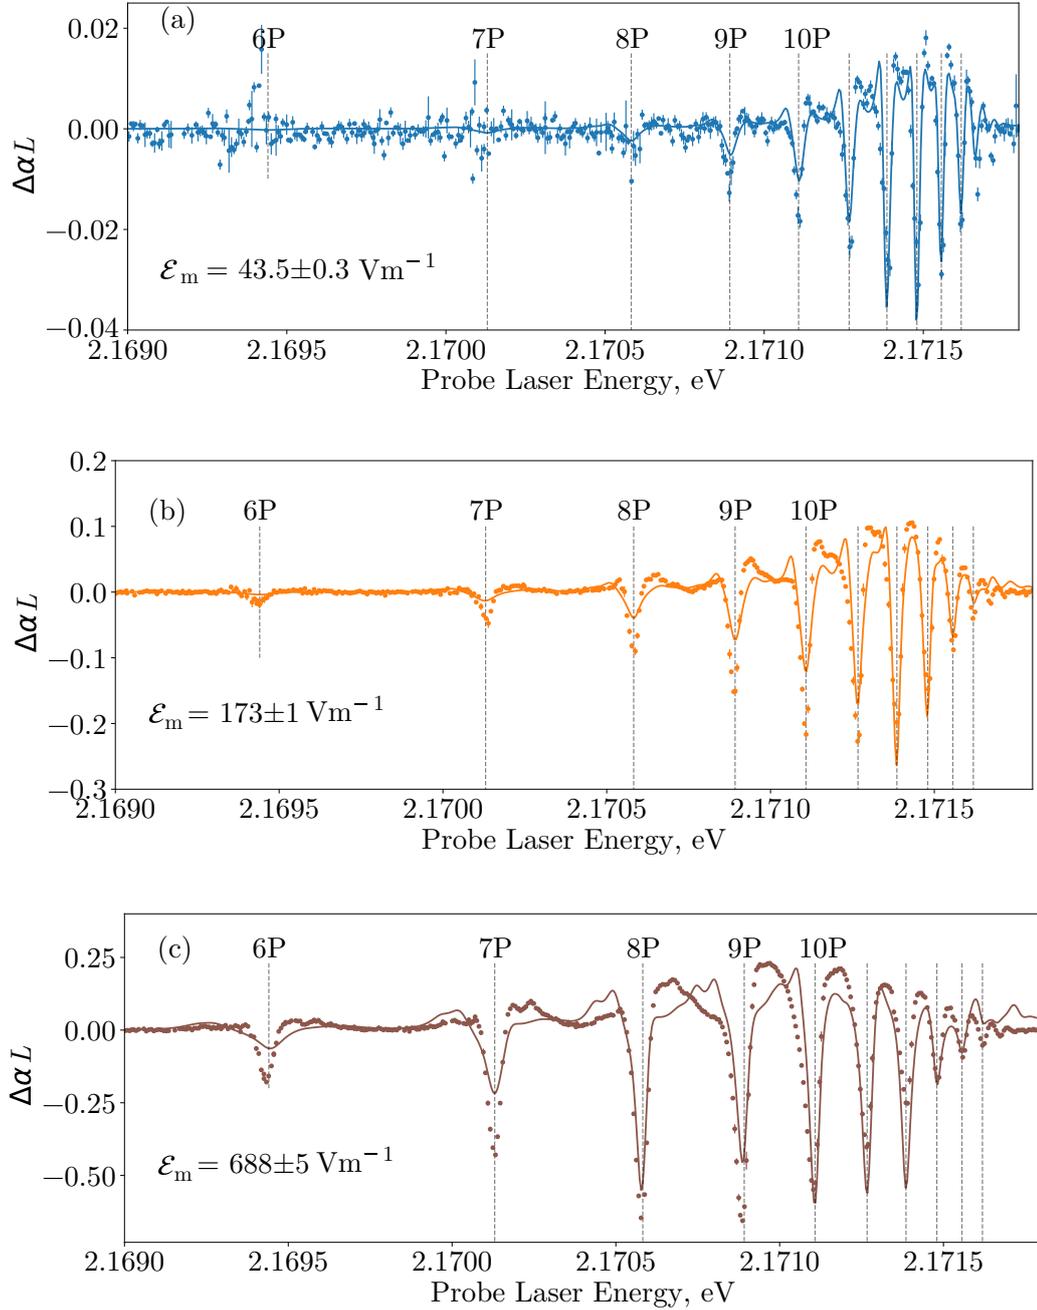


Figure 2.10: Comparison between the experimental (filled circles) and theoretical (solid lines) values of the change in absorption $\Delta\alpha L$ as a function of laser energy. Three different values of microwave field strength are shown spanning the perturbative (a), strong (b) and deep strong (c) driving regimes. The dashed vertical lines show the positions of the zero-field nP resonances.

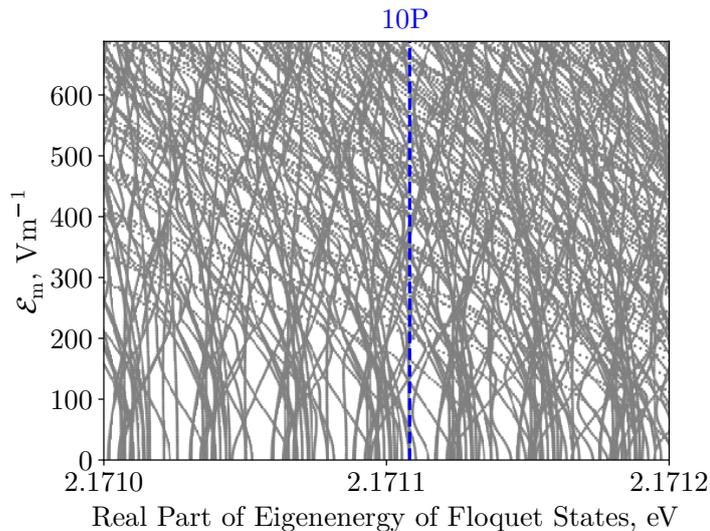


Figure 2.11: The real part of the eigenenergies of all the Floquet states in the region near $\hbar\omega_p = E_{10P}$. The density of states is very high due to the inclusion of $N \in \{-15, \dots, 15\}$. The periodicity that can be seen at $\mathcal{E}_m = 0 \text{ V m}^{-1}$ comes from the excitonic structure repeated at $\pm N\hbar\omega_m$. Above perturbative field strengths ($\mathcal{E}_m < 50 \text{ V m}^{-1}$) there is significant state mixing leading to large energy shifts and many state crossings. Only some of these states couple to the valence band and so contribute to probe laser absorption, as shown in figure 2.12.

field strength in the region near $\hbar\omega_p = E_{10P}$. At $\mathcal{E}_m = 0$, we reproduce the zero-field exciton, with copies of it periodically at $N\hbar\omega_m$ either side for each $N \in \{-15, \dots, 15\}$. As the field strength increases, the states mix together and no longer correspond to pure excitonic states. The perturbative region, where the change in state energies is mostly linear, is confined to $\mathcal{E}_m \lesssim 100 \text{ V m}^{-1}$ on the low energy side of the map, and $\mathcal{E}_m \lesssim 50 \text{ V m}^{-1}$ on the high energy side. This demonstrates how strong and ultrastrong driving, characterised here by non-linear response of the state energies and many state crossings, occurs at much lower field strengths for higher energies, where states are closer together and coupling strengths between them are greater.

Not all of the Floquet states couple to the valence band, however, so it is difficult to tell from figure 2.11 what the absorption spectrum would look like. From equation 2.43, we can see that only states with a non-zero $|nP\rangle$ component can contribute to the absorption. To visualise this, in figure 2.12 is plotted in grey the real part of the Floquet eigenenergies, $\text{Re}[E_q]$, against \mathcal{E}_m , with their opacity given by

$$a_q = \sum_{n,n'} R_{nP;0}^{(q)} [L_{n'P;0}^{(q)}]^*. \quad (2.59)$$

The choice here not to multiply a_q by the dipole matrix elements, as would be

natural looking at equation 2.43, is so that the high energy states are still visible, otherwise their small couplings to the valence band compared to the low energy excitons would make them impossible to see. The orange and green curves show where the nP to nD transitions enter the strong ($\Omega_{nP,nD} > \Gamma_{nP}$) and deep strong ($\Omega_{nP,nD} > \Gamma_{nP}, |E_{nP} - E_{nD}|$) driving regimes for each n . The theoretical absorption spectrum at $\mathcal{E}_m = 688 \text{ V m}^{-1}$ is shown above the Stark map. This plot gives immediate clarity to what is happening in the experiment. For low energy excitons (7P and below) the coupling to neighbouring states is small and the energy gap to them is large, so they do not mix significantly. One can see $|7S; +1\rangle$ and $|7F; 0\rangle$, the two closest visible states to $|7P; 0\rangle$, are gaining some 7P character at the highest field strengths, though not enough to produce peaks in the absorption spectrum. The 8P exciton is at the threshold of strong driving at the highest field strengths. It is isolated enough from nearby states that it is beginning to undergo Autler-Townes splitting, something that does not happen to any other field-free exciton states. Also, its mixing with nearby states leads to a reduction in its 8P character, which decreases the height of its absorption peak. The two states on the low-energy side of 8P are labelled $|8S; -1\rangle$ and $|8S; +1\rangle$ at zero field, and they couple strongly enough to $|8P; 0\rangle$ at the maximum field strength to produce peaks in the absorption spectrum. The 9P exciton enters the strong driving regime and is on the threshold of deep strong driving at the maximum field strength. It is exchanging several microwave photons before decaying, evidenced by the fact that we can see many nearby states that have significant 9P character. However, there is still a state that is mostly $|9P; 0\rangle$, giving rise to a large peak at $\hbar\omega_p = E_{9P}$. This is not true of the rest of the $|nP; 0\rangle$ states. Even at moderate field strengths, the coupling between them and states near and far is so strong that there are no states left with a dominant contribution from a single $|nP\rangle$ state. The resulting forest of states that all couple to the valence band is what gives rise to the absorption continuum observed in the experiment.

To give an example of how many states contribute to absorption at a given probe laser energy, the quantity

$$b_q = \text{Im} \left[\frac{1}{(\hbar\omega_p - E_q)} \sum_j R_{j;0}^{(q)} \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | j \rangle \sum_k [L_{k;0}^{(q)}]^* \langle k | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | 0 \rangle \right] \quad (2.60)$$

is shaded in red for laser energy $\hbar\omega_p = 2.171 \text{ eV}$, as indicated by the solid vertical line. Notice how many states contribute significantly to the absorption. This is due to the large energy width most of the states (which is not plotted, for readability) meaning they can interact with the laser even when far off resonance. Of course, the states that are closest to the laser in energy contribute the strongest, but because the contribution is mediated by the coupling to the valence band, even distant

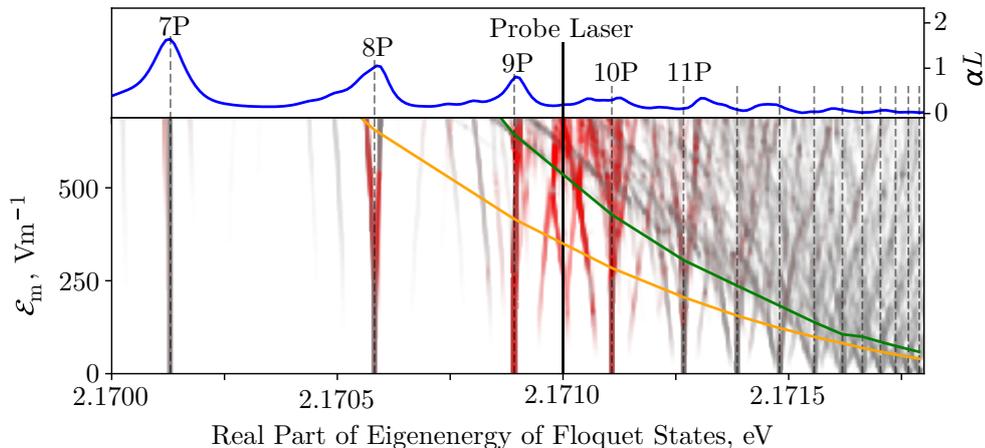


Figure 2.12: The real part of the eigenenergies of the Floquet states are plotted to show how the underlying structure of the exciton changes with microwave field strength. The Floquet states' opacity is given by their overlap with the zero-field $|nP\rangle$ states (eq. 2.59) to demonstrate how much they contribute to the laser absorption. The absorption spectrum at $\mathcal{E}_m = 688 \text{ V m}^{-1}$ is shown in the top plot for reference. The Floquet states also have a width to them, which is not shown. To give an idea of the importance of the width, the states which contribute to the absorption at the indicated laser energy ($\hbar\omega_p = 2.171 \text{ eV}$) are shaded in red, with deeper reds showing greater contribution. Above the orange curve, the nP to nD transitions enter the strong driving ($\Omega_{nP,nD} > \Gamma_{nP}$), and above the green curve the same transitions enter the deep strong driving regime ($\Omega_{nP,nD} > \Gamma_{nP}, |E_{nP} - E_{nD}|$).

states near 8P and 11P contribute. While the change in contribution is due mostly to states changing in coupling to the valence band, they are also changing in width, so can couple to the laser when further off-resonant.

Although we cannot treat this system as a two-level one, it is still useful to calculate the deep strong driving regime condition for some of the transitions to compare it to other strong coupling and driving systems. For example, we find that for the 11P to 11D transition we enter the deep strong driving regime for field strengths above $300 \pm 2 \text{ V m}^{-1}$. The point where all pairs of nP to nD transitions cross the strong and deep strong driving threshold is shown by the green curve in figure 2.12. At the highest fields applied, we can reach a deep strong driving parameter of up to $U = \Omega_{13P,13D}/\nu_{13P-13D} \sim 4$ for the 13P to 13D transition, significantly exceeding parameters of $U = 1$ [53] and $U = 2.1$ [66] in other strong driving works. The strong driving threshold, where the Rabi frequency is larger than the state linewidth, shown in orange in figure 2.12, is reached at a similar microwave field strength. This is because the state linewidth (FWHM of $40 \text{ } \mu\text{eV}$ for 11P) is comparable to the energy separation ($E_{11D} - E_{11P} = 58 \text{ } \mu\text{eV}$).

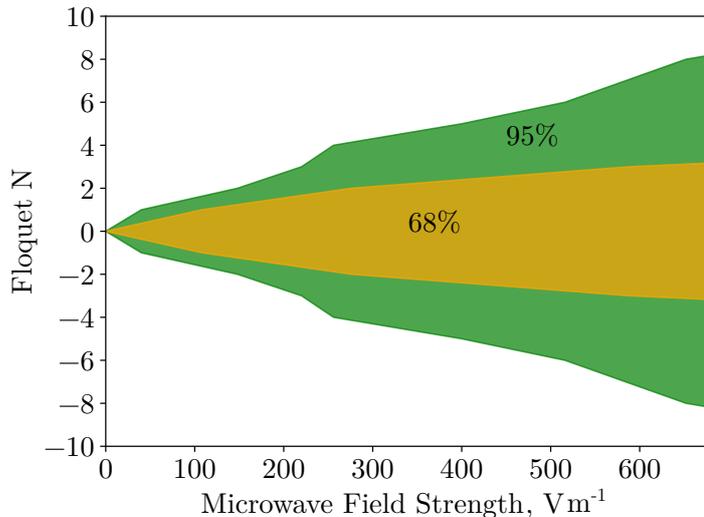


Figure 2.13: A heuristic for how much of the wavefunction oscillates at less than frequency $|N|\omega_m$, when $\hbar\omega_p = E_{12P}$. The yellow region shows where $X(N, \mathcal{E}_m) < 0.68$ and the green region is where $X(N, \mathcal{E}_m) < 0.95$, as defined in equation 2.61.

It is not clear from figure 2.12 how significant high- N states are in the wavefunction for a given laser energy, or in other words, how many microwave photons may be exchanged in the system. To investigate the extent of multiphoton transitions, in figure 2.13 I plot for $\hbar\omega_p = E_{12P}$ contours of the quantity

$$X(N, \mathcal{E}_m) = \frac{\sum_{-N < N' < N} \sum_{j > 0} |c_{j;N'}|^2}{\sum_{N'} \sum_{j > 0} |c_{j;N'}|^2}, \quad (2.61)$$

which measures how much of the wavefunction (excluding the valence band) oscillates at a frequency less than $|N|\omega_m$. This is a proxy for microwave photon transitions, since to populate a $|j, N\rangle$ state, at least N transitions must occur. We can see that even at the highest field strengths, most of the wavefunction consists of states $|j, N < 3\rangle$. However, a significant amount of the wavefunction is made up of states all the way out to $N = 8$, indicating at least 8 microwave photons are exchanged before the exciton decays, at the highest field strengths. The system must therefore be well beyond the validity of the rotating wave approximation, justifying our Floquet state-based approach.

With this many transitions involved, we may question whether ionisation could become significant, though it is not included in the model. Ionisation in excitons is analogous to ionisation in atoms since the conduction band acts like the continuum. Experiments on dissociating Rydberg excitons with static fields have shown the required ionisation field strength to follow an n^{-4} trend [93, 94, 95]. Applying these results to this system, the ionisation field at $n = 11$ would be 6000 V m^{-1} , an order

of magnitude larger than the largest field applied in our experiment. However, our field is not static, and so it is possible to ionise the exciton by absorbing multiple photons in a row [96]. Experiments on ionisation of hydrogen atoms in microwave fields reveal an n^{-5} scaling when the microwave frequency is less than the level spacing, as it is here, resulting in ionisation at much lower field strengths [97]. Comparing the binding energy (in frequency units) of the 12P exciton, 117 GHz, to the frequency of 8 microwave photons, 54 GHz, it is unlikely that the 12P exciton would be ionised by the field. At $\hbar\omega_p = E_{15P}$, the 95% contour of $X(N, \mathcal{E}_m)$ reaches out to $N = 17$ at $\mathcal{E}_m = 688 \text{ V m}^{-1}$. The binding energy of the 15P exciton is only 93 GHz, less than $17\omega_m = 119 \text{ GHz}$, implying that ionisation could have an effect on the very high energy part of the exciton series. Unfortunately, extending the model to include ionisation to a continuum via an arbitrary number of microwave photons is significantly non-trivial [98], and was beyond the scope of this thesis.

2.4.3 Phases of the dipole matrix elements

As discussed at the end of section 2.3.8, we do not have access to the phase of the dipole matrix elements $\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle$ between the valence band and the nP states, seen in equation 2.57. With 16 different nP states in the model ($n = 5$ to $n = 20$), the parameter space is too large to systematically explore. Therefore, to investigate how important a role these play in the laser absorption spectrum and the intensity of the predicted sidebands, 100 sets of matrix elements $\{\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle \exp(i\theta_n)\}$ were generated, with each θ_n chosen from a uniform random distribution on $[0, 2\pi)$.

Figure 2.14 shows the theoretical probe absorption spectra for two such sets of phases and how they vary with microwave field strength. These two sets were chosen here as they produced the two most distinct spectra of the phases tested. As one can see, little meaningful difference is made to the absorption spectrum by the choice of phases. The broad absorption continuum still forms at the same energies and at the same field strengths. The lack of response at the 6P and 7P energies is unaffected. Almost all of the new peaks that appear at high field strengths appear at the same energy with the same amplitude, like those between 7P & 8P and those between 8P & 9P. The biggest difference is near the peak corresponding to 9P, where the peak has moved to higher energy in the top plot, though this difference is still small. Given the relative lack of dependency of the absorption spectrum on the phases of the matrix elements, they were (arbitrarily) all chosen to be positive real in sections 2.3.10, 2.4.1 and 2.4.2. However, the choice of phase had a very strong effect on the predicted intensity of the sidebands.

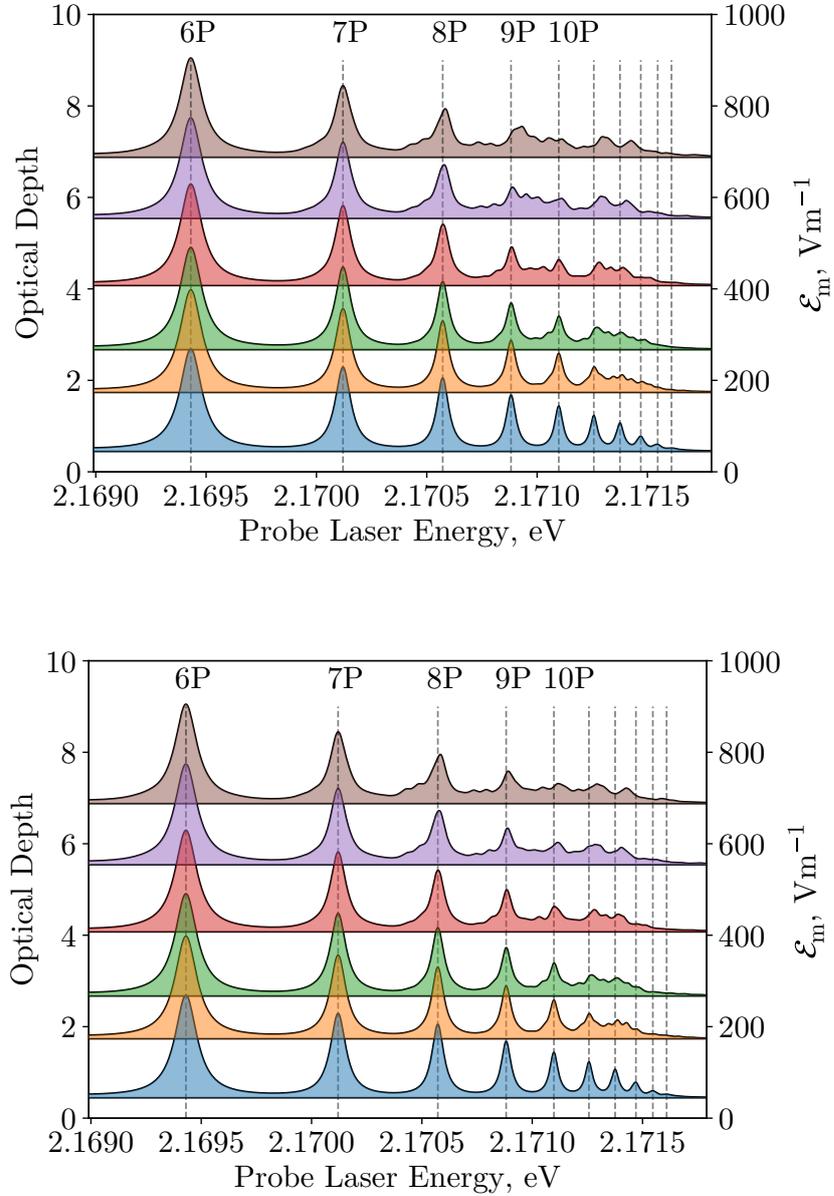


Figure 2.14: Plots of the theoretical absorption spectra at different microwave field strengths, \mathcal{E}_m , for two different sets of randomly chosen phases of $\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle$. These two choices of phases gave two most different spectra, indicating the absorption spectra is not significantly dependent on phase. Each spectrum has been offset in optical depth by $\mathcal{E}_m/100$ to indicate the microwave field strength it was taken at, shown on the righthand y-axis. The dashed vertical lines show the positions of the zero-field nP resonances.

2.4.4 Sidebands

In figure 2.15, the intensity of red and blue sidebands in the experiment at a selection of laser detunings are compared to the predictions of the model for each of the different sets of phases (see section 2.4.3). Almost all sets qualitatively reproduce the shape and quantitatively reproduce the onset of the sideband rollover over an order of magnitude in field strength and nearly three orders in intensity, for both the second- and fourth-order sidebands. It is also a significant improvement over the perturbative model used in [25], which could not predict rollover, nor fourth order sidebands.

The intensity is largely phase-independent near lower energy excitons, since they have relatively small couplings to neighbouring states and the energy gap to them is much larger than the microwave energy. This means that an exciton created through, say, the 8P state is overwhelmingly likely to decay through the 8P state, described by the $|\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | 8P \rangle|^2$ term in equation 2.44, which has no dependence on phase. Conversely, the intensity is very phase-dependent for higher energy laser detunings. As the exciton states get closer in energy it is more likely that an exciton created in one nP state decays through another, introducing cross-terms $\langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | nP \rangle \langle 0 | \hat{\mathbf{D}} \cdot \boldsymbol{\epsilon}_p | n'P \rangle \exp[i(\theta_n - \theta_{n'})]$ to equation 2.44, which are dependent on phase.

Looking at figure 2.15, one can see for any given energy some of the lines quantitatively reproduce the experimental data points, providing good evidence the model is capturing the essential physics of the system. It is therefore important to check if there is a common choice of phases, corresponding perhaps to the phases in reality, amongst the random sampling that quantitatively fit to many sidebands at once. Figure 2.16 shows an example of one choice of phases, namely, the set of phases that best-fit for fig. 2.16a, the second-order blue sidebands just below $\omega_p = E_{12P}$. These phases also fit well to the sideband data just below $\omega_p = E_{11P}$ (fig. 2.16b) for both the second- and fourth-order sidebands. However, the fit is poor for the corresponding red sidebands. Recall from equation 2.36 that η is a free parameter as well, and its value accounts for a rigid vertical shift in fig. 2.16. Unfortunately, looking at figure 2.15, we can see that for the required factor of 10 reduction in η for the red sidebands to fit as well as the blue sidebands, there are no curves that would fit the blue sidebands near $\omega_p = E_{12P}$. This is the same of all the choices of phase that were studied: no phases fit more than a few different sideband data sets at once. It is likely that due to the energy dependent absorption background, the phases of the matrix elements are energy dependent too, though in precisely what way is not obvious. Analysis of the phases in that way was considered beyond the

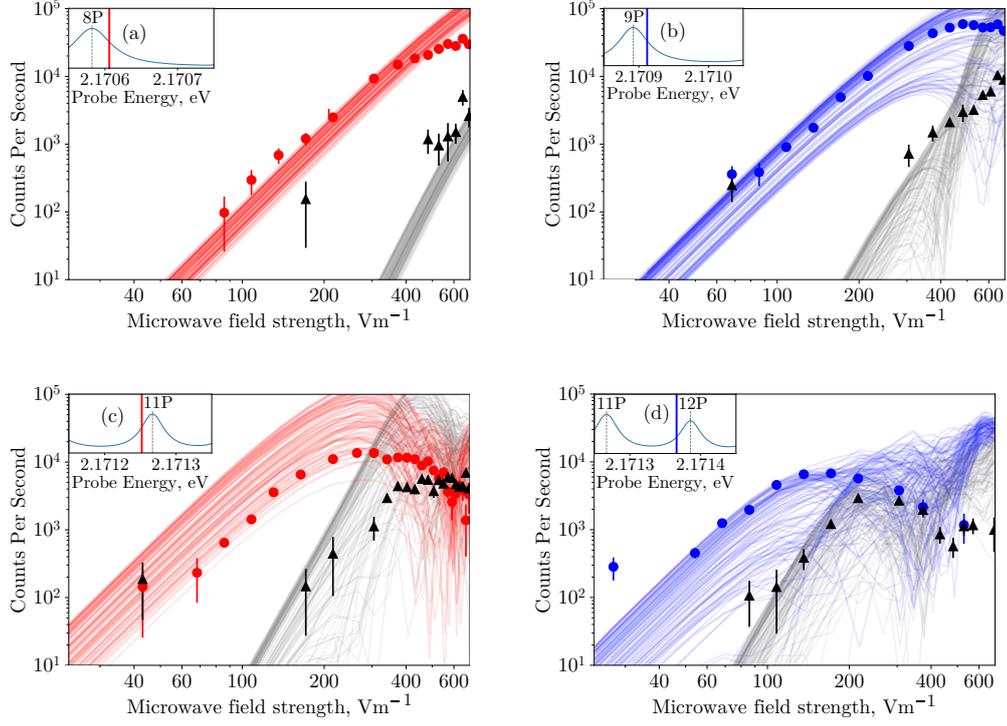


Figure 2.15: The intensity, expressed in counts per second, of second (colour) and fourth (black) order red- and blue-detuned sidebands for theory (lines) and experiment (points) at four different probe laser detunings. The vertical lines in the insets show where in the exciton spectra each set of sidebands was taken: (a) 2.17061 eV, (b) 2.17091 eV, (c) 2.17125 eV, (d) 2.17137 eV. Each theory curve is obtained with a different, randomly-chosen, set of phases for the valence band to nP dipole matrix elements, $\langle 0 | \hat{D} \cdot \epsilon_p | nP \rangle$, since these phases could not be determined from experiment.

scope of this thesis.

As mentioned in section 2.3.5, there is also an unknown prefactor η to the sideband intensity due to a lack of knowledge of the collection efficiency of the light in the experiment. Given that the data near 8P is largely phase independent, due to its relative isolation from other states, $\eta = 7.0 \times 10^{12} \text{ W m}^{-2}$ was obtained by fitting the average theory curve there to the data.

2.5 Extensions to the model

The model developed in this work does not include the effects of inter-exciton interactions. While the probe laser in the experiment was kept to an intensity below where interactions would be relevant, there is interesting work in the field

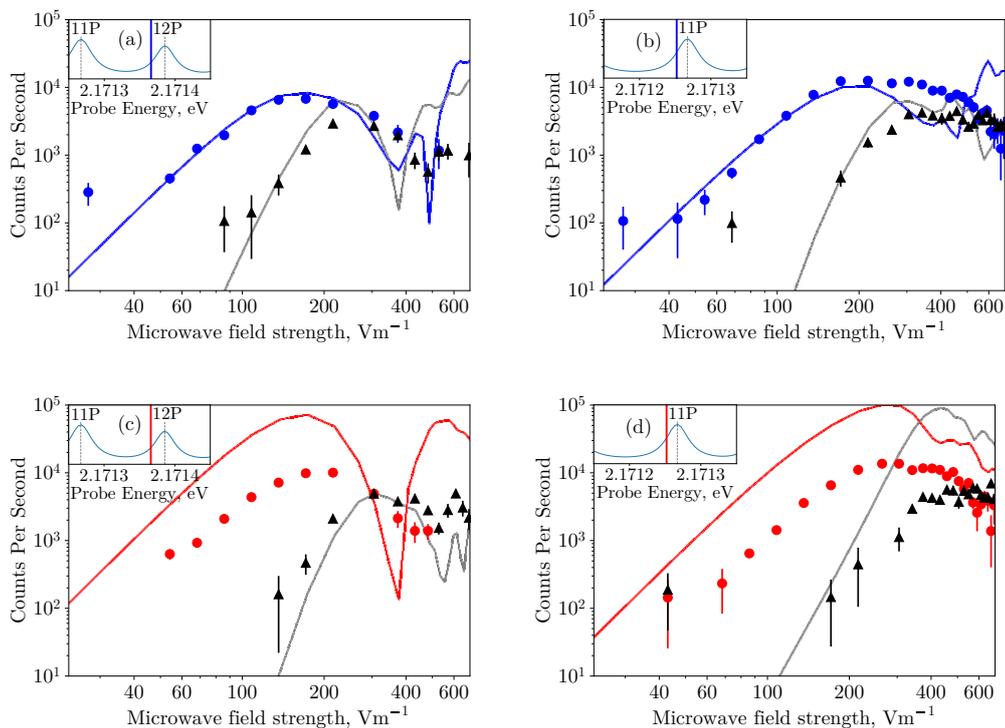


Figure 2.16: The intensity, expressed in counts per second, of second (colour) and fourth (black) order red- and blue-detuned sidebands for theory (lines) and experiment (points) at four different probe laser detunings. The vertical lines in the insets show where in the exciton spectra each set of sidebands was taken: (a) 2.17061 eV, (b) 2.17091 eV, (c) 2.17125 eV, (d) 2.17137 eV. The theory curves are obtained for the choice of phases for the valence band to nP dipole matrix elements, $\langle 0 | \hat{D} \cdot \epsilon_p | j \rangle$, that minimise the residuals in (a). This is an example of how I found no choice of phases that fit many sideband data sets at once.

studying dependence on Rydberg exciton density [99]. Since the van der Waals interaction energy between two excitons is in the microwave regime as well [16], this kind of experiment could be combined with a strong microwave field to explore tunable interactions between two Floquet systems, which could be understood by extending this model.

We do not include the possibility of ionisation in the model either. Extending the model to include ionisation in a complete and self-consistent way is unfortunately very complicated [98], as it requires a basis that can express both the discrete states of the exciton and the continuum of states in the conduction band. A simple method was discussed that adds a small number of ‘conduction band states’ with some effective coupling to the exciton states, but this would add many free parameters to the model in a poorly defined way, so it was not pursued.

The discrepancy in the relative polarisation dependence of the two fields between

theory and the experiment discussed in section 2.3.7 provides an interesting avenue to continue this work. A simple initial route could be to replace the spherical harmonics of the Wannier-Mott wavefunction with cubic harmonics, with the idea that they would better suit the symmetry of the crystal. The orientation of the crystal with respect to the two fields is also not included. It seems likely that the correct polarisation dependence will derive from a complete analysis of the exciton in these three different bases.

The phonon-mediated absorption background is only modelled phenomenologically in this work. If one were to find a way to model more completely the exciton's interaction with phonons, it would give a way of introducing the asymmetric width of the nP excitons [86, 87]. It would also have a significant effect on the decay widths of the other excitonic states, which are currently derived from 2-photon experiments that do not interact with the background [40]. Furthermore, it would introduce an energy-dependent phase to the valence band to nP couplings, which could be a solution to the phase problem in the sidebands. This would be a valuable addition to the model indeed.

2.5.1 A possible broadening of the excitonic states

At one point in the project data for the full-width-half-maximum of the excitonic states were misused as the half-width-half-maximum of those states, an increase in the experimentally measured widths by a factor of 2. Along with as some changes in the fitting of the 2-photon excitation data, the model was evaluated with the widths for the S, D, and F states 1.5 to 3 times too large. Interestingly, this led to significant improvement in comparisons with experimental data. Figure 2.17 shows the theoretical absorption spectrum for these artificially large widths, along with the experimental spectra for reference. Qualitatively it resembles the experimental spectra much more closely. The additional peaks that appear on the shoulders of the 8P and 9P peaks in the theoretical spectra in figure 2.9 do not appear here, just as they do not appear in the experiment. The absorption continuum is also much smoother, though still not as flat as in experiments. Figure 2.18 shows the change in absorption at $\mathcal{E}_m = 688 \pm 5 \text{ V m}^{-1}$ for the artificial widths against experiment. The quantitative agreement improves as well. The most dramatic improvement is that the change at 6P and 7P is the correct magnitude. There is also some improvement in the change at 9P. This could indicate that there is additional broadening of those states in the experiment above the widths measured in the 2-photon experiments [40].

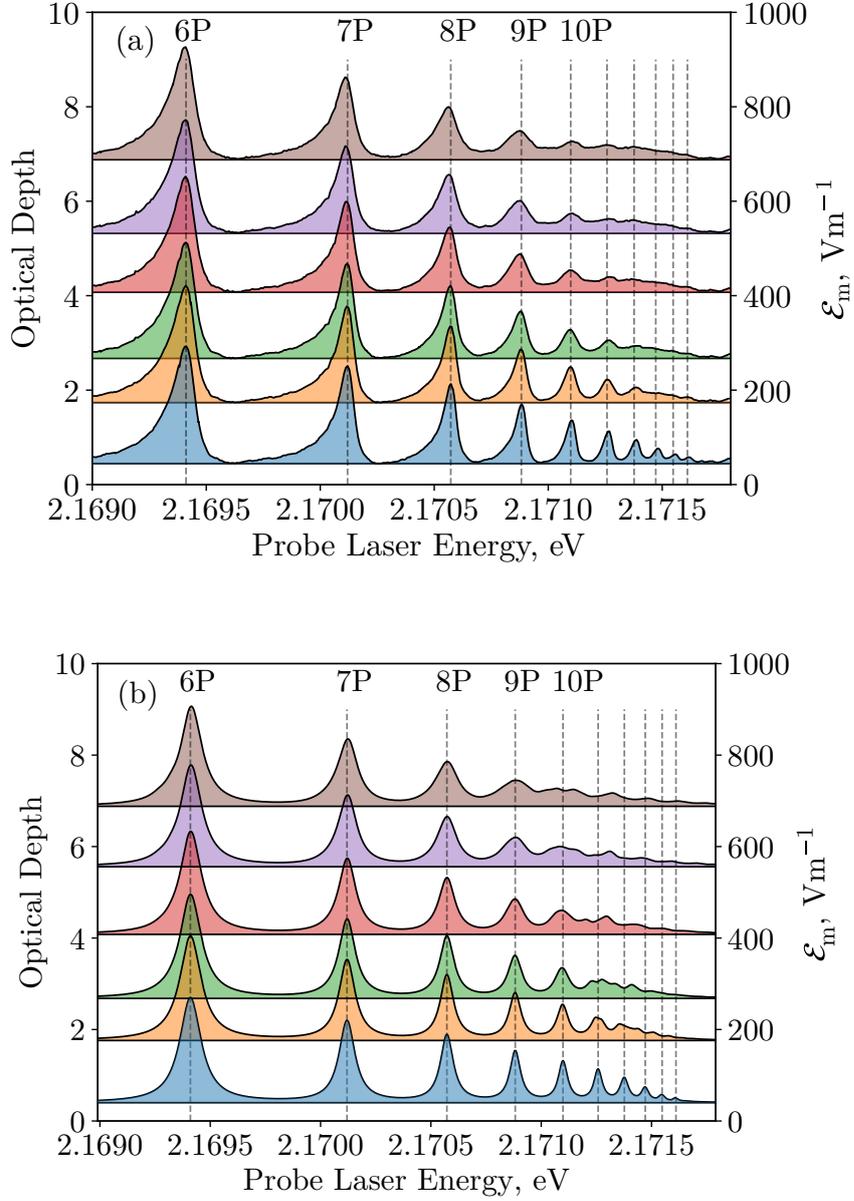


Figure 2.17: The probe absorption spectra at different microwave field strengths for (a) the experiment, and (b) the model with artificially large widths for the S, D, and F states. Each spectrum has been offset in optical depth by $\epsilon_m/100$ to indicate the microwave field strength it was taken at, shown on the righthand y-axis. The dashed vertical lines show the positions of the zero-field nP resonances.

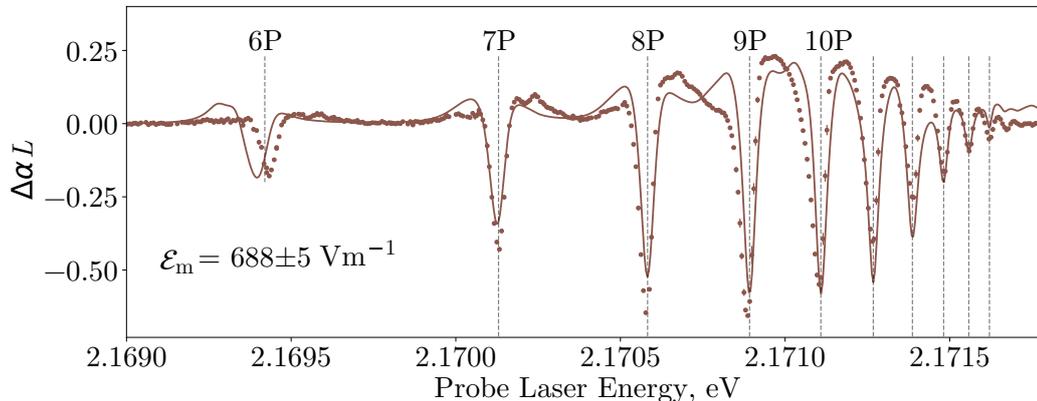


Figure 2.18: Comparison of the change in absorption $\Delta\alpha L$ as a function of laser energy between the experiment (filled circles) and the model with artificially large widths for the S, D, and F states (solid lines). The dashed vertical lines show the positions of the zero-field nP resonances.

2.5.2 Future experiments

Figure 2.19 shows the predicted absorption spectra for field strengths up to $\mathcal{E}_m = 1400 \text{ V m}^{-1}$. Unsurprisingly, the absorption continuum broadens, including states 8P and above. Features near 7P begin to couple strongly enough to the valence band to show as peaks in the absorption of the laser, indicating even states at those energies are participating in strong driving. Future experiments that can achieve higher field strengths, for example with superconducting resonators, could probe this extreme limit. Perhaps if ultrastrong driving (USD) and ultrastrong coupling (USC) physics can be done at lower Rydberg principal quantum numbers then poorer quality samples can be used, improving the scalability of technologies.

2.6 Conclusions

We have constructed a model for the effect of a strong microwave field on an ensemble of non-interacting Rydberg excitons in Cu_2O , probed by a weak laser field. Using Floquet theory, we have solved for the steady state of the system, something not previously done in excitons. The model qualitatively reproduces the absorption spectra of the probe laser and quantitatively reproduces the effect of the microwave field on the spectra up to field strengths 16 times larger than the perturbative limit [25], a remarkable result for a condensed matter system with no free parameters. We can be confident that, instead of ionising the excitons,

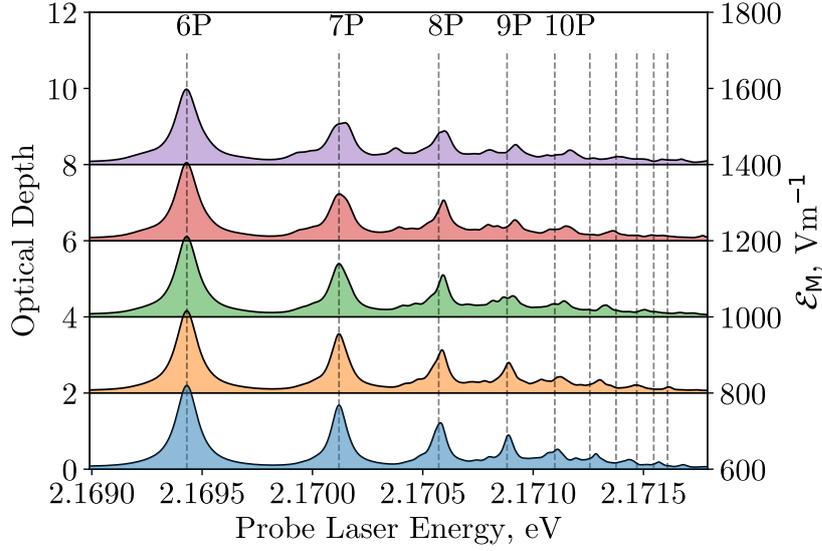


Figure 2.19: The theoretical laser absorption spectra for microwave field strengths above those of the experiment. Again, each spectrum has been offset in optical depth by $\mathcal{E}_m/100$ to indicate the microwave field strength it was taken at, shown on the righthand y-axis. The dashed vertical lines show the positions of the zero-field nP resonances.

high principal quantum number states are hybridised together into a dense forest of dressed ‘Floquet’ states, forming a broad absorption continuum. This does not occur in atomic systems [53] due to the non-radiative broadening of excitonic states allowing many states to be coupled with the same field. The continuum maintains a strong coupling to the microwave field, as shown by the measurements of the microwave-induced sidebands. Many excitonic transitions cross well into the deep strong driving regime, with couplings strengths in excess of 4 times their energy separation. At the highest field strengths there are significant processes involving the exchange of at least 8 microwave photons before the exciton decays.

The sidebands generated on the probe laser by the microwave field are also well reproduced, even quantitatively at lower laser energies. Despite a complicated dependence on unknown phases at high energies, the model qualitatively reproduces for both the second and fourth order red and blue sidebands the shape and intensity of the sideband rollover, and quantitatively predicts the field strength it occurs at, across the exciton spectrum. This is a significant result as it is at the core of optical non-linearities that can be used for microwave-optical conversion [25]. The near-quantitative agreement means that this model can be used to predict the behaviour of devices used for microwave-optical conversion well beyond the linear regime and deep into the ultrastrong driving regime, maximising the exploitation of the Rydberg mediated microwave-optical nonlinearity.

DFT study of defects in Cu₂O

3.1 Background

In the study of Rydberg excitons there is a problem: that the exciton spectra from different samples differ significantly in the heights of the absorption peaks and the number of resolvable peaks. Though they vary, the spectra which show the highest- n exciton peaks are all seen in natural samples, despite the unknown growth conditions. However, samples of such quality are rare. Work is currently being done on the growth of synthetic cuprite [95, 80, 26] to address this problem and improve potential scalability of cuprite-based quantum technologies.

To test the quality of a sample directly, its Rydberg exciton series must be measured, a lengthy experimental process, which affects turn-around time for feeding back to growth techniques. One method that has been devised to test for sample quality is by looking at the photoluminescence (PL) spectrum [100, 101, 80]. Photoluminescence is a technique where the crystal is illuminated by a laser with energy equal to the band gap, E_g , and the (off-axis) emitted light is spectroscopically resolved. It was noticed [80] that the PL spectrum in synthetic samples contained significant emission of energy less than E_g that the high-quality natural samples did not, leading to the hypothesis that whatever causes the emission also causes the degradation in the exciton spectra.

Naturally, the pure crystal cannot emit photons with less energy than E_g , so the cause must be defects in the crystal structure which introduce local electronic states in the band gap, called defect states. Electrons can then relax from the conduction band into the defect states, or out of them into the valence band, and emit the lost energy as photons. The defect states can also trap charge locally, and charged defects in Cu₂O are currently being studied for their effect on Rydberg

excitons [95, 89, 26, 27].

The photoluminescence spectrum of Cu_2O has been studied many times over its history, and 4 main peaks have been documented [27, 80, 89, 101, 100, 102, 28, 103, 104], at 1.2 eV, 1.35 eV, 1.5 eV, and 1.7 eV, as shown in the sketch in figure 3.1. It is commonly reported in the field that the peak at 1.35 eV is due to the copper vacancy, V_{Cu} , where a single copper atom is missing from the lattice. The peaks at 1.5 eV and 1.7 eV are commonly attributed respectively to V_{O}^+ and V_{O}^{2+} , the different charge states of the oxygen vacancy. Good justifications for these assignments could not be found. Many authors cite Ito et al. [103], a thorough early study of the effect of Cu_2O growth techniques on the PL spectrum. However, they do not offer any justification either.

Ito et al. in turn cite Bloem's study from 1958 [104] for the assignments, shown in figure 3.1. In Bloem's samples the heights of the peaks at 1.5 eV and 1.7 eV were always in the same ratio as each other, and disappeared in the samples grown in high oxygen pressure. Under the assumption that the only significant native defects were V_{Cu} and V_{O} , he therefore assigned the peaks at 1.5 eV and 1.7 eV to V_{O}^+ and V_{O}^{2+} , and the large peak at 1.35 eV to the remaining defect, V_{Cu} . He did not speculate on the origin of the peak at 1.2 eV, and few other authors pay it any mind either.

These arguments are no longer supported by the available data. For example, the samples in Ito et al. [103] have the heights of the 1.5 eV and 1.7 eV peaks in a variety of ratios, including their natural sample showing a bright peak for 1.7 eV and no peak at all at 1.5 eV, suggesting they come from different defects. Density functional theory studies [105, 106, 30] have shown that a host of native defects can form under various growth conditions, and so there is no longer a reason to assign the 'leftover' peak to V_{Cu} .

The initial aim of this project was to attempt to provide a theoretical basis for the defect assignments of the photoluminescence lines using density functional theory (DFT). DFT is an ab-initio method for calculating the electronic properties of matter and is ubiquitous in condensed matter physics and quantum chemistry. Standing on over 60 years of theory development [107], there are now many sophisticated out-of-the-box DFT codes built for high performance computing. In this work we make use of CASTEP [108], which specialises in periodic solids.

To keep the scope of the project manageable, we restricted ourselves to the study of native point defects in Cu_2O . Since the 4 PL peaks are consistent across many different samples, synthetic and natural, we did not study extrinsic defects, i.e. atoms other than copper or oxygen. The simplest point defects are the copper and

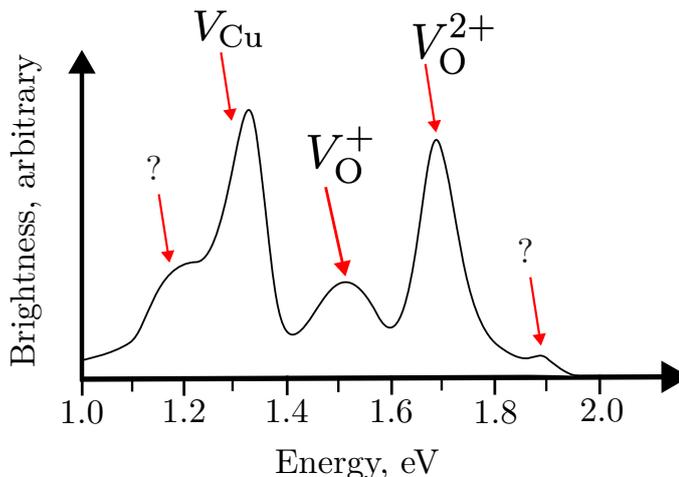


Figure 3.1: A sketch of an example PL spectrum showing the 4 main peaks commonly observed, with the defect assignments made by Bloem [104].

oxygen vacancies, V_{Cu} and V_{O} , where a single atom is missing from the lattice, and the anti-sites, Cu_{O} (O_{Cu}) where a copper (oxygen) atom is in the usual place of an oxygen (copper) atom. Next are the interstitials, Cu_i and O_i , where a copper or oxygen atom is inserted into the stoichiometric lattice. For each species of interstitial atom there are two different stable arrangements, labelled tetrahedral and octahedral, after the way they coordinate with neighbouring copper atoms. Finally, we also studied the split copper vacancy, $V_{\text{Cu}}^{\text{split}}$, a defect complex made of two copper vacancies and an interstitial copper atom half way in between. While other defect complexes have been studied in the context of some effects [109, 110], we concentrated this work on point defects for simplicity. However, we include $V_{\text{Cu}}^{\text{split}}$ as it has been the topic of much discussion around DFT and defect states in Cu_2O [30, 106, 31, 111].

While in the literature only one geometry for $V_{\text{Cu}}^{\text{split}}$ is mentioned, in the course of this work we studied two. While all individual copper atoms in the material are symmetry equivalent, nearest-neighbour pairs are not. If the two copper atoms removed were not bonded to the same oxygen atom, it forms a slightly different defect to one if they shared an oxygen atom. Only the former case is studied in the literature [30, 106, 31, 111]. We will denote the former defect $V_{\text{Cu}}^{\text{s},1}$ and the latter $V_{\text{Cu}}^{\text{s},2}$, which gives 10 different defects in total.

This is far from the first study of defects and defect states in Cu_2O . Defect states in the material have been of great interest to the semiconductor community mainly in the context of solar cells. In 2009, Soon et al. [106] modelled all the point defects as well as the split vacancy with PBE-DFT in $2 \times 2 \times 2$ supercells. They found that

the interstitials produced defect states in the band gap, and the copper vacancies gave some empty states right at the Fermi level, which can be seen as perturbations of the valence band in their band structures. Scanlon et al. [30] investigated the copper interstitials and vacancies in more detail under HSE-DFT, a more intensive method, to look for more complex exchange-driven behaviour, and this time found defect states well into the band gap for the copper vacancies. However, due to a lack of computing power, these results suffer from under-convergence with various DFT parameters. We expand on that in this work, and demonstrate issues in such studies. In 2013, Isseroff and Carter [31] revisited Scanlon et al.'s study of the copper vacancies with updated hardware. They found that the defect state attributed to V_{Cu} was indeed a convergence error causing the valence band maximum (VBM) to rise above the rest of the valence band, and that with proper k-point sampling it returned to its bulk behaviour. However, they found $V_{\text{Cu}}^{\text{s};1}$ still gave a state in the band gap.

This work builds on the accumulated knowledge of these fragmentary studies to provide a complete and thorough theoretical investigation of defect states in Cu_2O . In particular, we identify which features caused by the defects are robust and which are spurious effects of theoretical limitations by studying how the band structures change with supercell size and exchange-correlation approximations, an attention to detail not usually seen in the study of defects in general [32, 33, 34, 35, 36]. By the end, we are certain what *does not* cause photoluminescence, and cautiously make new assignments of the emission lines based on our results.

3.2 Theory

In this section I will give an overview of Density Functional Theory (DFT). This will not be a complete tutorial of DFT, as it is easily read about in many of other sources [29, 112], but I want to cover it in a way that makes accessible the parts that are relevant to this thesis.

To begin, if we want to model the properties of a solid, we need its many-body wavefunction, $\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M; t)$ (for N electrons and M nuclei, suppressing their spins), which is a function of all of its electron positions \mathbf{r}_i and nuclear positions \mathbf{R}_I . With the wavefunction, we can calculate everything there is to know about the solid, and to find it, we have to solve the many-body time-dependent Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M; t) = \hat{H}(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M) \Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M; t). \quad (3.1)$$

The non-relativistic Hamiltonian for all the electrons and nuclei in a solid can be written

$$\begin{aligned} \hat{H} = & - \sum_i \frac{\hbar^2}{2m_e} \nabla_i^2 + \frac{1}{2} \sum_{i \neq j} \frac{e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i,I} \frac{Z_I e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{R}_I|} \\ & - \sum_I \frac{\hbar^2}{2M_I} \nabla_I^2 + \frac{1}{2} \sum_{I \neq J} \frac{Z_I Z_J e^2}{4\pi\epsilon_0 |\mathbf{R}_I - \mathbf{R}_J|}. \end{aligned} \quad (3.2)$$

where m_e is the electron mass, e is the electron charge, M_I is the mass of each nucleus, and Z_I is the atomic number of each nucleus. Reading it from left to right, it is made of the sum of the kinetic energies of all the electrons (\hat{T}_e), the electron-electron Coulomb interaction (\hat{V}_{ee}), the electron-nuclear Coulomb interaction (\hat{V}_{en}), the kinetic energy of all the nuclei (\hat{T}_n), and the nuclear-nuclear Coulomb interaction (\hat{V}_{nn}). In this chapter, we will look only for ground state properties, and so consider only the time-independent Schrödinger equation

$$\hat{H}\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M) = \mathcal{E}\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M) \quad (3.3)$$

where the ground state is the eigenstate with the smallest eigenenergy, \mathcal{E} , which is the ground state energy.

3.2.1 Born-Oppenheimer Approximation

The wavefunction for a macroscopic solid involves $N = 10^{24}$ particles, each in 3 dimensions, so it is a $3^{10^{24}}$ -dimensional object. To make any progress, we are going to have to make some simplifications. The first simplification we can make, called the Born-Oppenheimer approximation, is that the nuclei are so much heavier than the electrons that they are effectively stationary on electronic time scales, so as the nuclei move the electrons remain in the instantaneous ground state of the nuclear potential. This makes the problem separable into electron and nuclear parts,

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N; \mathbf{R}_1, \dots, \mathbf{R}_M; t) \approx \Psi_{\mathbf{R}_1, \dots, \mathbf{R}_M}^{\text{elec}}(\mathbf{r}_1, \dots, \mathbf{r}_N; t) \chi(\mathbf{R}_1, \dots, \mathbf{R}_M; t). \quad (3.4)$$

The electronic Hamiltonian is therefore

$$\begin{aligned} \hat{H}_e = & -\frac{1}{2} \sum_i \nabla_i^2 + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i,I} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} \\ = & \hat{T}_e + \hat{V}_{ee} + \hat{V}_{en} \end{aligned} \quad (3.5)$$

where we have adopted Hartree units $e = \hbar = m_e = 4\pi\epsilon_0 = 1$ for simplicity. To find the electronic wavefunction we solve the electronic Schrödinger equation

$$\hat{H}_e \Psi_{\mathbf{R}_1, \dots, \mathbf{R}_M}^{\text{elec}}(\mathbf{r}_1, \dots, \mathbf{r}_N; t) = E(\mathbf{R}_1, \dots, \mathbf{R}_M) \Psi_{\mathbf{R}_1, \dots, \mathbf{R}_M}^{\text{elec}}(\mathbf{r}_1, \dots, \mathbf{r}_N; t) \quad (3.6)$$

for eigenenergy $E(\mathbf{R}_1, \dots, \mathbf{R}_M)$ as a function of nuclear positions (again, suppressing spin). The nuclear positions therefore lie on a potential energy surface,

$$V_{\text{Nuclear}}(\mathbf{R}_1, \dots, \mathbf{R}_M) = E(\mathbf{R}_1, \dots, \mathbf{R}_M) + V_{nn} \quad (3.7)$$

where

$$V_{nn} = \int_V d\mathbf{R}_1 \dots d\mathbf{R}_M \hat{V}_{nn} |\chi(\mathbf{R}_1, \dots, \mathbf{R}_M; t)|^2 \quad (3.8)$$

is the nuclear-nuclear energy term. The ground state geometries of solids are the minima of this potential energy surface.

3.2.2 Density Functional Theory

The wavefunction for the electrons is still a $\sim 3^{10^{24}}$ -dimensional object. If we study periodic systems, via Bloch's theorem we can reduce N to the number of particles in the unit cell, as we will see later. However, if we were to discretise even the unit cell into just 10 points along each axis, we would still need 1000^N floating point numbers to store the wavefunction, which at 4 bytes each exceeds the total amount of data storage in the world after 7 electrons. This is the problem that DFT tries to solve, by considering not the electron wavefunction but the electron density

$$\rho(\mathbf{r}) = N \sum_{\sigma_i} \int_V d\mathbf{r}_2 \dots d\mathbf{r}_N |\Psi(\mathbf{r}, \sigma_1, \mathbf{r}_2, \sigma_2, \dots, \mathbf{r}_N, \sigma_N)|^2, \quad (3.9)$$

where we explicitly write the spin of each electron, σ_i , have dropped the labels on Ψ for convenience. The normalisation of the density is therefore

$$\int_V d\mathbf{r} \rho(\mathbf{r}) = N. \quad (3.10)$$

Of course, we also lose a lot of information from the wavefunction by integrating over all-but-one of the variables, crucially including ways to calculate total energy and impose Pauli exclusion. Most of the work done in DFT is to replace everything that is lost in the transition from wavefunction to density.

3.2.3 Hohenberg-Kohn Theorems

Underpinning the mathematics of density functional theory are the theorems of Hohenberg and Kohn. Attentive readers may worry that one of the things lost from the many body electron wavefunction approach is uniqueness. It is not unreasonable to imagine two different potentials, that would produce two different ground state wavefunctions, leading to the same ground state density. Fortunately,

the first Hohenberg-Kohn theorem will dispel these fears. The first theorem states:

HK 1: *If two potentials, $v_1(\mathbf{r})$ and $v_2(\mathbf{r})$, lead to the same ground state density, then $v_1(\mathbf{r}) - v_2(\mathbf{r})$ must be a constant.*

Of course, only gradients of potentials are physically meaningful, so two potentials that differ by a constant define the same physical potential, or in other words, the external potential is uniquely determined by the ground state density (and vice versa). We can therefore write the potential as a functional of the density, $v[\rho](\mathbf{r})$, and the density as a functional of the potential, $\rho[v](\mathbf{r})$. The second HK theorem states:

HK 2: *There exists an energy functional $E[\rho]$, such that the density $\rho(\mathbf{r})$ that minimises $E[\rho(\mathbf{r})]$, subject to the constraint that $\int_V d\mathbf{r}\rho(\mathbf{r}) = N$, is the true ground state density of the many-body system.*

We can write the energy of the many-body electron system as a **density functional**,

$$\begin{aligned} E[\rho] &= F[\rho] + E_{en}[\rho] \\ &= F[\rho] + \int_V d\mathbf{r}\rho(\mathbf{r})v_{en}(\mathbf{r}) \end{aligned} \quad (3.11)$$

where

$$F[\rho] = \min_{\Psi \rightarrow \rho} \langle \Psi | \hat{T}_e + \hat{V}_{ee} | \Psi \rangle \quad (3.12)$$

is the minimum energy, over all wavefunctions that integrate to ρ , of the kinetic and electron-electron potential operators. Taking the minimising wavefunction to be $\Psi[\rho]$,

$$\begin{aligned} F[\rho] &= \langle \Psi[\rho] | \hat{T}_e + \hat{V}_{ee} | \Psi[\rho] \rangle \\ &= \langle \Psi[\rho] | \hat{T}_e | \Psi[\rho] \rangle + \langle \Psi[\rho] | \hat{V}_{ee} | \Psi[\rho] \rangle \\ &= T[\rho] + E_{ee}[\rho] \end{aligned} \quad (3.13)$$

which defines the kinetic energy functional, $T[\rho]$, and the electron-electron interaction energy functional, $E_{ee}[\rho]$. Taking the two HK theorems together, since $T[\rho]$ and $E_{ee}[\rho]$ have the same functional forms for all densities (known as universal functionals), and $v_{en}(\mathbf{r})$ uniquely determines the energy functional, the minimum of $E[\rho]$ must be unique for each $v_{en}(\mathbf{r})$, so $E[\rho]$ is all we need to find the exact ground state density.

Finding $T[\rho]$ and $E_{ee}[\rho]$ is the holy grail of DFT. However, they have not yet been found. The remainder of this theory section will be about the various strategies adopted to best approximate these functionals that are computable in reasonable time frames.

3.2.4 The Kohn-Sham auxiliary system

A great breakthrough in the hunt for approximations of both $T[\rho]$ and $E_{ee}[\rho]$ that enforce Pauli exclusion was the development of the Kohn-Sham auxiliary system. The idea is to recast the system of many *interacting* electrons in the potential of the nuclei into a system of the same number of *non-interacting* ‘electrons’ (called Kohn-Sham (KS) orbitals) in the effective potential of all the KS orbitals and nuclei [112, 29]. Crucially, this can be done in a way where the orbitals sum to the same density as the fully-interacting system. It is a matter of intense debate how closely KS orbitals and their eigenenergies map onto real electrons [113, 114], but in this thesis we will consider the KS orbitals as good approximations for electrons, and will refer to them as electrons interchangeably.

The wavefunction of the auxiliary system is constructed from a Slater determinant of all the orbitals to ensure it is antisymmetric. From that wavefunction, the electron density can be shown to be expressed in terms of the orbitals as

$$\rho(\mathbf{r}) = \sum_i f_i |\phi_i(\mathbf{r})|^2 \quad (3.14)$$

where $\phi_i(\mathbf{r})$ is the i -th Kohn-Sham orbital, and f_i is the orbital’s (fractional) occupancy, ranging from 0 if it is unoccupied to 2 if it is doubly occupied. The use of fractional occupancy is simply to smooth convergence.

In non-magnetic materials with an even number of electrons in the unit cell, a common, accurate approximation is that all occupied orbitals are doubly occupied, which makes accounting for spin much easier. Pure Cu_2O meets these criteria, as do most of the defects, so we will do our derivations here under that assumption, neglecting the spin labels on the orbitals (starting with equation 3.14). However, there are a few defects which introduce an odd number of electrons into the cell for which we must use spin-DFT, a good summary of which can be found here [115]. Spin-orbit coupling can be introduced via the pseudopotential, discussed later.

The density is used to calculate $E_{en}[\rho]$ as above, and can also be used to calculate the most common approximation to E_{ee} , called the Hartree energy,

$$E_{\text{H}} = \frac{1}{2} \iint_V d\mathbf{r}' d\mathbf{r} \frac{\rho(\mathbf{r}')\rho(\mathbf{r})}{|\mathbf{r}' - \mathbf{r}|}, \quad (3.15)$$

which is the classical Coloumb energy of a charge distribution. Note that in periodic systems it is more efficient to evaluate E_H in reciprocal space. Also, now that we are working with single particle states, it is simple to calculate the total kinetic energy of the KS system, $T_s[\rho]$, as simply the sum over the kinetic energies of each electron,

$$T_s[\rho] = -\frac{1}{2} \sum_i f_i \int_V d\mathbf{r} \phi_i^*(\mathbf{r}) \nabla^2 \phi_i(\mathbf{r}). \quad (3.16)$$

Putting all the above pieces together would not yield the same energy as the interacting system, and so would give the wrong ground state density. Namely, the Hartree approximation to the interaction energy neglects the exchange interaction (and includes a self-interaction error), and we have lost terms known as correlation that come from electron density in one place affecting how the electrons interact with the density at another. All this missing energy is bundled up into the exchange-correlation (xc) energy functional, $E_{xc}[\rho]$, which is defined as

$$\begin{aligned} E_{xc}[\rho] &= E[\rho] - (T_s[\rho] + E_H[\rho] + E_{en}[\rho]) \\ &= T[\rho] + E_{ee}[\rho] - (T_s[\rho] + E_H[\rho]). \end{aligned} \quad (3.17)$$

Thus the exact total energy can be expressed in terms of the KS orbitals as

$$E[\rho] = T_s[\rho] + E_{ee}[\rho] + E_{en}[\rho] + E_{xc}[\rho] \quad (3.18)$$

$$\begin{aligned} &= -\frac{1}{2} \sum_i f_i \int_V d\mathbf{r} \phi_i^*(\mathbf{r}) \nabla^2 \phi_i(\mathbf{r}) + \frac{1}{2} \iint_V d\mathbf{r}' d\mathbf{r} \frac{\rho(\mathbf{r}')\rho(\mathbf{r})}{|\mathbf{r}' - \mathbf{r}|} \\ &\quad + \int_V d\mathbf{r} v_{en}(\mathbf{r})\rho(\mathbf{r}) + E_{xc}[\rho]. \end{aligned} \quad (3.19)$$

However, knowing the form of $E_{xc}[\rho]$ would mean one could calculate the exact total energy, so it is believed to be at least as hard as solving the full interacting system. The form of the exact exchange functional is known, though computationally expensive, however the correlation functional is unknown, therefore clever approximations must be made for the xc-functional to keep computational costs down without sacrificing too much accuracy [116] (a huge research field in its own right). Many xc-functionals have been developed that span the range of computational cost and accuracy. The ones used for this work are discussed in section 3.2.6.

Finally, we must find the orbitals which minimise the total energy under the constraint that each orbital is normalised such that the density integrates to N . We

can do this by introducing Lagrange multipliers ϵ_i to $E[\rho]$, and minimising instead

$$G[\rho] = E[\rho] - \sum_i \epsilon_i f_i \int_V d\mathbf{r} |\phi_i(\mathbf{r})|^2. \quad (3.20)$$

Taking the functional derivative of $G[\rho]$ with respect to each orbital and setting it to zero,

$$\begin{aligned} 0 &= \frac{\delta G[\rho]}{\delta \phi_i^*(\mathbf{r})} = \frac{\delta E[\rho]}{\delta \phi_i^*(\mathbf{r})} - \epsilon_i \phi_i(\mathbf{r}) \\ &= -\frac{1}{2} \nabla^2 \phi_i(\mathbf{r}) + \left[\int_V d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|} \right] \phi_i(\mathbf{r}) + v_{en}(\mathbf{r}) \phi_i(\mathbf{r}) + v_{xc}(\mathbf{r}) \phi_i(\mathbf{r}) - \epsilon_i \phi_i(\mathbf{r}) \\ &= -\frac{1}{2} \nabla^2 \phi_i(\mathbf{r}) + \left[v_H(\mathbf{r}) + v_{en}(\mathbf{r}) + v_{xc}(\mathbf{r}) \right] \phi_i(\mathbf{r}) - \epsilon_i \phi_i(\mathbf{r}), \end{aligned} \quad (3.21)$$

where

$$v_H(\mathbf{r}) = \int_V d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|}, \quad (3.22)$$

and

$$v_{xc}(\mathbf{r}) = \frac{\delta E_{xc}[\rho]}{\delta \rho(\mathbf{r})} \quad (3.23)$$

define the Hartree and exchange correlation local potentials respectively. Note that this derivation holds only for approximate exchange-correlation energies that are explicit functionals of the density, as opposed to, for example, functionals of the orbitals. Rearranging, and collecting the potential terms together into the Kohn-Sham potential, $v_{KS}(\mathbf{r})$, we obtain the single particle Kohn-Sham equations,

$$\left[-\frac{1}{2} \nabla^2 + v_{KS}(\mathbf{r}) \right] \phi_i(\mathbf{r}) = \epsilon_i \phi_i(\mathbf{r}), \quad (3.24)$$

from which the KS orbitals are found. Since v_{KS} depends on all the orbitals via the density, the solutions to this equation must be found *self-consistently*, meaning that the true solutions are the orbitals which sum to a density which produces a KS potential whose solutions are the same Kohn-Sham orbitals.

We can now naïvely imagine the arc of an algorithm for KS-DFT:

1. Make an educated guess* at the ground state density $\rho(\mathbf{r})$
2. Calculate the potential $v_{KS}(\mathbf{r})$
3. Find the KS orbitals for $v_{KS}(\mathbf{r})$ from the one particle Schrödinger equation

*There are many ways of doing this, and they are all beyond the scope of this thesis.

4. Calculate the ground state density from the orbitals
5. If the change in the density over each cycle is smaller than some threshold, stop. Otherwise, return to step 2

This algorithm is fine in principle, but would in practice take a long time and may not converge. Hence, we must be cleverer and make use of the total energy. The actual procedure followed in DFT codes like CASTEP is much more complicated and involves taking the functional derivative of the total energy with respect to the orbitals, so that one can take decreasing steps in energy to the ground state much more efficiently. As one might expect, there are also many choices of algorithm which trade on speed and accuracy [117], the discussion of which is beyond the scope of this work. Also note that the energy minimum is unique only if there are no stable magnetic states. This is not a problem if we assumed double occupancy of orbitals, but in cases where we cannot (such as defect cells with odd numbers of electrons), it is wise to solve the system several times from random starting points to explore the different minima the system can fall into.

3.2.5 Periodic systems and Bloch's theorem

Up until this point the theory we have discussed can be applied to any system of electrons. In this work, however, we are particularly interested in cuprous oxide, which is a periodic system, and we can use this to make significant progress in solving for the KS orbitals. The first step we can make is in recognising that a sufficiently large crystal is very well approximated by an infinite crystal, where the potential v_{KS} must be periodic, such that

$$v_{\text{KS}}(\mathbf{r}) = v_{\text{KS}}(\mathbf{r} + \mathbf{L}) \quad (3.25)$$

for any real space lattice vector \mathbf{L} of the crystal lattice, given by

$$\mathbf{L} = n_1 \mathbf{a} + n_2 \mathbf{b} + n_3 \mathbf{c}, \quad (3.26)$$

where for the Cu_2O cubic cell of side-length a the vectors $\mathbf{a} = a\hat{x}$, $\mathbf{b} = a\hat{y}$, and $\mathbf{c} = a\hat{z}$ and n_i are integers. We can also define the spatial frequencies, called reciprocal lattice vectors \mathbf{G} , over which V_{KS} is periodic, as

$$\mathbf{G} = m_1 \mathbf{a}^* + m_2 \mathbf{b}^* + m_3 \mathbf{c}^*, \quad (3.27)$$

where m_i are integers and the vectors \mathbf{a}^* , \mathbf{b}^* , and \mathbf{c}^* are for the cubic cell simply

$$\mathbf{a}^* = \frac{2\pi}{a} \hat{\mathbf{x}}, \quad (3.28)$$

$$\mathbf{b}^* = \frac{2\pi}{a} \hat{\mathbf{y}}, \quad (3.29)$$

$$\mathbf{c}^* = \frac{2\pi}{a} \hat{\mathbf{z}}. \quad (3.30)$$

We have already seen in chapter 2 that solutions to periodic Hamiltonians take a particular form. While in time-periodic systems the result is called Floquet's theorem, in space-periodic systems it is called Bloch's theorem, and it tells us that the solutions to equation 3.24 take the form

$$\phi_i^k(\mathbf{r}) = \exp(i\mathbf{k} \cdot \mathbf{r}) u_i^k(\mathbf{r}) \quad (3.31)$$

where $u_i^k(\mathbf{r})$ is periodic with the same periodicity as V_{KS} and \mathbf{k} is called the *quasi-momentum*, also known as the crystal momentum, or *k-point*. These Bloch wavefunctions obey the orthonormality condition

$$\int_V d\mathbf{r} \phi_i^k(\mathbf{r}) \phi_i^{k'*}(\mathbf{r}) = \delta(\mathbf{k} - \mathbf{k}') \quad (3.32)$$

such that orbitals at different k -points are orthogonal to each other. Since the density is made of the squared norm of the orbitals, it too will have the periodicity of the potential. We use these facts to impose periodic boundary conditions on $u_i^k(\mathbf{r})$ and $\rho(\mathbf{r})$, reducing our problem greatly from infinite size and infinite electrons to one unit cell and only the electrons in the unit cell.

Again, like the quasi-energy, the quasi-momentum has units of momentum but is not conserved (only the momentum of the electrons plus the ions is). Unlike in chapter 2, where there were discrete eigenenergies corresponding to the discrete exciton spectrum, the infinite number of electrons in an infinite crystal correspond to an effective continuum of allowed quasi-momenta when condensed down into the periodic unit cell (addressed in more detail in section 3.3).

As before, we can show that the k -points are not uniquely defined. Since $u_i^k(\mathbf{r})$ is periodic in \mathbf{L} , we can expand it as a Fourier series,

$$u_i^k(\mathbf{r}) = \sum_{\mathbf{G}} c_{ik}^{\mathbf{G}} \exp(i\mathbf{G} \cdot \mathbf{r}). \quad (3.33)$$

Choosing to represent the orbitals with the basis of plane waves in this way is actually a significant choice amongst DFT codes. Many molecular DFT codes use a basis of atomic orbitals, but most crystal codes including CASTEP prefer plane waves, as they better describe highly delocalised states and are not biased

to any regions of space. Using 3.33, we could recast the orbitals, for any choice of reciprocal lattice vector \mathbf{K} , as

$$\begin{aligned}\phi_i^k(\mathbf{r}) &= \exp(i[\mathbf{k} + \mathbf{K}] \cdot \mathbf{r}) \exp(-i\mathbf{K} \cdot \mathbf{r}) u_i^k(\mathbf{r}) \\ &= \exp(i[\mathbf{k} + \mathbf{K}] \cdot \mathbf{r}) \sum_{\mathbf{G}} c_{ik}^{\mathbf{G}} \exp(i[\mathbf{G} - \mathbf{K}] \cdot \mathbf{r}) \\ &= \exp(i[\mathbf{k} + \mathbf{K}] \cdot \mathbf{r}) \sum_{\mathbf{G}'} c_{i,k+\mathbf{K}}^{\mathbf{G}'} \exp(i\mathbf{G}' \cdot \mathbf{r}) = \phi_i^{k+\mathbf{K}}(\mathbf{r})\end{aligned}\quad (3.34)$$

where $\mathbf{G}' = \mathbf{G} - \mathbf{K}$ is simply a re-indexing of the sum by moving the origin of the infinite crystal lattice. Therefore, the orbitals ϕ_i^k are only unique in \mathbf{k} modulo $2\pi/a$ in each direction; all possible electrons are contained within a reciprocal-space box $-\pi/a < k_{x,y,z} < \pi/a$, called the *first Brillouin zone**. Finally, it follows that the construction of the density from the orbitals in equation 3.14 becomes

$$\rho(\mathbf{r}) = \sum_{i,k} f_i^k |\phi_i^k(\mathbf{r})|^2 \quad (3.35)$$

where k indexes the quasi-momenta in the first Brillouin zone, i is the band index, and f_i^k is the occupancy of the i -th band at the k -th k-point.

3.2.6 Exchange-Correlation Functionals

There are many choices in Kohn-Sham DFT for the approximate exchange-correlation energy functional, $E_{\text{xc}}[\rho]$, which vary greatly in computational cost and the results they produce. The simplest are the local functionals, whose potentials at \mathbf{r} are only a function of the density at \mathbf{r} , making them very fast to compute. For example the local density approximation (LDA) is defined as

$$E_{\text{xc}}[\rho] = \int_V d\mathbf{r} \epsilon_{\text{xc}}^{\text{HEG}}[\rho(\mathbf{r})] \rho(\mathbf{r}), \quad (3.36)$$

where $\epsilon_{\text{xc}}^{\text{HEG}}[\rho(\mathbf{r})]$ is the exchange-correlation energy per particle of the homogeneous electron gas (HEG). The LDA performs reasonably well in infinite crystals, matching phonon frequencies and bond lengths to within a few percent [118] (usually a few percent too short), and works best in systems with smoothly varying densities which are most like the HEG.

A small step above strictly local functionals is a category known as the generalised gradient approximation (GGA), defined by

$$E_{\text{xc}}[\rho] = \int_V d\mathbf{r} \epsilon_{\text{xc}}[\rho(\mathbf{r}), \nabla\rho(\mathbf{r})] \rho(\mathbf{r}), \quad (3.37)$$

*The other Brillouin zones are all the other possible boxes

for a choice of function ϵ_{xc} which depends not only on the local density but also on the gradient of the density. Many options for ϵ_{xc} have been developed, but by far the most popular for solids is that of Perdew, Burke and Ernzerhof (PBE) [119]. This is the one we make the most use of in this work. Without excessive computational cost, it gives good structures to within a few percent, and more accurate binding energies than LDA, which will be useful for studying the binding of defect states to defect sites.

Both functionals also significantly underestimate the band gap. This is expected, as even the exact KS potential would underestimate the band gap due to the derivative discontinuity [120, 121, 122], which causes the eigenvalues of the unoccupied bands to differ from their true values by a constant, even for the exact KS potential. In fact, there is no mathematical guarantee that the eigenenergies of the orbitals have any physical meaning at all aside from the highest-occupied orbital. Nevertheless, quite surprisingly, the KS eigenenergies provide a good approximation to actual excitation energies of the system [123], which is why band structures from KS-DFT are used all over solid-state physics (including in this work).

While PBE provides an excellent baseline for calculations, exchange and correlation are inherently non-local effects, so it is unlikely that physics that depends strongly on exchange and correlation can be captured in a (semi)local functional of the density. For this reason, KS-DFT is often extended to what is called generalised Kohn-Sham (GKS) DFT. Without wishing to dive too deep into details, in this scheme, instead of the exchange-correlation energy being an explicit functional of the density, it is a functional of the orbitals. Then, when you take the functional derivative of the energy to find the KS equations,

$$\frac{\delta E_{xc}[\{\phi\}]}{\delta \phi_i^*(\mathbf{r})} = \hat{V}_{xc}\phi_i(\mathbf{r}) \quad (3.38)$$

defines the operator xc-potential that acts on the orbitals, as opposed to a local potential that multiplies the orbitals. Allowing for operator potentials widens the scope of possible potentials but often comes at considerable computational cost.

The relevant GKS scheme to this work is called Hartree-Fock. First, recall that in the KS scheme the wavefunction $|\Phi\rangle$ is constructed from a Slater determinant of the orbitals $\{\phi_i\}$. Instead of using the classical approximation for E_{ee} , the full quantum mechanical electron-electron interaction energy then expands as

$$\begin{aligned} E_{ee} &= \langle \Phi | \frac{1}{|\mathbf{r}' - \mathbf{r}|} | \Phi \rangle \\ &= \frac{1}{2} \iint_V d\mathbf{r}' d\mathbf{r} \frac{\rho(\mathbf{r}')\rho(\mathbf{r})}{|\mathbf{r}' - \mathbf{r}|} - \frac{1}{2} \iint_V d\mathbf{r}' d\mathbf{r} \frac{\rho(\mathbf{r}', \mathbf{r})\rho'(\mathbf{r}, \mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|}, \end{aligned} \quad (3.39)$$

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_i f_i \phi_i(\mathbf{r})\phi_i(\mathbf{r}') \quad (3.40)$$

where $\rho(\mathbf{r}, \mathbf{r}')$ is the 1-body reduced density matrix. The first term is the familiar Hartree energy, E_H , and the second term is called the Hartree-Fock exact exchange energy, E_x^{HF} . Note that the density-density term in the Hartree energy, $\rho(\mathbf{r})\rho(\mathbf{r}') = \sum_{i,j} \phi_i^*(\mathbf{r})\phi_i(\mathbf{r})\phi_j^*(\mathbf{r}')\phi_j(\mathbf{r}')$, resembles the numerator in the HF exchange energy with the coordinates of two orbitals exchanged. Compared to just E_H in regular KS-DFT, one can see that this form of E_{ee} corrects for the self-interaction error, and, as the name implies, includes the exact exchange interaction energy. From the functional derivative of E_x^{HF} with respect to the orbitals, we find the Fock exchange operator is given by

$$\hat{V}_x \phi_i(\mathbf{r}) = - \int_V d\mathbf{r}' \frac{\rho(\mathbf{r}, \mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|} \phi_i(\mathbf{r}'). \quad (3.41)$$

It is tempting, then, not to include a correlation correction at all, an approach called Hartree-Fock (HF) theory. It turns out that the correlation effects missing in HF theory are important. Hartree-Fock greatly overestimates the band gap of all materials, and in fact cannot describe metals at all, since it encounters a divergence in the derivative of the single particle energy when a single particle is at the Fermi energy, and gives zero density of states at the Fermi energy for the homogeneous electron gas.

The underestimation of band gaps by local functionals and the overestimation by Hartree-Fock led to the development of hybrid functionals, which try to combine the successes of both methods. Care should be taken in choosing a hybrid functional, as parameterisations of them are often done to fit a particular observable in a narrow class of systems, making them no longer *ab initio*. It is therefore possible to do much more computationally expensive calculations for less reliable results. All hybrid functionals split the exchange energy into a fraction α of HF exchange and a fraction $1 - \alpha$ of a density functional approximation, and the correlation comes entirely from the density functional approximation. For example, E_{xc} for the PBE0 functional is

$$E_{xc}^{\text{PBE0}} = \alpha E_x^{\text{HF}} + (1 - \alpha) E_x^{\text{PBE}} + E_c^{\text{PBE}}. \quad (3.42)$$

The functional we employ in this work is that of Heyd, Scuseria, and Ernzerhof, called HSE06 [124, 125]*. To arrive at HSE06, the exchange energies in PBE0 are first split into short- and long-range parts, as

$$E_{xc}^{\text{PBE0}} = \alpha E_x^{\text{HF;SR}}(\omega) + \alpha E_x^{\text{HF;LR}}(\omega) + (1 - \alpha) E_x^{\text{PBE;SR}}(\omega) + (1 - \alpha) E_x^{\text{PBE;LR}}(\omega) + E_c^{\text{PBE}}. \quad (3.43)$$

*HSE03 is the same functional but the original 2003 paper [124] had a mistake in the screening length parameter, so it should not be used.

where $1/\omega$ is the length at which the transition from short- to long-range occurs. Numerically it was observed that the long-range exchange energies $E_x^{\text{HF};\text{LR}}(\omega)$ and $E_x^{\text{PBE};\text{LR}}(\omega)$ are small and almost exactly the same size, so the long range HF contribution cancels with the $-\alpha$ contribution from PBE. The HSE xc energy is therefore defined as

$$E_{\text{xc}}^{\text{HSE}} = \alpha E_x^{\text{HF};\text{SR}}(\omega) + (1 - \alpha) E_x^{\text{PBE};\text{SR}}(\omega) + E_x^{\text{PBE};\text{LR}}(\omega) + E_c^{\text{PBE}}, \quad (3.44)$$

where the so-called screened Hartree-Fock energy takes the form

$$E_x^{\text{HF};\text{SR}}(\omega) = -\frac{1}{2} \iint_V d\mathbf{r}' d\mathbf{r} \rho(\mathbf{r}', \mathbf{r}) \frac{\text{erf}(\omega|\mathbf{r}' - \mathbf{r}|)}{|\mathbf{r}' - \mathbf{r}|} \rho'(\mathbf{r}, \mathbf{r}') \quad (3.45)$$

where erf is the error function. Helpfully, the screening of the exact exchange dampens the singularity from pure HF at the Fermi level, which allows it to treat states near the Fermi level more accurately. The mixing fraction $\alpha = 1/4$ was determined *ab initio* from perturbation theory [124], making this functional more generally applicable than some other hybrids, and the parameter $\omega = 0.2 \text{ \AA}^{-1}$ was found to be a good fit to several observables across many different semiconductors [124, 125].

3.2.7 Pseudopotentials

Another computational trick used in DFT that I will quickly discuss is pseudopotentials. One could perform the calculation including every electron in the unit cell, known as an all-electron calculation. However, from chemistry we know that only the valence electronic states will involve themselves in bonding, and what we call the ‘core’ electronic states will be, to a very good approximation, atomic. For this reason, most DFT codes work under the frozen core approximation, where they solve for the states of an isolated atom of each atomic species in the cell and reuse the solutions for the core electrons inside some cutoff radius r_c in the full solution. In this work, a core of helium was used for oxygen, and a core of argon was used for copper.

However, the speed up for plane-wave basis codes (such as CASTEP) is much more significant than just using fewer electrons. We can use the idea of valence and core electrons to calculate an effective potential for the valence electrons which reproduces the effect of the core electrons together with the nuclear potential. This will not result in exactly the same orbitals for the valence electrons as in an all electron code, but it can be done in a way where the orbitals are exactly the same outside the cutoff radius r_c , where all the bonding and interesting physics occurs. When we do this, the typical rapid oscillations that electron orbitals do near the

nuclei smooth out significantly, requiring far fewer plane waves in equation 3.33 to describe them. Figure 3.2 shows an example of this. The great reduction in the size of the basis set we have to use for the calculation comes with a speed up of several orders of magnitude. This method had lots of issues in early DFT calculations, but was only really fixed by 2016 when they became as accurate as all-electron methods [126].

There are two dominant schemes in CASTEP for generating pseudopotentials. Norm-conserving pseudopotentials (NCP) [127] require that the charge inside r_c integrates to the same value as in the all-electron case. Ultrasoft pseudopotentials [128] relax this constraint in order to achieve, as the name suggests, smoother wavefunctions still. Ultrasoft pseudopotentials are much more complicated to construct and to apply, but the time spent on doing so is worth it for the reduced basis set *in the case of local functionals*. For non-local functionals, it is much faster to use NCPs. In this work we used CASTEP's internal on-the-fly pseudopotential generator for both ultrasoft and norm-conserving pseudopotentials, using PBE for the atomic calculation.

3.2.8 Spin polarisation and spin-orbit coupling

Most DFT codes including CASTEP have the option whether to consider the effect of spin. However, as mentioned previously, since Cu_2O is not a magnetic material, the two spin channels are expected to be degenerate, so we can save computational resources and perform the calculation for half the number of electrons with the understanding that each band will be doubly occupied.

However, when there is an odd number of electrons (such as for a copper vacancy), or we wish to include spin-orbit coupling, the calculation must be spin-polarised. However, very briefly, if the calculation needs to be spin-polarised (such as if there is an odd number of electrons) the orbitals are also labelled by their spin and the system is solved twice, once for each spin. The KS potential then also includes a contribution dependent on the spin density. A good introduction to the details can be found here [115].

Under PBE, this always resulted in one of the doubly occupied bands straddling the Fermi level and becoming half occupied. However, the HF exchange energy is divergent at the Fermi level, causing significant spin splitting, which can be unphysical [129]. This plays a role in several of the defects, and must be taken into account when drawing conclusions from the HSE results.

Most of the calculations in this work were performed without spin-orbit coupling. The initial aim of the project was to establish whether or not a particular defect

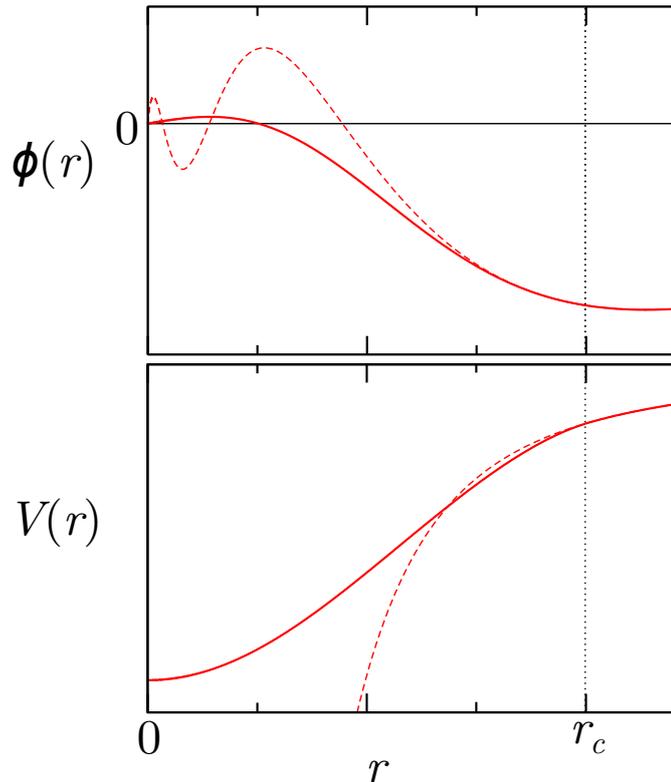


Figure 3.2: An example of a pseudopotential and its pseudo-wavefunction. The dashed line shows the rapidly oscillating valence orbital (top) resulting from the $1/r$ nuclear potential (bottom), and the solid line shows the smoothed valence orbital (top) that comes from the pseudopotential (bottom). The pseudopotential is designed so that both are the same outside the cutoff radius r_c .

produced a defect state in the band gap, with little emphasis placed on its exact energy. Since spin-orbit coupling makes a correction of ~ 0.1 eV to orbital energies, it was considered unlikely that it would affect the existence of a defect state, and so to save on computational cost it was not included.

To understand the effect it could have on the specific energies of the defect states, some of the calculations in this work include spin-orbit coupling. There are two methods for introducing spin-orbit coupling to calculations. The first is through relativistic DFT, and the second is with an appropriate pseudopotential. Fortunately, the second method is simpler, much faster to compute, and just as accurate [130] for our applications. For the spin-orbit calculations in this work, we use the CASTEP on-the-fly spin-orbit pseudopotential generator, using PBE for the atomic calculation.

3.3 Convergence with parameters

As discussed in section 3.2.5, the KS orbitals each have an associated k-point \mathbf{k} , and their periodic part can be decomposed into a Fourier series with wavevectors for each reciprocal lattice vector \mathbf{G} . While in principle there is a continuum of k-points and an infinite number of reciprocal lattice vectors, in practice we can only compute finitely many things. It is important that we include enough of each that our calculation is accurate but no more or it will take too long. The canonical way that the Brillouin zone is sampled for integration is with a Monkhorst-Pack (MP) grid [131], which in the cubic lattice is a regularly spaced cubic grid. Mathematically,

$$\rho(\mathbf{r}) = \int_{\text{BZ}} \Psi_{\mathbf{k}}^*(\mathbf{r})\Psi_{\mathbf{k}}(\mathbf{r})d\mathbf{k} \approx \sum_{\{\mathbf{k}\}} \Psi_{\mathbf{k}}^*(\mathbf{r})\Psi_{\mathbf{k}}(\mathbf{r})\Delta\mathbf{k} \quad (3.46)$$

where the sum that approximates the integral is over the k-points in the MP grid, and $\Delta\mathbf{k}$ is the spacing between them. Once the self-consistent ground state density and potential have been found with a sufficient sample of k-points, the potential can be used to obtain the ground state orbitals at any desired \mathbf{k} . The largest G-vector, \mathbf{G}_{cut} , included in the Fourier series is traditionally referred to by its energy,

$$E_{\text{cut}} = \frac{\hbar^2}{2m_e}G_{\text{cut}}^2, \quad (3.47)$$

known as the cutoff energy.

Figure 3.3 shows how the total final energy of a DFT calculation for the unit cell of Cu_2O in CASTEP changes with MP k-point grid density and cutoff energy. The change in energy of the Cu_2O unit cell is of order meV above a $7 \times 7 \times 7$ MP k-point grid, which for this work we consider well converged. The total energy is similarly converged for $E_{\text{cut}} = 1100$ eV for the ultrasoft pseudopotentials used with PBE. The norm-conserving pseudopotentials that must be used for HSE06 require a higher cutoff of $E_{\text{cut}} = 1600$ eV to be converged to meV total energy because they are less smooth than ultrasoft pseudopotentials. Unfortunately, due the available computational resources, such a high cutoff under a non-local functional was not feasible, so we had to settle for $E_{\text{cut}} = 1000$ eV, as well as a reduced k-point grid sampling of $4 \times 4 \times 4$. This converged the energy difference between the VBM and the conduction band minimum (CBM) to within 1 meV, indicating (but unfortunately not guaranteeing) that the band structure has stopped changing in a material way.

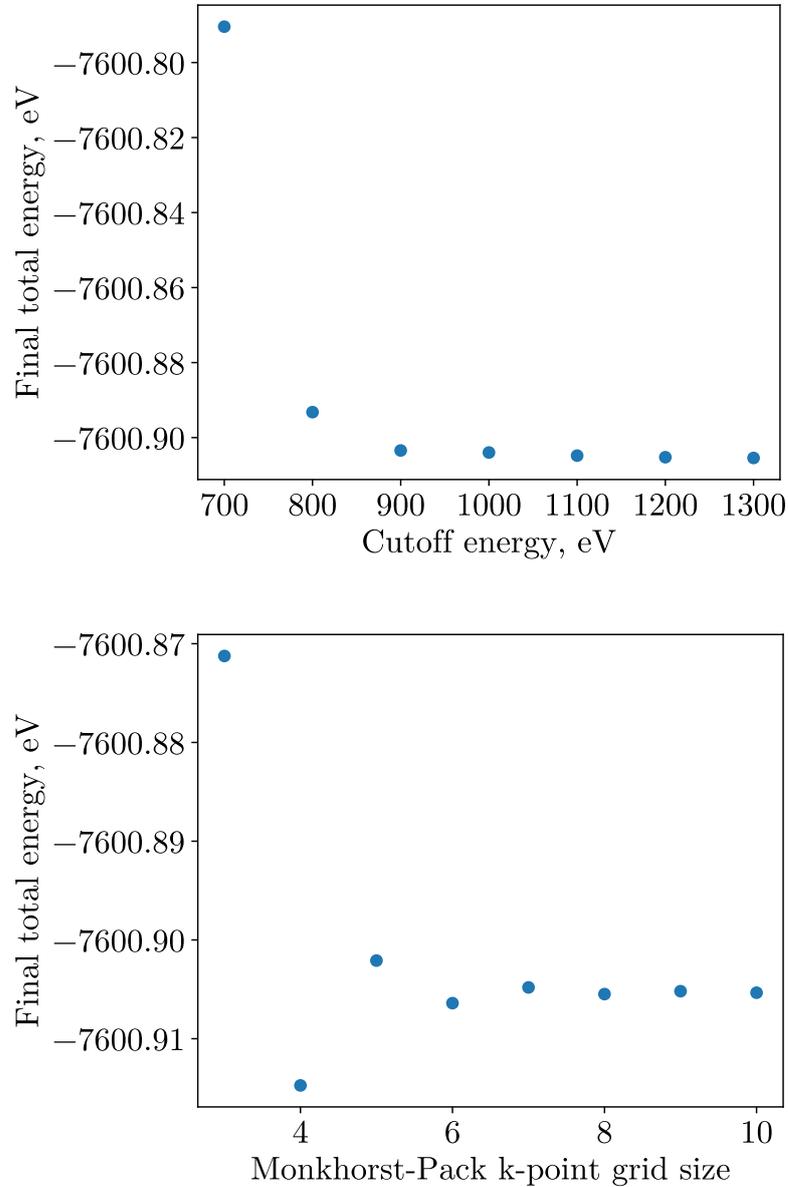


Figure 3.3: The convergence of the final energy of the Cu_2O unit cell for ultra-soft pseudopotentials. **Top:** plot shows how the final calculated energy changes with the length of the Fourier series used to represent the orbitals, which is traditionally measured with the energy of the highest frequency component, called the cutoff energy. **Bottom:** the dependence on the density of k-point sampling. The Monkhorst-Pack grid method samples the Brillouin zone in a regular cubic grid, and the x -axis shows the side length of the cube.

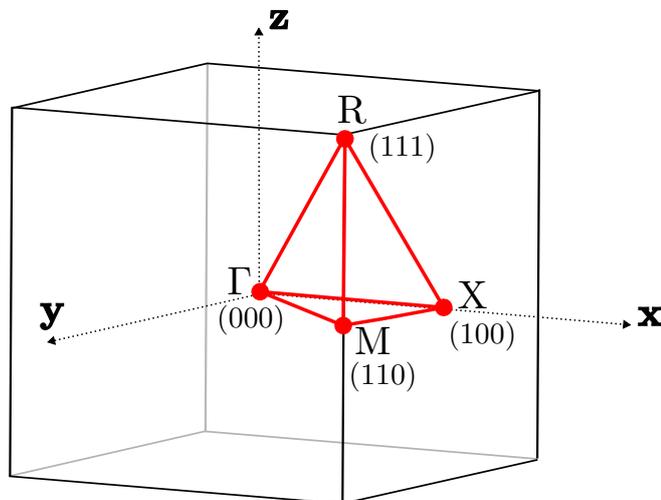


Figure 3.4: The high symmetry point labels in the first Brillouin zone for the primitive cubic lattice (the lattice for cuprous oxide). The Γ point is at $\mathbf{k} = (0, 0, 0)$. Band structures in this thesis follow paths along the red lines. The points R, M, and X represent $\mathbf{k} = (1, 1, 1) \pi/a$, $\mathbf{k} = (1, 1, 0) \pi/a$, and $\mathbf{k} = (1, 0, 0) \pi/a$, respectively.

3.4 DFT on pure Cu_2O

3.4.1 Band Structures

One of the most central tools used in condensed matter physics for understanding the electronic structure of a material is the band structure. Most of the results in this chapter are presented in that form, so it is important to know how to read one. The general idea is that we would like to show how the eigenenergies, $\epsilon_i^{\mathbf{k}}$, of the electrons, $\phi_i^{\mathbf{k}}(\mathbf{r})$, vary as a function of quasi-momentum, \mathbf{k} , and band, i . The accepted solution is to plot the eigenenergies along a path through \mathbf{k} -space. This path is found to be most illuminating if it follows piecewise some of the highest symmetry lines in the crystal, taking corners at the high symmetry points. Figure 3.4 shows the points of interest for a primitive cubic lattice such as Cu_2O . The point Γ always represents $\mathbf{k} = 0$ in all space groups, and is the most important point to study for us because in Cu_2O it is where the band gap occurs, so is where exciton physics is determined. The points R, M, and X represent $\mathbf{k} = (1, 1, 1) \pi/a$, $\mathbf{k} = (1, 1, 0) \pi/a$, and $\mathbf{k} = (1, 0, 0) \pi/a$, respectively. By symmetry, all the other high symmetry points are equivalent to these ones.

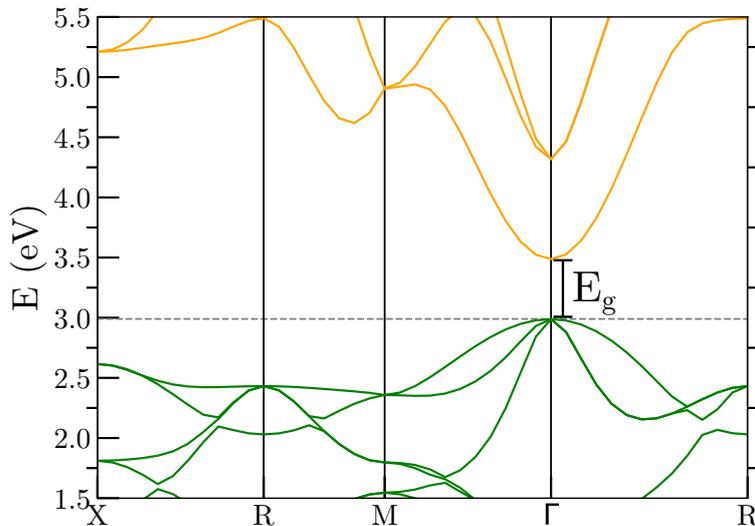


Figure 3.5: The band structure of a single Cu₂O unit cell under the PBE xc-functional. Occupied bulk states are shown in green, while unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The band gap, $E_g = 0.499$ eV, is at the Γ point, making Cu₂O a direct-gap semiconductor. Note that E_g as given by PBE-DFT is less than the experimental gap.

3.4.2 Perdew–Burke–Ernzerhof

Figure 3.5 shows an example of a band structure of Cu₂O produced with CASTEP under the Perdew–Burke–Ernzerhof (PBE) xc-functional, through the path X-R-M- Γ -R. Each Kohn-Sham orbital forms a different curve, sometimes also called a band. The occupied orbitals making up the valence band are shown in green, while the empty conduction band states are in yellow, with the Fermi energy marked by the dashed horizontal line. The smallest energy separation between an occupied and unoccupied orbital, called the band gap, $E_g = 0.499$ eV, occurs between two points at the same \mathbf{k} , making cuprous oxide a direct band gap semiconductor. It is important to note that the band gap given by PBE-DFT is less than the experimental band gap of 2.18 eV. Inaccurate band gaps are a problem inherent in Kohn-Sham DFT, as mentioned in 3.2.6. Such a significant underestimation of the band gap by PBE-DFT indicates that Cu₂O is a strongly correlated material. Despite this, PBE-DFT is known to reproduce well the shapes and orderings of bands, and so these are the things we are trying to extract from our simulations.

We can also learn a lot from the curvature of the bands. For a free electron,

$$E = \frac{\hbar^2 k^2}{2m}, \quad (3.48)$$

so we can identify the second derivative with the mass,

$$m = \hbar^2 \left(\frac{d^2 E}{dk^2} \right)^{-1}. \quad (3.49)$$

This is called effective mass. Since it is proportional to the inverse of the curvature, the flatter a band, the higher the effective mass of the electron, and the more localised the state. The effective mass can even be negative, as it is at the top of the valence band. This is interpreted as a positively charged, positive mass quasiparticle, i.e. the hole. Bands that are parabolic around Γ are most like free electrons, and their orbitals are delocalised across the crystal. This immediately provides an intuitive understanding for why Cu_2O is such a good platform for Rydberg excitons: the conduction electron and valence hole are nearly-free almost until the Brillouin zone edge in multiple directions, unusual for common semiconductors [132, 133], allowing for long range interactions akin to a free space hydrogen atom.

3.4.3 HSE06

Figure 3.6 shows the band structure for the pure Cu_2O unit cell under HSE06. As expected from a hybrid functional, the biggest difference from PBE is the size of the band gap. Here, $E_g^{\text{HSE}} = 1.60$ eV, comparable to the experimental gap $E_g = 2.18$ eV, but unsurprisingly not exact. The increase in band gap mostly comes from a constant increase in energy of the CBM, but without a change to its shape, giving confidence that we can still draw conclusions about exciton physics.

Scanlon et al. [30], the first to study Cu_2O defects with HSE-DFT, found a band gap of $E_g^{\text{HSE}} = 2.12$ eV by changing the fraction of HF exchange from $\alpha = 0.25$ to $\alpha = 0.275$. One must be careful when tuning functionals to produce experimental results. Given the systemic problems DFT often has with band gaps, there is no guarantee that a more accurate band gap will lead to better accuracy in any other property of the material. For this reason, we choose to keep the mixing fraction and screening length at their default values, and not put excessive weight on the quantitative results we achieve with HSE06 and instead focus on their qualitative nature.

3.5 Supercells

In order to study *defects* in Cu_2O in DFT, we need to construct larger cells that which contain multiple copies of the Cu_2O unit cell, called supercells. A $j \times k \times l$

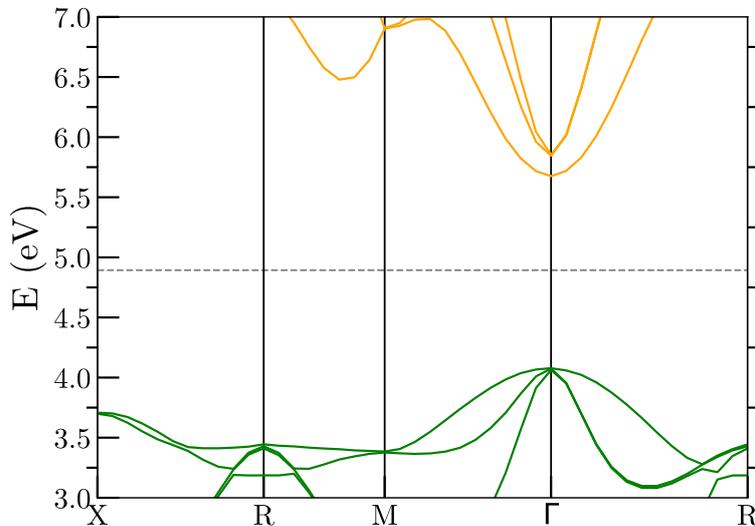


Figure 3.6: The Cu_2O unit cell band structure under HSE06. Occupied bands are in green and unoccupied bands are in yellow. The band gap is 1.60 eV.

supercell is a cuboid cell j unit cells wide, k unit cells deep, and l unit cells tall. If we were to naively introduce a defect into the 6-atom Cu_2O unit cell, for example a copper vacancy, we would introduce that vacancy into every periodic copy of the unit cell and create a crystal of Cu_3O_2 , but placing the defect in a supercell dilutes it into a bulk crystal structure.

For a baseline of comparison, figure 3.7 shows the band structures for pure Cu_2O under PBE in a $2 \times 2 \times 2$ and $3 \times 3 \times 3$ supercell, and under HSE06 in a $2 \times 2 \times 2$ supercell. There are several differences from the single unit cell to pay attention to. With increasing supercell size, the number of electrons increases (48 atoms worth in $2 \times 2 \times 2$, 162 atoms worth in $3 \times 3 \times 3$), creating many more bands. At the same time, the bands become much flatter and undergo an effect called band folding. Since the edge of the Brillouin zone occurs at π/a , when we make the cell twice as big on each side, we make the Brillouin zone twice as small on each side. One could imagine that the R point in the $2 \times 2 \times 2$ supercell occurs halfway along the $\Gamma - R$ line in the single unit cell. This makes the bands appear flatter, but the curvature is unaffected since the change in k along the k -point path is half as much as well. Also, due to the symmetry of k -space, bands are reflected from the new boundaries back into the Brillouin zone, adding yet more bands to the band structure.

Due to the condensing of k -space in supercells, the k -point sampling required for the same accuracy can be reduced as well. The scaling is actually very simple: for a supercell of $N \times N \times N$ unit cells, we need only a $K/N \times K/N \times K/N$ k -point

MP grid. Therefore, the k-point grids for the $2 \times 2 \times 2$ and $3 \times 3 \times 3$ PBE supercells were a $4 \times 4 \times 4$ MP grid and a $2 \times 2 \times 2$ MP grid respectively. For the HSE $2 \times 2 \times 2$ supercell we used $2 \times 2 \times 2$ k-points, which coincides with the k-point density that Isseroff and Carter [31] needed to find a defect state in the band gap for the split copper vacancy under HSE06.

3.6 The Defects

3.6.1 Searching for defect states

We are now ready to model the defects. Experimentally, the fractional deviation from stoichiometry in a sample is of the order 10^{-5} [134], which is equivalent to 1 point defect per $\sim 10^4$ unit cells. This level of dilution is not achievable with the resources available for this project. However, for the correct physics, the defects need only be dilute enough to not ‘feel’ the effects of neighbouring defects. To investigate where this dilute limit lies, we model each type of defect under PBE in a $2 \times 2 \times 2$ and $3 \times 3 \times 3$ supercell, and under HSE06 in a $2 \times 2 \times 2$ supercell. As will be seen, the decay on this finite-size effect can be inferred from the difference between the PBE results for the two sizes of supercell with and without the defect, and varies significantly between defects. Unfortunately, since HSE06 is a hybrid functional, it is much more expensive to compute, and so we were not able to run calculations under HSE06 for $3 \times 3 \times 3$ supercells, but the convergence of the PBE band structure with supercell size can often be used to infer the behaviour of HSE in a larger supercell.

As a reminder, the goal of this modelling is to look for defects that introduce states in the band gap, so that we can identify candidates for the photoluminescence lines in experiments. As discussed in section 3.2.6, however, DFT does not accurately reproduce band gaps, so we cannot directly compare the energies of the states with the energies of the emissions. Furthermore, differences in band energies are only a zero-th order approximation for excitation energies, as discussed in chapter 4. Therefore, we are looking for any state which falls within the band gap as calculated for the given functional in the bulk crystal.

In all the defect supercell bandstructures the states deep into the valence and conduction bands will look much denser than in the pure crystal. This is because in the pure crystal the high degree of symmetry causes many of the bands to be degenerate. However, the defect breaks symmetries in the crystal, leading to small splittings of most bands into several states. These kinds of perturbations will not

make a significant difference to the behaviour of Cu_2O , especially not from the perspective of excitons and photoluminescence.

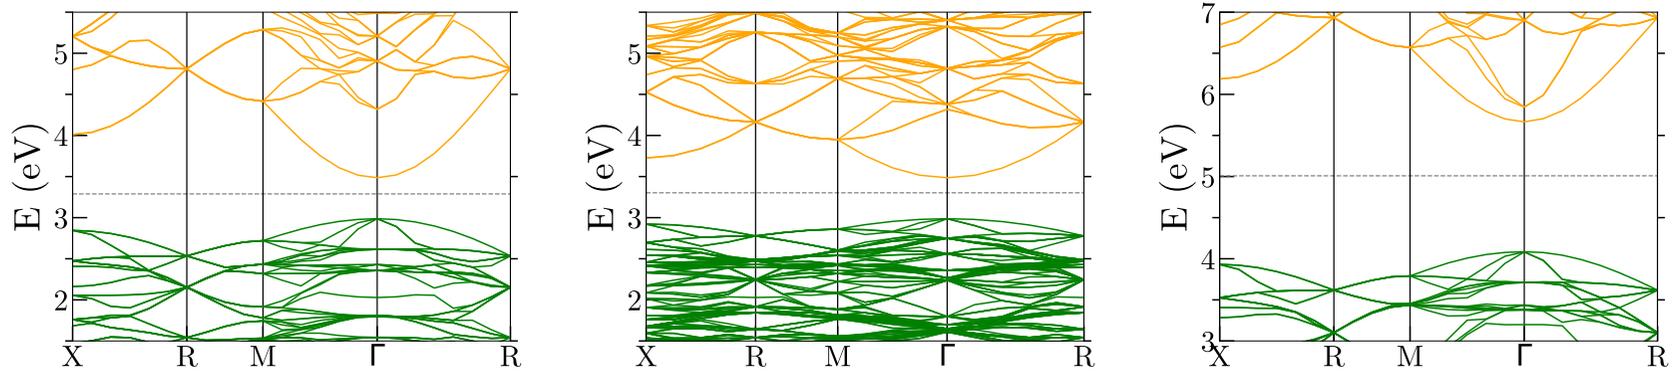


Figure 3.7: Band structures for pure Cu_2O in supercells. From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied states are shown in green and unoccupied states are shown in yellow. The dotted line denotes the Fermi level.

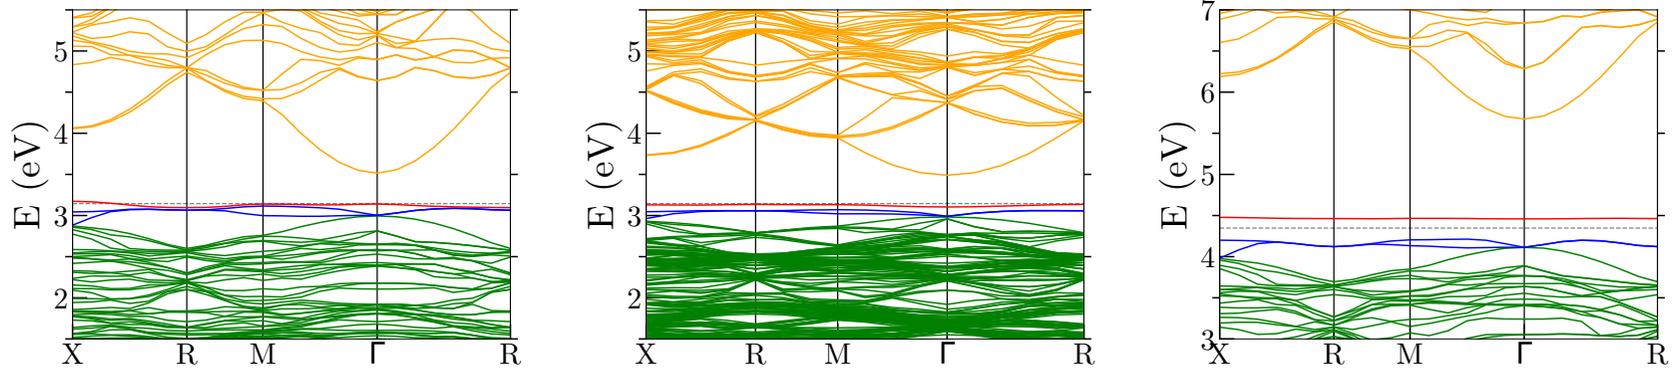


Figure 3.8: Band structures for the octahedral form of the oxygen interstitial, O_i^{oct} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The occupied defect states are shown in blue and the unoccupied ones are shown in red.

3.6.2 Oxygen interstitials

To begin with the most positive results, the band structures for the two geometries of the oxygen interstitials are shown in figures 3.8 and 3.9. The occupied bulk states are shown in green and the unoccupied bulk states are shown in yellow. The occupied defect states are shown in blue and the unoccupied defect states are shown in red. Both geometries for these interstitials show clear defect states in the band gap across both supercell sizes and for both PBE and HSE xc-functionals. The robustness of these results is stark when compared to the other defects. Both geometries also produce 3 well localised spin-degenerate states in the gap, 2 occupied and 1 unoccupied, which are easily understood as the 2P electrons of the additional oxygen atom.

As can be seen for O_i^{oct} (fig. 3.8), the unoccupied defect state remains deep in the band gap for the $3 \times 3 \times 3$ supercell. Along with the fact that it is so flat, meaning it is very well localised, we can be confident it is not a consequence of neighbouring defects interacting with each other but a genuine, additional, local electronic state that would exist in the dilute defect limit. The two occupied defect states for O_i^{oct} are degenerate with the VBM at Γ . Across all three calculations the defect had no effect on the band gap, lending further evidence that the states would not move in relation to the band gap in the dilute limit.

In contrast, figure 3.9 shows that all three defect states introduced by O_i^{tet} exist firmly inside the bulk band gap and would produce two distinct photoluminescence lines in experiments. All three states localise further from the $2 \times 2 \times 2$ supercell to the $3 \times 3 \times 3$ supercell under PBE, and their similarity under HSE suggests they would do the same. The band gaps in the $2 \times 2 \times 2$ supercells are slightly larger than their bulk values, but in the $3 \times 3 \times 3$ supercell the band gap returns to the bulk PBE value, so we can assume the same for HSE.

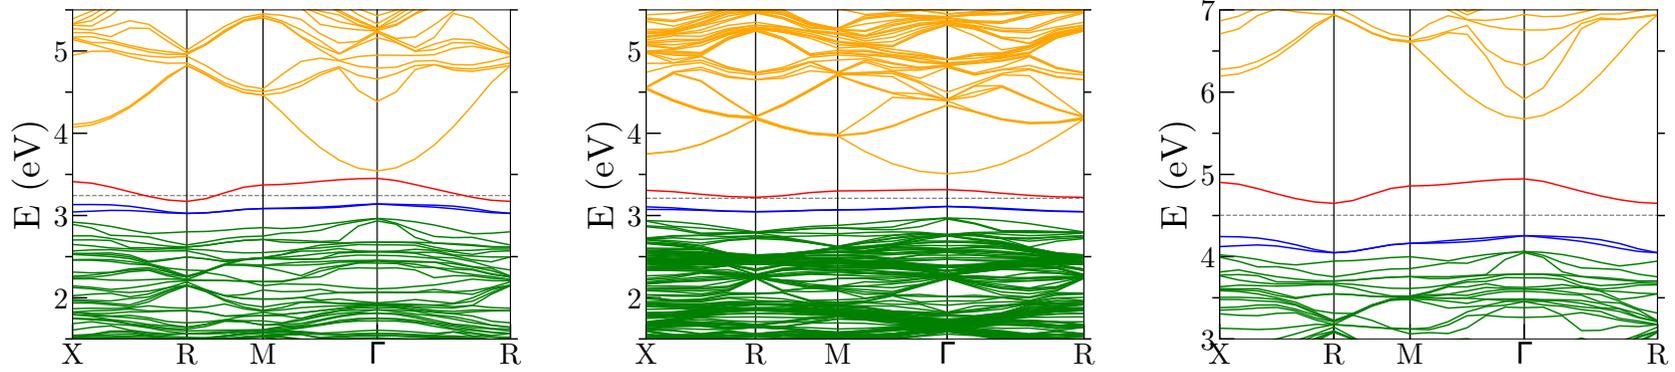


Figure 3.9: Band structures for the tetrahedral form of the oxygen interstitial, O_i^{tet} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The occupied defect states are shown in blue and the unoccupied ones are shown in red.

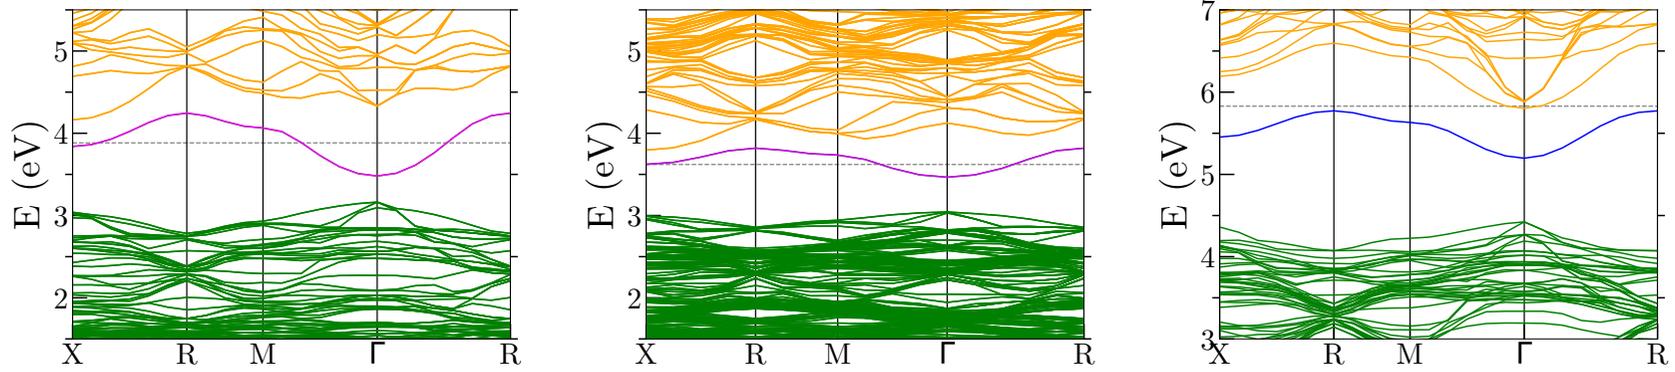


Figure 3.10: Band structures for the octahedral form of the copper interstitial, Cu_1^{oct} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. Occupied defect states are shown in blue and half-occupied ones are in purple. The dotted line denotes the Fermi level.

3.6.3 Copper interstitials

The copper interstitials, both in the octahedral (figure 3.10) and tetrahedral (figure 3.11) arrangements, produce similar distortions from the bulk structure, so we will discuss them together.

Under PBE, for the $2 \times 2 \times 2$ supercell, a half-occupied state appears at $E_g = 0.32$ eV for the octahedral arrangement and $E_g = 0.34$ eV for the tetrahedral arrangement, well within the bulk band gap at Γ . The state has a very wide dispersion, indicating a large extent in real space, and so a large finite-size effect. Comparing these band structures to the bulk crystal, we also see that the CBM is missing, so we could think that these defect states are simply a perturbation of the CBM that would return to bulk behaviour in the dilute limit.

Interestingly, the defect states are somewhat robust to increasing the supercell size, indicating that this state could exist even in the dilute defect limit. The band increases to $E_g = 0.42$ eV for octahedral and $E_g = 0.40$ eV for tetrahedral for the $3 \times 3 \times 3$ PBE supercell, still less than but closer to the bulk band gap. Comparing again to the bulk $3 \times 3 \times 3$ PBE supercell band structure, a band appears at the bottom of the conduction band which could be the bulk CBM returning to its bulk behaviour. One can imagine two scenarios for the dilute limit under PBE: the defect state rejoins the conduction band as the bulk CBM; or the state remains in the band gap and the bulk CBM emerges from the conduction band. A further study of the $4 \times 4 \times 4$ supercell would be needed to distinguish between these two.

The HSE analysis gives us a different angle to understand these defects. The half-occupied state predicted by PBE is split by Hartree-Fock exchange into a lower occupied spin state and an upper unoccupied spin state. The occupied spin states are 0.93 eV and 0.98 eV above the VBM for Cu_{oct}^i and Cu_{tet}^i respectively, and the unoccupied ones not only resemble the bulk CBM but are exactly the HSE bulk gap above the VBM.

Our interpretation of these results is that the interstitial copper atoms provide an additional, local binding of one of the CBM states, leading to a local state in the band gap, near in energy to the CBM, which extends over several unit cells in real space.

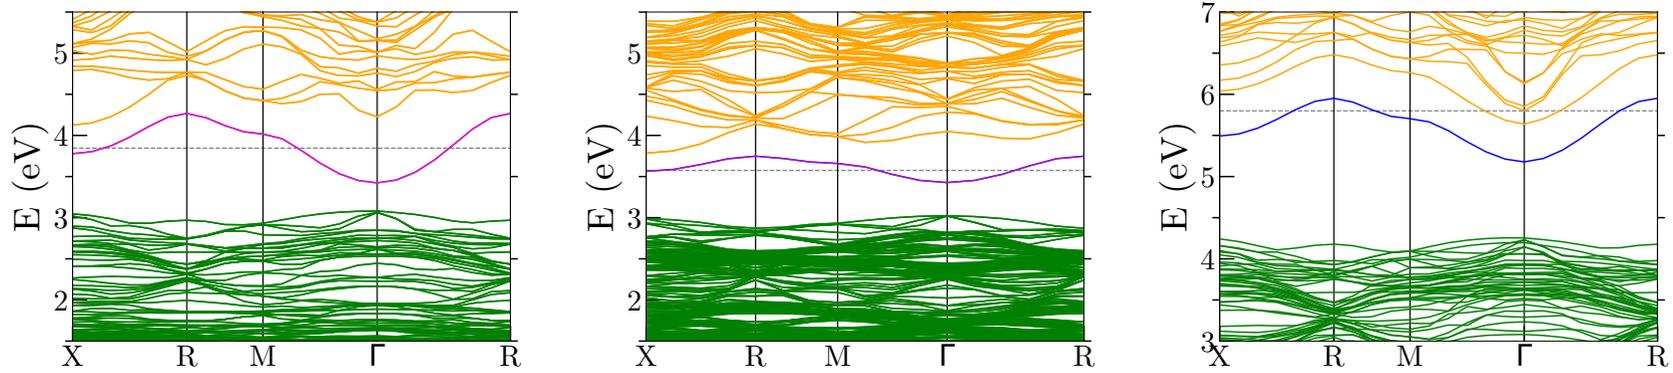


Figure 3.11: Band structures for the tetrahedral form of the copper interstitial, Cu_i^{tet} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level.

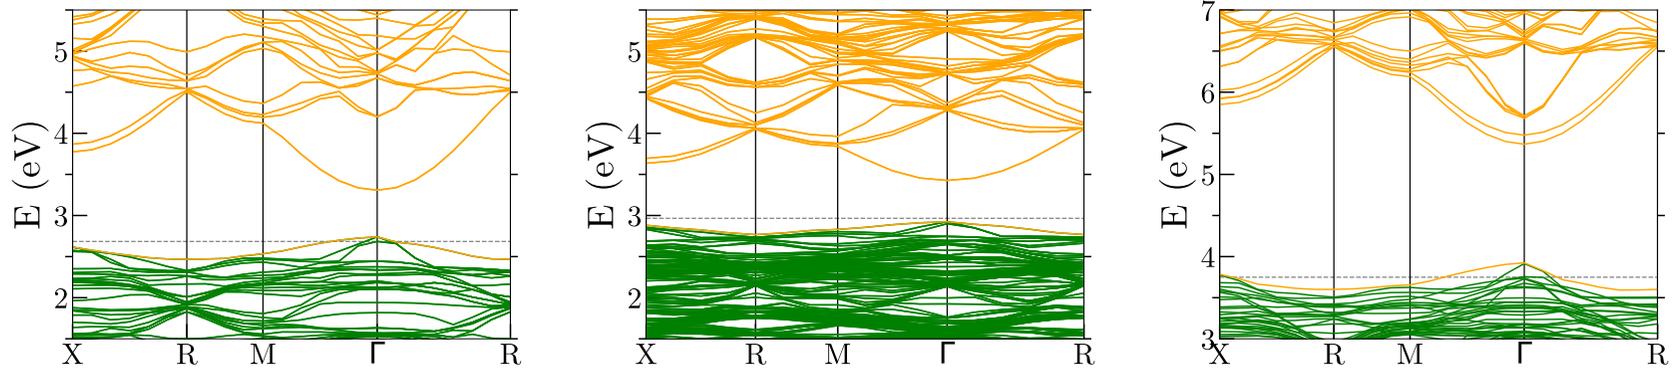


Figure 3.12: Band structures for the simple copper vacancy, V_{Cu} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The simple copper vacancy does not exhibit defect states in the band gap.

3.6.4 Copper vacancy

The band structures for the simple copper vacancy for $2 \times 2 \times 2$ and $3 \times 3 \times 3$ supercells under PBE and for the $2 \times 2 \times 2$ supercell under HSE06 are seen in figure 3.12. The occupied states are shown in green and the unoccupied states are shown in yellow. Across all three band structures, we see V_{Cu} does not introduce a clear electronic state in the band gap. In the smaller supercell under both PBE and HSE06 we see that V_{Cu} perturbs the valence band such that the VBM becomes slightly peaked. Were we only to look at the eigenvalues at the Γ point, it would be easy to believe that a defect state has appeared in the band gap just above the valence band. However, when we dilute the defect further in the $3 \times 3 \times 3$ supercell under PBE, we see the structure return to be identical to the valence and conduction bands in the bulk crystal (fig. 3.7), demonstrating that the peaking of the valence band is merely a finite-size effect caused by an unphysical effective density of defects. Given the similarity between the PBE and HSE upper valence bands in the $2 \times 2 \times 2$ supercell, we can be confident that HSE is not introducing new physics which is absent from the PBE level of analysis, and that the peak in the valence band under HSE is also a finite-size effect.

Across all three band structures we also see that the top of the valence band is partially unoccupied. This is expected, since the removal of a copper atom leaves an odd number of electrons. Under PBE the VBM is spin degenerate and half occupied, but under HSE the top state is actually only of one spin channel, a consequence of the spin splitting caused by Hartree-Fock exchange near the Fermi level. Neither effect would produce photoluminescence lines of energy less than E_g , but they are responsible for much of the p-type conductivity observed in the semiconductor literature on cuprous oxide [135].

3.6.5 Split copper vacancies

Aside from the simple copper vacancy, there are two important copper vacancy complexes known as *split vacancies*. They are formed of two nearest-neighbour copper vacancies, with an interstitial copper atom placed at the midpoint between the two vacancies. Up to symmetry, there are two distinct choices to place the nearest-neighbouring copper vacancies: two copper atoms bonded to different oxygen atoms, or two copper atoms bonded to the same oxygen atom.

Under PBE, the band structure for the first split copper vacancy (where the copper atoms were bonded to different oxygen atoms), seen in figure 3.13, looks very similar to the simple copper vacancy and the second split copper vacancy. The

only difference is that the VBM for $V_{\text{Cu}}^{s;1}$ is raised at R to the same energy as at Γ in the $2 \times 2 \times 2$ supercell. The VBM mostly restored to that of the bulk in the $3 \times 3 \times 3$, but even if this persists in the dilute limit it will not have an effect on the photoluminescence. Like the other copper vacancies, the VBM also becomes partially unoccupied, contributing to native p-type conductivity, but not to photoluminescence.

Under HSE, a defect state appears in the band gap 0.35 eV above the valence band maximum. The band is very flat and well localised, and only exists for one of the spin channels. This is in excellent agreement with Isseroff and Carter [31], who found a defect state with only one allowed spin with a peak in the density of states 0.57 eV above the VBM. The difference in energy is perhaps due to the lack of spin orbit coupling in this work. The fact that this state only exists for one spin channel suggests a similar concern as for the copper interstitials. As discussed in section 3.2.6, hybrid functionals like HSE struggle to accurately treat partially-occupied states, since they encounter divergences in the derivative of the energy at the Fermi level, so this defect state could be an artefact. However, the other two copper vacancies have a partially-occupied VBM and do not exhibit such an artefact under HSE, which encourages us to believe the defect state for $V_{\text{Cu}}^{s;1}$ is physical and due to some physics captured by HSE not present under PBE.

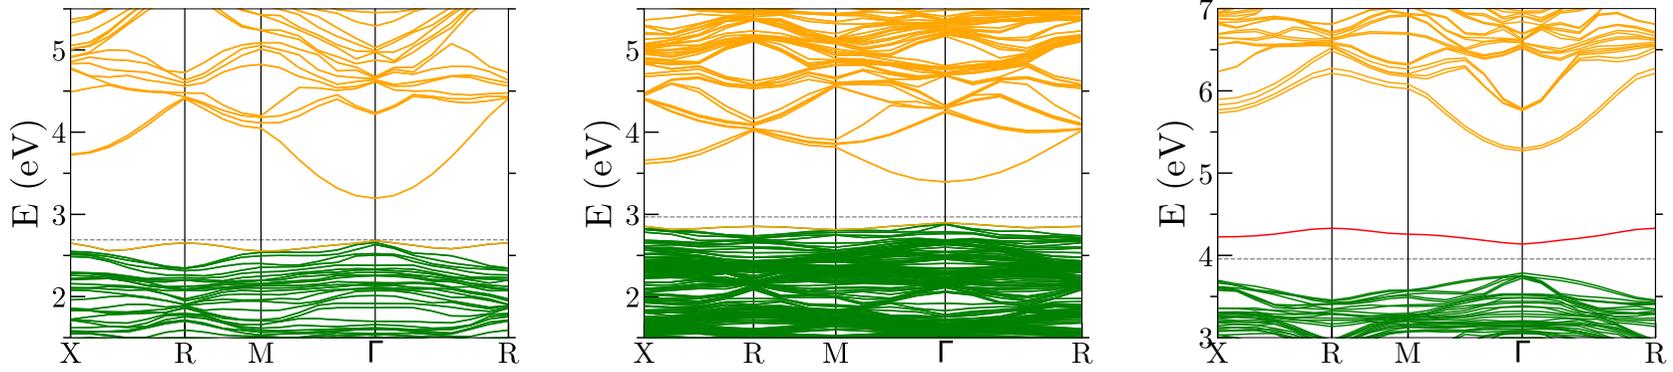


Figure 3.13: Band structures for the first form of the split copper vacancy, $V_{\text{Cu}}^{\text{s},1}$. From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green, unoccupied bulk states are shown in yellow, and unoccupied defect states are shown in red. The dotted line denotes the Fermi level.

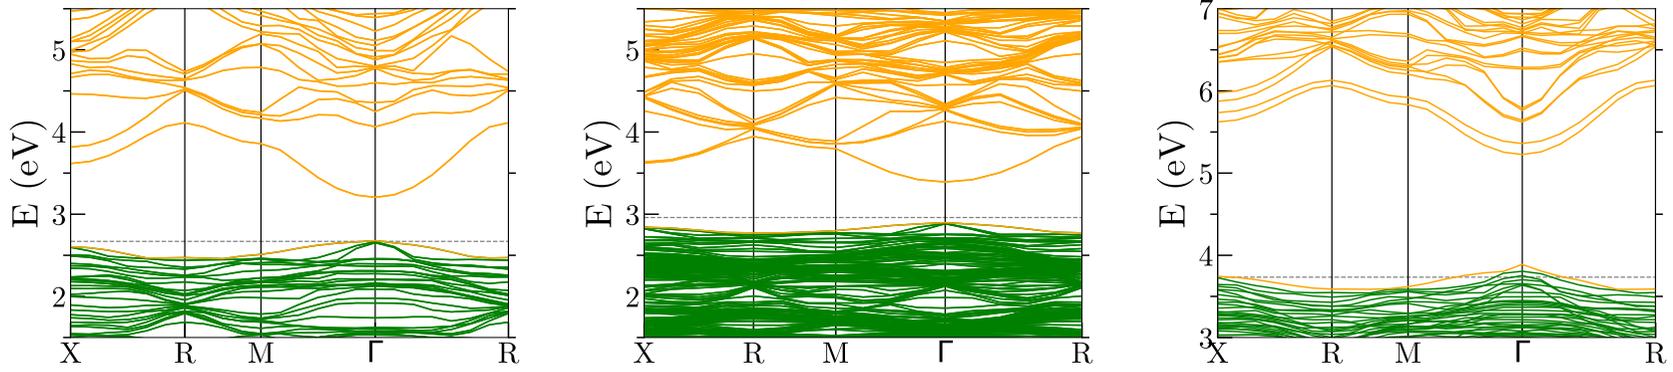


Figure 3.14: Band structures for the second form of the split copper vacancy, $V_{\text{Cu}}^{\text{s},2}$. From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The second form of the split copper vacancy does not exhibit defect states in the band gap.

The band structures for the second form of the split copper vacancy, $V_{\text{Cu}}^{\text{s},2}$ (where the two copper atoms were bonded to the same oxygen atom), for $2 \times 2 \times 2$ and $3 \times 3 \times 3$ supercells under PBE and for the $2 \times 2 \times 2$ supercell under HSE06 are seen in figure 3.12. Across all three band structures, as with V_{Cu} , we see $V_{\text{Cu}}^{\text{s},2}$ does not introduce a clear electronic state in the band gap. In fact, the effect of the defect on the bulk structure is similar to that of V_{Cu} , but here the peak in the VBM is only seen under HSE. One difference from the simple copper vacancy is that the CBM has separated from the conduction band, breaking its degeneracy at the high symmetry points with the second-lowest conduction state. However, since the degeneracy is almost completely restored in the $3 \times 3 \times 3$ supercell, we can safely conclude that this is a finite size effect. Given the similarity in the HSE case, we can also conclude that the split CBM would rejoin the conduction band proper in the dilute defect limit. Just like the simple copper vacancy, $V_{\text{Cu}}^{\text{s},2}$ partially depopulates the valence band, which does not result in photoluminescence peaks but does cause p-type conductivity.

3.6.6 Oxygen vacancy

Figure 3.15 shows the supercell band structures for the oxygen vacancy, V_{O} . This defect presents no defect states in the band gap across the three levels of analysis, and so could not produce the photoluminescence lines seen in experiment. In fact, V_{O} significantly opens the band gap to $E_g = 0.78$ eV (PBE, $2 \times 2 \times 2$) by flattening the CBM, however, it does return to $E_g = 0.56$ eV in the $3 \times 3 \times 3$ supercell. While it is hard to be certain how dilute the defect would need to be for the gap to return to its bulk value, we can be certain that larger supercells will not introduce states inside the band gap if they are not seen in these artificially high defect concentrations.

The behaviour of the defect is qualitatively similar under HSE, with a greater flattening of the CBM leading to an enlargement of the gap to $E_g = 1.83$ eV. With a shape very similar to that of the PBE CBM in the $2 \times 2 \times 2$ supercell, we expect this effect also to lessen with increasing supercell size. Regardless, as for PBE, given that there is no additional physics that would give rise to photoluminescence at extreme defect densities, none will arise in the dilute limit.

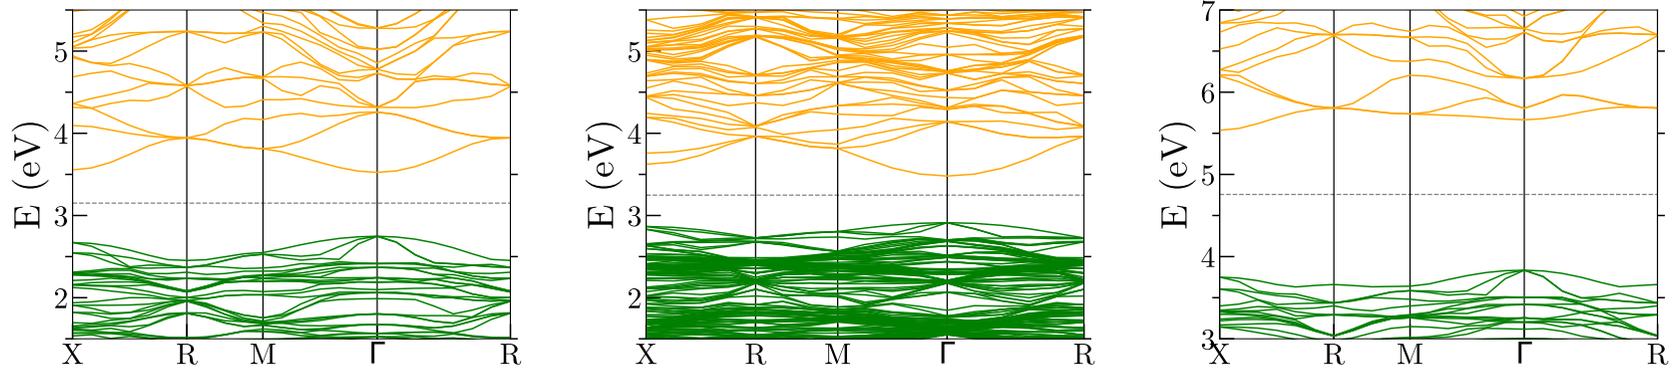


Figure 3.15: Band structures for the oxygen vacancy, V_{O} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The oxygen vacancy does not exhibit defect states in the band gap.

3.6.7 Anti-sites

3.6.7.1 Copper replacing oxygen

The copper anti-site Cu_O introduces a myriad of defect states in the band gap, as seen in figure 3.16, though most of them are degenerate with the VBM or CBM at Γ . At the PBE level, there is one (doubly spin degenerate) defect state inside the band gap at the Γ point, and it persists from the $2 \times 2 \times 2$ supercell to the $3 \times 3 \times 3$ supercell. It also localises significantly in the $3 \times 3 \times 3$ supercell, which suggests that the state would persist in the dilute limit.

However, Cu_O also greatly increases the band gap under PBE to 1.05 eV in the $2 \times 2 \times 2$ supercell. In the $3 \times 3 \times 3$ supercell the band gap relaxes to 0.72 eV, suggesting a very slow convergence to the dilute limit, even slower than the copper interstitials. Since the defect state goes from 0.17 eV below the CBM at $2 \times 2 \times 2$ to 0.12 eV below at $3 \times 3 \times 3$ supercell size, it is not clear whether in the dilute limit it would become essentially degenerate with the CBM. At the HSE level, the inclusion of Hartree-Fock exchange splits the states by spin, such that 3 states could produce meaningful photoemission. However, given how slowly this defect converges to the dilute limit in the PBE data, it is not possible to draw meaningful conclusions about how much of the HSE behaviour is a finite size effect.

The perturbation of the VBM that has been labelled as a defect state in the $2 \times 2 \times 2$ PBE band structure is seen to return to bulk behaviour in the $3 \times 3 \times 3$ supercell. The perturbation of the VBM is larger under HSE06, and has a component of spin-splitting at the edges of the Brillouin zone, but it is very similar in shape to the perturbation under PBE, so there is no strong reason to believe it would not return to bulk behaviour in the dilute limit.

3.6.7.2 Oxygen replacing copper

The oxygen anti-site O_Cu , seen in figure 3.17, produces remarkably similar results to $V_\text{Cu}^{\text{s};2}$. The rest of the conduction band separates slightly from the CBM while retaining its delocalisation, and like $V_\text{Cu}^{\text{s};2}$ it returns to bulk behaviour in the $3 \times 3 \times 3$ supercell. The separation of the CBM looks very similar at the HSE level, and we expect it to return to bulk behaviour in the dilute limit as well.

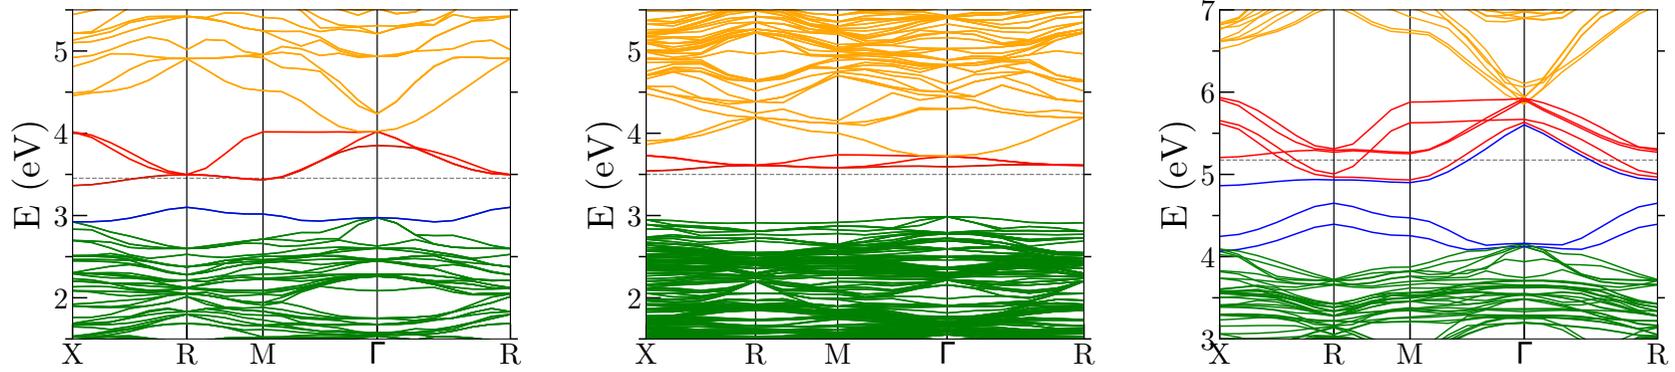


Figure 3.16: Band structures for the copper anti-site, Cu_O . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The occupied defect states are shown in blue and the unoccupied ones are shown in red.

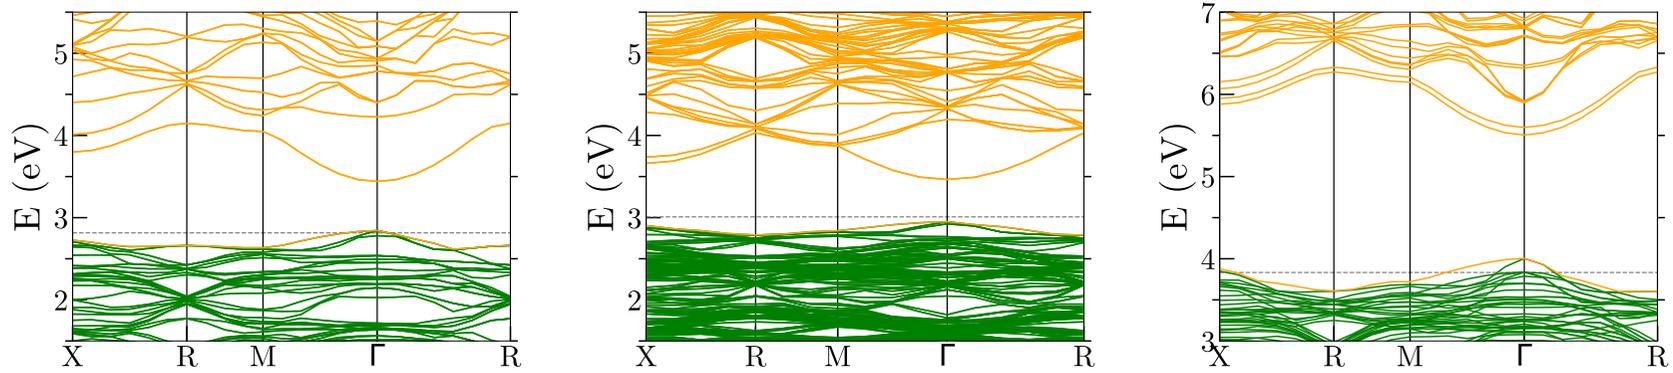


Figure 3.17: Band structures for the oxygen anti-site, O_{Cu} . From left to right: PBE band structure in a $2 \times 2 \times 2$ supercell; PBE band structure in a $3 \times 3 \times 3$ supercell; HSE06 band structure in a $2 \times 2 \times 2$ supercell. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The oxygen anti-site does not produce defect states in the band gap.

Also like $V_{\text{Cu}}^{s;2}$ and V_{Cu} , the removal of a copper atom removes one electron from the VBM. Copper has an odd number of electrons, whereas oxygen has an even number, so O_{Cu} leaves one electron unpaired. The absence of 10 copper d electrons allows space for the new oxygen valence electrons to sit below the VBM, unlike in the oxygen interstitials, where they are the highest energy occupied orbitals. The band gap in the PBE $2 \times 2 \times 2$ supercell opens slightly to $E_g = 0.59$ eV, but returns to the bulk value in the $3 \times 3 \times 3$ supercell. The band gap under HSE opens to $E_g = 1.66$ eV, so even if the peak in the VBM were to persist in the dilute limit it would not produce photoluminescence lines below the band gap energy in experiment.

3.6.8 The effect of charge

Since much work is being done at the moment on the effect of charged defects on Rydberg excitons [26, 27], it is important to understand the stability of these defect states to the effect of adding or subtracting electrons. We computed the band structures of each defect under PBE in both supercell sizes for $\{N - 2, \dots, N + 2\}$ electrons, where N is the neutral number of electrons. For the oxygen interstitials, we found all 3 defect levels shifted rigidly up for each additional electron and down for each removed electron, indicating that the different charge states of these defects could show distinct peaks in the PL spectrum. The shift was ~ 0.5 eV for O_i^{tet} (pushing the unoccupied defect level above the CBM) and ~ 0.2 eV for O_i^{oct} . For every other defect, we found a change of ± 2 in the number of electrons had no effect on the band structure. Unfortunately, due to the significantly increased computational cost and unstable convergence of HSE, we did not have time to investigate the charge states under HSE. This would be important for future work especially in the case of the copper interstitials and $V_{\text{Cu}}^{s;1}$, where the defect level predicted by HSE could be dependent on HSE's treatment of the half-occupied defect state.

3.6.9 Spin-orbit coupling

As mentioned in section 3.2.8, most calculations in this work were performed without spin-orbit coupling. This was to save on a 4x increase in computational cost, as it was not expected to have an effect on the existence of any defect states. It does, however, have an effect on the placement of the defect states within the band gap.

Figure 3.18 shows the band structures with spin-orbit coupling for pure Cu_2O in the unit cell, O_{tet}^i ($2 \times 2 \times 2$), and Cu_{oct}^i ($2 \times 2 \times 2$), compared to their band structures

without spin-orbit coupling. One can see the effect of spin-orbit coupling on the unit cell is to split the degeneracy of the upper valence band, resulting in a band gap ~ 0.2 eV smaller than that without spin-orbit coupling. In the defect supercells, the change in the band gap changes the energy at which the defect state appears within the band gap, raising the defect level by ~ 0.1 eV for O_{tet}^i and lowering the defect level of Cu_{oct}^i by ~ 0.2 eV. However, it has almost no effect on the defect state itself. The fact that the upper defect state for O_{tet}^i crosses the CBM at Γ is not a cause for concern. Since the curvature of this band is high in the $2 \times 2 \times 2$, it will flatten out closer to its average energy in the dilute limit (as seen in figure 3.9). Given the size of the shift in defect level energy, an additional error bar of ± 0.1 eV should be applied to calculations without spin-orbit coupling.

All three calculations are performed under PBE and only up to the $2 \times 2 \times 2$ supercell, since the calculations are prohibitively expensive otherwise. Since spin-orbit calculations involve different pseudopotentials, the zero-energies are not the same as for calculations without spin-orbit coupling. Therefore, both sets of band structures are plotted here as their energy differences from the Fermi energy, to make them easily comparable.

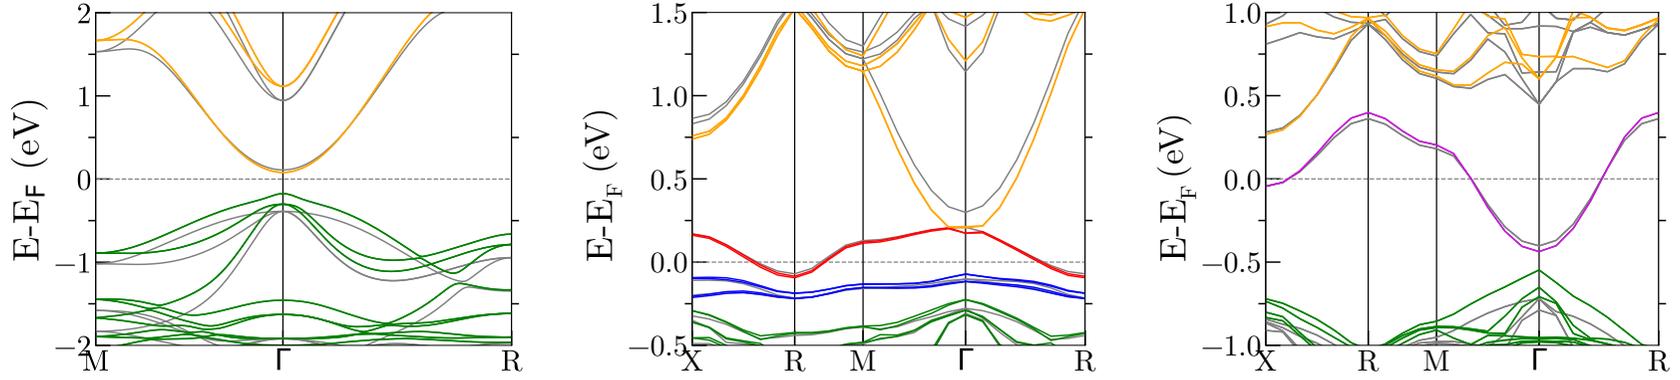


Figure 3.18: The effect of spin-orbit coupling on band structures. From left to right: Pure Cu_2O in the unit cell; O_{tet}^i in a $2 \times 2 \times 2$ supercell; Cu_{oct}^i in a $2 \times 2 \times 2$ supercell. All four calculations were under the PBE xc-functional. Occupied bulk states are shown in green and unoccupied bulk states are shown in yellow. The dotted line denotes the Fermi level. The occupied defect states are shown in blue, the unoccupied ones are shown in red, and half-occupied defect states are shown in purple. The grey states depict the band structure of the cell without spin-orbit coupling, for comparison.

	V_{Cu}	$V_{\text{Cu}}^{s;1}$	$V_{\text{Cu}}^{s;2}$	V_{O}	Cu_i^{oct}	Cu_i^{tet}	O_i^{oct}	O_i^{tet}	CuO	O_{Cu}
ΔH_F , eV	0.48	0.80	1.59	2.47	2.30	2.67	0.20	-0.24	6.17	2.07
PBE $2 \times 2 \times 2$	No	No	No	No	Yes	Yes	Yes	Yes	Yes	No
PBE $3 \times 3 \times 3$	No	No	No	No	Maybe	Maybe	Yes	Yes	Yes	No
HSE06 $2 \times 2 \times 2$	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes	No

Table 3.1: The enthalpy of formation of the defects studied in this work, along with whether or not the given level of theory predicts a defect state in the band gap. The energies are calculated from the energies of the pure and defective $3 \times 3 \times 3$ supercells under PBE.

3.6.10 Formation energies

To assess how bright the photoluminescence from each defect would be, we need to know their relative concentration in the crystal. The formation enthalpy, $\Delta H_F(D)$, of a neutral defect D , tells us how easy or hard it is to form the defect in the perfect crystal, where $\Delta H_F(D) < 0$ indicates a preference to form the defect and $\Delta H_F(D) > 0$ indicates a barrier, and so should give us a heuristic for concentration. In general, it is given by

$$\Delta H_F(D) = E_D - E_P + \sum_i n_i (\mu_i^{\text{elem}} + \Delta\mu_i), \quad (3.50)$$

where E_D is the energy of the supercell with a defect, E_P is the energy of the perfect supercell, and n_i is the number of added (negative n_i) or subtracted (positive n_i) atoms of species i and chemical potential μ_i^{elem} . The additional term $\Delta\mu_i$ changes the chemical potentials depending on growth conditions, and other authors take this into account in their analysis [30, 106, 31]. For simplicity, in this work we do not include this effect, and set $\Delta\mu_{\text{Cu}} = \Delta\mu_{\text{O}} = 0$. The elemental chemical potentials $\mu_{\text{Cu}}^{\text{elem}}$ and $\mu_{\text{O}}^{\text{elem}}$ are obtained from the energy per atom of the ‘natural’ state of the elements, in this case O_2 and metallic Cu. For consistency, their geometries were first optimised under PBE, giving a lattice constant of 3.63 Å for Cu and a bond length of 1.21 nm for O_2 , resulting in $\mu_{\text{Cu}}^{\text{elem}} = 1680.93$ eV/atom and $\mu_{\text{O}}^{\text{elem}} = 436.80$ eV/atom.

Table 3.1 shows the formation enthalpies as calculated from the $3 \times 3 \times 3$ supercells under PBE. They are in reasonable agreement with literature [106, 30], keeping in mind the lack of charge and growth conditions analysis. Many of the formation enthalpies can be intuitively described from a chemistry perspective. Since CuO is a more stable copper oxide than Cu_2O , it is easiest to add oxygen atoms or

remove copper atoms, so the oxygen interstitials and copper vacancies are the lowest energies and will occur most commonly in a sample. The second split vacancy is slightly higher in energy than the other two because it is harder to remove the second copper atom from a given oxygen atom than the first. The oxygen vacancy and the copper interstitials are a harder class of defect to form as they increase the ratio of copper to oxygen, moving further away from CuO, and so will be less common but still noticeable. The anti-sites are significantly more expensive than their interstitial counterparts as they involve the formation of a defect and the insertion of an atom into an unfavourable location. The combination makes Cu_O prohibitively expensive to form.

3.7 Discussion and conclusions

There is no evidence within our DFT calculations that the oxygen vacancy V_{O} , simple copper vacancy V_{Cu} , second split copper vacancy $V_{\text{Cu}}^{s;2}$, or oxygen anti-site O_{Cu} give rise to any states within the bulk band gap, nor would produce any signal in photoluminescence experiments. In fact, given their small perturbations in $2 \times 2 \times 2$ supercells for both xc-functionals and lack of effect on the band structure of the $3 \times 3 \times 3$ supercell, it is unlikely they would have any effect on excitons at all in the dilute limit. Therefore, there appears to be no reason to label the photoluminescence peaks observed in experiments [104, 103, 100, 101, 102, 28] with either V_{Cu} or with V_{O} .

DFT provides strong evidence that the two possible geometries of the oxygen interstitial would result in one (octahedral) or more (tetrahedral) electronic states midway into the bulk band gap. Their low or even negative enthalpy of formation indicates they would produce strong emission lines in photoluminescence in defective samples grown in most conditions. Given how localised these states are, it is also likely that their energies will be sensitive to the charge state of the defect. Further analysis of the effect of charge could be fruitful.

We are also confident that the $V_{\text{Cu}}^{s;1}$ defect would produce a defect state in the band gap, and therefore a peak in the photoluminescence. It's low formation energy would make it a common defect, especially in samples grown in copper-deficient conditions. However, the defect state only supports one electron per defect, as opposed to the two (spin up and spin down) electrons supported by each level from the oxygen interstitials, so the $V_{\text{Cu}}^{s;1}$ peak may be somewhat suppressed.

Despite being well under-converged in supercell size in this study, it is quite possible that the copper anti-site could introduce a defect state into the band gap in the

dilute limit. Nevertheless, its prohibitively high formation energy would prevent it from contributing any signal to the photoluminescence spectrum.

It is difficult to draw concrete conclusions regarding the copper interstitials. The effect of the defect is too long range to be properly converged in the $3 \times 3 \times 3$ PBE supercell, so it is difficult to say whether the half-occupied band is a perturbation of the CBM or a poorly localised defect state. The much stronger exchange treatment in HSE06 splits this band by spin and populates fully one of the spin channels, but hybrid functionals are known to do poorly with metal-like states such as this one [129]. Given how strongly these defects interact with neighbouring defects, it is also hard to say whether this splitting would decay to zero in the dilute limit. If the defect did have a state in the band gap, it is safe to say it would lie near the bottom of the conduction band and produce photoemission of energy similar to the band gap energy, ~ 2 eV.

Now comes the difficult task in science of not just saying what something isn't, but saying what it is. Given the trouble DFT has with accurately predicting band gaps, it is hard to directly relate the energy of the defect to that of a photoemission line, especially given the error bar of ± 0.1 eV from the lack of spin-orbit coupling. Even if we were to extend the band gap artificially to its experimental value, it is not clear whether we should keep the defects at their gap above the VBM, at their gap below the CBM, or something else entirely. Having said that, PBE and HSE both seem to agree roughly on the fractional position of a given defect state within the band gap. To achieve more accurate estimates of state energies, one would have to employ more computationally expensive corrections. Spin-orbit coupling is a step in the right direction, but ultimately photoluminescence is an excited state phenomenon, so the most accurate calculations would involve beyond DFT methods like GW (discussed in chapter 4).

There are 4 peaks commonly observed across authors and samples [104, 103, 100, 101, 102, 28] at ~ 1.2 eV, ~ 1.35 eV, ~ 1.5 eV, and ~ 1.7 eV, and a fifth small peak observed by Frazer et al [101] at ~ 1.9 eV. No one has yet identified clear trends in peak sizes against growth conditions. Natural samples seem to be more often dominant in 1.7 eV emission [101, 103], Li et al. [28] and Koirala et al. [102] synthesise samples that are dominant in 1.7 eV emission, and Frazer et al. [101], Ito et al. [103] and Bloem [104, 134] synthesise samples with all kinds of emissions with no obvious pattern. The emissions at 1.35 eV and the emission at 1.7 eV are usually stronger than the other 3.

Since the oxygen interstitials have the lowest formation energies and the clearest defect states, we assign the largest two peaks to them. The unoccupied defect state due to O_i^{oct} is lower in energy than that of O_i^{tet} , so we will assign the peak at 1.35

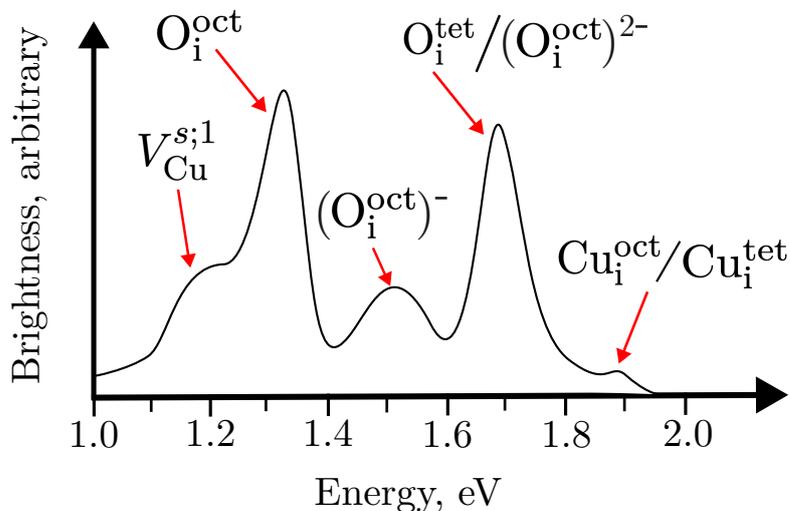


Figure 3.19: A sketch of all the commonly observed peaks in photoluminescence experiments on Cu_2O , and a plausible assignment of the defects which cause them.

eV to O_i^{oct} and the peak at 1.7 eV to O_i^{tet} . Since the defect level for O_i^{oct} rises by about 0.2 eV for each added electron, this assignment also admits an interpretation of the peak at 1.5 eV as $(\text{O}_i^{\text{oct}})^-$, and the peak at 1.7 eV as having a contribution from $(\text{O}_i^{\text{oct}})^{2-}$. Why different growth conditions would favour one geometry of an oxygen interstitial over another is left for future work.

The peak at 1.9 eV, as the highest energy observed peak, we assign to the copper interstitials, the only neutral defects studied that could produce a peak that high. That peak is always the weakest of the 5 (where it is even visible), reflecting the higher formation energy. While there is a slight difference between the octahedral and tetrahedral arrangements in the spin-splitting under HSE, the average energy of the defect band is very similar and the behaviour of these two defects under PBE is very similar, so we assign them to the same emission line.

The peak at 1.2 eV we assign to the first split copper vacancy, $V_{\text{Cu}}^{s;1}$. The average energy of its defect state in the band gap under HSE would suggest it would give the lowest energy PL peak, and similar to that of the O_i^{oct} . Additionally, there is a low energy barrier for this defect to move and become instead V_{Cu} or $V_{\text{Cu}}^{s;2}$, which produce no photoluminescence, and the defect state itself supports only one electron per defect site, hence why this peak is weak amongst all samples and growth conditions. The same is not easily said of any of the other defect-emission candidates.

It is important to stress that this is by no means a conclusive assignment of the

photoluminescence lines, only a plausible one, especially given the lack of spin-orbit coupling. Much stronger conclusions could be drawn from excited state calculations with HSE, as the difference between KS orbital eigenenergies are only a zero-th order approximation to the actual excitation energies.

Excitons in DFT

4.1 Background

In chapter 2 we used a phenomenological model to study excitons in cuprous oxide, and how they behave in strong driving fields, and then in chapter 3 we used DFT to study *ab initio* the ground state properties of cuprous oxide and its various native defects. The obvious questions then arise: can we use DFT to study excitons in cuprous oxide *ab initio*? How do they interact with defects? Unfortunately, DFT is only formulated for ground states of materials, and as mentioned in the previous chapter, Kohn-Sham orbitals and their energy differences only provide a zero-th order approximation to excitations. To study excited states like excitons, we must extend beyond DFT.

Beyond DFT methods have been used to calculate the binding energies of many bulk semiconductors [136] and 2D materials [137] from first principles, as well as for more complex phenomena like dispersion relations [138]. There have also been several studies of Cu_2O with these techniques [139, 140, 141], all aimed at the binding energy of the 1S exciton. However, there have been no studies of the excitations of defects in cuprous oxide, nor do the studies make use of supercells of Cu_2O .

In this chapter, we want to take the first steps into this corner of exciton research, to explore what kinds of things we could learn from *ab initio* studies of excitons in Cu_2O and how they interact with defects. The work is by no means a complete study, and is designed to open avenues for future research in this area. We present a method for estimating the radius of the exciton by exciting supercells of the material, as well as analyse real-space plots of the electronic densities, something rarely seen in literature.

The rungs on the ladder of increasing theoretical and computational complexity

above DFT are time-dependent density functional theory (TDDFT) [142], the GW approximation [143], and the Bethe-Salpeter equation [144]. The latter two, while they strive for greater accuracy and allow for more detailed analysis, are computationally much more expensive than the first. TDDFT is computationally the simplest and so is most appropriate for exploratory work like this. Furthermore, we will be working with linear response TDDFT, a kind of first order perturbation theory in the time-dependent potential that is exciting the system. Such approximations are necessary to make working with excited states in large supercells computationally tractable.

4.2 TDDFT Theory

In this section I will introduce only the concepts of TDDFT which are necessary for understanding the results in this chapter. A thorough and mathematically rigorous tour of TDDFT can be found elsewhere [142].

4.2.1 Runge-Gross theorem

The Hamiltonian for the problem that TDDFT is trying to solve is the same as DFT in equation 3.5, but this time for a time-dependent external potential,

$$\hat{H} = \hat{T}_e + \hat{V}_{ee} + \hat{V}_{\text{ext}}(t), \quad (4.1)$$

where the external potential $\hat{V}_{\text{ext}}(t) = \sum_i v(\mathbf{r}_i, t)$ is a scalar potential acting on the various electrons. This potential can be thought of as the same as for the ground state system before some time t_0 , at which point the perturbation is switched on, such that

$$v(\mathbf{r}, t) = v_{\text{en}}(\mathbf{r}) + \theta(t - t_0)\delta v(\mathbf{r}, t) \quad (4.2)$$

where $\delta v(\mathbf{r}, t)$ is the perturbing potential and $\theta(t - t_0)$ is the step function. As with DFT, we would like to work with the time-dependent density, $\rho(\mathbf{r}, t)$, formed from the wavefunction, rather than directly with the wavefunction. Luckily, the Runge-Gross theorem states [142]:

RG: Two densities, $\rho(\mathbf{r}, t)$ and $\rho'(\mathbf{r}, t)$, evolving from a common initial many-body state $\Psi_0(t_0)$ under the influence of two different potentials $v(\mathbf{r}, t)$ and $v'(\mathbf{r}, t) \neq v(\mathbf{r}, t) + c(t) \forall t > t_0$, will start to become different infinitesimally later than t_0 .

Therefore, there is a one-to-one correspondence between densities and potentials, allowing us to write the time-dependent potential $v[\rho](\mathbf{r}, t)$ as a functional of the time-dependent density, and so construct a time-dependent density functional theory. It is possible to take the time dependent density and potential and solve self-consistently for the time dependence of the Kohn-Sham orbitals. This scheme is known as real-time TDDFT [145], and is usually reserved for electron dynamics in molecules because of its high computational cost. Instead, we will employ a method called linear response, which enables us to look for resonant frequencies for the exciting potential and the resulting effect on the density purely in terms of ground state properties of the material.

4.2.2 Linear response

Consider the functional Taylor expansion of the density $\rho(\mathbf{r}, t)$ in powers of the perturbing potential δv :

$$\rho(\mathbf{r}, t) = \rho_0(\mathbf{r}) + \delta\rho(\mathbf{r}, t) + \delta^2\rho(\mathbf{r}, t) + \dots \quad (4.3)$$

For a weak perturbing potential, such as the probing laser field used for exciton spectroscopy in Cu_2O , we can consider only the first order change in density, given by

$$\delta\rho(\mathbf{r}, t) = \int_{-\infty}^{\infty} dt' \int_V d^3\mathbf{r}' \chi(\mathbf{r}, t, \mathbf{r}', t') \delta v(\mathbf{r}', t') \quad (4.4)$$

where

$$\chi(\mathbf{r}, t, \mathbf{r}', t') = \left. \frac{\delta\rho(\mathbf{r}, t)}{\delta v(\mathbf{r}', t')} \right|_{v_0(\mathbf{r})} \quad (4.5)$$

is called the density-density response function. This is the key quantity in linear response theory. A physically intuitive way to think about the response function is it tells you how the density at one point and time, (\mathbf{r}, t) , changes due to a small perturbation to the density by a potential at another point and time, (\mathbf{r}', t') . We find the full response of the density by adding up all the perturbations over all space and time. Naturally, to respect causality, $\chi(\mathbf{r}, t, \mathbf{r}', t') = 0$ if $t < t'$.

Taking the Fourier transform of $\chi(\mathbf{r}, t, \mathbf{r}', t')$, the response function can be expressed in the basis $\{|\Psi_n\rangle\}$, the many-body eigenstates of \hat{H} , as [142]

$$\chi(\mathbf{r}, \mathbf{r}', \omega) = \lim_{\eta \rightarrow 0^+} \sum_n \left[\frac{\langle \Psi_0 | \hat{n}(\mathbf{r}) | \Psi_n \rangle \langle \Psi_n | \hat{n}(\mathbf{r}') | \Psi_0 \rangle}{\omega - \Omega_n + i\eta} - \frac{\langle \Psi_0 | \hat{n}(\mathbf{r}') | \Psi_n \rangle \langle \Psi_n | \hat{n}(\mathbf{r}) | \Psi_0 \rangle}{\omega + \Omega_n + i\eta} \right] \quad (4.6)$$

where $\hat{n}(\mathbf{r})$ is the density operator and Ω_n are the eigenenergies of the eigenstates $|\Psi_n\rangle$. Of course, we cannot find the states $|\Psi_n\rangle$ or their energies, but the point of this result is to show that the density-density response function has poles at the exact excitation energies of the many body system, so if we can find $\chi(\omega)$ some other way, we can find the resonant frequencies of the system.

4.2.3 KS linear response

We can do exactly the same as the interacting system for the Kohn-Sham system. Since $v_{\text{KS}}[\rho](\mathbf{r}, t)$ can be expressed as a functional of the density, and vice versa, the first order change in $v_{\text{KS}}[\rho](\mathbf{r}, t)$ can be expressed as

$$\delta v_{\text{KS}}[\rho](\mathbf{r}, t) = \delta v(\mathbf{r}, t) + \int d^3\mathbf{r}' \frac{\delta\rho(\mathbf{r}', t)}{|\mathbf{r} - \mathbf{r}'|} + \delta v_{\text{xc}}(\mathbf{r}, t). \quad (4.7)$$

This allows us to define the KS density-density response function

$$\chi_{\text{s}}(\mathbf{r}, t, \mathbf{r}', t') = \left. \frac{\delta\rho(\mathbf{r}, t)}{\delta v_{\text{KS}}(\mathbf{r}', t')} \right|_{v_{\text{KS}}[\rho_0](\mathbf{r})} \quad (4.8)$$

such that the first order density difference can also be calculated from

$$\delta\rho(\mathbf{r}, t) = \int dt' \int d^3\mathbf{r}' \chi_{\text{s}}(\mathbf{r}, t, \mathbf{r}', t') \delta v_{\text{KS}}(\mathbf{r}', t'). \quad (4.9)$$

Substituting equation 4.7 into equation 4.9, we can find a recursive expression for the change in density,

$$\begin{aligned} \delta\rho(\mathbf{r}, t) = & \int dt' \int d^3\mathbf{r}' \chi_{\text{s}}(\mathbf{r}, t, \mathbf{r}', t') \left[\delta v(\mathbf{r}', t') \right. \\ & \left. + \int dt_1 \int d^3\mathbf{r}_1 \left\{ \frac{\delta(t' - t_1)}{|\mathbf{r}' - \mathbf{r}_1|} + f_{\text{xc}}(\mathbf{r}', t', \mathbf{r}_1, t_1) \right\} \delta\rho(\mathbf{r}_1, t_1) \right] \end{aligned} \quad (4.10)$$

where

$$f_{\text{xc}}(\mathbf{r}, t, \mathbf{r}', t') = \left. \frac{\delta v_{\text{xc}}(\mathbf{r}, t)}{\delta\rho(\mathbf{r}', t')} \right|_{\rho_0(\mathbf{r})} \quad (4.11)$$

is called the xc-kernel, which we will discuss later. This means that the change in density must be calculated self-consistently, as the change in KS potential depends on the change in density.

The same change in density can also be calculated with equation 4.4. When we substitute in eq. 4.4, we note that the equality of the two methods for calculating

$\delta\rho$ must hold for all δv , so we can equate the integrands of both sides. Doing so gives

$$\begin{aligned} \chi(\mathbf{r}, t, \mathbf{r}', t') &= \chi_s(\mathbf{r}, t, \mathbf{r}', t') \\ &+ \int_{-\infty}^{\infty} dt_1 \int_V d^3\mathbf{r}_1 \int_{-\infty}^{\infty} dt_2 \int_V d^3\mathbf{r}_2 \chi_s(\mathbf{r}, t, \mathbf{r}_1, t_1) \left\{ \frac{\delta(t_1 - t_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right. \\ &\quad \left. + f_{xc}(\mathbf{r}_1, t_1, \mathbf{r}_2, t_2) \right\} \chi(\mathbf{r}_2, t_2, \mathbf{r}', t') \end{aligned} \quad (4.12)$$

which relates the interacting response function to the KS response function recursively. In Fourier space the relation simplifies greatly to

$$\chi^{-1}(\mathbf{r}, \mathbf{r}', \omega) = \chi_s^{-1}(\mathbf{r}, \mathbf{r}', \omega) - \frac{1}{|\mathbf{r} - \mathbf{r}'|} - f_{xc}(\mathbf{r}, \mathbf{r}', \omega). \quad (4.13)$$

Details of this step can be found in appendix A. Analogous to the interacting response function, the KS response function can be found from the *orbitals* as [142]

$$\chi_s(\mathbf{r}, \mathbf{r}', \omega) = \lim_{\eta \rightarrow 0^+} \sum_{j,k} (f_k - f_j) \frac{\phi_j(\mathbf{r}) \phi_k^*(\mathbf{r}) \phi_j^*(\mathbf{r}') \phi_k(\mathbf{r}')}{\omega - \omega_{jk} + i\eta} \quad (4.14)$$

where f_j is the occupancy of orbital j and $\omega_{jk} = \epsilon_j - \epsilon_k$ is the difference between the eigenenergies of the j -th and k -th orbitals.

This gives us a recipe for finding the resonant frequencies and the resulting density change from ground state properties:

1. Find the ground state orbitals $\phi_i(\mathbf{r})$ of the system with DFT,
2. Compute the KS density-density response function from equation 4.14,
3. Use χ_s and equation 4.13 to self-consistently find the interacting response function,
4. The poles of your interacting response function in frequency space are the resonant frequencies, and the effect on the density is given by equation 4.4.

Just like with DFT, the above algorithm is a tool to help think about how TDDFT works. It would take a long time to converge the interacting response function and find its poles manually. There are also numerical problems introduced when trying to calculate $\delta\rho$ from the product of a divergent χ and an infinitesimal δv . In practice, TDDFT codes construct and solve the Cassida equation (under the Tamm-Dancoff approximation in the case of CASTEP) to avoid these issues and find the poles directly [142].

Briefly, while we have seen that external perturbations at the resonant frequencies of the system cause divergences in the density-density response function, the external perturbation is, in fact, not required. If the resonant oscillation is set in motion, the system will sustain it even in the absence of the external field. Using this fact, equation 4.10 can be rewritten without the perturbing potential, in frequency space, as

$$\delta\rho(\mathbf{r},\omega) = \int d^3\mathbf{r}' \chi_s(\mathbf{r},\mathbf{r}',\omega) \int d^3\mathbf{r}'' \left\{ \frac{1}{|\mathbf{r}' - \mathbf{r}''|} + f_{xc}(\mathbf{r}',\mathbf{r}'',\omega) \right\} \delta\rho(\mathbf{r}'',\omega). \quad (4.15)$$

We can see that this resembles an eigenvalue problem, where $\delta\rho(\mathbf{r},\omega)$ is an eigenvector of the frequency-dependent integral operator acting on the right hand side. The frequencies for which its eigenvalue is 1 are the solutions and therefore the resonant frequencies of the system. This equation, after the definition of many more intermediate mathematical objects, can be morphed into the Casida equation, which is solved for these frequencies. The Tamm-Dancoff approximation neglects to find the negative solutions which correspond to de-excitations of the system, resulting in significant computational speed up for little reduction in accuracy. Further details can be found here [142].

4.2.4 XC kernel

As mentioned above, just as DFT relies on approximations to the xc-potential, TDDFT relies on approximations to the xc-kernel, which requires a time-dependent xc-potential first. While these can reach arbitrary complexity in search of ever greater accuracy, in keeping with the exploratory theme of this chapter, we will use some quite simple ones. The most helpful is the adiabatic approximation, which takes the xc-kernel to be frequency (and therefore time) independent, and is simply given by the ground state xc-kernel,

$$f_{xc}(\mathbf{r},\mathbf{r}',\omega) = f_{xc}(\mathbf{r},\mathbf{r}',0) = \frac{\delta v_{xc}[\rho_0](\mathbf{r})}{\delta\rho_0(\mathbf{r}')} = \frac{\delta E_{xc}[\rho_0]}{\delta\rho_0(\mathbf{r})\delta\rho_0(\mathbf{r}')}, \quad (4.16)$$

where $\rho_0(\mathbf{r})$ is the ground state density. This approximation, like the adiabatic theorem of quantum mechanics, is exact in the case where the perturbation occurs slowly enough for the system to be in its instantaneous ground state. There have been developments in approximate xc-kernels beyond this approximation, but this is the most common starting place [142]. From here, it is natural to use the PBE approximation for E_{xc} , in keeping with the use of PBE throughout the previous chapter.

Supercell size	$2 \times 2 \times 2$	$3 \times 3 \times 3$	$4 \times 4 \times 4$
Triplet excitation energy, eV	0.20	0.23	0.24
Triplet binding energy, eV	0.30	0.27	0.26
Singlet excitation energy, eV	0.35	0.27	/

Table 4.1: The convergence of the first excitation energies with supercell size for the pure crystal, and their implied binding energies, under PBE. Note the excitation energies are much smaller than experiment because PBE underestimates the band gap. We only managed to compute the lowest excitation energy for the $4 \times 4 \times 4$ supercell.

4.2.5 Computational details

To make sure that the excited states are accurate, we must first start with a well-converged ground state. Since, to zeroth order, we will be exciting an electron from one orbital to another, we will end up with some half-filled bands, so the ground state calculation must be spin-polarised. We did not perform spin polarised calculations for the pure crystal and the oxygen interstitials in the previous chapter, so before the performing TDDFT calculations we computed these ground states.

The momentum of a photon is many orders of magnitude less than that of an electron, so excitations do not change the k -point of an electron. Since the band-gap is at Γ , and hence the lowest excitation, we perform all the following TDDFT calculations at Γ as well.

4.3 Results and discussion

4.3.1 Pure Cu_2O

The first direction we chose to explore is the simplest: we wanted to see what excitation energies and binding energies linear response TDDFT under PBE would predict for the perfect crystal. Also, in the same way as the defects, there is a finite size effect that occurs when exciting a cell which is too small. Exciting one electron for each unit cell has a strong effect on the rest of the electrons, and also on the excited state, and would represent an unphysical density of excitons. In this way, we can use the supercell size to estimate the radius of the 1S exciton.

Table 4.1 shows how the energy of the first few excitations of the pure crystal converges with supercell size. In the $4 \times 4 \times 4$ supercell the lowest excitation energy, E_{ex}^1 , has converged to within 10 meV, which is roughly the point at which the periodic reflection of the excitation becomes negligible. This gives an implied upper

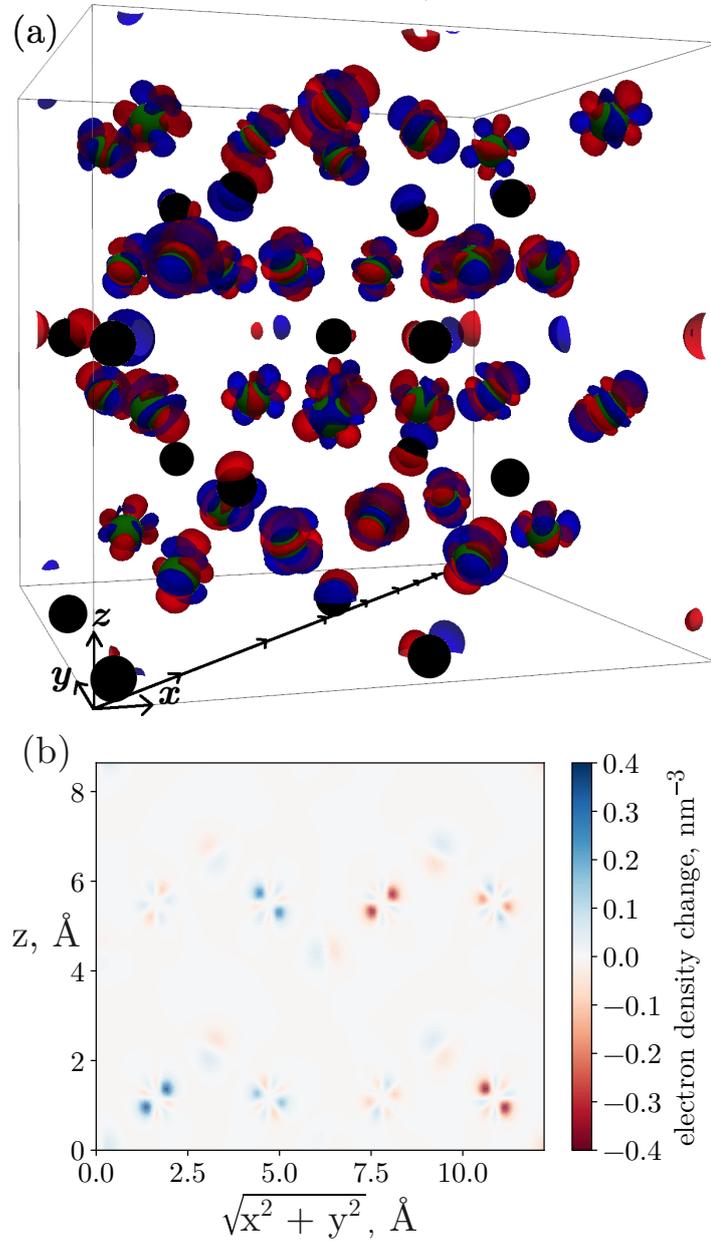


Figure 4.1: The first excited state of the pure Cu_2O crystal under PBE in the $2 \times 2 \times 2$ supercell. (a) Isosurfaces of $\delta\rho(\mathbf{r})$. The blue (red) surface contains 90% of the positive (negative) part of $\delta\rho(\mathbf{r})$, showing where electron density has moved toward (away from) as a result of the excitation. The green spheres are the copper atoms and the black spheres are the oxygen atoms. (b) A slice through $\delta\rho(\mathbf{r})$ in the $[110]$ direction (indicated by the black line), showing the value of $\delta\rho(\mathbf{r})$.

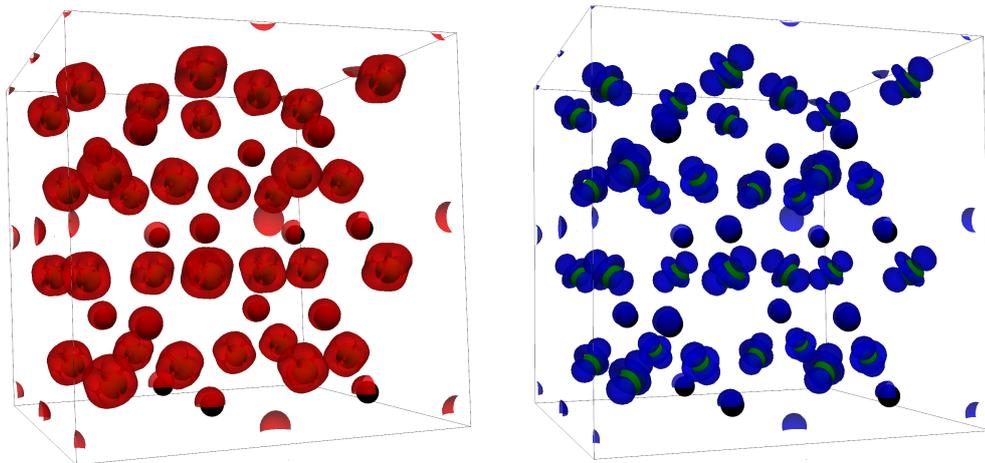


Figure 4.2: Isosurfaces of the electron density, $|\phi_i^{\mathbf{k}=\mathbf{0}}(\mathbf{r})|^2$, of the VBM (left, red) and CBM (right, blue) of Cu_2O in the $2 \times 2 \times 2$ supercell. The isosurface contains 90% of the electron density of the band. The green spheres are the copper atoms and the black spheres are the oxygen atoms.

bound for the 1S exciton Bohr radius of 8.6 \AA (half the width of the supercell), in good agreement with $a_0 = 5.3 \text{ \AA}$ from Wannier-Mott style calculations and experimental measurements [146]. The excitation energies are much lower than the experimental values ($\sim 2 \text{ eV}$) due to PBE underestimating the band gap. To account for this, we also show the binding energy, $E_g^{\text{PBE}} - E_{\text{ex}}^1$, changes with supercell size. It appears to converge to $\sim 0.25 \text{ eV}$, an *overestimation* of the experimental 98 meV [146] binding energy of the 1S exciton, meaning LR-TDDFT with PBE overestimates the attraction between the electron and the hole. Contrary to experiment [147], this method predicts the lowest singlet excitation (the 1S para-exciton) energy to be higher than for the triplet, though it does decrease significantly with energy to 40 meV above the triplet, as opposed to the observed 10 meV below the triplet. It is possible that the para-exciton has a more severe finite-size error, but in the large supercell limit converges to a lower energy than the ortho-exciton; more work is needed on that front.

Figure 4.1a shows isosurfaces of the change in electron density $\delta\rho(\mathbf{r})$ in response to the excitation in the pure $2 \times 2 \times 2$ supercell. The blue surfaces contain 90% of the increase in electron density, and the red surfaces contain 90% of the decrease in electron density, with the copper atoms shown as green balls and the oxygen atoms as black balls. Figure 4.1b shows a slice of $\delta\rho(\mathbf{r})$ in the $[110]$ direction (along the $x = y$ diagonal). The first thing to notice is that the excitation is delocalised throughout the entire cell. This runs counter to the typical intuition of the exciton as like a hydrogen atom, radiating out from a central point. However, this is not

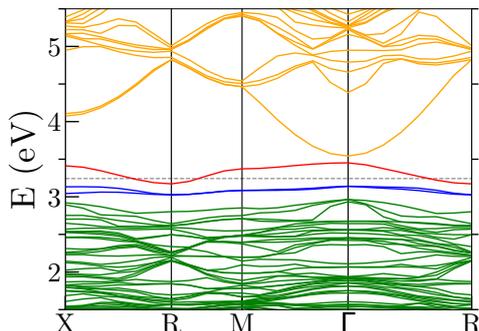


Figure 4.3: The band structure of the tetrahedral oxygen interstitial under PBE in the $2 \times 2 \times 2$ supercell, reprinted here to help with the interpretation of the excited states.

unexpected, as in a periodic crystal there is no preferred region of space to localise the exciton around. Also note that $\delta\rho(\mathbf{r})$ has obvious orbital symmetry around each species of atom. The excitations on the oxygen atoms resemble a P orbital gaining a bias in one direction along its axis, and the excitations on the copper atoms resemble changes between various D orbitals.

These can also be understood from plotting the bands which the excitation occurs between. Figure 4.2 shows isosurfaces of the electron density for the VBM in red and the CBM in blue. The CBM closely resembles the positive part of $\delta\rho(\mathbf{r})$, supporting our interpretation that the first excitation promotes one electron into the lowest conduction band. However, the VBM does not conform to the same simple mapping. This is because there are 3 degenerate valence bands at the Γ -point (where the excitation takes place), so there are actually 3 degenerate lowest energy excitations from a linear combination of the three bands to the CBM. For a similar reason, the valence band plotted in figure 4.2 does not neatly resemble atomic orbitals like the CBM does, since at Γ the three bands have the same eigenvalue, so the bands can be any linear combination of the three eigenvectors.

It is well known that the degeneracy at the top of the valence band is lifted when introducing spin-orbit coupling, which leads to the energy difference between the ‘yellow’ and ‘green’ exciton series [146]. It also plays a role in the energy difference between the ortho- and para-exciton [146, 148, 39]. While spin-orbit coupling was not included in these calculations for simplicity, its inclusion is an obvious next step to take and could improve results in regards to the singlet excitation and the interpretability of the bands involved in the excitation.

We also notice that the change in electron density is quite small. Since charge is

Excited state	Orbitals	Triplet or singlet	Energy, eV
1	Lower defect \rightarrow upper defect	Triplet	0.18
2	Lower defect \rightarrow upper defect	Triplet	0.24
3	Lower defect \rightarrow upper defect	Triplet	0.24
4	Lower defect \rightarrow upper defect	Triplet	0.28
5	VBM \rightarrow upper defect	Triplet	0.30
6	Lower defect \rightarrow CBM	Triplet	0.32
7	Deep valence \rightarrow upper defect	Triplet	0.35
8	Deep valence \rightarrow upper defect	Triplet	0.38
9	Lower defect \rightarrow CBM	Triplet	0.38
10	VBM \rightarrow upper defect	Triplet	0.43

Table 4.2: The first 10 excited states of O_{tet}^i in the $2 \times 2 \times 2$ supercell.

conserved,

$$\int_{\text{unit cell}} d^3\mathbf{r} \delta\rho(\mathbf{r}) = 0, \quad (4.17)$$

but we can still define a metric for the total amount of displaced electron compared to the ground state,

$$\Delta n = \frac{1}{2} \int_{\text{unit cell}} d^3\mathbf{r} |\delta\rho(\mathbf{r})| \quad (4.18)$$

, which in general is quite different to zero. One might expect, since we are exciting an electron from one orbital to an orthogonal one, that $\Delta n \sim 1$, but in practice they seem to be a lot less than one. This is because orthogonal orbitals will have different phase in the same regions, but can still have similar amplitudes, and so density. For the first excited state of the pure $2 \times 2 \times 2$ supercell, $\Delta n = 6 \times 10^{-4}$, because the valence and conduction bands share much of the same space.

4.3.2 Tetrahedral oxygen interstitial

With all the work done in the previous chapter modelling defects in CASTEP, and in this chapter studying excited states in CASTEP, it makes sense to see the effect of a defect on the excited states. With limited time, we could only study in sufficient detail one defect, so the tetrahedral oxygen interstitial was chosen as it has 2 clear defect levels in the band gap (one to trap holes and one to trap electrons) which are predicted under PBE.

Table 4.2 shows the first 10 excited states in the $2 \times 2 \times 2$ supercell. Looking at the band structure for this defect (reproduced in figure 4.3) the first few excited states are naturally between the occupied and unoccupied defect levels, with the VBM

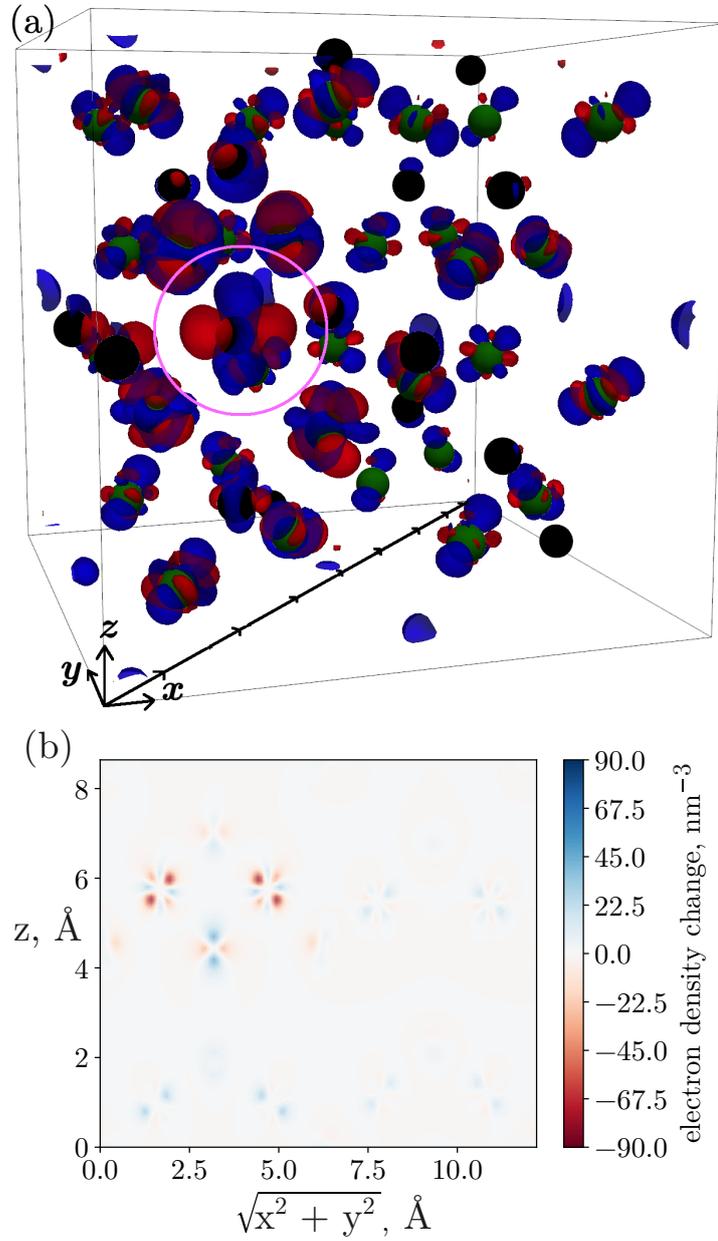


Figure 4.4: The first excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell. It is between one of the lower defect states and the upper defect state. (a) Isosurfaces of $\delta\rho(\mathbf{r})$. The blue (red) surface contains 90% of the positive (negative) part of $\delta\rho(\mathbf{r})$, showing where electron density has moved toward (away from) as a result of the excitation. The green spheres are the copper atoms and the black spheres are the oxygen atoms, and the pink circle highlights the interstitial oxygen. (b) A slice through $\delta\rho(\mathbf{r})$ in the $[110]$ direction (indicated by the black line), showing the value of $\delta\rho(\mathbf{r})$.

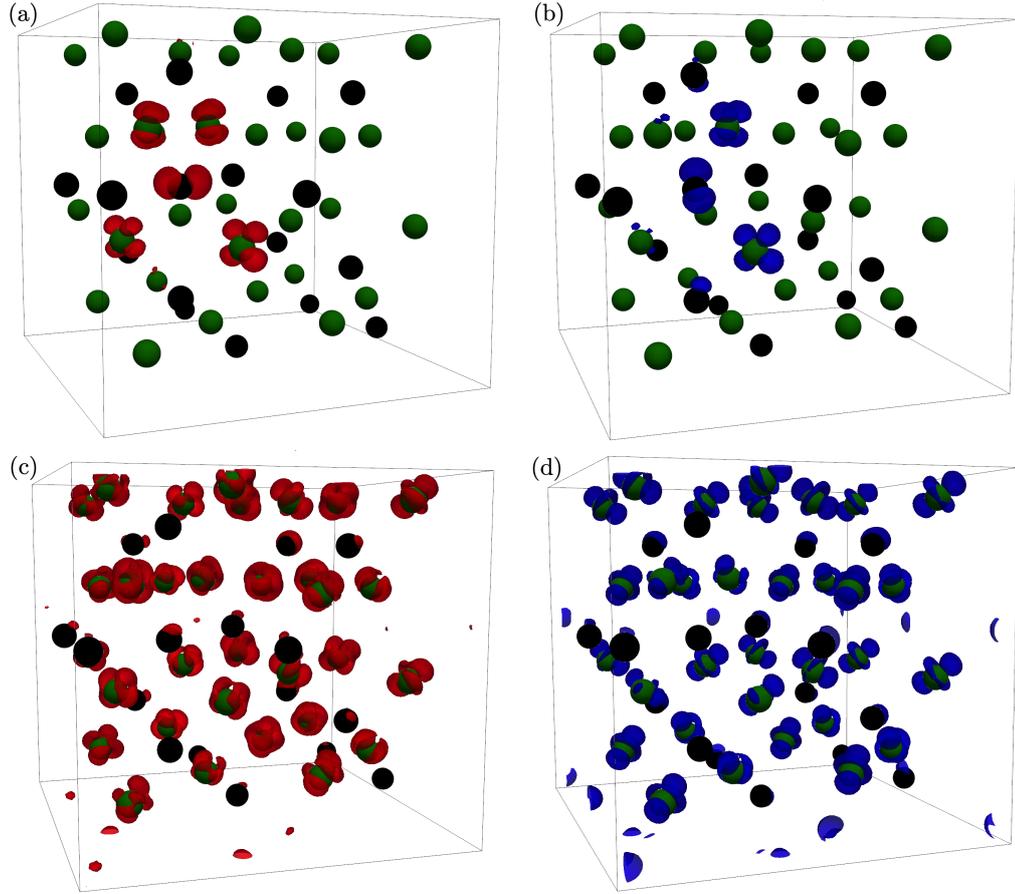


Figure 4.5: Isosurfaces of the electron density, $|\phi_i^{\mathbf{k}=\mathbf{0}}(\mathbf{r})|^2$, of (a) the upper defect state; (b) a lower defect state; (c) a VBM; (d) the CBM of O_{tet}^i in the $2 \times 2 \times 2$ supercell. The isosurfaces contain 90% of the electron density of the band. The green spheres are the copper atoms and the black spheres are the oxygen atoms.

\rightarrow defect and defect \rightarrow CBM above them. The degeneracy between the two lower defect states is broken because TDDFT calculations must be spin-polarised (not shown in the band structure), causing the first 4 excited states to have slightly different energies. Interestingly, even though the band gap ($E_g^{\text{PBE}} = 0.5$ eV) is unchanged by the defect, the VBM \rightarrow CBM exciton was not one of the 10 calculated excited states. Its absence means that the defect has pushed the energy of the exciton from 0.23 eV to greater than 0.43 eV, significantly decreasing the binding energy.

Figure 4.4 shows the first excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell, between one of the lower defect states and the upper defect state, along with a slice through the [110] direction. Compared to the defect-free crystal, this excited state is strongly localised, with most of $\delta\rho(\mathbf{r})$ concentrated on the interstitial oxygen and the tetrahedrally coordinated copper atoms. This is unsurprising, given the localisation

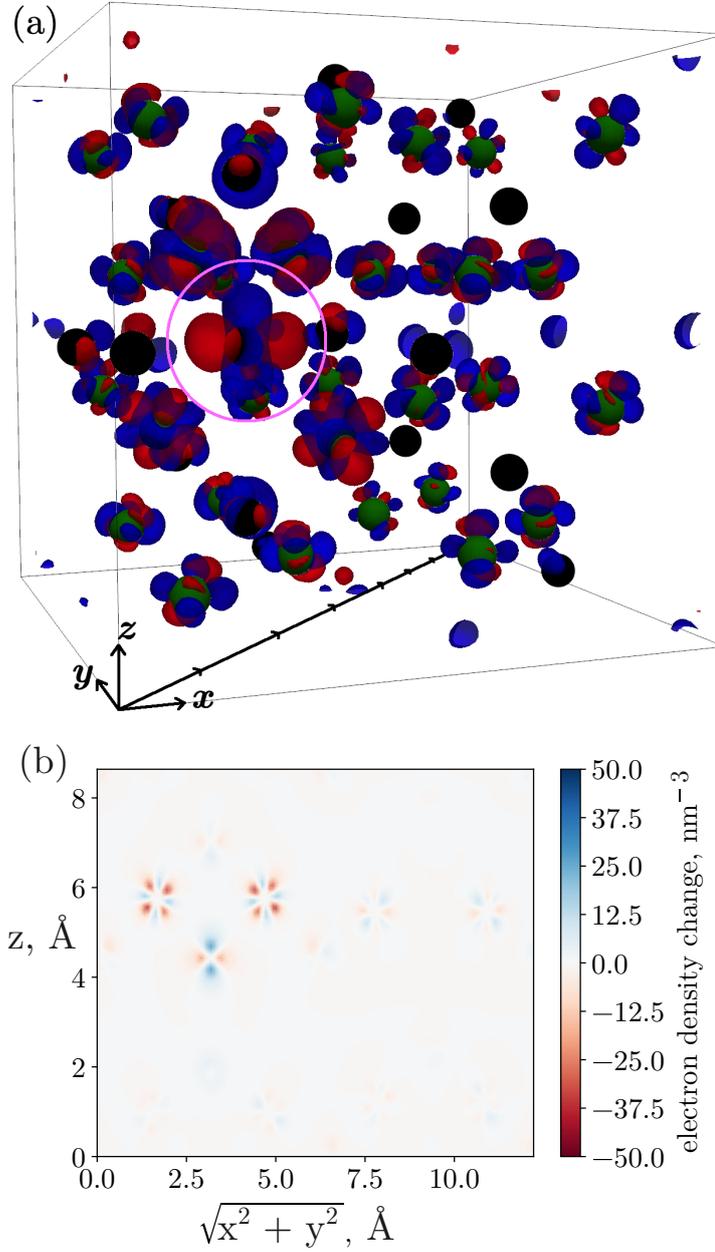


Figure 4.6: The fifth excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell. It is between the VBM and the upper defect state. (a) Isosurfaces of $\delta\rho(\mathbf{r})$. The blue (red) surface contains 90% of the positive (negative) part of $\delta\rho(\mathbf{r})$, showing where electron density has moved toward (away from) as a result of the excitation. The green spheres are the copper atoms and the black spheres are the oxygen atoms, and the pink circle highlights the interstitial oxygen. (b) A slice through $\delta\rho(\mathbf{r})$ in the $[110]$ direction (indicated by the black line), showing the value of $\delta\rho(\mathbf{r})$.

of the orbitals involved, shown in figure 4.5. In fact, what is surprising is how much the electron density changes away from the defect given how local the two defect orbitals are. This is an example of the many-body effects that resist a simple KS band-to-band interpretation of excitations in solids. Among other things, the polarisation created by the movement of the excited electron can have long range effects on other electrons. Due to the high degree of localisation of both bands, for this state $\Delta n = 0.13$, much higher than in the pure crystal, which is a theme with the excited states of the defect cell.

Figures 4.6 and 4.7 show the fifth and sixth excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell. The fifth excitation is between the VBM and the upper defect state, and is a state the exciton could end up in if the electron became trapped in the upper defect state. It is highly localised around the interstitial oxygen, which could indicate that an exciton would be pinned to the defect site. Strangely, it is very similar to the first excitation state, with slightly lower $\Delta n = 0.08$ implying it is less localised. The sixth excited state, between one of the lower defect levels and the CBM, represents the hole from an exciton becoming trapped in the lower defect state of O_{tet}^i . It resembles the first and fifth excited states, with its principle difference being an even greater degree of localisation at $\Delta n = 0.2$. Again, an exciton which found itself in this state would be pinned to the defect site.

4.4 Conclusions

We have demonstrated a novel method for estimating the radius of excitons based on supercell convergence, which achieved good agreement with experimental measurements and Mott-Wannier models. In doing so, we have also shown the importance of appropriate supercell choice for excited state calculations, especially when it comes to excited states above the first. Another novel result is that we can understand some of the properties of the excited states by studying real-space plots of the bands that are involved in the excitations, something rarely seen in the excited state literature.

We have shown how the localised defect states of oxygen interstitials can cause a localisation of the exciton, perhaps opening the way for experiments that pin excitons to defect sites. Surprisingly, the presence of the oxygen interstitial has a strong effect on the energy of the 1S exciton, even though the defect does not significantly disturb the valence band maximum or the conduction band minimum, which could have consequences for an exciton that encompasses such a defect even if not trapped by it.

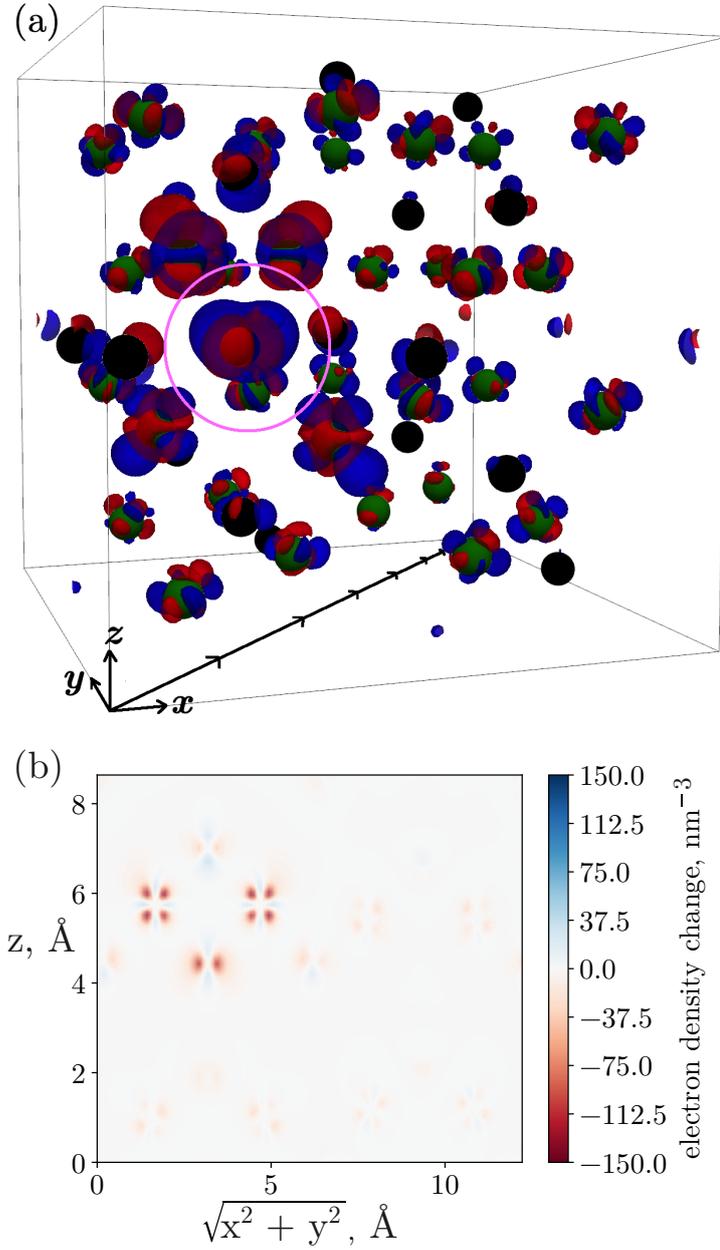


Figure 4.7: The sixth excited state of O_{tet}^i in the $2 \times 2 \times 2$ supercell. It is between one of the lower defect states and the CBM. (a) Isosurfaces of $\delta\rho(\mathbf{r})$. The blue (red) surface contains 90% of the positive (negative) part of $\delta\rho(\mathbf{r})$, showing where electron density has moved toward (away from) as a result of the excitation. The green spheres are the copper atoms and the black spheres are the oxygen atoms, and the pink circle highlights the interstitial oxygen. (b) A slice through $\delta\rho(\mathbf{r})$ in the $[110]$ direction (indicated by the black line), showing the value of $\delta\rho(\mathbf{r})$.

Lots of exciting work can be done in this area. It is clearly important that the O_{tet}^i should be studied in a $3 \times 3 \times 3$ supercell to see the full extent of the trapped exciton and the convergence of the excitation energies. All the other defects that exhibit defect states in the band gap could also have varied and interesting effects on the exciton. Indeed, it may be true that even defects that do not produce photoluminescence have an effect on the exciton series, as is suspected from annealing results [80].

The inclusion of spin-orbit coupling could have significant effects on excitations involving the valence band. Moreover, other authors have had success with TDDFT using hybrid xc-kernels like HSE06 [141], and the improved band gap of such functionals could enable more direct numerical interpretation of results. Further up the computational ladder, it would eventually be interesting to see what results GW or Bethe-Salpeter calculations would give for defect-exciton interactions.

Conclusion

In this thesis, we have studied Rydberg excitons and defects in Cu_2O with a variety of theoretical and computational techniques. In chapter 2, we pushed the limits of the usual ‘top down’ theory of Mott-Wannier excitons by applying it, in conjunction with Floquet theory, to the ultrastrong driving of Rydberg excitons by a microwave field. The model performed well, reproducing many experimental results even at the highest experimentally-applied field strengths. We have come to understand the flattening of the $n\text{P}$ absorption peaks at high n not as ionisation of the excitons, but as the formation of a quasi-continuum of hybridised excitonic states. Here, many states enter ultrastrong and deep strong driving all with the same microwave field, something not observed in atoms due to the large energy widths of Rydberg excitons. The continuum maintains a strong coupling to the microwave and laser field, as evidenced by the generation of strong sidebands on the probe laser, opening avenues to microwave-to-optical conversion using cuprous oxide.

The strength of the sidebands is also reproduced well, capturing the onset and shape of the rollover in intensity across the exciton spectrum for both the second and fourth order red and blue sidebands, and for low n the agreement is quantitative. This is a significant result as it is at the core of optical non-linearities that can be used for microwave-optical conversion. However, the sidebands are one area we start to see shortcomings of the Mott-Wannier approach. Without a detailed theory for the excitation process of the excitons, we do not have access to the phases of the couplings between the $n\text{P}$ excitons and the valence band, which limits the model’s predictive power for the sidebands. A theory of the interaction between the excitation of the phonon background and the exciton series would also allow for treatment of the asymmetry of the $n\text{P}$ absorption peaks, another limit of the top down theory approach. In the future, it would be interesting to see an ab initio condensed matter approach to this problem.

The other significant part of the Mott-Wannier model that is limited by a neglect of the underlying crystal is in the dependence on the polarisation of the fields. The spherical harmonics predicted by Mott-Wannier theory were shown to give a strong dependence of the effect of the microwave field on the polarisation angle, however in experiment the dependence was close to zero. Future work could make use of replacements for spherical harmonics, such as cubic harmonics, that are more compatible with the symmetries of the crystal, but it is likely that the symmetries will demand a more radical change to the exciton wavefunctions. Progress on this front would also have positive effects on the theory of interactions between Rydberg excitons.

In chapter 3 we studied Cu_2O from the bottom up using a careful and thorough application of DFT, in order to predict which defects cause which spectral lines in photoluminescence experiments, a key technique in characterising Cu_2O sample quality for Rydberg excitons. This project was broadly successful, firstly in determining that the simple copper and oxygen vacancies, to which all the PL peaks are assigned throughout the literature, are almost certainly not responsible for photoemissions below the band gap. Furthermore, we conclude that oxygen interstitials must be responsible for some of the PL peaks, and make tentative assignments of which ones they are.

Beyond these findings, our conclusions are less concrete. The copper interstitials seem like they could produce higher energy PL peaks below the band gap, and so one peak is assigned to them. However, the states associated with these defects extend over several unit cells such that they are clearly not in the dilute limit even in $3 \times 3 \times 3$ supercells, something to be investigated in future work. The first split copper vacancy is assigned to one of the PL peaks, but only displays a defect state in the band gap under the HSE06 xc-functional. This could be as a result of the problems HSE has treating the semi-metallic valence band introduced by the defect. Future work might investigate charge states or excited states of this defect under HSE, to check whether it remains in the band gap when occupied by an electron.

Finally, in chapter 4, we began to make connections between the atomic and solid state approaches to Rydberg exciton theory by studying excited states with time-dependent DFT. We found this approach to be very promising. The excitation energy of the 1S exciton, like the defect states, was sensitive to supercell size, and using this fact gave a good estimate for the radius of the exciton. We also found it illuminating to see 3-dimensional representations of the excited states, and how they related to the representations of the valence and conduction bands predicted by ground state DFT.

Following the excitations of the pure crystal, we then explored how a defect with a localised defect state (O_{tet}^i) would effect the excited states of Cu_2O . We saw that excitations to or from the defect states localised the excited states, potentially enabling the engineering of defects to pin excitons to desired sites. Further studies could expand this analysis to the rest of the native defects, more excited states, larger supercells, and higher level approximations. Quantitative assignment could even be made of the photoluminescence peaks through excitations of the defect states.

Rydberg excitons provide a unique opportunity for the marriage of atomic physics and quantum condensed matter in one system. We have learnt a lot about the successes and limits of each approach in cuprous oxide, and seen ways in which they can each benefit from advances in the other. Going forward, I hope new avenues for quantum technologies emerge from a more complete understanding of this system built of a union of these two fields.

Transforming the Dyson-type equation into frequency space

Here we will transform the Dyson-type equation (eq. 4.13) into a relation between the response function inverses (eq. 4.14). The first step is to use the fact that the density-density response functions and xc-kernel depend only on the time differences, $t - t'$. The response functions and the xc-kernel can then be replaced by their Fourier transforms,

$$\begin{aligned}\chi_{(s)}(\mathbf{r}, \mathbf{r}', t - t') &= \int_{-\infty}^{\infty} d\omega \chi_{(s)}(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')}, \\ f_{xc}(\mathbf{r}, \mathbf{r}', t - t') &= \int_{-\infty}^{\infty} d\omega \chi_{(s)}(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')}.\end{aligned}\quad (\text{A.1})$$

The Dyson-type equation then becomes, after some rearranging,

$$\begin{aligned}\int_{-\infty}^{\infty} d\omega \chi(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')} &= \int_{-\infty}^{\infty} d\omega \chi_s(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')} \\ &+ \left\{ \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \int_{-\infty}^{\infty} d\omega \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \int_{-\infty}^{\infty} d\Omega \chi(\mathbf{r}_2, \mathbf{r}', \Omega) \right. \\ &\left. \int_{-\infty}^{\infty} d\theta f_{xc}(\mathbf{r}_1, \mathbf{r}_2, \theta) e^{i\omega t} e^{-i\Omega t'} \int_{-\infty}^{\infty} dt_1 e^{it_1(\theta-\omega)} \int_{-\infty}^{\infty} dt_2 e^{it_2(\Omega-\theta)} \right\} \\ &+ \left\{ \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \int_{-\infty}^{\infty} d\omega \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \int_{-\infty}^{\infty} d\Omega \chi(\mathbf{r}_2, \mathbf{r}', \Omega) \right. \\ &\left. \int_{-\infty}^{\infty} dt_1 e^{i\omega(t-t_1)} \int_{-\infty}^{\infty} dt_2 e^{i\Omega(t_2-t')} \frac{\delta(t_1 - t_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right\}.\end{aligned}\quad (\text{A.2})$$

Using the fact that

$$\int_{-\infty}^{\infty} dt e^{i(\omega-\Omega)t} = \delta(\omega - \Omega), \quad (\text{A.3})$$

we can reduce greatly to

$$\begin{aligned} \int_{-\infty}^{\infty} d\omega \chi(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')} &= \int_{-\infty}^{\infty} d\omega \chi_s(\mathbf{r}, \mathbf{r}', \omega) e^{i\omega(t-t')} \\ &+ \left\{ \int_{-\infty}^{\infty} d\omega \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \chi(\mathbf{r}_2, \mathbf{r}', \omega) f_{xc}(\mathbf{r}_1, \mathbf{r}_2, \omega) e^{i\omega(t-t')} \right\} \\ &+ \left\{ \int_{-\infty}^{\infty} d\omega \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \chi(\mathbf{r}_2, \mathbf{r}', \omega) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} e^{i\omega(t-t')} \right\}. \end{aligned} \quad (\text{A.4})$$

We can see that the same Fourier transform is being applied to the left- and right-hand sides of equation A.4. Therefore, we can equate the Fourier components, and find that

$$\begin{aligned} \chi(\mathbf{r}, \mathbf{r}', \omega) &= \chi_s(\mathbf{r}, \mathbf{r}', \omega) \\ &+ \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \chi(\mathbf{r}_2, \mathbf{r}', \omega) \left\{ f_{xc}(\mathbf{r}_1, \mathbf{r}_2, \omega) + \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \right\} \end{aligned} \quad (\text{A.5})$$

To remove the integral operators we make use of the inverse density-density response function, which has the property

$$\int_V d^3\mathbf{r}' \chi(\mathbf{r}, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}'', \omega) = \delta(\mathbf{r} - \mathbf{r}''). \quad (\text{A.6})$$

Multiplying both sides first by $\chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega)$ and integrating over \mathbf{r}' , we get

$$\begin{aligned} \int_V d^3\mathbf{r}' \chi(\mathbf{r}, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega) &= \int_V d^3\mathbf{r}' \chi_s(\mathbf{r}, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega) \\ &+ \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r}_2 \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \left\{ f_{xc}(\mathbf{r}_1, \mathbf{r}_2, \omega) + \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \right\} \\ &\int_V d^3\mathbf{r}' \chi(\mathbf{r}_2, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega), \end{aligned} \quad (\text{A.7})$$

and after simplifying,

$$\begin{aligned} \delta(\mathbf{r} - \mathbf{r}_3) &= \int_V d^3\mathbf{r}' \chi_s(\mathbf{r}, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega) \\ &+ \int_V d^3\mathbf{r}_1 \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \left\{ f_{xc}(\mathbf{r}_1, \mathbf{r}_3, \omega) + \frac{1}{|\mathbf{r}_1 - \mathbf{r}_3|} \right\}. \end{aligned} \quad (\text{A.8})$$

Next, multiplying both sides by $\chi_s^{-1}(\mathbf{r}_4, \mathbf{r}, \omega)$ and integrating over \mathbf{r} , we get

$$\begin{aligned} \int_V d^3\mathbf{r} \chi_s^{-1}(\mathbf{r}_4, \mathbf{r}, \omega) \delta(\mathbf{r} - \mathbf{r}_3) &= \\ &\int_V d^3\mathbf{r}' \int_V d^3\mathbf{r} \chi_s^{-1}(\mathbf{r}_4, \mathbf{r}, \omega) \chi_s(\mathbf{r}, \mathbf{r}', \omega) \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega) \\ &+ \int_V d^3\mathbf{r}_1 \int_V d^3\mathbf{r} \chi_s^{-1}(\mathbf{r}_4, \mathbf{r}, \omega) \chi_s(\mathbf{r}, \mathbf{r}_1, \omega) \left\{ f_{xc}(\mathbf{r}_1, \mathbf{r}_3, \omega) + \frac{1}{|\mathbf{r}_1 - \mathbf{r}_3|} \right\}. \end{aligned} \quad (\text{A.9})$$

After simplifying once,

$$\begin{aligned} \chi_s^{-1}(\mathbf{r}_4, \mathbf{r}_3, \omega) &= \int_V d^3\mathbf{r}' \delta(\mathbf{r}_4 - \mathbf{r}') \chi^{-1}(\mathbf{r}', \mathbf{r}_3, \omega) \\ &\quad + \int_V d^3\mathbf{r}_1 \delta(\mathbf{r}_4 - \mathbf{r}_1) \left\{ f_{xc}(\mathbf{r}_1, \mathbf{r}_3, \omega) + \frac{1}{|\mathbf{r}_1 - \mathbf{r}_3|} \right\}, \end{aligned} \quad (\text{A.10})$$

and after simplifying again,

$$\chi_s^{-1}(\mathbf{r}_4, \mathbf{r}_3, \omega) = \chi^{-1}(\mathbf{r}_4, \mathbf{r}_3, \omega) + f_{xc}(\mathbf{r}_4, \mathbf{r}_3, \omega) + \frac{1}{|\mathbf{r}_4 - \mathbf{r}_3|}. \quad (\text{A.11})$$

Simply rearranging and renaming the position variables arrives us at equation 4.14, or

$$\chi^{-1}(\mathbf{r}, \mathbf{r}', \omega) = \chi_s^{-1}(\mathbf{r}, \mathbf{r}', \omega) - \frac{1}{|\mathbf{r} - \mathbf{r}'|} - f_{xc}(\mathbf{r}, \mathbf{r}', \omega). \quad (\text{A.12})$$

Bibliography

- [1] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. Yavuz, T. Walker, and M. Saffman. Observation of Rydberg blockade between two atoms. *Nature Physics*, 5(2):110–114, 2009.
- [2] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys. Entanglement of two individual neutral atoms using Rydberg blockade. *Phys. Rev. Lett.*, 104:010502, Jan 2010. doi: 10.1103/PhysRevLett.104.010502. URL <https://link.aps.org/doi/10.1103/PhysRevLett.104.010502>.
- [3] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Physical Review Letters*, 85(10):2208, 2000.
- [4] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Reviews of modern physics*, 82(3):2313–2363, 2010.
- [5] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler. A Rydberg quantum simulator. *Nature Physics*, 6(5):382–388, 2010.
- [6] M. Morgado and S. Whitlock. Quantum simulation and computing with Rydberg-interacting qubits. *AVS Quantum Science*, 3(2), 2021.
- [7] M. M. Müller, A. Kölle, R. Löw, T. Pfau, T. Calarco, and S. Montangero. Room-temperature Rydberg single-photon source. *Phys. Rev. A*, 87:053412, May 2013. doi: 10.1103/PhysRevA.87.053412. URL <https://link.aps.org/doi/10.1103/PhysRevA.87.053412>.
- [8] D. P. Ornelas-Huerta, A. N. Craddock, E. A. Goldschmidt, A. J. Hachtel, Y. Wang, P. Bienias, A. V. Gorshkov, S. L. Rolston, and J. V. Porto. On-

- demand indistinguishable single photons from an efficient and pure source based on a Rydberg ensemble. *Optica*, 7(7):813–819, 2020.
- [9] C. Tresp, C. Zimmer, I. Mirgorodskiy, H. Gorniaczyk, A. Paris-Mandoki, and S. Hofferberth. Single-photon absorber based on strongly interacting Rydberg atoms. *Phys. Rev. Lett.*, 117:223001, Nov 2016. doi: 10.1103/PhysRevLett.117.223001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.117.223001>.
- [10] H. Gorniaczyk, C. Tresp, J. Schmidt, H. Fedder, and S. Hofferberth. Single-photon transistor mediated by interstate Rydberg interactions. *Phys. Rev. Lett.*, 113:053601, Jul 2014. doi: 10.1103/PhysRevLett.113.053601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.113.053601>.
- [11] Y. Hao, G. Lin, X. Lin, Y. Niu, and S. Gong. Single-photon transistor based on cavity electromagnetically induced transparency with Rydberg atomic ensemble. *Scientific reports*, 9(1):4723, 2019.
- [12] M. Brune, S. Haroche, V. Lefevre, J. M. Raimond, and N. Zagury. Quantum nondemolition measurement of small photon numbers by Rydberg-atom phase-sensitive detection. *Phys. Rev. Lett.*, 65:976–979, Aug 1990. doi: 10.1103/PhysRevLett.65.976. URL <https://link.aps.org/doi/10.1103/PhysRevLett.65.976>.
- [13] H. Hattermann, D. Bothner, L. Ley, B. Ferdinand, D. Wiedmaier, L. Sárkány, R. Kleiner, D. Koelle, and J. Fortágh. Coupling ultracold atoms to a superconducting coplanar waveguide resonator. *Nature communications*, 8(1):2254, 2017.
- [14] S. D. Hogan, J. A. Agner, F. Merkt, T. Thiele, S. Filipp, and A. Wallraff. Driving Rydberg-Rydberg transitions from a coplanar microwave waveguide. *Physical Review Letters*, 108(6):063004, 2012.
- [15] T. Kazimierczuk, D. Fröhlich, S. Scheel, H. Stolz, and M. Bayer. Giant Rydberg excitons in the copper oxide Cu_2O . *Nature*, 514(7522):343–347, 2014. ISSN 14764687. doi: 10.1038/nature13832. URL <https://www.nature.com/articles/nature13832>.
- [16] V. Walther, S. O. Krüger, S. Scheel, and T. Pohl. Interactions between Rydberg excitons in Cu_2O . *Phys. Rev. B*, 98:165201, Oct 2018. doi: 10.1103/PhysRevB.98.165201. URL <https://link.aps.org/doi/10.1103/PhysRevB.98.165201>.

- [17] J. Heckötter, M. Aßmann, and M. Bayer. Rydberg excitons and quantum sensing. *MRS Bulletin*, 49(9):948–957, 2024.
- [18] M. A. M. Versteegh, S. Steinhauer, J. Bajo, T. Lettner, A. Soro, A. Romanova, S. Gyger, L. Schweickert, A. Mysyrowicz, and V. Zwiller. Giant Rydberg excitons in Cu_2O probed by photoluminescence excitation spectroscopy. *Phys. Rev. B*, 104:245206, Dec 2021. doi: 10.1103/PhysRevB.104.245206. URL <https://link.aps.org/doi/10.1103/PhysRevB.104.245206>.
- [19] J. Frenkel. On the transformation of light into heat in solids. i. *Physical Review*, 37(1):17, 1931.
- [20] G. H. Wannier. The structure of electronic excitation levels in insulating crystals. *Physical Review*, 52(3):191, 1937.
- [21] H. Fidler, J. Terpstra, and D. A. Wiersma. Dynamics of frenkel excitons in disordered molecular aggregates. *The Journal of chemical physics*, 94(10):6895–6907, 1991.
- [22] M. Hayashi and K. Katsuki. Absorption spectrum of cuprous oxide. *Journal of the Physical Society of Japan*, 5(5):380B–381, 1950.
- [23] S. B. Nam, D. C. Reynolds, C. W. Litton, R. J. Almassy, T. C. Collins, and C. M. Wolfe. Free-exciton energy spectrum in GaAs. *Phys. Rev. B*, 13:761–767, Jan 1976. doi: 10.1103/PhysRevB.13.761. URL <https://link.aps.org/doi/10.1103/PhysRevB.13.761>.
- [24] V. Walther, R. Johne, and T. Pohl. Giant optical nonlinearities from Rydberg excitons in semiconductor microcavities. *Nature communications*, 9(1):1309, 2018.
- [25] J. D. Pritchett, L. A. P. Gallagher, A. Brewin, H. Q. X. Wong, W. Langbein, S. A. Lynch, C. S. Adams, and M. P. A. Jones. Giant microwave–optical Kerr nonlinearity via Rydberg excitons in cuprous oxide. *APL Photonics*, 9(3):031303, 03 2024. ISSN 2378-0967. doi: 10.1063/5.0192710. URL <https://doi.org/10.1063/5.0192710>.
- [26] J. Heckötter, B. Panda, K. Brägelmann, M. Harati, and M. Aßmann. Neutralization of impurities by continuous-wave excitation of high- n Rydberg excitons. *Physical Review B*, 108(23):235212, 2023.
- [27] M. Bergen, V. Walther, B. Panda, M. Harati, S. Siegeroth, J. Heckötter, and M. Aßmann. Taming charged defects: Large scale purification in semiconductors using Rydberg excitons. *arXiv preprint arXiv:2310.11726*, 2023.

- [28] J. Li, Z. Mei, D. Ye, H. Liang, L. Liu, Y. Liu, A. Galeckas, A. Y. Kuznetsov, and X. Du. Engineering of optically defect free cu₂o enabling exciton luminescence at room temperature. *Optical Materials Express*, 3(12):2072–2077, 2013.
- [29] R. M. Martin. *Electronic structure: basic theory and practical methods*. Cambridge university press, 2020.
- [30] D. O. Scanlon, B. J. Morgan, G. W. Watson, and A. Walsh. Acceptor levels in p-type cu₂o: rationalizing theory and experiment. *Physical review letters*, 103(9):096405, 2009.
- [31] L. Y. Isseroff and E. A. Carter. Electronic structure of pure and doped cuprous oxide with copper vacancies: suppression of trap states. *Chemistry of Materials*, 25(3):253–265, 2013.
- [32] A. Shafqat, T. Iqbal, and A. Majid. A dft study of intrinsic point defects in monolayer mose₂. *Aip Advances*, 7(10), 2017.
- [33] A. Kohan, G. Ceder, D. Morgan, and C. G. Van de Walle. First-principles study of native point defects in zno. *Physical Review B*, 61(22):15019, 2000.
- [34] A. Janotti and C. G. Van de Walle. Native point defects in zno. *Physical Review B—Condensed Matter and Materials Physics*, 76(16):165202, 2007.
- [35] J. Lyons, A. Janotti, and C. Van de Walle. The role of oxygen-related defects and hydrogen impurities in hfo₂ and zro₂. *Microelectronic engineering*, 88(7):1452–1456, 2011.
- [36] K. Xiong, J. Robertson, M. Gibson, and S. Clark. Defect energy levels in hfo₂ high-dielectric-constant gate oxide. *Applied physics letters*, 87(18), 2005.
- [37] A. Ramasubramaniam. Large excitonic effects in monolayers of molybdenum and tungsten dichalcogenides. *Physical Review B—Condensed Matter and Materials Physics*, 86(11):115409, 2012.
- [38] A. Gali. Time-dependent density functional study on the excitation spectrum of point defects in semiconductors. *physica status solidi (b)*, 248(6):1337–1346, 2011.
- [39] F. Schweiner. *Theory of excitons in cuprous oxide*. PhD thesis, Universitat Stuttgart, 2017.
- [40] J. P. Rogers, L. A. P. Gallagher, D. Pizzey, J. D. Pritchett, C. S. Adams, M. P. A. Jones, C. Hodges, W. Langbein, and S. A. Lynch. High-resolution

- nanosecond spectroscopy of even-parity Rydberg excitons in Cu_2O . *Phys. Rev. B*, 105:115206, Mar 2022. doi: 10.1103/PhysRevB.105.115206. URL <https://link.aps.org/doi/10.1103/PhysRevB.105.115206>.
- [41] F. Schöne, S.-O. Krüger, P. Grünwald, H. Stolz, S. Scheel, M. Aßmann, J. Heckötter, J. Thewes, D. Fröhlich, and M. Bayer. Deviations of the exciton level spectrum in Cu_2O from the hydrogen series. *Phys. Rev. B*, 93:075203, Feb 2016. doi: 10.1103/PhysRevB.93.075203. URL <https://link.aps.org/doi/10.1103/PhysRevB.93.075203>.
- [42] I.-C. Hoi, C. M. Wilson, G. Johansson, T. Palomaki, B. Peropadre, and P. Delsing. Demonstration of a single-photon router in the microwave regime. *Phys. Rev. Lett.*, 107:073601, Aug 2011. doi: 10.1103/PhysRevLett.107.073601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.073601>.
- [43] W.-B. Yan and H. Fan. Single-photon quantum router with multiple output ports. *Scientific Reports*, 4(11):4820, Apr. 2014. ISSN 2045-2322. doi: 10.1038/srep04820. URL <https://www.nature.com/articles/srep04820>.
- [44] P. Forn-Díaz, J. J. García-Ripoll, B. Peropadre, J.-L. Orgiazzi, M. A. Yurtalan, R. Belyansky, C. M. Wilson, and A. Lupascu. Ultrastrong coupling of a single artificial atom to an electromagnetic continuum in the nonperturbative regime. *Nature Physics*, 13(11):39–43, Jan. 2017. doi: 10.1038/nphys3905. URL <https://www.nature.com/articles/nphys3905>.
- [45] L. García-Álvarez, J. Casanova, A. Mezzacapo, I. L. Egusquiza, L. Lamata, G. Romero, and E. Solano. Fermion-fermion scattering in quantum field theory with superconducting circuits. *Phys. Rev. Lett.*, 114:070502, Feb 2015. doi: 10.1103/PhysRevLett.114.070502. URL <https://link.aps.org/doi/10.1103/PhysRevLett.114.070502>.
- [46] J. Braumüller, M. Marthaler, A. Schneider, A. Stehli, H. Rotzinger, M. Weides, and A. V. Ustinov. Analog quantum simulation of the Rabi model in the ultra-strong coupling regime. *Nature Communications*, 8(11):779, Oct. 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-00894-w. URL <https://www.nature.com/articles/s41467-017-00894-w>.
- [47] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020. doi: 10.1146/annurev-conmatphys-031119-050605. URL <https://doi.org/10.1146/annurev-conmatphys-031119-050605>.

- [48] L. A. Martínez-Martínez, R. F. Ribeiro, J. Campos-González-Angulo, and J. Yuen-Zhou. Can ultrastrong coupling change ground-state chemical reactions? *ACS Photonics*, 5(1):167–176, Jan. 2018. doi: 10.1021/acsp Photonics.7b00610. URL <https://doi.org/10.1021/acsp Photonics.7b00610>.
- [49] J. Flick, N. Rivera, and P. Narang. Strong light-matter coupling in quantum chemistry and quantum photonics. *Nanophotonics*, 7(9):1479–1501, Sept. 2018. ISSN 2192-8614. doi: 10.1515/nanoph-2018-0067. URL <https://www.degruyter.com/document/doi/10.1515/nanoph-2018-0067/html>.
- [50] T. Niemczyk, F. Deppe, H. Huebl, E. P. Menzel, F. Hocke, M. J. Schwarz, J. J. Garcia-Ripoll, D. Zueco, T. Hümmer, E. Solano, A. Marx, and R. Gross. Circuit quantum electrodynamics in the ultrastrong-coupling regime. *Nature Physics*, 6(10):772–776, Oct 2010. ISSN 1745-2481. doi: 10.1038/nphys1730. URL <https://doi.org/10.1038/nphys1730>.
- [51] U. Bhattacharya, T. Lamprou, A. S. Maxwell, A. Ordóñez, E. Pisanty, J. Rivera-Dean, P. Stammer, M. F. Ciappina, M. Lewenstein, and P. Tzallas. Strong-laser-field physics, non-classical light states and quantum information science. *Reports on Progress in Physics*, 86(9):094401, aug 2023. doi: 10.1088/1361-6633/acea31. URL <https://dx.doi.org/10.1088/1361-6633/acea31>.
- [52] G. A. Peterson, S. Kotler, F. Lecocq, K. Cicak, X. Y. Jin, R. W. Simmonds, J. Aumentado, and J. D. Teufel. Ultrastrong parametric coupling between a superconducting cavity and a mechanical resonator. *Phys. Rev. Lett.*, 123:247701, Dec 2019. doi: 10.1103/PhysRevLett.123.247701. URL <https://link.aps.org/doi/10.1103/PhysRevLett.123.247701>.
- [53] J. D. R. Tommey and S. D. Hogan. Resonant Rydberg-atom-microwave-field interactions in the ultrastrong-driving regime: Beyond the rotating-wave approximation. *Phys. Rev. A*, 100:053417, Nov 2019. doi: 10.1103/PhysRevA.100.053417. URL <https://link.aps.org/doi/10.1103/PhysRevA.100.053417>.
- [54] L. A. P. Gallagher, J. P. Rogers, J. D. Pritchett, R. A. Mistry, D. Pizzey, C. S. Adams, M. P. A. Jones, P. Grünwald, V. Walther, C. Hodges, W. Langbein, and S. A. Lynch. Microwave-optical coupling via Rydberg excitons in cuprous oxide. *Phys. Rev. Res.*, 4:013031, Jan 2022. doi: 10.1103/PhysRevResearch.4.013031. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.4.013031>.

- [55] T. Byrnes, N. Y. Kim, and Y. Yamamoto. Exciton–polariton condensates. *Nature Physics*, 10(11):803–813, 2014.
- [56] D. N. Basov, A. Asenjo-Garcia, P. J. Schuck, X. Zhu, and A. Rubio. Polariton panorama. *Nanophotonics*, 10(1):549–577, 2020.
- [57] M. D. Fraser, S. Höfling, and Y. Yamamoto. Physics and applications of exciton–polariton lasers. *Nature materials*, 15(10):1049–1052, 2016.
- [58] S. Klemmt, T. Harder, O. Egorov, K. Winkler, R. Ge, M. Bandres, M. Emmerling, L. Worschech, T. Liew, M. Segev, et al. Exciton-polariton topological insulator. *Nature*, 562(7728):552–556, 2018.
- [59] M. Makhonin, A. Delphan, K. W. Song, P. Walker, T. Isoniemi, P. Claronino, K. Orfanakis, S. K. Rajendran, H. Ohadi, J. Heckötter, et al. Nonlinear Rydberg exciton-polaritons in cu₂o microcavities. *Light: Science & Applications*, 13(1):47, 2024.
- [60] K. Orfanakis, S. K. Rajendran, V. Walther, T. Volz, T. Pohl, and H. Ohadi. Rydberg exciton–polaritons in a cu₂o microcavity. *Nature Materials*, 21(77):767–772, July 2022. ISSN 1476-4660. doi: 10.1038/s41563-022-01230-4. URL <https://www.nature.com/articles/s41563-022-01230-4>.
- [61] W. Bao, X. Liu, F. Xue, F. Zheng, R. Tao, S. Wang, Y. Xia, M. Zhao, J. Kim, S. Yang, et al. Observation of Rydberg exciton polaritons and their condensate in a perovskite cavity. *Proceedings of the National Academy of Sciences*, 116(41):20274–20279, 2019.
- [62] P. Forn-Díaz, L. Lamata, E. Rico, J. Kono, and E. Solano. Ultrastrong coupling regimes of light-matter interaction. *Reviews of Modern Physics*, 91(2), June 2019. ISSN 1539-0756. doi: 10.1103/revmodphys.91.025005. URL <http://dx.doi.org/10.1103/RevModPhys.91.025005>.
- [63] W. Qin, A. F. Kockum, C. S. Muñoz, A. Miranowicz, and F. Nori. Quantum amplification and simulation of strong and ultrastrong coupling of light and matter. *Physics Reports*, 1078:1–59, Aug. 2024. ISSN 0370-1573. doi: 10.1016/j.physrep.2024.05.003. URL <https://www.sciencedirect.com/science/article/pii/S0370157324001571>.
- [64] Y. Kaluzny, P. Goy, M. Gross, J. M. Raimond, and S. Haroche. Observation of self-induced rabi oscillations in two-level atoms excited inside a resonant cavity: The ringing regime of superradiance. *Phys. Rev. Lett.*, 51:1175–1178, Sep 1983. doi: 10.1103/PhysRevLett.51.1175. URL <https://link.aps.org/doi/10.1103/PhysRevLett.51.1175>.

- [65] A. Frisk Kockum, A. Miranowicz, S. De Liberato, S. Savasta, and F. Nori. Ultrastrong coupling between light and matter. *Nature Reviews Physics*, 1(11):19–40, Jan. 2019. ISSN 2522-5820. doi: 10.1038/s42254-018-0006-2. URL <https://www.nature.com/articles/s42254-018-0006-2>.
- [66] C. Deng, J.-L. Orgiazzi, F. Shen, S. Ashhab, and A. Lupascu. Observation of floquet states in a strongly driven artificial atom. *Phys. Rev. Lett.*, 115:133601, Sep 2015. doi: 10.1103/PhysRevLett.115.133601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.133601>.
- [67] P. Nataf and C. Ciuti. Protected quantum computation with multiple resonators in ultrastrong coupling circuit qed. *Phys. Rev. Lett.*, 107:190402, Nov 2011. doi: 10.1103/PhysRevLett.107.190402. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.190402>.
- [68] R. Stassi, M. Cirio, and F. Nori. Scalable quantum computer with superconducting circuits in the ultrastrong coupling regime. *npj Quantum Information*, 6(1):67, 2020.
- [69] F. J. Garcia-Vidal, C. Ciuti, and T. W. Ebbesen. Manipulating matter by strong coupling to vacuum fields. *Science*, 373(6551):eabd0336, 2021.
- [70] N. J. Lambert, A. Rueda, F. Sedlmeir, and H. G. L. Schwefel. Coherent conversion between microwave and optical photons—an overview of physical implementations. *Advanced Quantum Technologies*, 3(1):1900077, 2020. doi: <https://doi.org/10.1002/qute.201900077>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qute.201900077>.
- [71] X. Han, W. Fu, C.-L. Zou, L. Jiang, and H. X. Tang. Microwave-optical quantum frequency conversion. *Optica*, 8(8):1050–1064, Aug 2021. doi: 10.1364/OPTICA.425414. URL <https://opg.optica.org/optica/abstract.cfm?URI=optica-8-8-1050>.
- [72] A. Brewin, L. A. Gallagher, J. D. Pritchett, H. Q. Wong, R. M. Potvliege, S. J. Clark, and M. P. Jones. Microwave-optical spectroscopy of Rydberg excitons in the ultrastrong driving regime. *New Journal of Physics*, 26(11):113018, 2024.
- [73] J. H. Shirley. Solution of the schrödinger equation with a hamiltonian periodic in time. *Physical Review*, 138(4B):B979, 1965.
- [74] J. Pritchett. *Microwave induced optical nonlinearities in cuprous oxide*. PhD thesis, Durham University, 2023. URL <https://etheses.dur.ac.uk/15291/>.

- [75] J. Heckötter, P. Rommel, J. Main, M. Aßmann, and M. Bayer. Analysis of the fine structure of the D -exciton shell in cuprous oxide. *physica status solidi (RRL) – Rapid Research Letters*, 15(n/a):2100335, 2021. doi: <https://doi.org/10.1002/pssr.202100335>.
- [76] T. Dittrich, P. Hänggi, G.-L. Ingold, B. Kramer, G. Schön, and W. Zwerger. *Quantum transport and dissipation*, volume 3. Wiley-Vch Weinheim, 1998.
- [77] T. Mori. Floquet states in open quantum systems. *Annual Review of Condensed Matter Physics*, 14(1):35–56, 2023.
- [78] T. Oka and S. Kitamura. Floquet engineering of quantum materials. *Annual Review of Condensed Matter Physics*, 10(1):387–408, 2019.
- [79] R. M. Potvliege. Non-hermitian floquet dynamics in absorption spectroscopy. *arXiv:2406.16559*, 2024.
- [80] S. A. Lynch, C. Hodges, S. Mandal, W. Langbein, R. P. Singh, L. A. P. Gallagher, J. D. Pritchett, D. Pizzey, J. P. Rogers, C. S. Adams, and M. P. A. Jones. Rydberg excitons in synthetic cuprous oxide Cu_2O . *Phys. Rev. Mater.*, 5:084602, Aug 2021. doi: 10.1103/PhysRevMaterials.5.084602. URL <https://link.aps.org/doi/10.1103/PhysRevMaterials.5.084602>.
- [81] J. Mund, D. Fröhlich, D. R. Yakovlev, and M. Bayer. High-resolution second harmonic generation spectroscopy with femtosecond laser pulses on excitons in Cu_2O . *Physical Review B*, 98(8):085203, 2018.
- [82] A. Brewin. Microwave-optical spectroscopy of Rydberg excitons in the ultrastrong driving regime. URL <https://doi.org/10.15128/r25q47rn767>.
- [83] R. Loudon. *The quantum theory of light*. OUP Oxford, 2000.
- [84] P. Yu and M. Cardona. *Fundamentals of Semiconductors: Physics and Materials Properties*. Graduate Texts in Physics. Springer Berlin Heidelberg, 2010. ISBN 978-3-642-00710-1. URL https://books.google.co.uk/books?id=5aBuKYBT_hsC.
- [85] E. Wigner. *Group theory: and its application to the quantum mechanics of atomic spectra*. Academic Press, New York, 1959.
- [86] F. Schweiner, J. Main, and G. Wunner. Linewidths in excitonic absorption spectra of cuprous oxide. *Physical Review B*, 93(8):085203, Feb. 2016. doi: 10.1103/PhysRevB.93.085203. URL <https://link.aps.org/doi/10.1103/PhysRevB.93.085203>.

- [87] H. Stolz, F. Schöne, and D. Semkat. Interaction of Rydberg excitons in cuprous oxide with phonons and photons: optical linewidth and polariton effect. *New Journal of Physics*, 20(2):023019, Feb. 2018. ISSN 1367-2630. doi: 10.1088/1367-2630/aaa396. URL <https://dx.doi.org/10.1088/1367-2630/aaa396>.
- [88] J. Heckötter. *Strongly interacting Rydberg excitons in Cu₂O*. PhD thesis, TU Dortmund University, 2020. URL <https://eldorado.tu-dortmund.de/handle/2003/39805>.
- [89] S. O. Krüger, H. Stolz, and S. Scheel. Interaction of charged impurities and Rydberg excitons in cuprous oxide. *Phys. Rev. B*, 101:235204, Jun 2020. doi: 10.1103/PhysRevB.101.235204. URL <https://link.aps.org/doi/10.1103/PhysRevB.101.235204>.
- [90] R. W. Boyd, A. L. Gaeta, and E. Giese. Nonlinear optics. In *Springer Handbook of Atomic, Molecular, and Optical Physics*, pages 1097–1110. Springer, 2008.
- [91] L. Gallagher. *Optical and microwave spectroscopy of Rydberg excitons in Cu₂O*. PhD thesis, Durham University, 2022.
- [92] N. B. Delone and V. P. Krainov. Ac stark shift of atomic energy levels. *Physics-Uspekhi*, 42(7):669, July 1999. ISSN 1063-7869. doi: 10.1070/PU1999v042n07ABEH000557. URL <https://iopscience.iop.org/article/10.1070/PU1999v042n07ABEH000557/meta>.
- [93] J. Heckötter, M. Freitag, D. Fröhlich, M. Aßmann, M. Bayer, M. A. Semina, and M. M. Glazov. Dissociation of excitons in cu₂O by an electric field. *Phys. Rev. B*, 98:035150, Jul 2018. doi: 10.1103/PhysRevB.98.035150. URL <https://link.aps.org/doi/10.1103/PhysRevB.98.035150>.
- [94] J. Heckötter, M. Freitag, D. Fröhlich, M. Aßmann, M. Bayer, M. A. Semina, and M. M. Glazov. Influence of the wavefunction distribution on exciton dissociation in electric field. *Physics of the Solid State*, 60(8):1506–1509, Aug 2018. ISSN 1090-6460. doi: 10.1134/S1063783418080115. URL <https://doi.org/10.1134/S1063783418080115>.
- [95] J. Heckötter, D. Janas, R. Schwartz, M. Aßmann, and M. Bayer. Experimental limitation in extending the exciton series in Cu₂O towards higher principal quantum numbers. *Phys. Rev. B*, 101:235207, Jun 2020. doi: 10.1103/PhysRevB.101.235207. URL <https://link.aps.org/doi/10.1103/PhysRevB.101.235207>.

- [96] T. F. Gallagher. *Microwave excitation and ionization*, page 162–194. Cambridge Monographs on Atomic, Molecular and Chemical Physics. Cambridge University Press, 1994.
- [97] T. Gallagher, A. Dalgarno, C. U. Press, F. Read, and R. Zare. *Rydberg Atoms*. Cambridge Monographs on Atomic, Molecular and Chemical Physics. Cambridge University Press, 1994. ISBN 9780521385312. URL <https://books.google.co.uk/books?id=8JIPEhHWT-cC>.
- [98] R. Shakeshaft. Sturmian expansion of green’s function and its application to multiphoton ionization of hydrogen. *Phys. Rev. A*, 34:244–252, Jul 1986. doi: 10.1103/PhysRevA.34.244. URL <https://link.aps.org/doi/10.1103/PhysRevA.34.244>.
- [99] J. Heckötter, V. Walther, S. Scheel, M. Bayer, T. Pohl, and M. Aßmann. Asymmetric Rydberg blockade of giant excitons in cuprous oxide. *Nature Communications*, 12(1):3556, Jun 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-23852-z. URL <https://doi.org/10.1038/s41467-021-23852-z>.
- [100] L. Frazer, E. J. Lenferink, K. B. Chang, K. R. Poeppelmeier, N. P. Stern, and J. B. Ketterson. Evaluation of defects in cuprous oxide through exciton luminescence imaging. *Journal of Luminescence*, 159:294–302, 2015.
- [101] L. Frazer, K. B. Chang, R. D. Schaller, K. R. Poeppelmeier, and J. B. Ketterson. Vacancy relaxation in cuprous oxide ($\text{Cu}_2\text{O}_{1-y}$). *Journal of Luminescence*, 183:281–290, 2017.
- [102] S. Koirala, N. Naka, and K. Tanaka. Correlated lifetimes of free paraexcitons and excitons trapped at oxygen vacancies in cuprous oxide. *Journal of luminescence*, 134:524–527, 2013.
- [103] T. Ito and T. Masumi. Detailed examination of relaxation processes of excitons in photoluminescence spectra of Cu_2O . *Journal of the Physical Society of Japan*, 66(7):2185–2193, 1997.
- [104] J. Bloem. Discussion of some optical and electrical properties of Cu_2O . *Philips Research Reports*, 13:167–193, 1958.
- [105] H. Raebiger, S. Lany, and A. Zunger. Origins of the p-type nature and cation deficiency in Cu_2O and related materials. *Physical Review B—Condensed Matter and Materials Physics*, 76(4):045209, 2007.

- [106] A. Soon, X.-Y. Cui, B. Delley, S.-H. Wei, and C. Stampfl. Native defect-induced multifarious magnetism in nonstoichiometric cuprous oxide: First-principles study of bulk and surface properties of $\text{Cu}_{2-\delta}\text{O}$. *Physical Review B—Condensed Matter and Materials Physics*, 79(3):035205, 2009.
- [107] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical review*, 136(3B):B864, 1964.
- [108] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. Probert, K. Refson, and M. C. Payne. First principles methods using castep. *Zeitschrift für kristallographie-crystalline materials*, 220(5-6):567–570, 2005.
- [109] A. Mittiga, F. Biccari, and C. Malerba. Intrinsic defects and metastability effects in Cu_2O . *Thin Solid Films*, 517(7):2469–2472, 2009.
- [110] D. O. Scanlon and G. W. Watson. Undoped n-type Cu_2O : fact or fiction? *The Journal of Physical Chemistry Letters*, 1(17):2582–2585, 2010.
- [111] N. Dongfang, Y. S. Al-Hamdani, and M. Iannuzzi. Understanding the role of oxygen-vacancy defects in Cu_2O (111) from first-principle calculations. *Electronic Structure*, 5(3):035001, 2023.
- [112] N. Ashcroft and D. Mermin. *Solid State Physics*. Saunders College Publishing, 1976.
- [113] R. Stowasser and R. Hoffmann. What do the kohn-sham orbitals and eigenvalues mean? *Journal of the american chemical society*, 121(14):3414–3420, 1999.
- [114] A. I. Krylov. From orbitals to observables and back. *The Journal of Chemical Physics*, 153(8), 2020.
- [115] C. R. Jacob and M. Reiher. Spin in density-functional theory. *International Journal of Quantum Chemistry*, 112(23):3661–3684, 2012.
- [116] A. D. Kaplan, M. Levy, and J. P. Perdew. The predictive power of exact constraints and appropriate norms in density functional theory. *Annual Review of Physical Chemistry*, 74(1):193–218, 2023.
- [117] N. D. Woods, M. Payne, and P. Hasnip. Computing the self-consistent field in kohn-sham density functional theory. *Journal of Physics: Condensed Matter*, 31(45):453001, 2019.
- [118] J. Kohanoff and N. Gidopoulos. Density functional theory: basics, new trends and applications. *Handbook of molecular physics and quantum chemistry*, 2 (Part 5):532–568, 2003.

- [119] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized gradient approximation made simple. *Physical review letters*, 77(18):3865, 1996.
- [120] J. P. Perdew, R. G. Parr, M. Levy, and J. L. Balduz Jr. Density-functional theory for fractional particle number: derivative discontinuities of the energy. *Physical Review Letters*, 49(23):1691, 1982.
- [121] W. Yang, A. J. Cohen, and P. Mori-Sanchez. Derivative discontinuity, bandgap and lowest unoccupied molecular orbital in density functional theory. *The Journal of chemical physics*, 136(20), 2012.
- [122] J. P. Perdew, W. Yang, K. Burke, Z. Yang, E. K. Gross, M. Scheffler, G. E. Scuseria, T. M. Henderson, I. Y. Zhang, A. Ruzsinszky, et al. Understanding band gaps of solids in generalized kohn–sham theory. *Proceedings of the national academy of sciences*, 114(11):2801–2806, 2017.
- [123] K. Capelle. A bird’s-eye view of density-functional theory. *Brazilian journal of physics*, 36:1318–1343, 2006.
- [124] J. Heyd, G. E. Scuseria, and M. Ernzerhof. Hybrid functionals based on a screened coulomb potential. *The Journal of chemical physics*, 118(18):8207–8215, 2003.
- [125] J. E. Peralta, J. Heyd, G. E. Scuseria, and R. L. Martin. Spin-orbit splittings and energy band gaps calculated with the heyd-scuseria-ernzerhof screened hybrid functional. *Physical Review B—Condensed Matter and Materials Physics*, 74(7):073101, 2006.
- [126] K. Lejaeghere, G. Bihlmayer, T. Björkman, P. Blaha, S. Blügel, V. Blum, D. Caliste, I. E. Castelli, S. J. Clark, A. Dal Corso, et al. Reproducibility in density functional theory calculations of solids. *Science*, 351(6280):aad3000, 2016.
- [127] D. Hamann, M. Schlüter, and C. Chiang. Norm-conserving pseudopotentials. *Physical review letters*, 43(20):1494, 1979.
- [128] D. Vanderbilt. Soft self-consistent pseudopotentials in a generalized eigenvalue formalism. *Physical review B*, 41(11):7892, 1990.
- [129] W. Gao, T. A. Abtey, T. Cai, Y.-Y. Sun, S. Zhang, and P. Zhang. On the applicability of hybrid functionals for predicting fundamental properties of metals. *Solid State Communications*, 234:10–13, 2016.
- [130] M. Dolg and X. Cao. Relativistic pseudopotentials: their development and scope of applications. *Chemical reviews*, 112(1):403–480, 2012.

- [131] H. J. Monkhorst and J. D. Pack. Special points for brillouin-zone integrations. *Physical review B*, 13(12):5188, 1976.
- [132] I. Remediakis and E. Kaxiras. Band-structure calculations for semiconductors within generalized-density-functional theory. *Physical Review B*, 59(8):5536, 1999.
- [133] C. Ekuma, M. Jarrell, J. Moreno, and D. Bagayoko. Re-examining the electronic structure of germanium: A first-principle study. *Physics Letters A*, 377(34-36):2172–2176, 2013.
- [134] F. Biccari. *Defects and doping in Cu₂O*. PhD thesis, Sapienza, Universita di Roma, 2012.
- [135] H. Al-Jawhari. A review of recent advances in transparent p-type cu₂o-based thin film transistors. *Materials Science in Semiconductor Processing*, 40: 241–252, 2015.
- [136] M. Grunert, M. Großmann, and E. Runge. Predicting exciton binding energies from ground-state properties. *Physical Review B*, 110(7):075204, 2024.
- [137] L. C. Gomes, P. E. Trevisanutto, A. Carvalho, A. S. Rodin, and A. H. Castro Neto. Strongly bound mott-wannier excitons in ges and gese monolayers. *Phys. Rev. B*, 94:155428, Oct 2016. doi: 10.1103/PhysRevB.94.155428. URL <https://link.aps.org/doi/10.1103/PhysRevB.94.155428>.
- [138] M. Gatti and F. Sottile. Exciton dispersion from first principles. *Physical Review B—Condensed Matter and Materials Physics*, 88(15):155113, 2013.
- [139] Y. Wang, S. Lany, J. Ghanbaja, Y. Fagot-Revurat, Y. P. Chen, F. Soldera, D. Horwat, F. Mücklich, and J.-F. Pierson. Electronic structures of c u 2 o, c u 4 o 3, and cuo: A joint experimental and theoretical study. *Physical Review B*, 94(24):245418, 2016.
- [140] L. Dash, F. Bruneval, V. Trinité, N. Vast, and L. Reining. Electronic excitations: Ab initio calculations of electronic spectra and application to zirconia zro₂, titania tio₂ and cuprous oxide cu₂o. *Computational materials science*, 38(3):482–493, 2007.
- [141] A. Zivkovic, N. H. de Leeuw, B. G. Searle, and L. Bernasconi. Electronic excitations in copper oxides: Time-dependent density functional theory calculations with a self-consistent hybrid kernel. *The Journal of Physical Chemistry C*, 124(45):24995–25003, 2020.

- [142] C. A. Ullrich. *Time-dependent density-functional theory: concepts and applications*. OUP Oxford, 2011.
- [143] F. Aryasetiawan and O. Gunnarsson. The gw method. *Reports on progress in Physics*, 61(3):237, 1998.
- [144] X. Blase, I. Duchemin, D. Jacquemin, and P.-F. Loos. The bethe–salpeter equation formalism: From physics to chemistry. *The Journal of Physical Chemistry Letters*, 11(17):7371–7382, 2020.
- [145] M. R. Provorse and C. M. Isborn. Electron dynamics with real-time time-dependent density functional theory. *International Journal of Quantum Chemistry*, 116(10):739–749, 2016.
- [146] G. M. Kavoulakis, Y.-C. Chang, and G. Baym. Fine structure of excitons in cu₂o. *Phys. Rev. B*, 55:7593–7599, Mar 1997. doi: 10.1103/PhysRevB.55.7593. URL <https://link.aps.org/doi/10.1103/PhysRevB.55.7593>.
- [147] K. Karpinska, P. van Loosdrecht, I. Handayani, and A. Revcolevschi. Para-excitons in cu₂o—a new approach. *Journal of luminescence*, 112(1-4):17–20, 2005.
- [148] F. Schweiner, J. Main, G. Wunner, and C. Uihlein. Even exciton series in cu₂o. *Physical Review B*, 95(19):195201, 2017.

Colophon

This thesis is based on a template developed by Matthew Townson and Andrew Reeves. It was typeset with L^AT_EX 2_ε. It was created using the *memoir* package, maintained by Lars Madsen, with the *madsen* chapter style. The font used is Latin Modern, derived from fonts designed by Donald E. Kunth.